

Lower bound for the communication volume required for an optically interconnected array of points

Haldun M. Ozaktas and Joseph W. Goodman

Information Systems Laboratory, Durand Building, Department of Electrical Engineering, Stanford University, Stanford, California 94305

Received July 12, 1989; accepted April 24, 1990

The information-carrying capacity of optical fields is usually stated in terms of an area density as being related to communication through a surface. We render these well-understood results in a form such that they can be interpreted as a volume-density limit, applicable to an arbitrary array of points communicating with one another. An important example of such a situation is an optically interconnected computing system. We show that regardless of their actual spread or mutual overlap, optical communication links may be viewed as solid wires of minimum cross section $\lambda^2/2\pi$ for the purpose of calculating bounds on volumes and cross sections. Thus the results of area-volume complexity theory for solid wires are also applicable to optically communicating systems. The maximum number of binary pulses that may be in transit in an optical communication network occupying volume V is found to be $\rho 2\pi V/\lambda^3$, ρ denoting the modulation bandwidth normalized by the carrier frequency. Previously suggested optical-interconnection schemes are discussed in this context.

1. INTRODUCTION

A convenient unit for measuring information is the bit. The capacity of an information channel can then 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 Gabor¹ and later was elaborated by Winthrop,² among others. Winthrop gave the following result for the number of degrees of freedom F associated with a quasi-monochromatic optical wave field over a given surface S :

$$F = \int_S \chi(P) dS, \quad (1)$$

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

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

with θ being the angle between the element of solid angle and a unit vector perpendicular to the surface in question (Fig. 1). (It might have been noticed that we are omitting certain factors of 2 that 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 that bring in an additional factor of 2 is given in Appendix A.) 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 that 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 Eq. (1) we may write

$$A(P)\chi(P) = 1. \quad (3)$$

This relationship is guaranteed to be preserved with free propagation² 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 invariant³) and is also closely connected to radiometric and thermodynamic considerations.⁴ 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.

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 Fig. 2. Several authors⁵⁻⁷ have adapted results from area-volume complexity theory (for instance, see Refs. 8 and 9) 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.

In this paper, 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




Fig. 1. 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 \hat{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 Eq. (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 Winthrop²).




Fig. 2. Communication through a finite surface.

$$V = \frac{\lambda^2}{2\pi} L_{total} \tag{4}$$

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 that each independent optical information channel has a minimum cross section of $\lambda^2/2\pi$ as if it were a solid wire.

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 propagation as a medium of communication. Since electromagnetic phenomena underlie all currently observed and utilized modalities of information transfer, these results may be considered to have universal significance.

In Section 2 we state our assumptions and derive the above result. In Section 3 we further discuss its meaning and derive another basic result relating the signal delay to the number of optical connections established. In Section 4 we give simple examples to illustrate how our results may be

used to discuss to what extent previously proposed optical-interconnection schemes approach fundamental limitations.

2. LOWER BOUND FOR THE COMMUNICATION VOLUME FOR AN OPTICALLY INTERCONNECTED ARRAY OF POINTS

Our discussions are based on a scalar theory of light. We will assume that all sources emit spatially coherent quasi-monochromatic radiation of a given center frequency f . 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 attributed mainly to the former effect.

We will consider the following model as illustrated in Fig. 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 these signals are 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 B). 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 signal delay along that interconnection with the velocity of propagation in the medium. Obviously, the communication length is greater than or 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 V , as given in many physics texts¹⁰:

$$\mathcal{D}(f) = \frac{8\pi V}{c\lambda^2} \tag{5}$$

c being the speed of light in the medium of propagation. As written, the above equation includes a factor of 2 in order to




Fig. 3. Volume allocated for optical communication. The unshaded region, of volume V , is allocated for establishing optical communication between transducers located on the islands and/or the enclosing surface.

account for polarization. Also, in the derivation of this equation it was assumed that standing waves exist throughout the volume in all three dimensions, whereas in our discussion we do not permit double sidedness along the directions of propagation (see Appendix A). So we divide by 4 to maintain consistency with our discussion:

$$\mathcal{D}(f) = \frac{2\pi V}{c\lambda^2}. \quad (6)$$

The above equation states that a volume V can support $(2\pi V/c\lambda^2)\Delta f$ degrees of freedom, where $\Delta f \ll f$ is the temporal bandwidth of modulation. We may also write this in an alternative form assuming that the temporal bandwidth is a fraction ρ of the optical carrier frequency; i.e., $\Delta f = \rho f$ with $\rho \ll 1$. $\rho 2\pi V/\lambda^3$ is then the maximum number of binary pulses that can be in transit at a given time in an optical communication network occupying volume V . This is sometimes also termed the population capacity of a communication network.¹¹

