*Preprint of a paper to appear in Applied Optics, Vol. 33, No. 11, November 1994*

# A Scalable Optical Hypercube-based Interconnection Network for Massively Parallel Computing

**Ahmed Louri and Hongki Sung**
**Department of Electrical and Computer Engineering**
**The University of Arizona**
**Tucson, AZ 85721**
**Phone: (602) 621 – 2318, Fax: (602) 621 – 8076**

**E-mail: louri@ece.arizona.edu**

## Abstract

Two important parameters of a network for massively parallel computers are scalability and modularity. Scalability has two aspects; size and time (or generation). Size-scalability refers to the property that the size of the network can be increased with nominal effect on the existing configuration. Also, the increase in size is expected to result in a linear increase in performance. Time-scalability implies that the communication capabilities of a network should be large enough to support the evolution of processing elements through generations. A modular network enables the construction of a large network out of many smaller ones. The lack of these two important parameters has limited the use of certain types of interconnection networks in the area of massively parallel computers. This paper presents a new modular optical interconnection network, called an Optical Multi-Mesh Hypercube (OMMH), which is both size- and time-scalable. The OMMH combines positive features of both the hypercube (small diameter, high connectivity, symmetry, simple routing, and fault tolerance) and the torus (constant node degree and size-scalability) networks. Also presented is a three-dimensional optical implementation of the OMMH network. A basic building block of the OMMH network is a hypercube module which is constructed with free-space optics to provide compact and high-density localized hypercube connections. The OMMH network is then constructed by putting together such basic building blocks with multi-wavelength optical fibers which realize torus connections. The proposed implementation methodology is intended to fully exploit the advantages of both space-invariant free-space and multiwavelength fiber-based optical interconnects technologies. The analysis of the proposed implementation shows that such a network is optically feasible in terms of the physical size and the optical power budget.

*Key words:* optical interconnects, scalability, parallel processing, space-invariant interconnects, multi-wavelength multiplexing.

# Contents


**1 Introduction** **3**

**2 Topology of Optical Multi-Mesh Hypercube Network** **5**

2.1 Definition of OMMH Network . . . . . 5

2.2 Message Routing in OMMH . . . . . 6

2.3 Diameter and Link Complexity . . . . . 6

2.4 Fault Tolerance of the OMMH Network . . . . . 7

**3 Modular and Scalable Optical Interconnection Architecture of OMMH Network and Its Optical Implementation** **8**

3.1 Scalable Interconnection Architecture of OMMH Network . . . . . 8

3.1.1 Size-scalability of the OMMH Network . . . . . 8

3.1.2 Generation-scalability of the OMMH Network . . . . . 9

3.2 Optical Implementation of OMMH Network . . . . . 9

3.2.1 Optical Implementation of Space-invariant Hypercubes Using Binary Phase Gratings 9

3.2.2 Design of Torus Links to Connect Hypercube Modules . . . . . 12

3.2.3 Optical Hardware Required for Torus Links . . . . . 14

3.2.4 Rationale for the Two-level Design Approach . . . . . 15

3.2.5 Evaluation of the OMMH Implementation . . . . . 16

**4 Conclusions** **17**

# 1 Introduction

The quest for Teraflops ($10^{12}$ floating point operations per second) supercomputers combined with the launching of the High Performance Computing and Communication (HPCC) initiative is putting major emphasis on exploiting massive parallelism with greater than one thousand processing elements networked to form massively parallel computers (Ultracomputers)[1, 2]. A key element, and deciding factor in terms of performance and cost of these computers is the interconnection network[3]. The interconnection network for massively parallel computers must not only be adequate in terms of communication bandwidth, latency, and connectivity but it must also be modular and scalable.

Scalability of a network consists of two aspects; size-scalability and generation-scalability (or time-scalability)[2]. Size-scalability refers to the property that the size of the network (e.g., the number of communicating nodes) can be increased with nominal change in the existing configuration. Also, the increase in system size is expected to result in an increase in performance comparable to the increasing size. A generation-scalable network can be implemented in a new technology, and interconnection bandwidth of the network should grow at the same rate as processing speed and memory. Without increasing interconnection bandwidth, we cannot fully exploit the increased speed of evolutionary processing elements. Modular networks enable the construction of a large network out of smaller networks.

Numerous topologies have been explored for parallel computers[4, 5, 6, 7]. However, the lack of size-scalability and modularity of some of these networks have limited their use in massively parallel computing systems despite their many other advantages. For example, one of the most popular network for parallel computers is the binary $n$-cube, also called as a hypercube. The attractiveness of the hypercube topology is its small diameter, which is the maximum number of links (or hops) a message has to travel to reach its final destination between any two nodes. A binary $n$-cube network has $2^n$ nodes and the diameter is $n$. Each node is numbered in such a way that there is only one binary digit difference between any node and its $n$ neighbors (node degree) that are directly connected to it. This property greatly facilitates the routing of messages through the network. In addition, the regular and symmetric nature of the network provides fault-tolerance. Despite its small diameter, high connectivity, simple routing scheme, and fault tolerance, the hypercube topology is rarely adopted in the most recent projects for massively parallel computers such as Intel Paragon, Cray Research MPP Model, Caltech Mosaic C, MasPar MP-1, Tera Computer Tera Multiprocessor, Stanford Dash Multiprocessor which

are based on the torus/mesh topology[1, 8]. One major reason is its lack of size-scalability. As the dimension of the hypercube is increased by one, one additional link needs to be added to every node in the network. In addition to the changes in the node configuration, at least a doubling of the number of existing nodes is required for the regular hypercube network to expand and to remain as a hypercube.

