Limitations on Optical Free-Space Crossbar-Like Interconnection Networks

Dror G. Feitelson Larry Rudolph Eugen Schenfeld

Department of Computer Science, The Hebrew University of Jerusalem
Givat Ram, 91904 Jerusalem, Israel

ABSTRACT

Providing the required interconnections between the processors of a parallel computer is a difficult problem: latency, switching control, cost, and crosstalk effects have to be taken into account. It is widely believed that the design might be simplified if optical technology is used. However, even optical interconnections cannot cater for an unlimited number of processors. For systems with up to about $10^5 - 10^6$ processors, it seems that the arrangement of the processors in space so that each can "see" the others is the limiting factor; A large part of the volume involved must stay empty for the information-bearing light beams. above that, diffraction effects and power requirements place a limit on the number of processors.

1. INTRODUCTION

A basic problem in the design of parallel computers is that of connecting the processing elements (PEs) to each other. A minimal requirement is to enable data transfers between any pair of PEs, maybe through intermediaries; this saves hardware, but complicates and delays the routing. A maximal requirement is to have a direct channel between every pair; this is fast and simple, but very wasteful. A practical compromise is to create an interconnection mechanism that allows each PE to connect to only one other PE at a time, but with dynamic reconfiguration. Crossbars and multistage networks are based on this idea.

When a reconfigurable network is designed, the questions of layout in 3-D space, routing between the various elements, and control over the routing switches must be settled. Different implementations differ in the above respects. They can be evaluated and compared according to the following criteria:

- Density: the number of PEs that may be connected within a given volume.
- Degree of locality supported by the network, i.e. the minimal, average, and maximal path lengths.
- Possible interconnection patterns: point-to-point, multicast, broadcast, etc.
- Achievable data rates.
- Achievable reconfiguration rates.
- Power requirements.

These properties are determined by the technology that is used, by the layout of the network, and by the properties of the switching device. This paper is about networks using optical free-space technology, based on the "DynaSwitch" switching element [1]; possible designs that use this switch are outlined in the next section. Section 3 investigate how the network layout affects the number of PEs that may be connected within a given volume, based on geometrical considerations. Diffraction considerations and power requirements are taken into account in sections 4 and 5, respectively.

2. NETWORK DESIGN

2.1. General considerations

We identify three possible basic designs (fig. 1):

1. There is a single central switching device, to which all the PEs are connected. This central device is responsible for performing the routing, according to the specified configuration.
2. The switches are distributed throughout the network. Each switch is simpler than before, but a number of switches must be passed on the way from the source to the destination.

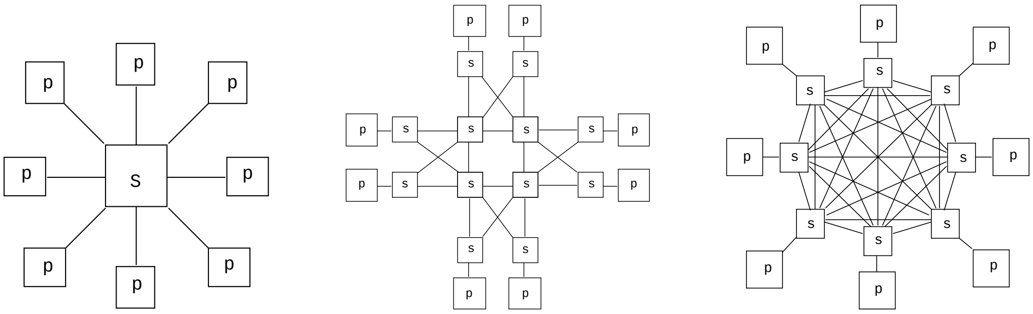

Figure 1: *Three possible relations between processors and switches in a network.*

3. The switches are associated with the PEs, with each PE responsible for the routing of the data that it transmits. The switches are fully connected.

The first option is problematic when large systems are considered. This is seen from the following observations:

- If the surface area of the switch does not increase with the number of PEs, the wiring complexity for both data and control lines becomes prohibitive. This might limit the number of elements that are connected. A number of proposed devices suffer from this drawback [2, 3, 4].
- If only one central device is used, the whole configuration might have to be reset even if only a small part of it changes. The new setup is computed by a single controller, instead of exploiting the possible parallelism inherent in the routing problem. This happens, for example, with monolithic routing holograms [5, 2]
- Physical limits on single devices (e.g. limited space-bandwidth product) limit the number of channels that can be supported. Examples include the crossbar based on an optical matrix-vector multiplier [6], and multiwavelength star couplers [7].