In order to derive our main result let us assume that the temporal information modulation bandwidth associated with the output transducers is Δf . This quantity will cancel out in our final result. Let each connection be numbered with the index $i = 1, 2, \dots, n$, where it is assumed that there are 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

$$\frac{\Delta f}{1/\tau_i} = \frac{\Delta f}{c/L_i} = \frac{\Delta f L_i}{c}, \quad (7)$$

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

$$\sum_{i=1}^n \frac{\Delta f L_i}{c} = \frac{\Delta f}{c} \sum_{i=1}^n L_i. \quad (8)$$

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

$$\frac{2\pi V}{c\lambda^2} \Delta f \geq \frac{\Delta f}{c} \sum_{i=1}^n L_i = \frac{\Delta f}{c} L_{\text{total}} \quad (9)$$

or

$$V \geq \frac{\lambda^2}{2\pi} L_{\text{total}}. \quad (10)$$

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

If 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 $\lambda L_{\text{total}}/\pi$.

3. DISCUSSION

We repeat the main conclusion of Section 2: The minimum communication volume required for an optical-interconnection network with total communication length L_{total} is $\lambda^2 L_{\text{total}}/2\pi$. We state this result 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 that each light path occupies constant-width tubes with no further overlap permitted, i.e., by treating them as solid wires: The minimum effective nonoverlapping cross section required for each independent 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 Eq. (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 apertures, such that the guide cross sections are $\simeq \lambda^2$. Clearly, it is, in principle, 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 communication 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 n as a parameter, we require that the volume be given by the above-stated results, within a constant that is of the order of unity and independent of n . This will be clarified in the following and through examples in Section 4.

The reader may have noticed that until now we refrained from discussing how the communication lengths are specified. This will be determined mainly by the communication requirements of the computing system. For a so-called mesh architecture all interconnections will be to nearest neighbors, so that L_{total} , hence the communication volume required, can be small. Certain architectures will require longer interconnections. Without attempting a detailed

discussion of such issues, here we simply illustrate the major mechanism that determines L_{total} for a communication-volume-limited optical-interconnection network. For concreteness, let us assume a more or less cubic array of n pairs of points that are to be connected (i.e., a total of $2n$ points). We would like to pack this large array of points as closely as possible so as to reduce the communication delay between distant points of the array. From relation (10) we may write

$$\frac{\lambda^2}{2\pi} nL_{\text{ave}} \leq (\text{total volume}), \quad (11)$$

where L_{ave} denotes the average interconnection length. For the purposes of this discussion L_{ave} will be assumed to be proportional to the linear extent of the system, which will in turn be assumed to be proportional to $(\text{total volume})^{1/3}$. Thus we will write $L_{\text{ave}} = \zeta(\text{total volume})^{1/3}$, where ζ is a constant. This is not of course generally true; the average interconnection length need not be proportional to the linear extent of the system. However, a more elaborate model is beyond the scope of the present paper. It can be shown to be true in particular for highly interconnected architectures such as binary hypercubes of high dimension. We can now solve for the total volume and the average interconnection length as

$$(\text{total volume}) \geq \left(\frac{\lambda^2}{2\pi} \zeta n \right)^{3/2}, \quad (12)$$

$$L_{\text{ave}} \geq \frac{\zeta^{3/2} n^{1/2} \lambda}{\sqrt{2\pi}}. \quad (13)$$

The average signal delay may then be found as

$$\tau = \frac{L_{\text{ave}}}{c} \geq \left(\frac{\zeta^{3/2}}{\sqrt{2\pi}} \right) \frac{n^{1/2}}{f}, \quad (14)$$

which we write as

$$\frac{n^{1/2}}{\tau} \leq \left(\frac{\sqrt{2\pi}}{\zeta^{3/2}} \right) f, \quad (15)$$

where f is the frequency of the optical sources. The last relation represents a tradeoff between the signal delay and the number of optical connections established. Our result may be considered to be a generalization of that given by Shamir for communication between two planes⁵ to an arbitrarily overlapping pattern of interconnections between points laid out on a three-dimensional grid. Shamir has already noted that this result constitutes a fundamental limit for parallel processing involving global communication. Similar results have been previously established for communication using solid wires.¹¹

4. EXAMPLES

In the remainder of this paper we will no longer concern ourselves with numerical factors such as 2π . Although architectures that are communication-volume limited in the sense previously defined can exist in principle, squeezing out the last factor of 2 or so would require considerable ingenuity and would usually not be practical.