Torus networks (henceforth, the mesh is referred to as a *torus* if the mesh has wraparound connections in the rows and columns) are easily implemented because of the simple regular connection and small number of links (four) per node. Due to its constant node degree, the torus network is highly size-scalable. With a network size of $N$ nodes, the minimal incremental size is approximately $N^{1/2}$ for a perfectly balanced network. However, the torus network also suffers from a major limitation which is its large diameter ($N^{1/2}$ for an $N$-node network) along with its limited connectivity. Despite the fact that the mesh/torus topology has limited connectivity and a large diameter, many recent projects for massively parallel computers targeting Teraflops use this topology for the interconnection network.

Motivated by these limitations, we have explored a novel topology for optical interconnection networks, called *Optical Multi-Mesh Hypercube* (OMMH), which combines the advantages of both the hypercube (small diameter, high connectivity, symmetry, simple control and routing, fault tolerance, etc.) and the mesh (constant node degree and size-scalability) topologies, while circumventing their disadvantages (lack of size-scalability of the hypercube, and large diameter of the mesh/torus). The topology of the OMMH network is size-scalable. Time-scalability is provided by the optics-based interconnection architecture. We have developed a three-dimensional (3-D) optical design methodology which exploits the advantage of both space-invariant free-space and multiwavelength fiber-based optical interconnects technologies. The proposed implementation is also analyzed for an example OMMH network with ten-cube as a basic building block. The analysis includes examination of power flow, efficiency, and system volume. It is shown that a ten-cube (1024 nodes) module is containable within a $25.4mm \times 25.4mm \times 203.2mm$ volume with power efficiency of 16% and the torus subnetwork has power efficiency of 23.6%.

The distinctive advantages of the proposed design methodology include: (1) an efficient and scalable interconnection network; (2) better utilization of the space-bandwidth product (SBWP) of optical imaging systems; (3) full exploitation of the parallelism of free-space optics and high bandwidth of fiber-optics; and (4) compatibility with the emerging two-dimensional (2-D) optical logic and switching, and opto-electronic integrated circuit technologies.

## 2 Topology of Optical Multi-Mesh Hypercube Network

In this section, we define the structure of the OMMH. We then compare and contrast structural properties of the OMMH with the hypercube network.

### 2.1 Definition of OMMH Network

An OMMH is characterized by a triplet $(l, m, n)$, where $l$ represents the row dimension of a torus, $m$ the column dimension of the torus, and $n$ the dimension of a binary hypercube.

An $(l, m, n)$-OMMH network is constructed as follows. For two nodes $(i_1, j_1, k_1)$ and $(i_2, j_2, k_2)$, where $0 \leq i_1 < l$, $0 \leq i_2 < l$, $0 \leq j_1 < m$, $0 \leq j_2 < m$, $0 \leq k_1 < 2^n$, and $0 \leq k_2 < 2^n$:

1. There is a link (called a torus link) between two nodes if (1) $k_1 = k_2$ and (2) two components, $i$ and $j$, differ by one in one component while the other component is identical.

2. There also exists a torus link for the wraparound connection in the row if (1) $k_1 = k_2$ and (2) $i_1 = i_2$, $j_1 = 0$, $j_2 = m - 1$, or for the wraparound connection in the column if (1) $k_1 = k_2$ and (2) $j_1 = j_2$, $i_1 = 0$, $i_2 = l - 1$.

3. There is also a link (called a hypercube link) between two nodes if and only if (1) $i_1 = i_2$, and (2) $j_1 = j_2$, and (3) $k_1$ and $k_2$ differ by one bit position in their binary representation (Hamming distance of one).

Fig. 1 shows a $(4, 4, 3)$-OMMH interconnection where solid lines represent hypercube links and dashed lines represent torus links. A $(4, 4, 3)$-OMMH consists of $4 \times 4 \times 2^3 = 128$ nodes. Small black circles represent nodes of the OMMH network which are, in this paper, abstractions of PEs which consist of electronic processing modules for computation and optical sources/detectors for communication. Both ends of torus links, dashed lines, are connected for wraparound connections of the torus if they have the same labels. The size of the OMMH can grow without altering the number of links per node by expanding the size of the torus; for example, by inserting 3-cubes in the row or column of the torus in Fig. 1. This feature allows the OMMH to be size scalable. More discussion on the scalability issue will follow in subsection 3.1.

An interesting isomorphic network is shown in Fig. 2. The same network is redrawn as a $4 \times 4$ torus-clustered 3-cube. It can be viewed as 8 concurrent toruses where 8 nodes having identical torus addresses

form one 3-cube. It can also be viewed as 16 concurrent 3-cubes in which 16 nodes having identical hypercube addresses form a $4 \times 4$ torus. The $(4, 4, 3)$-OMMH in Fig.1 looks like a 3-cube-clustered $4 \times 4$ torus. Depending on the problem at hand, the OMMH can be configured as torus-clustered hypercubes or hypercube-clustered toruses.

### 2.2 Message Routing in OMMH

Due to the regularity of the structure, a distributed routing scheme can be implemented without global information. Since the OMMH is a point-to-point network, packet communication is assumed in the message routing scheme. For an $(l, m, n)$-OMMH network, let the addresses of two arbitrary nodes $S$ and $T$ be $(i_s, j_s, k_s)$ and $(i_t, j_t, k_t)$, respectively, where $0 \leq i_s < l$, $0 \leq i_t < l$, $0 \leq j_s < m$, $0 \leq j_t < m$, $0 \leq k_s < 2^n$, and $0 \leq k_t < 2^n$. The message routing scheme from $S$ to $T$ is that of an $n$-cube network or that of an $l \times m$ torus network or a combination of the two depending upon the relative locations of the nodes.