The second scheme overcomes these difficulties at least to some degree, by employing simpler switches. However, it suffers from the following drawbacks:

- A distributed control scheme is needed to coordinate all the switches. The simplest solution is to use a packet-switching scheme, but then the switches need to have processing and buffering capabilities, and control information must be transmitted with the data.
- The latency could be higher, because data has to pass more active elements. In packet switching, for example, the data might be buffered in each node it passes.

We therefore consider networks in which the switching is distributed, with one switch associated with each PE; direct links are provided between every pair of switches. Such networks are totally unpractical with electronic technology, because they require full interconnection. They are plausible, however, with optical free-space interconnections. The networks we study are based on the "DynaSwitch" switching element [1]. Technically, these networks use a beam steering approach rather than a beam spreading and masking approach [8, 9]. In other words, the switches are not fully connected on a permanent basis, but rather each can potentially form a direct connection with all the others. The switch is characterized by the following properties:

- One-to-N functionality, with an option to connect to k of the outputs in parallel ($k \leq N$).
- Simple electrical control that defines which of the N outputs are connected to the input.
- Maximal data rate of 10 Gb/s.
- Reconfiguration time of 1μs–1ms, depending on the power supply.
- External optical power supply.

It is implemented by a dynamic hologram written in a photorefractive crystal. This hologram deflects the output beam from the PE to the desired direction, and splits it if necessary.

The question is how to use a set of such switches, one per PE, to create a crossbar-like network. Each PE/switch combination has a volume v that has to be taken into account. The PE has one input and one output; the output

Table 1: *Parameters of the analysis.*

switch	x^2	area allocated per PE/switch for transmission
	y^2	area allocated per PE/switch for detection
	z^2	total area allocated per PE/switch
	α	Maximal deflection angle
	v	PE/switch volume
light	ℓ	length of light's path segment
	s_0	radius of light beam at transmission
	s_ℓ	radius of light beam after distance ℓ
	λ	wavelength of light
	E	transmission power
sphere	R	radius of sphere
cube	f^2	size of facet in routing hologram in cube
	β	board form factor (see fig. 5)
general	a^2	number of PEs in array or board
	N	total number of PEs
	V	total volume

may be directed by the switch to k other PEs. If a collision occurs at the input, as a result of multiple PEs that send messages to the same PE, this is detected by the destination PE. We require that all the PEs be able to transmit in parallel, and the network should handle all this activity without causing any limitations.

2.2. The compared designs

We study three basic designs for a network based on the "DynaSwitch": in the first two are closely related: in both the PEs are located on the surface of a certain volume that is used for communication. The difference is that in the first design the communication volume is cubical, and in the second it is spherical. In the third design the PEs are located within the volume, with interconnection channels passing between them.

The communication in the three designs is done as follows:

The Array: The PEs are arranged in a square 2-D array on one face of a cube. The opposite face is a big mirror. Communication is performed by directing a beam of light at a certain point on the mirror so that it is reflected to the destination PE.

The Hemisphere: This is an extension of the array design: instead of a cube we have a spherical volume, with one half of the sphere containing the PEs and the other half a reflecting surface. If the maximal deflection angle of the switch is too small, the whole surface cannot be utilized. In such cases, only a pair of spherical caps are used: one with the processors, the other reflecting.

The Cube: The PEs are arranged in small arrays on boards. In addition to the PEs, each board contains a set of routing holograms, which include a facet for each destination PE. The boards are stacked in a cube structure. In order to transmit to a certain destination PE, the transmitter directs a light beam at the facet associated with the destination, in the routing hologram of the board above it (the output beam is emitted from the top of the PE). The routing holograms direct the beam from one board to the next, until it reaches the bottom of the destination PE. A detailed description of this scheme is given in a separate paper [10].

The parameters that are used throughout the analysis are defined in table 1.

3. GEOMETRICAL CONSIDERATIONS

Geometrical limitations are the result of the spatial layout of the network. The basic limiting factor is the deflection angle of the switch, α (this is measured from the normal, so the total coverage is 2α). In this section, we assume the communication is mediated by perfect rays of light. The next section will deal with diffraction effects.

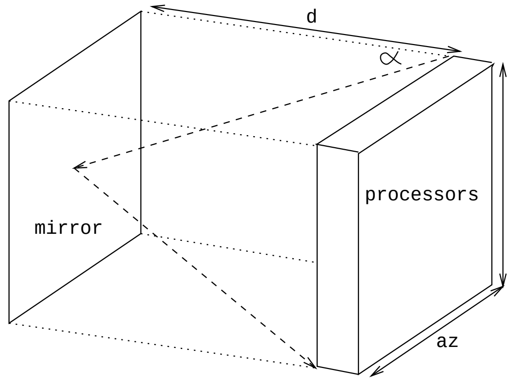

Figure 2: *Figure to aid the derivation for the array, showing the route that requires the maximal d.*