To illustrate a case that is not communication-volume limited, let us assume that the transducers among which connections are to be established are constrained to be on

the surface of a sphere. Assume that there are a total of $2n$ transducers with n connections to be made between them. We wish to determine the minimum radius R and volume V of the sphere as a function of the number of connections n . Also, suppose that the transducer areas are given as $m^2\lambda^2$ with $m^2 > (1/\pi)$.² The surface area, the radius, and, hence, the volume of the sphere are constrained to minimum values of

$$4\pi R^2 = 2nm^2\lambda^2, \quad (16)$$

$$R = \left(\frac{n}{2\pi} \right)^{1/2} m\lambda, \quad (17)$$

$$V = \frac{4}{3} \pi R^3 = \frac{1}{3} \left(\frac{2}{\pi} \right)^{1/2} n^{3/2} m^3 \lambda^3 \sim n^{3/2} m^3 \lambda^3. \quad (18)$$

The lower bound of the communication volume is, using $V \geq \lambda^2 nL_{\text{ave}}/2\pi$,

$$V_{\text{lb}} = \frac{\lambda^2}{2\pi} n\zeta 2R = \frac{1}{\pi} \left(\frac{1}{2\pi} \right)^{1/2} \zeta n^{3/2} m\lambda^3 \sim \zeta n^{3/2} m\lambda^3, \quad (19)$$

where we used $L_{\text{ave}} = \zeta 2R$ since $2R$ is the linear extent of the system. The ratio between the two previous equations for volume is approximately

$$\frac{V}{V_{\text{lb}}} \sim \frac{m^2}{\zeta}. \quad (20)$$

It is seen that for large m (i.e., transducer areas $\gg \lambda^2$) or small ζ (i.e., $L_{\text{ave}} \ll$ system linear extent) we are doing considerably worse than as predicted by our lower bounds. 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.

The distinction between being communication-volume limited or transducer-surface limited arises from the fact that independently excited optical wave fields can noninterferingly share the same volume. Hence, regardless of how large the transducer sizes and areas associated with each independent channel of information are [Eq. (3)], it is always, in principle, possible to achieve globally an effective cross section $\sim \lambda^2$. As the transducer areas approach this size, being communication-volume limited or transducer-surface limited becomes one and the same thing.

We will further discuss these ideas by using two more examples. Figure 4 depicts holographic imaging configurations that have been suggested for intrachip and chip-to-chip optical communication for very-large-scale integration circuits.^{12,13} The scheme shown in Fig. 4(a) is a space-invariant design that utilizes the full aperture for each multiplexed volume hologram. Figure 4(b) shows a space-variant multifacet implementation for which one facet is allocated for each connection. The scheme of Fig. 4(c) is a modification of Fig. 4(b) in which collimating lenses are used above each transducer. This permits smaller transducer sizes (important in terms of energy considerations) and also has been shown to be more robust for alignment inaccuracies and drifts in parameter values.¹⁴ In Figs. 4(a) and 4(b) we will denote the transducer areas as a . In Fig. 4(c) a will denote the area of the collimating holographic lenslets. It is shown