1. *Routing within a hypercube:* if $i_s = i_t$ and $j_s = j_t$, then $S$ and $T$ are within the same hypercube. The routing scheme for this case is exactly the same as that of the $n$-cube network.

2. *Routing within a torus:* if $k_s = k_t$, then $S$ and $T$ are within the same torus. The routing scheme for this case is exactly the same as that of the $l \times m$ torus network [9].

3. *Routing through toruses and hypercubes:* if none of the above two cases are true, $S$ and $T$ share neither a hypercube nor a torus. There are several options available for this case. One option uses the hypercube routing scheme until the message arrives at the same torus where $T$ resides, and then uses the torus routing scheme for the message to arrive at $T$. In another option, the torus routing scheme can first be applied to forward the message to the same hypercube where $T$ resides, and then the message can reach $T$ using the hypercube routing scheme. We can also mix the hypercube and the torus routing until the message is forwarded to the same hypercube or to the same torus where $T$ resides, and then we can forward the message to $T$ using the hypercube or the torus routing scheme, respectively.

### 2.3 Diameter and Link Complexity

The distance between two nodes in a network is defined as the number of links connecting these two nodes. The diameter of a network is defined as the maximum of all the shortest distances between

any two nodes. The diameter of the network is of great importance since it determines the maximum number of hops that a message may have to take. For two extreme cases, the diameter of a linear array with $N$ nodes is $(N-1)$ while that of a completely connected network is unity. An $l \times m$ torus has diameter $(\lfloor l/2 \rfloor + \lfloor m/2 \rfloor)$. The diameter of a hypercube with $N$ nodes is $\log_2 N$. Thus, the diameter of $(l, m, n)$-OMMH is $(\lfloor l/2 \rfloor + \lfloor m/2 \rfloor + n)$.

Link complexity or node degree is defined as the number of links per node. The higher the link complexity, the greater is the hardware complexity and, consequently, the cost of the network. The node degree of a hypercube with $N$ nodes is $\log_2 N$ and that of $(l, m, n)$-OMMH is $(n+4)$. $N$ is equal to $(l \times m \times 2^n)$ if the hypercube and the OMMH have the same network size. A comparison of diameters should be accompanied by a comparison of link complexity, because a higher connectivity resulting from a higher link complexity is expected to lead to smaller diameters. Fig. 3(a) compares the diameters of the hypercube and the OMMH, where $(16, 16, n)$-OMMH means the size of the torus in the OMMH is fixed and the size of the hypercube in the OMMH is changed to have the same network size for comparison purposes. Similarly, $(l, m, 4)$-OMMH implies the size of the hypercube in the OMMH is fixed and that of the torus is changed. Fig. 3(b) compares link complexities or node degrees of the hypercube and that of the OMMH. Fig. 3(c) depicts the growth of the total number of links in the network as the network size increases. For a network size of one million nodes, the hypercube network contains about 10.5 million links while the $(l, m, 4)$-OMMH has about 4.2 million links and $(16, 16, n)$-OMMH has approximately 8.4 million links. Since one link implies one physical path, electrical or optical, between two nodes, the OMMH network is cost-efficient compared to the hypercube network in terms of hardware requirement.

### 2.4 Fault Tolerance of the OMMH Network

As the number of components in a system grows, the probability of the existence of faulty components increases. For a large-scale system, we cannot always expect that all components in such a system are free from any failures. However, we need to expect such a system to continue to operate correctly in the presence of a reasonable number of failures.

Due to the concurrent presence of toruses and hypercubes in the OMMH network, rerouting of messages in the presence of a single faulty link or a single faulty node can easily be done with little modification of existing fault-free routing algorithms. In the OMMH network, any single faulty link or any single faulty node can be bypassed by only two additional hops as long as that particular node

is not involved in the communication, namely, the node is neither the source nor the destination for any message. This can briefly be proven as follows. When torus routing scheme is being applied in the presence of a faulty link or a node, one additional hop is needed to forward the message to a neighboring torus subnetwork through a hypercube link ($n$ such neighboring toruses exist in $(l, m, n)$-OMMH), and another hop is needed to return the message to the orignal torus subnetwork. Similarly, when hypercube routing scheme is being applied, a message can detour a faulty link or node with two additional hops; one to forward the message to a neighboring hypercube subnetwork and another to return the message to the original hypercube subnetwork.

# 3 Modular and Scalable Optical Interconnection Architecture of OMMH Network and Its Optical Implementation

In this section, we first discuss modularity and scalability issues of the OMMH interconnection architecture, and present an optical implementation of the OMMH network. We then discuss the rationale and performance of the proposed OMMH implementation.

## 3.1 Scalable Interconnection Architecture of OMMH Network

### 3.1.1 Size-scalability of the OMMH Network

Size-scalable networks have the property that the size of the system (e.g., the number of communicating nodes) can be increased with nominal changes in the existing configuration. Also, the increase in system size is expected to result in an increase in performance to the extent of the increase in size. As the dimension of the hypercube is increased by one, one more link needs to be added to every node in the network. In addition to the changes in the node configuration, at least a doubling of the size is required for the hypercube network to expand. This implies that the hypercube does not allow an incremental expansion of small sizes. Thus, the hypercube network is not scalable according to the above definition. We should note that the hypercube network may be scalable at a greater cost. Moreover, it is not modular[6]. The lack of size-scalability and modularity have limited the application of the hypercube topology to large-scale massively parallel systems.