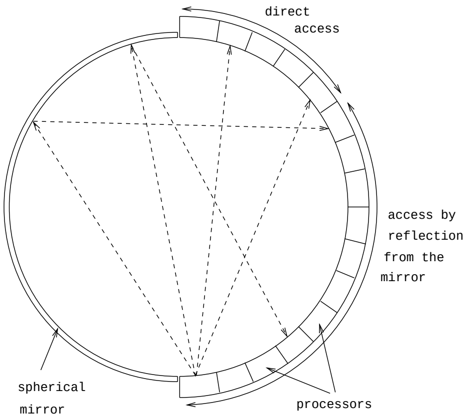

Figure 3: *Routing in the sphere when $\alpha \geq 30°$ (cross section). PEs near the rim use both direct beams and reflected beams.*

3.1. The Array

The maximal communication distance in this design is from one corner to the opposite corner, using a reflection off the center of the mirror (fig. 2). The distance d between the array and the mirror needed for this is

$$d = \frac{az}{\sqrt{2}\tan\alpha}$$

The total volume is then

$$V = a^2z^2d + Nv = \frac{(az)^3}{\sqrt{2}\tan\alpha} + Nv$$

And as $N = a^2$ we get

$$V = \frac{z^3}{\sqrt{2}\tan\alpha}N^{3/2} + Nv \tag{1}$$

3.2. The Hemisphere

As noted above, the hemisphere displays a discontinuity when α is changed. Below $30°$, the PEs and mirror occupy two spherical caps facing each other. Above $30°$ the spherical caps turn into touching hemispheres, so that all the surface is utilized; additional increase of α has no effect (fig. 3 shows how the routing is done in this case). We therefore divide the discussion into two cases.

Below $30°$

The arrangement is shown in fig. 4. The number of PEs in this case is determined by the surface area of the cap, which has a central angle of 4α. Ignoring problems associated with the shape and arrangement of the PEs, it is

$$N = \frac{2\pi R^2(1-\cos 2\alpha)}{z^2} \tag{2}$$

The volume that is used includes both caps and the cylinder between them, plus the volume of the PEs themselves. This is given by

$$\begin{aligned} V &= \frac{4\pi}{3}R^3(1-\cos^3 2\alpha) + Nv \\ &= \frac{\sqrt{2}}{3\sqrt{\pi}}\frac{(1-\cos^3 2\alpha)z^3}{(1-\cos 2\alpha)^{3/2}}N^{3/2} + Nv \end{aligned} \tag{3}$$

where equation (2) is used to express R in terms of N and z.

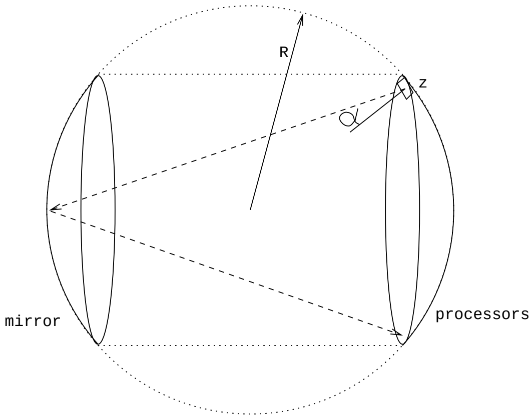

Figure 4: *Figure to aid the derivation for the hemisphere under 30°, showing the path using the maximal α and thus limiting the size of the spherical caps that are used.*

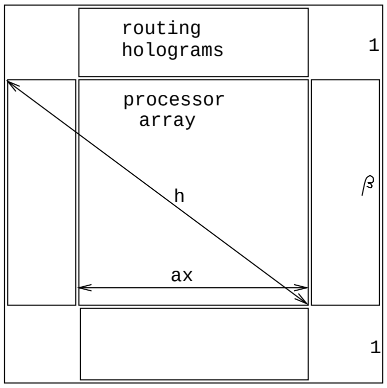

Figure 5: *Configuration of a board, showing the maximal horizontal distance h. The board form factor, β, is defined as the ratio of the width of the PE array to the width of the routing hologram.*

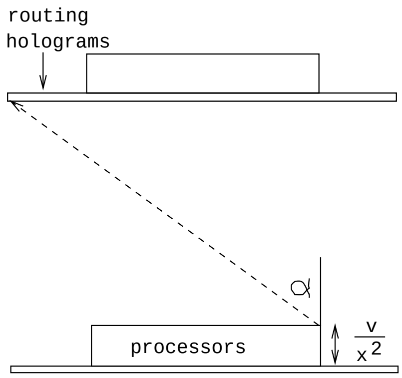

Figure 6: *The vertical distance between boards depends on the deflection angle and on the height of the PEs.*