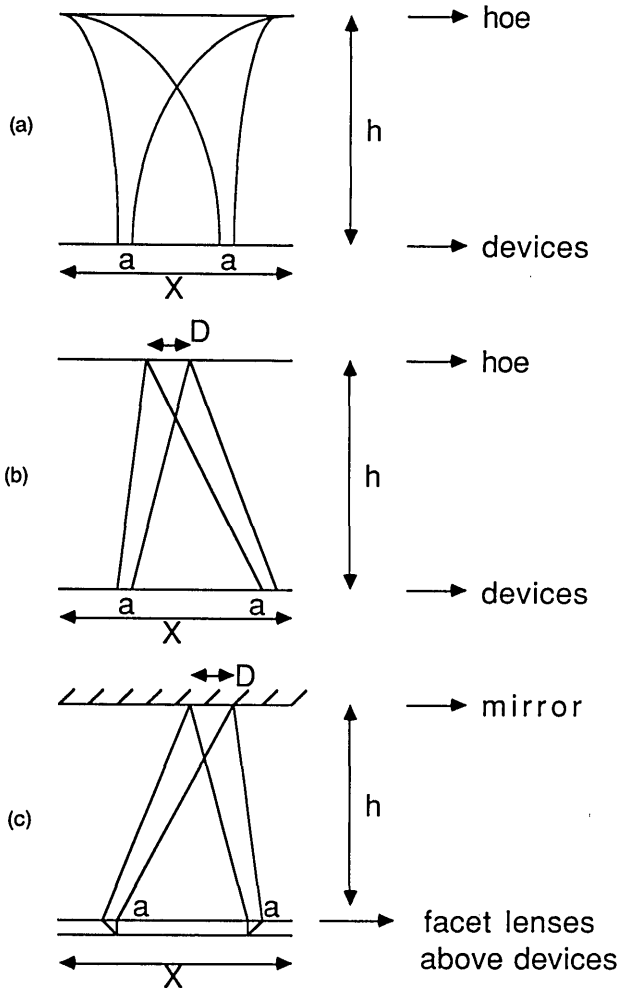


Fig. 4. Holographic-optical-interconnection schemes. a denotes the areas of the transducers [the area of the facet lenses in (c)]. X is the linear extent of the devices to be connected. h is the height of the holographic element from the device plane. (a) Space-invariant imaging configuration for which volume multiplexed holograms are used. (b) Multifacet space-variant design. (c) Modification of this scheme with transmission rather than reflection holograms.

in Appendix C that, for the multifacet designs, owing to diffraction considerations one can approximate

$$n \sim \frac{a}{\lambda^2}, \tag{21}$$

where we have assumed that there are n connections to be formed between n pairs of transducers. If we are to be able to form connections between distant parts of the planar layout, the height of the holographic optical element from the layout plane should be approximately equal to the maximum length that is to be connected. Although this latter condition is true also for the scheme depicted in Fig. 4(a), no such condition as relation (21) holds and the transducer areas can, in principle, be reduced down to the order of the wavelength squared.

We first consider the space-invariant design shown in Fig. 4(a). With the remarks of the previous paragraph in mind, note that $2n$ transducers forming n links will occupy an area of $X^2 \simeq 2n\lambda^2$. Since $h \sim X$, a typical communication length is approximately $2h \simeq n^{1/2}\lambda$. Using relation (10) we arrive at a lower bound for the volume as

$$V_{lb} = \frac{\lambda^2}{2\pi} n 2h \simeq \frac{\lambda^2}{2\pi} n^{3/2} \lambda \sim n^{3/2} \lambda^3. \tag{22}$$

The actual communication volume is found simply, again with $h \sim X$, as

$$V = X^2 h \simeq X^3 \sim n^{3/2} \lambda^3, \tag{23}$$

which, within a numerical factor, is the same as relation (22). However, notice that, if the transducers are made considerably larger than a wavelength, this design becomes transducer-surface limited. The high $f\#$ beams will only occupy a fraction of the available modal space.

Now we turn to the space-variant designs. Again $X^2 = 2na$, with a being given by relation (21). So again with $h \sim X$ we obtain the volume $X^2 h \simeq X^3 \sim n^{3/2} \lambda^3$. The typical communication distance is now $2h \simeq n\lambda$. Dividing the total volume by this distance and by n , we find that the average effective cross section required for communication is

$$n\lambda^2. \tag{24}$$

This implies that not only is the typical interconnection distance larger by a factor of $n^{1/2}$ over the communication-volume-limited case, but also the effective cross section is not a constant but rather a linear function of n . In this respect, especially for large values of n , this design is far from approaching fundamental limits. Since the effective cross section is a growing function of n , it is even worse than transducer-surface-limited designs. We noted that high-numerical-aperture single-mode waveguides of cross section $\sim \lambda^2$ serve as an existence proof that arbitrary space-variant connections may be achieved with communication-volume-limited systems. A problem of great practical interest is to devise space-variant free-space architectures that are also communication-volume limited. Such an architecture may involve a three-dimensional layout of the points to be connected.