As can be seen in Fig. 3(b), the OMMH with a constant cube as a basic building block (e.g., $(l, m, 4)$-OMMH) has a constant node degree, which means that the size of the OMMH is ready to be scaled up by expanding the size of the torus without affecting the link complexity of existing nodes as is the case

in expanding the size of the hypercube network. However, we cannot just add one node to the OMMH. For an $(l, m, n)$-OMMH, we need to add at least $l \times 2^n$ nodes (if $l < m$) when the torus subnetwork needs to be balanced.

An OMMH network is constructed from simple building blocks (hypercubes) in a modular and incremental fashion. These building blocks, once constructed, are left undisturbed when the network grows in size. The OMMH can be viewed as a two-level interconnection network: high-density, local connections for hypercube links (within a basic module), and high bit rate, low-density and long connections for the torus links connecting the basic building blocks. Thus, the size of the OMMH can be increased by adding hypercube modules, which provides modularity and size-scalability.

#### 3.1.2 Generation-scalability of the OMMH Network

Generation-scalable architecture are designed with consideration of what future implementation might be like. Such architectures will survive throughout generations. A generation scalable network can be implemented in a new technology, and interconnection bandwidth of the network should grow at the same rate as processing speed and memory. Without increasing interconnection bandwidth, we cannot fully exploit the increased speed of evolutionary processing elements. Generation-scalability in the OMMH interconnection architecture is provided by the use of high-bandwidth optics which would match communications bandwidth requirements of future processing elements.

### 3.2 Optical Implementation of OMMH Network

In Subsection, we first present an optical implementation of the hypercube networks for constructing basic building blocks, and then show how to design torus links to connect such hypercube modules to construct the OMMH network.

#### 3.2.1 Optical Implementation of Space-invariant Hypercubes Using Binary Phase Gratings

We discuss an optical implementation of the 3-D space-invariant hypercube network of which design methodology is proposed in Ref. [10]. The design methodology is based on an observation that nodes in an interconnection network can be partitioned into two sets of nodes such that any two nodes in a set do not have a direct link (except for completely connected networks). This is a well-known problem

of bipartitioning a graph if the interconnection network is represented as a graph. For a binary $n$-cube, nodes whose addresses differ by more than one in Hamming distance can be in the same partition, since no link exists between two nodes if their Hamming distance is greater than one. Besides bipartitioning the graph, we arrange the nodes in each partition onto the plane such that interconnection between two planes becomes space-invariant. Two partitions of nodes are called $Plane_L$ and $Plane_R$, respectively, in Ref. [10]. The methodology uses free-space optical multiple imaging technique, to replicate and spatially shift the image of one partition of nodes. Multiple images are then simultaneously incident on the other partition. The locations of multiple image spots on the receiving partition are determined by the required connection patterns[10].

There is a wide variety of optical means to generate multiple images. These include phase gratings[11, 12], beam splitters[13], multi-split lenses[14], lenslet arrays[15], arrays of mirrors[16], and holographic techniques[17]. A hypercube-based architecture for cellular image processing has been demonstrated with holograms[18] and a new concept of grid patterns for the layout of optoelectronic integrated circuit chip has been proposed[19].

In the following, we discuss the design of binary phase gratings (BPGs) for the five-cube implementation as an example. Fig. 4 describes a hardware arrangement of optical components for a space-invariant five-cube network. For clarity, only a 2-D view is shown. A BPG is added at the pupil plane between two imaging lenses to provide necessary beam steering operations. This type of arrangement was first proposed in Ref. [20]. We extended it here for the implementation of space-invariant hypercube networks. Since the interconnection patterns are space-invariant, any beam-steering operation performed on one of the beams must be performed on all of the beams that pass through the BPG. The beam-steering operation of the BPG is dictated by the grating equation shown in Eq. 1 which describes the relationship between the angle of the incident beam($\theta_i$), the period of the grating($p$), the wavelength of the light($\lambda$), the grating order($m$), and the angle of the $m$-th order's diffracted beam($\theta_{d_m}$).

$$p(\sin\theta_{d_m} - \sin\theta_i) = m\lambda \tag{1}$$

Assume that the size of a node in one-dimension be $d$ and the focal length of each lens be $f$. Let $h^m_{PE0}$ be the distance of an image spot in $Plane_R$ from the optical axis made by $m$-th order diffracted beam from PE0. Then,

$$h^m_{PE0} = f \times \tan\theta_{d_m} \tag{2}$$

Given that $\theta_i = \tan^{-1}(1.5d/f)$, Eq. 2 can be rewritten as:

$$h^m_{PE0} = f \times \tan\{\sin^{-1}[\frac{m\lambda}{p} + \sin(\tan^{-1}(\frac{1.5d}{f}))]\} \tag{3}$$

We assume that the structure of the grating is designed such that the power of the incident beam is equally distributed into 0th-order, $\pm$1st order, and $\pm$3rd order of diffracted beams and others are suppressed. We can have different amounts of optical power from the original beam to be routed into the different orders by changing the periodic structure of the grating. To have different angular spacings, we should change the period of the grating[17]. Since PE0 is supposed to be connected with PE1, PE4, and PE16 for the five-cube network, the following conditions should be satisfied.

$$\begin{aligned} h^0_{PE0} &= 1.5 \times d \\ h^{+1}_{PE0} &= 0.5 \times d \\ h^{+3}_{PE0} &= -1.5 \times d \\ h^{-1}_{PE0} &> 2.0 \times d \\ h^{-3}_{PE0} &> 2.0 \times d \end{aligned} \tag{4}$$

Note that conditions for $h^{-1}_{PE0}$ and $h^{-3}_{PE0}$ make $-$1st and $-$3rd order diffracted beams fall outside $Plane_R$ to avoid unwanted connections.