Above 30°

Above 30° the equations are simpler, because the full hemisphere is utilized and α no longer has any effect. The number of PEs is

$$N = \frac{1}{2}\frac{4\pi R^2}{z^2} \tag{4}$$

and the volume needed for N PEs is

$$V = \frac{4\pi}{3}R^3 + Nv = \frac{\sqrt{2}z^3}{3\sqrt{\pi}}N^{3/2} + Nv \tag{5}$$

3.3. The Cube

In the cube, the transmission and detection are done from opposite sides, so the area needed per PE[1] is x^2 rather than z^2. The PEs are arranged in a square array[2], with four rectangular routing holograms around them (fig. 5). The size of the holograms is expressed by the ratio β of their width to that of the PE array. The board area, including the PEs, the routing holograms, and area for support structures, is

$$\frac{(\beta+2)^2 a^2 x^2}{\beta^2}$$

[1] There is a hidden assumption here that $y \leq x$.
[2] Note that now a^2 is the number of PEs per board, not the total number

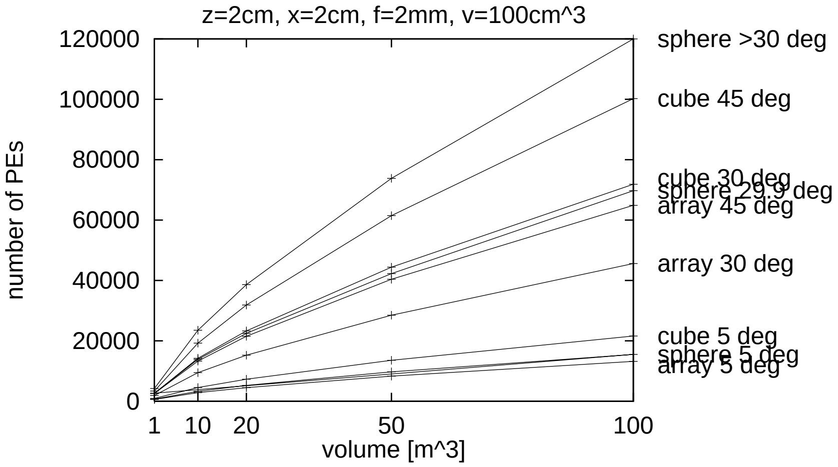

Figure 7: *A comparison of the number of PEs accommodated within a given volume in the three schemes. In the cube, β is set for the maximal number.*

The maximal horizontal distance is from a corner of the PE array on the board to the far corner of the hologram. This is

$$h = \frac{\sqrt{2\beta^2 + 2\beta + 1}\, ax}{\beta}$$

With a maximal deflection angle of α, the vertical distance to the hologram is $h/\tan\alpha$. The distance between boards, however, must also take into account the volume of the PEs. Thus the total distance between boards is (fig. 6)

$$\frac{\sqrt{2\beta^2 + 2\beta + 1}\, ax}{\beta \tan\alpha} + \frac{v}{x^2}$$

An additional constraint is imposed by the holograms: the four holograms per board must be large enough to contain N facets. This gives

$$N f^2 = \frac{4a^2x^2}{\beta} \tag{6}$$

which actually means that once x, f, and β are given, the number of boards is determined: it is simply $N/a^2 = (4x^2)/(\beta f^2)$.

From all the above, the following expression for the volume is derived[3]:

$$\begin{aligned} V &= \#\text{ boards} \times \text{board area} \times \text{board distance} \\ &= \frac{N}{a^2}\frac{(\beta+2)^2a^2x^2}{\beta^2}\left(\frac{\sqrt{2\beta^2+2\beta+1}\, ax}{\beta\tan\alpha} + \frac{v}{x^2}\right) \\ &= \frac{(\beta+2)^2\sqrt{2\beta^2+2\beta+1}}{2\beta^{5/2}\tan\alpha}N^{3/2}x^2 f + \frac{(\beta+2)^2}{\beta^2}Nv \end{aligned} \tag{7}$$

3.4. Discussion

Using a simple geometrical consideration, we find that the volume needed for a network of N PEs always has two terms: the first is proportional to $N^{3/2}$, which means that in some sense the PEs are always situated on a surface enclosing the communication volume, even in the cube. The second is proportional to Nv, and describes the volume required by the PEs themselves.

When typical values are assigned to the parameters, the performance of the three designs can be compared numerically. The results are plotted in fig. 7. It turns out that the full hemisphere is the most volume efficient when α is

[3] This expression ignores extra layers for crossing from one column to another.

large. If $\alpha < 30°$, however, the cube takes the lead. The efficiency of the cube is further improved if the facet size f is reduced. The array is always the least efficient of all. The volume needed for the interconnections is very apparent: in fact, most of the volume is left empty. All this empty space is needed for the imformation-carrying beams of light.