5. CONCLUSIONS

We saw that for an arbitrary array of points communicating with one another, 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 permitted to overlap in coordinate space, or in Fourier space, but not in 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 $\rho 2\pi V/\lambda^3$, where ρ is the modulation bandwidth of the output transducers normalized by the carrier frequency.

The utility of our global viewpoint is that it enables one to model the basic mechanisms that limit how closely one can pack an array of optically interconnected primitive computing elements to form a larger computing system. The advantage that is to be gained by using free-space alternatives

(i.e., overlap of independently excited wave fields possible) over guided-wave alternatives (with no overlap permitted 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 permitted. As the $f\#$'s in question increase, however, it becomes more and more important to be able to noninterferingly overlap independent wave fields so as to make better utilization of the Fourier space and hence the available modal volume.

Most suggested optical global interconnection architectures consist of a planar array of devices. We saw that the communication volume required by such systems are transducer-surface limited unless the transducers are reduced to sizes of the order of one wavelength. The problem of minimizing the communication volume will increase in importance as smaller transducer areas and reduced gate 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 that enable efficient utilization of both coordinate and Fourier spaces, even with transducer areas greater than of the order of one wavelength.

Another important conclusion may be drawn from relation (10). Observe that the minimum communication volume is linear only 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 $\text{jinc}(r)$,^{15,16} functions whose self-convolutions are identical to themselves. Thus, apart from the effects of aberrations, the diffraction-limited spot size does not increase on concatenating several identical imaging systems in order 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 interconnections is superlinear in distance. This is because longer interconnections 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. This advantage of optical interconnections in terms of interconnect density is independent of and in addition to those discussed by Feldman *et al.*⁷

Throughout our exposition we only briefly discussed how the communication lengths are specified in terms of the total volume of the computing system. More accurate specification of communication lengths will require introduction of a computational model and is beyond the scope of this paper.

APPENDIX A. ADDITIONAL FACTORS OF 2 CONTRIBUTING TO THE NUMBER OF DEGREES OF FREEDOM

There are several factors of 2 that can be included in Eq. (1) and in our major results, such as relation (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^{17,18} 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 that is sensitive to the temporal phase of the optical carrier. With the understanding that they may be easily reintroduced whenever appropriate, we did not include either.

Bidirectional (i.e., full-duplex) utilization of the channels does not necessitate that bidirectional links be formed between every transmitting and receiving transducer pair. It is only necessary that one efficiently utilizes given cross-sectional areas by making use of these additional available degrees of freedom, as depicted in Fig. 5.

APPENDIX B. EXTENSION TO FAN IN AND FAN OUT

The extension of our main result to fan in and fan out is straightforward. Consider, for instance, the fan-out situation symbolically depicted in Fig. 6. Given the locations of the source and detectors, one arranges power division to occur at such points so that the total distance required, $L_1 + L_2 + L_3$, is minimal. Then the contribution of this fan-out link to the total volume required is just $\lambda^2(L_1 + L_2 + L_3)/2\pi$.




Fig. 5. Bidirectional utilization of available channels. Bidirectional utilization of the channels does not require each transducer pair to form bidirectional links. Here it is symbolically shown how this last factor of 2 may be squeezed out. Whenever one of the counterpropagating signals in a certain path diverts out into another direction, another enters the now-empty propagation mode.




Fig. 6. Fan-out situation. A multiple-signal net is shown. 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.

APPENDIX C. MINIMUM FACET AREA FOR THE MULTIFACET-HOLOGRAPHIC-INTERCONNECTION SCHEME

We refer to Fig. 4(b). Assume that we want to form interconnections between pairs of transducers that, in the worst case, may be separated from one another by a distance that is comparable with the total extent of the devices, which are laid out on a planar surface. We also assume that the total area of the holographic optical element is approximately equal to the total area occupied by these devices. Since n connections are being formed between n pairs of transducers, we write

$$2na \simeq nD^2, \quad (25)$$

where D is the width of a facet. We also have

$$a \simeq \frac{h^2}{D^2} \lambda^2, \quad (26)$$

since h/D corresponds to the $f\#$. There also exists a limit to how large X/h can be. Detailed theoretical and experimental analyses of these issues were previously given in Refs. 19 and 20. We require only an approximate expression, however; we will take

$$X \leq 2h. \quad (27)$$

Since $X^2 \simeq 2na$, combining the above, we obtain that the minimum value of a satisfies

$$\frac{a}{\lambda^2} \sim n \quad (28)$$

within a numerical factor of the order of unity. A similar relation can be shown to hold for Fig. 4(c). An engineering analysis of this and similar configurations may be found in Ref. 21.