Similarly, the beam from PE5 generates multiple spots in $Plane_R$ for which the distances from the optical axis are:

$$h^m_{PE5} = f \times \tan\{\sin^{-1}[\frac{m\lambda}{p} + \sin(\tan^{-1}(\frac{0.5d}{f}))]\} \tag{5}$$

To make connections from PE5 to PE1, PE4, and PE21, the following set of conditions should hold:

$$\begin{aligned} h^0_{PE5} &= 0.5 \times d \\ h^{+1}_{PE5} &= -0.5 \times d \\ h^{-1}_{PE5} &= 1.5 \times d \\ h^{+3}_{PE5} &< -2.0 \times d \\ h^{-3}_{PE5} &> 2.0 \times d \end{aligned} \tag{6}$$

Note that conditions for $h^{+3}_{PE5}$ and $h^{-3}_{PE5}$ make +3rd and $-$3rd diffracted beams fall outside $Plane_R$. Since PE0 and PE5 are symmetrically placed, with respect to the optical axis, with PE17 and PE20,

we can determine the period of grating($p$) to provide the required connections for the 5-cube network by solving Eqs. 4 and 6 given the size of a node, the focal length of the lens, and the wavelength of the light source. However, we cannot have an exact solution since image spots generated by both PE0 and PE5 cannot be placed on uniform spacings in $Plane_R$. An approximate solution could be determined by a computer program which optimizes conditions in Eqs. 4 and 6. By optimization we mean minimization of errors in each condition. For example, given that the node size in one dimension is $5mm$, the wavelength of the light source $960nm$, and the focal length of the lens $50.8mm$, the optimum period of the grating is computed to be $19.6\mu m$ which causes maximum misalignment of $9.0\mu m$ at PE21 from the PE5 connection. The feature size of a node for the construction of massively parallel computers will mainly be dependent on the area required for processing elements and memories for a single node, not on the size of the light source/detector. With wafer scale integration, the size of a node would be small enough (assuming fine grain or medium grain parallelism) so as not to make imaging lenses impractically large[21].

The size of a basic $n$-cube module that can be implemented is primarily determined by the number of fanouts that can be managed by the BPG since an $n$-cube implementation requires $2n-1$ fanouts. The BPG must be able to generate $2n-1$ beams of equal power. We should note that a hypercube of a relatively small size could be easily implemented with electronics if the bandwidth requirement is not large, but as the network size grows, optics is more advantageous than electronics both in design complexity and bandwidth

### 3.2.2 Design of Torus Links to Connect Hypercube Modules

An $(l, m, n)$-OMMH can be constructed as follows:

**(1)** $l \times m$ $n$-cube modules as described in Sec. 3.2.1 are placed in a $l \times m$ matrix form.

**(2)** $l \times m$ nodes, each of which is from the same location of the $n$-cube modules, are connected to form a torus of dimension $l \times m$ (A node consists of an $n$-cube module with torus link interface. An example of a five-cube module with torus interface is shown in Figs. 6 and 7.).

**(3)** Step 2 is repeated until every node is connected, resulting in $2^n$ toruses of size $l \times m$.

Since two adjacent $n$-cube modules are connected by $2^n$ torus links, the number of optical fibers required grows exponentially as $n$ increases. A possible solution for reducing the number of optical

fibers required is the use of a wavelength multiplexing (WM) technique. However, a straightforward use of the WM also requires a prohibitively large number of different wavelengths. For example, to connect two ten-cube modules, we need $2^{10} = 1024$ different wavelengths. A wavelength-node assignment technique[22] can alleviate this problem as follows.

Referring to Ref. [10], an $n$-cube layout ($Plane_L$ or $Plane_R$) consists of $2^{\lfloor (n-1)/2 \rfloor}$ non-empty rows and $2^{\lceil (n-1)/2 \rceil}$ non-empty columns. For $Plane_L$ and $Plane_R$, we assign the following wavelengths to the nodes in the first row; $\lambda_1$, $\lambda_2$, $\cdots$, $\lambda_{2^{\lceil (n-1)/2 \rceil}}$. Then, we assign $\lambda_2$, $\cdots$, $\lambda_{2^{\lceil (n-1)/2 \rceil}}$, $\lambda_1$, as wavelengths to the nodes in the second row. In general, wavelength-assignment in a row is achieved by rotating the wavelength-assignment of the previous row by one column. This wavelength assignment results in no two nodes in the same row or column having an identical wavelength. Figure 5 shows a wavelength-assignment for a 5-cube module. We then use a $2^{\lceil (n-1)/2 \rceil}$-channel wavelength multiplexed fiber to connect two rows in the adjacent two $n$-cube modules. Similarly, a $2^{\lfloor (n-1)/2 \rfloor}$-channel fiber is used to connect two columns in the adjacent two $n$-cube modules. Thus, an implementation of $(l, m, n)$-OMMH using the above wavelength assignment method requires no more than $2^{\lceil (n-1)/2 \rceil}$ different wavelengths. In addition, no more than $2^{\lceil (n-1)/2 \rceil}$ optical fibers will be required for the connections between any two adjacent $n$-cube modules.