As for path lengths, the cube allows us to arrange the boards as we wish, so we can always get a more or less cubical shape; in the array and hemisphere small αs cause the shape to be long and narrow, so all the paths are long.

4. DIFFRACTION CONSIDERATIONS

Modeling a system with rays of light rather than beams of light is sometimes useful, but in our case it is an over-simplification that may create overly optimistic results. In real-life situations, one must take beam-expansion resulting from diffraction into account. We start with the following formula that describes the spot size for a beam with Gaussian cross section and initial radius s_0 after a distance ℓ [11]:

$$s_\ell = \sqrt{s_0^2 + \frac{\ell^2 \lambda^2}{\pi^2 n^2 s_0^2}}$$

where n is the index of refraction; in our case, $n = 1$. In addition, due to deflection angles that can be as much as α, the effective source radius is $s_0 \cos\alpha$, and the spot size is $\frac{1}{\cos\alpha}$ of the size given by the formula. Therefore we use the following version:

$$s_\ell = \sqrt{s_0^2 + \frac{\ell^2 \lambda^2}{\pi^2 s_0^2 \cos^4\alpha}} \tag{8}$$

4.1. The Array

The basic equations are the same as in the previous section, with the added constraint that $2s_\ell \leq z$. Using equation (8) and squaring on both sides, we get

$$4s_0^2 + \frac{4\ell^2\lambda^2}{\pi^2 s_0^2 \cos^4\alpha} \leq z^2$$

The distance in the maximal path is

$$\ell = 2\frac{d}{\cos\alpha} = \frac{\sqrt{2}\,az}{\sin\alpha}$$

Using this, the inequality may be turned into the following quadratic equation in s_0^2:

$$4s_0^4 - z^2 s_0^2 + \frac{8a^2 z^2 \lambda^2}{\pi^2 \sin^2\alpha \cos^4\alpha} \leq 0$$

From this we can derive a limit on the size of the array as a function of the ratio of z to λ. To do so, we solve for s_0^2:

$$s_0^2 = \frac{z^2 \pm \sqrt{z^4 - \frac{128a^2z^2\lambda^2}{\pi^2\sin^2\alpha\cos^4\alpha}}}{8}$$

A solution exists only if

$$z^4 - \frac{128a^2z^2\lambda^2}{\pi^2\sin^2\alpha\cos^4\alpha} \geq 0$$

Dividing by z^2 and noting that $a^2 = N$, this leads to

$$N \leq \frac{\pi^2 \sin^2\alpha \cos^4\alpha \; z^2}{128\lambda^2} \tag{9}$$

Note that this limit does not depend on the volume. However, it does depend on z, so it is not a global limit.

Next, we can derive an expression for the volume needed for a certain number of PEs. We use the limit (9) to get a lower bound on z, and plug it into the basic equation for the total volume (1), $V = \frac{1}{\sqrt{2}\tan\alpha} N^{3/2} z^3 + Nv$. The result simplifies to

$$V \geq \frac{1024\lambda^3 N^3}{\pi^3 \sin^4\alpha \cos^5\alpha} + Nv \tag{10}$$

4.2. The Hemisphere

When we address beam expansion in the hemisphere, we must also take into account the focusing properties of a spherical mirror. Luckily it turns out that the size of a spot that the beam creates on the sphere surface is the same as the size it would have created if it continued without reflection.

The maximal distance a beam might have to traverse is $\ell = 4R$, so the requirement on the spot size is

$$4s_0^2 + \frac{64R^2\lambda^2}{\pi^2 s_0^2 \cos^4\alpha} \leq z^2$$

where a value of $\alpha = 30°$ should be used whenever the maximal deflection angle exceeds $30°$. This again reduces to a quadratic inequality

$$4s_0^4 - z^2 s_0^2 + \frac{64R^2\lambda^2}{\pi^2\cos^4\alpha} \leq 0$$

The condition for the existence of a solution again provides an upper limit on the number of PEs that depends only on the ratio of z to λ. For $\alpha < 30°$ we use the basic equation for the number of PEs (2), $N = 2\pi R^2(1-\cos 2\alpha)/z^2$, and find that this limit is

$$N \leq \frac{\pi^3 \sin^2\alpha \cos^4\alpha \; z^2}{256\lambda^2} \tag{11}$$

As for the array, we can also find the volume required for any number of PEs. We derive a lower bound on z from (11) and plug it into the equation for the volume (3), $V = \frac{\sqrt{2}(1-\cos^3 2\alpha)}{3\sqrt{\pi}(1-\cos 2\alpha)^{3/2}} z^3 N^{3/2} + Nv$, giving (after simplification)

$$V \geq \frac{2048(1-\cos^3 2\alpha)\lambda^3}{3\pi^5 \sin^6\alpha \cos^6\alpha} N^3 + Nv$$