ACKNOWLEDGMENTS

The authors thank Yaakov Amitai for helpful discussions. This research was supported by the U.S. Air Force Office of Scientific Research under Grant AFOSR-88-0024.

REFERENCES

1. D. Gabor, "Light and information," in *Progress in Optics*, E. Wolf, ed. (North-Holland, Amsterdam, 1961), Vol. 1, Chap. 4.
2. J. T. Winthrop, "Propagation of structural information in optical wave fields," *J. Opt. Soc. Am.* **61**, 15-30 (1971).
3. M. Born and E. Wolf, *Principles of Optics*, 6th ed. (Pergamon, Oxford, 1980), pp. 164-169.
4. K. D. Moller, *Optics* (University Science, Mill Valley, Calif., 1988), p. 388.
5. J. Shamir, "Fundamental speed limitations on parallel processing," *Appl. Opt.* **26**, 1567 (1987).
6. R. Barakat and J. Reif, "Lower bound on the computational efficiency of optical computing systems," *Appl. Opt.* **26**, 1015-1018 (1987).
7. M. R. Feldman, C. C. Guest, T. J. Drabik, and S. C. Esener, "Comparison between electrical and free space optical interconnects for fine grain processor arrays based on interconnect density capabilities," *Appl. Opt.* **28**, 3820-3829 (1989).
8. C. D. Thompson, "Area-time complexity for VLSI," in *Proceedings of the 11th Annual ACM Symposium on the Theory of Computing* (Association for Computing Machinery, New York, 1979), pp. 81-88.
9. J. D. Ullman, *Computational Aspects of VLSI* (Computer Science Press, Rockville, Md., 1984).
10. R. P. Feynman, R. B. Leighton, and M. Sands, *The Feynman Lectures on Physics* (Addison-Wesley, Reading, Mass., 1965), Vol. 3, pp. 4-11.
11. A. C. Hartmann and J. D. Ullman, "Model categories for theories of parallel systems," in *Parallel Computing: Theory and Experience*, G. J. Lipovski and M. Malek, eds. (Wiley, New York, 1986), App. H.
12. J. W. Goodman, F. J. Leonberger, S-Y. Kung, and R. Athale, "Optical interconnections for VLSI systems," *Proc. IEEE* **72**, 850-866 (1984).
13. R. K. Kostuk, J. W. Goodman, and L. Hesselink, "Optical imaging applied to microelectronic chip-to-chip interconnections," *Appl. Opt.* **24**, 2851-2858 (1985).
14. E. Bradley, P. Kit Lai Yu, and A. R. Johnston, "System issues relating to laser diode requirements for VLSI holographic optical interconnects," *Opt. Eng.* **28**, 201-211 (1989).
15. J. W. Goodman, *Introduction to Fourier Optics* (McGraw-Hill, New York, 1968), p. 62.
16. R. N. Bracewell, *The Fourier Transform and Its Applications*, 2nd ed. (McGraw-Hill, New York, 1986).
17. A. V. Oppenheim and R. W. Shafer, *Digital Signal Processing* (Prentice-Hall International, London, 1975), p. 26.
18. F. T. S. Yu, *Optics and Information Theory* (Wiley, New York, 1976).
19. R. K. Kostuk, J. W. Goodman, and L. Hesselink, "Design considerations for holographic optical interconnects," *Appl. Opt.* **26**, 3947-3953 (1987).
20. R. K. Kostuk, "Multiple grating volume reflection holograms with application to optical interconnects," Ph.D. dissertation (Stanford University, Stanford, Calif., 1986).
21. M. R. Feldman and C. C. Guest, "Interconnect density capabilities of computer generated holograms for optical interconnection of very large scale integrated circuits," *Appl. Opt.* **28**, 3134-3137 (1989).