Now, we consider an optical implementation of the $(l, m, n)$-OMMH network. We assume the availability of two optical components: a quadrant beam splitter (QBS) which splits a single beam into four beams and an $i$-channel wavelength multiplexor (WMUX) which multiplexes beams with $i$ different wavelengths into a single beam (also demultiplexes since it is bi-directional). The realization of these two components with current technology will be discussed in detail in the following subsection. Figues 6 and 7 shows an example of a five-cube basic module construction including the torus link interface. We assume that each node has two light sources; one source, $S_h$, illuminates the BPG to generate the required hypercube links and the second source, $S_t$, is coupled with an optical fiber for the torus links. The number of detectors per node is dependent on the incoming channel distinction technique[23]. If we use the space division multiplexing technique, we need four detectors per node for the torus connection, and $2^n$ detectors per node for the $n$-cube connection (the number of detectors in the $n$-cube module can be reduced greatly if a channel encoding technique is used[23, 24]). A QBS is attached to every $S_t$ to provide the four fanouts, $S_{t_N}$, $S_{t_S}$, $S_{t_E}$, and $S_{t_W}$ (north, south, east, and west). A WMUX is located at both ends of each row and each column. Let each WMUX at the right end of a row be $\text{WMUX}_E$, each

WMUX at the left end of a row be $\text{WMUX}_W$, each WMUX at the top of a column be $\text{WMUX}_N$, and each WMUX at the bottom of a column be $\text{WMUX}_S$. In a given row, a $\text{WMUX}_E$ multiplexes light from the $S_{t_E}$ sources of that row into a single fiber which is then connected to a $\text{WMUX}_W$ in the neighboring $n$-cube module. Similarly, $S_{t_N}$s $S_{t_S}$s, and $S_{t_W}$s sources are multiplexed by $\text{WMUX}_N$, $\text{WMUX}_S$, and $\text{WMUX}_W$, respectively. In the receiving module, those signals are demultiplexed by the WMUX and routed to the corresponding nodes. Figure 6 illustrates a side view of a five-cube module with torus link interface. For clarity, only the 2-D view is shown and, thus, only two fanouts by a QBS is given. Figure 7 show the corresponding top view of the module.

### 3.2.3 Optical Hardware Required for Torus Links

In this subsection, we discuss the functionality and limits of two optical components used in the implementation of torus links.

*Quadrant Beam Splitter(QBS):*

The function of the QBS is to split one beam into four beams. An optical arrangement of the QBS using graded index (GRIN) lenses[25] is illustrated in Fig. 8.a. Four small GRIN lenses are placed on the end facet of the large GRIN lens. The large lens is used to collimate a beam from a single trunk fiber and the aperture of the collimated beam is divided into four by the smaller lenses. The small lenses then focus the beams onto fibers. Beam combination or merging is performed but in the opposite direction. Figure 8.b illustrates the geometry of the QBS with GRIN lenses for the purpose of calculating power loss occurring at the connection between the large GRIN lens and small GRIN lenses. Since four small GRIN lenses do not cover the entire end-facet area of the large GRIN lens, some portion of beam aperture from the large GRIN lens cannot be captured by four small GRIN lenses, resulting in power loss. Suppose that a radius of a small lens is $r$. The smallest possible radius of the large lens that can cover four small lenses is then $r + \sqrt{2}r$. Thus, $4\pi r^2/[\pi(1+\sqrt{2})^2 r^2] = 68.6\%$ of the end-facet area of the large GRIN lens is covered by the four small lenses. Therefore, approximately 31.4% of power is lost from the large GRIN lens to the four small lenses during the beam splitting process.

A more power efficient (less than 20% power loss) QBS has been reported in Ref. [26] using substrate mode holograms to reduce mechanical alignment and chromatic sensitivity. The QBS design with substrate mode holograms is better than the design with GRIN lenses in terms of power efficiency,

feature size, alignment, and fiber coupling efficiency. However, substrate mode multiplexed holograms for the QBSs are not commercially available at this time.

*Wavelength Multiplexor(WMUX):*

A wavelength multiplexor/demultiplxor (WMUX) using a GRIN lens and a blazed grating is discussed in Ref. [27]. WMUXs of this type allow more of the total bandwidth of the optical fiber to be used and more than ten channels are currently available. Typical values of the insertion loss and the crosstalk in available WMUXs are generally $1 \sim 2$ dB and less than $-30$ dB, respectively. Since $(l, m, n)$-OMMH requires $2^{\lceil (n-1)/2 \rceil}$-channel WMUXs, with 8-channel WMUXs, it is possible to implement any size of OMMH networks if $n \leq 7$.

### 3.2.4 Rationale for the Two-level Design Approach

As discussed in the previous subsection, the optical implementation of the OMMH network consists of two levels: free-space space-invariant optics for the construction of basic building blocks, and multiwavelength fibers for the torus links. The rationale for the two-level design approach is as follows; the use of space-invariant free-space optics would result in compact and simple building blocks that can be easily reproduced[23, 28]. However, it would not be easy to implement scalable optical interconnects with totally space-invariant optics only, since a single space-invariant optical component such as a hologram is used to image multiple nodes for totally space-invariant interconnects. Thus, it would be necessary to redesign the component in order to increase the number of nodes. However, since the minimum incremental size of the OMMH is one hypercube module (a basic building block), the use of space-invariant optics within the basic building block will not limit the scalability of the OMMH. We use multiwavelength fiber optics to connect the basic building blocks because fiber optics would provide affordable scalable interconnects and the wavelength multiplexing technique would make a better utilization of the transmission capacity of an optical fiber[29, 30, 31, 32]. The breakdown of functional requirements for the OMMH network is consistent with the advantages of free-space and optical fiber technologies.