The whole process can be repeated for $\alpha \geq 30°$. Using the equation for the number of PEs in this case (4), $N = 2\pi R^2/z^2$, and the equality $\cos 30° = \frac{\sqrt{3}}{2}$, we find that the limit on N is

$$N \leq \frac{9\pi^3 z^2}{8192\lambda^2} \tag{12}$$

To find the volume required for a certain number of PEs we derive a lower bound on z from this equation and plug it into the equation for the volume of the sphere (5), $V = \frac{\sqrt{2}}{3\sqrt{\pi}} z^3 N^{3/2} + Nv$, giving (after simplification)

$$V \geq \frac{2^{20}\lambda^3}{81\pi^5} N^3 + Nv$$

4.3. The Cube

In this case the requirement is that the spot created by the expansion of the beam not be larger than a facet of the hologram:

$$4s_0^2 + \frac{4(2\beta^2+2\beta+1)a^2x^2\lambda^2}{\beta^2\pi^2 s_0^2 \sin^2\alpha \cos^4\alpha} \leq f^2$$

Once again we get a quadratic expression for s_0^2, and in this case the condition for the existence of a solution is

$$f^4 \geq \frac{64(2\beta^2+2\beta+1)a^2x^2\lambda^2}{\beta^2\pi^2\sin^2\alpha\cos^4\alpha}$$

Using the constraint that the hologram must be large enough to contain all the required facets (6), $Nf^2 = 4a^2x^2/\beta$, we can express a^2x^2 by the other parameters. We then get

$$N \leq \frac{\pi^2\beta \sin^2\alpha \cos^4\alpha f^2}{16(2\beta^2+2\beta+1)\lambda^2} \tag{13}$$

Table 2: *Asymptotic behavior and crossover points with and without diffraction considerations.*

scheme	*geometry*	*diffraction*	*crossover*
array	$V \approx 3.7 \times 10^{-6} N^{3/2}$	$V \approx 1.1 \times 10^{-14} N^3$	$N \approx 4.9 \times 10^5$
sphere	$V \approx 2.1 \times 10^{-6} N^{3/2}$	$V \approx 4.2 \times 10^{-17} N^3$	$N \approx 1.4 \times 10^7$
cube	$V \approx 4.6 \times 10^{-6} N^{3/2}$	$V^* \approx 2.5 \times 10^{-14} N^3$	$N^* \approx 3.2 \times 10^5$

$\alpha = 30°$, $z = x = 2$cm, $f = 2$mm, $\lambda = 1\mu$m, $\beta = 7$

* not a tight bound

As before, we would like to derive a lower bound on the volume needed for any number of PEs. Note, however, that the bound in equation (7) is expressed in terms of both x and f, while equation (13) can only provide a lower bound for f. We are therefore forced to make the following crude approximation: as a^2 is the number of PEs on *one* board, $a^2 \leq N$. Using this in equation (6), we get a bound on x in terms of f: it is $x^2 \geq f^2\beta/4$. This allows us to derive the bound

$$V \geq \frac{8(\beta+2)^2(2\beta^2+2\beta+1)^2\lambda^3}{\pi^3\beta^3\sin^4\alpha\cos^5\alpha} N^3 + \frac{(\beta+2)^2}{\beta^2} N v$$

Note, however, that this bound is not tight. If $a^2 \ll N$ it may predict volumes that are too large by orders of magnitude.

4.4. Discussion

The above results indicate that asymptotically diffraction poses a more severe constraint on the number of PEs that can be connected than geometrical considerations alone. We find that for a fixed z (f in that case of the cube) there is an upper limit to the number of PEs that may be connected (regardless of the volume!). However, if we make z (f) larger, more PEs can be connected. The price is that enlarging z (f) causes the volume to grow faster, and indeed we find that asymptotically the volume grows as the *third* power of the number of PEs. This is much faster than the $N^{3/2}$ growth rate due to geometrical considerations.

Numerical calculations show, however, that in the reasonable range of volumes (up to and exceeding 100 cubic meters), and with reasonable values for the design parameters (e.g. $z = 2$cm, $f = 2$mm), the diffraction effects do not come into play. The reason is that the constants in the equations that take diffraction into account are much smaller than the constants in the earlier equations that only use geometry. Taking the array as an example, we see that in equation (1) $V \propto N^{3/2}$, while in equation (10) we have $V \propto \lambda^3 N^3$. As λ is very small, the second expression is initially smaller. The crossover point occurs only when N is half a million. Therefore we can say that practical systems are not affected by diffraction, and are only limited by geometrical constraints. The constants and crossover points are tabulated in table 2.

5. POWER CONSIDERATIONS

In this section we consider limitations that may arise because not all the light that is emitted actually reaches the detector. The caculations do not take fan-out into consideration. If multicasts are used, the numbers given here should be divided by the maximal fan-out.

5.1. The Array

We assume that the mirror is 100% reflective. Given that the area for each PE, z^2, is determined so as to be just larger than the spot size, and assuming uniform illumination, the power that hits the detector is simply

$$d = E\frac{y^2}{z^2}$$

Using the limit on the number of processors (9) to derive a bound on z^2, we get

$$d \leq \frac{\pi^2\sin^2\alpha\cos^4\alpha\; Ey^2}{128\lambda^2 N}$$

This must be larger than the threshold of the detector. Denoting this threshold by D, we get a limit on the number of PEs:

$$N \leq \frac{\pi^2 \sin^2 \alpha \cos^4 \alpha \; E y^2}{128 \lambda^2 D}$$

An interesting question is whether it is possible to improve this by focusing the light on the detector. If so, we can simply use the diameter of the focusing lens for y, instead of the diameter of the detector itself. But this idea is problematic, for two reasons:

- If the maximal deflection angle α is large, the detector must also accept light coming from large angles. A lens will focus light coming at such angles to a point that is up to a distance of the focal length away from the optical axis. Thus it would require a detector with a diameter twice the focal length of the lens.
- The lens might obstruct the beam that is transmitted from the same processor. To prevent this, the area allocated for the processor would have to be increased.

5.2. The Hemisphere

This case is nearly identical to the array, except for the limit on the number of processors. For $\alpha < 30°$ this limit is given by (11). The final result is then

$$d \leq \frac{\pi^3 \sin^2 \alpha \cos^4 \alpha \; E y^2}{256 \lambda^2 N}$$

and limit on the number of the PEs is

$$N \leq \frac{\pi^3 \sin^2 \alpha \cos^4 \alpha \; E y^2}{256 \lambda^2 D}$$

For $\alpha \geq 30°$, the limit is given by (12). The final limit on the number of the PEs is

$$N \leq \frac{9 \pi^3 E y^2}{8192 \lambda^2 D}$$

5.3. The Cube

In this case the limits are not caused by the limited area of the detector: the detector always receives its light from the same direction, so all the light may be focused on it. However, the holograms along the way cause energy losses. The total loss is minimized if the efficiencies of the holograms along the path change. These changes are concerted so that the fraction of the energy that arrives at the destination is the same for all the transmitters, regardless of how far away they are. Thus holograms in the far branches of a path will need high efficiency, while holograms near the destination will have very low efficiencies. Counting from the farthest transmitter, a hologram m steps along the path has efficiency $1/m$. The maximal path effectively collects light from $\frac{N}{a^2}$ sources. At each one there are a^2 first-order beams, one of which coincides with the desired path. Thus the power that arrives at the end is

$$\begin{aligned} d &= \frac{E}{a^2} \prod_{m=2}^{N/a^2} \left(1 - \frac{1}{m}\right) \\ &= \frac{E}{a^2} \frac{1}{\frac{N}{a^2}} = \frac{E}{N} \end{aligned}$$

Hence the limit on the number of processors is simply $N \leq E/D$. This agrees with the arguments presented by Goodman [12] about fan-in limitations when multiple beams are combined; it shows that the scheme with changing efficiencies is optimal. Note that the above derivation is based on the assumption that no additional losses occur as a result of diffraction in the propagation between boards. This assumption is based on the use of perfect telescopic arrays [13] where necessary.

Table 3: *Limits imposed by power requirements, with no fan-out. E/D is of the order of 10^5.*

scheme	*bound*
array	$N \leq 4.3 \times 10^6 E/D$
sphere	$N \leq 1.4 \times 10^7 E/D$
cube	$N \leq E/D$

$\alpha = 30°$, $y = 2\mathrm{mm}$, $\lambda = 1\mu\mathrm{m}$

5.4. Discussion

Power considerations limit the cube more than the other designs, because of the loss due to redundant beams that are created by the holograms. The quantum limit on the power threshold needed by a detector operating at 10 GHz with an error rate of 10^{-9} is about 21 photons, which corresponds to $D = 2.1 \times 10^{-8}$ Watt on average for $\lambda = 1\mu$m [14]. As semiconductor lasers with an output power of $E =\sim$ 100mW are feasible, this would indicate that power requirements only become a limiting factor when more than a million processors are involved.

Practical detectors, however, are much less efficient than the theoretical bound predicted by the quantum limit. The minimum number of photons required per bit is about an order of magnitude more than that implied by the quantum limit. Using the same numbers as above, this means that no more than 500,000 PEs may be connected by the cube design. If a fan-out capability is desired, i.e. if multicast interconnections are also supported, this number must be divided by the maximal fan-out.