### 3.2.5 Evaluation of the OMMH Implementation

To demonstrate the viability of the OMMH implementation, an evaluation of the power efficiency and system volume was performed with the proposed design based on practical component sizes. The design consists of ten-cube modules using $25.4mm$ diameter lenses, each with a focal length of $50.8mm$. A detector diameter of $50\mu m$ was calculated to provide the connection density to produce a ten-cube while allowing for space division techniques to be incorporated to avoid signal overlap.

With this system design, the system volume of the hypercube module is $25.4mm \times 25.4mm \times 203.2mm$. The efficiency of the hypercube was calculated to be approximately 16%. This efficiency is due mostly to the phase hologram (which is theoretically 33% maximum) and to the unwanted fan-out beams that result from the space-invariant nature of the design. The unwanted fan-out contributes an efficiency ($\eta_h$) given by

$$\eta_h = \frac{n}{2n-1}, \tag{7}$$

where $n$ is the hypercube dimension. For the ten-cube, this is 53%.

As for the OMMH mesh subnetwork, a power analysis of the fiber optic system was performed. Figure 9 shows a single unidirectional link of the mesh. For this link, a $-1dB$ loss occurs from the insertion of the laser signal into the fiber. The QBS suffers a $-0.97dB$ loss, while each WMUX loses $-1dB$ of power (this calculation is for total system efficiency rather than per channel power flow so the QBS loss does not include fan-out). Furthermore, the fiber is assumed to be at most $1m$ in length and a mean operating wavelength of $960nm$ is assumed for the loss calculation. At this wavelength, the fiber has an attenuation of $-3.5dB/km$. Thus, the fiber loss for the system is $-0.0035dB$. Furthermore, the detector loss is $-1dB$. Connection losses in the system were calculated based upon reflection at the fiber interface. For the QBS, each connection suffers a $-0.45dB$ loss. As for the WMUXs, an index matching oil was assumed to be used between the fiber and the GRIN lens of the WMUX to ease the index transition. Using this setup, the connection losses around the WMUXs are $-0.1dB$ each. The final result is a total of $-6.27dB$ of system losses. This equates to a system efficiency of 23.6%.

To gain a sense of the bandwidth of the proposed implementation, a leading edge VCSEL source is capable of producing $6.5Gbps$ with an output power of $2.2mW$[33]. When used with this system, a detector must have a sensitivity of $0.13mW$ for the same bandwidth. This sensitivity can be reached with a detector having a bit error rate of $10^{-17}$[34].

# 4 Conclusions

Scalable networks and architectures are becoming more and more desirable for massively parallel computers since they can grow in size without major changes of the existing system configuration (size-scalability), as well as they are able to employ new evolving technologies (generation-scalability). In fact, scalable network topologies are becoming the preferred choice for the computer industry despite their inherently limited topological characteristics such as low connectivity, large diameters, long average distances, and lack of fault tolerance. For example, many recent projects for the development of ultracomputers (Intel Paragon, Cray Research MPP Model, Caltech Mosaic C, MasPar MP-1, Kendall Square Research KSR-1, Stanford Dash Multiprocessor, Tera Computer Tera Multiprocessor, and Thinking Machine Corporation CM-5, etc) are based on the scalable topologies such as the mesh/torus, ring or tree. Interconnection networks which are not only scalable, but also possess good topological characteristics such as small diameter, high connectivity, constant node degree, simple routing scheme, and fault tolerance, would greatly enhance the performance of massively parallel computers.

We presented in this paper a new interconnection network, called the Optical Multi-Mesh Hypercube (OMMH), for massively parallel computers. The distinctive features of the OMMH network are its scalability, both in size and generation, and modularity while retaining positive features of both the hypercube (high connectivity, small diameter, simple message routing, and fault tolerance) and the mesh (constant node degree and scalability) topologies. We also proposed an optical interconnection architecture of the OMMH and its three-dimensional implementation. The proposed implementation is divided into two levels; space-invariant free-space optical interconnects for localized high-density hypercube modules and high bandwidth multiwavelength optical fiber links for global low-density torus connections. This breakdown of functional requirements for the OMMH implementation is intended to fully exploit the advantages of free-space space-invariant optics (parallelism, simple and compact design, high connectivity, and cost efficiency) as well as wavelength multiplexed fiber-based optics (full utilization of transmission bandwidth and scalability). In addition, the breakdown is intended to provide modularity and scalability both in size and generation. The two-level design methodology enables the construction of the OMMH network in a modular and incremental fashion (size-scalability) and the use of high bandwidth wavelength multiplexed optics in the OMMH can satisfy communication bandwidth requirements of the current or near future processing elements (generation-scalability). We have also

analyzed the proposed implementation. The implementation demonstrates good feasibility by showing a reasonable optical power efficiency and a volume size capable for inclusion within the case of a massively parallel computer.

This research was supported by an NSF grant No. MIP 9310082 and a grant from USWest. The authors thank Dr. Matthew Derstine, the topical editor, and anonymous reviewers for their comments to improve the quality of this paper. They also express their gratitude to Mr. Michael Major, at the University of Arizona, for his help in analyzing performance of the proposed implementation.

## References


[1] K. Hwang, *Advanced Computer Architecture: Parallelism, Scalability, Programmability*. New York, NY: McGraw-Hill, 1993.

[2] G. Bell, "Ultracomputers: A Teraflop Before Its Time," *Communications of the ACM*, vol. 35, pp. 27–47, Aug 1992.

[3] H. S. Stone and J. Cocke, "Computer Architecture in the 1990s," *Computer*, pp. 30–38, Sep. 1991.