The above limit is a result of losses in the holographic routing elements. The other designs, which do not use holographic routing, are therefore not subject to such limits. Ideally, these designs would not be limited at all by power considerations (provided that $E > D$). However, some losses do occur due to the limited size of the detectors, thus limiting the number of processors after all. Table 3 contains a numerical comparison. If we substitute the assumption of uniform illumination by an assumption of a Gaussian beam that is centered on the detector, the constants in the array and hemisphere designs become even larger.

6. CONCLUSIONS

Numerous suggestions for optical interconnection networks and switches have been made in recent years. Most of them present a small working prototype, but do not include a study of the limitations and scalability of the design. As optical technology is typically introduced to cope with large parallel system, the lack of such a study undermines the value of these proposals. General studies like the one reported here are not only important for the evaluation of a certain design, but also give insight into the overall capablities and limitations of optical interconnections. They may also be relevant to other models, e.g. to optical interconnections for VLSI and to optical neural networks.

Our results are of course not exact, as many approximations were made in the derivations. However, four main conclusions may be drawn:

1. All the different schemes have the same asymptotic behavior. The details depend strongly on the choice of various parameters, such as the surface area and volume of each PE.
2. The geometrical considerations involved in organizing the processors and routing elements in 3-D space are the main factor that limits the number of processors that can be accommodated. In reasonable volumes, and under a certain set of parameter values, this number is less than 100,000.
3. Most of the volume of a parallel computer using optical free-space interconnections will be empty space traversed by light beams, and only a small fraction of the volume will be used for the PEs themselves.
4. Schemes using independent paths that only converge on the detector are more energy efficient than schemes that combine paths using holograms or semi-transparent mirrors, because of the N-fold loss due to the fan-in [12].

When a specific network is evaluated, general comments like the above are insufficient; a detailed study should be made. Specifically, the constants in the various expressions should be evaluated, and manufacturability and alignment considerations must be taken into account. The cube design, for example, is found to perform well in terms of volume use, but the number of processors it can accommodate is somewhat limited by power considerations. This limit is offset by the fact that it does not contain any global elements that have critical placement requirements, and by the fact that the path lengths are shorter and display some degree of locality.

7. ACKNOWLEDGEMENT

The authors would like to thank Joseph Goodman for his help in assessing the validity of these results.

8. REFERENCES

[1] L. Rudolph and E. Schenfeld: *DynaSwitch.* Patent application, 1989.

[2] E. Marom and N. Konforti: "Programmable optical interconnects". In *Proc. Intl. Optical Computing Conf.,* SPIE vol. **700** pp. 209–213, July 1986

[3] J. B. McManus, R. S. Putnam, and H. J. Caulfield: "Switched holograms for reconfigurable optical interconnections: demonstration of a prototype device". *Appl. Opt.* **27(20)** 4244–4250, Oct 1988

[4] D. Psaltis, D. Brady, and K. Wagner: "Adaptive optical networks using photorefractive crystals". *Appl. Opt.* **27(9)** 1752–1759, May 1988

[5] J. W. Goodman, F. J. Leonberger, S-Y. Kung, and R. A. Athale: "Optical interconnections for VLSI systems". *Proc. IEEE* **72(7)** 850–865, July 1984

[6] A. A. Sawchuk, B. J. Jenkins, C. S. Raghavendra, and A. Varma: "Optical crossbar networks". *Computer* **20(6)** 50–60, June 1987

[7] M. S. Goodman: "Multiwavelength networks and new approaches to packet switching". *IEEE Comm. Mag.* **27(10)** 27–35, Oct 1989

[8] D. O. Harris and A. VanderLugt: "Acousto-optic photonic switch". *Opt. Lett.* **14(21)** 1177–1179, Nov 1989

[9] A. Hartmann and S. Redfield: "Design sketches for optical crossbar switches intended for large-scale parallel processing applications". *Opt. Eng.* **28(4)** 315–327, Apr 1989

[10] L. Rudolph, D. G. Feitelson, and E. Schenfeld: "An optical interconnection network with 3-D layout and distributed control". This proceedings.

[11] A. Yariv: *Quantum Electronics.* Wiley, 2nd ed., 1975

[12] J. W. Goodman: "Fan-in and fan-out with optical interconnections". *Opt. Acta* **32(12)** 1489–1496, Dec 1985

[13] A. W. Lohmann and F. Sauer: "Holographic telescope arrays". *Appl. Opt.* **27(14)** 3003–3007, July 1988

[14] S. D. Personick: "Fundamental limits in optical communication". *Proc. IEEE* **69(2)** 262–266, Feb 1981