[4] L. N. Bhuyan and D. P. Agrawal, "Generalized Hypercube and Hyperbus Structures Constructing Massively Parallel Computers," *IEEE Trans. Comput.*, vol. C-33, pp. 323–333, 1984.

[5] N.-F. Tzeng and S. Wei, "Enhanced Hypercubes," *IEEE Trans. Comput.*, vol. 40, pp. 284–294, Mar 1991.

[6] J. R. Goodman and C. H. Sequin, "Hypertree: a Multiprocessor Interconnection Topology," *IEEE Trans. Comput.*, vol. C-30, pp. 923–933, 1981.

[7] C. L. Seitz, "The Cosmic Cube," *Communications of the ACM*, vol. 28, pp. 22–33, 1985.

[8] D. Lenoski, J. Laudon, K. Gharachorloo, W. D. Weber, A. Gupta, J. Hennessy, M. Horowitz, and M. Lam, "The Stanford Dash Multiprocessor," *IEEE Computer*, pp. 63–79, Mar 1992.

[9] D. Nassimi and S. Sahni, "An Optimal Routing Algorithm for Mesh Connected Parallel Computers," *Journal of the ACM*, vol. 27, pp. 6–29, January 1980.

[10] A. Louri and H. Sung, "A Design Methodology for three-dimensional space-invariant hypercube networks using graph bipartitioning," *Optics Letters*, vol. 18, pp. 2050–2052, Dec 1993.

[11] L. P. Boivin, "Multiple Imaging Using Various Types of Simple Phase Gratings," *Applied Optics*, vol. 11, pp. 1782–1792, 1972.

[12] F. B. McCormick, "Generation of Large Spot Arrays from a Single Laser Beam by Multiple Imaging with Binary Phase Gratings," *Optical Engineering*, vol. 28, pp. 299–304, Apr. 1989.

[13] K. H. Brenner and A. Huang, "Optical Implementation of the Perfect Shuffle Interconnection," *Applied Optics*, vol. 27, pp. 135–137, 1988.

[14] A. S. Kumar and R. M. Vasu, "Multiple Imaging and Multichannel Optical Processing with Split Lenses," *Applied Optics*, vol. 26, pp. 5345–5349, 1987.

[15] N. Streibl, U. Nolscher, J. Jahns, and S. Walker, "Array Generation with Lenslet Arrays," *Applied Optics*, vol. 30, pp. 2739–2742, 1991.

[16] Y. Sheng, "Space Invariant Multiple Imaging for Hypercube Interconnections," *Applied Optics*, vol. 29, pp. 1101–1105, 1990.

[17] K. S. Urquhart, S. H. Lee, C. C. Guest, M. R. Feldman, and H. Farhoosh, "Computer Aided Design of Computer Generated Holograms for Electron Beam Fabrication," *Applied Optics*, vol. 28, pp. 3387–3396, 1989.

[18] K.-S. Huang, A. A. Swachuk, B. K. Jenkins, P. Chavel, J.-M. Wang, and A. G. Weber, "Digital Optical Cellular Image Processor (DOCIP): Experimental Implementation," *Applied Optics*, vol. 32, pp. 166–173, Jan. 1993.

[19] M. W. Haney, "Self-similar Grid Patterns in Free-space Shuffle-exchange Networks," *Optics Letters*, vol. 18, pp. 2047–2049, Dec 1993.

[20] T. J. Cloonan and M. J. Herron, "Optical Implementation and Performance of One-dimensional and Two-dimensional Trimmed Inverse Augmented Data Manipulator Networks for Multiprocessor Computer Systems," *Optical Engineering*, vol. 28, pp. 305–314, 1989.

[21] C. M. Habiger and R. M. Lea, "Hybrid-WSI: a Massively Parallel Computing Technology?," *Computer*, vol. 26, pp. 50–61, Apr 1993.

[22] Y. Li, A. W. Lohmann, and S. B. Rao, "Free-space Optical Mesh-connected Bus Networks Using Wavelength-division Mutiple Access," *Applied Optics*, vol. 32, pp. 6425–6437, 1993.

[23] A. Louri and H. Sung, "Efficient Implementation Methodology for Three-dimensional Space-invariant Hypercube-based Free-space Optical Interconnection Networks," *Applied Optics*, vol. 32, pp. 7200–7209, Dec. 1993.

[24] A. Louri, H. Sung, Y. Moon, and B. P. Zeigler, "An Efficient Signal Distinction Scheme for Large-scale Free-space Optical Networks Using Genetic Algorithms," in *To appear in Photonics in Switching: Topical Meeting, OSA*, Salt Lake City, Utah, March 12–17 1995.

[25] J. P. E. Green, *Fiber Optic Networks*. Englewood Cliffs, NJ: Prentice Hall, 1993.

[26] M. Kato, Y.-T. Huang, and R. K. Kostuk, "Multiplexed Substrate-mode Holograms," *J. Opt. Soc. Am. A*, vol. 7, pp. 1441–1447, 1990.

[27] B. D. Metcalf and J. F. Providakes, "High-capacity Wavelength Demultiplexer with a Large-diameter GRIN Rod Lens," *Applied Optics*, vol. 21, pp. 794–796, 1982.

[28] G. E. Lohman and K. H. Brenner, "Space-invariance in Optical Computing Systems," *Optik*, vol. 89, pp. 123–134, 1992.

[29] R. J. Vetter and D. H. C. Du, "Distributed Computing with High-Speed Optical Networks," *IEEE Comput.*, pp. 8–18, Feb. 1993.

[30] M. G. Hluchyj and M. J. Karol, "ShuffleNet: An Application of Generalized Perfect Shuffles to Multihop Lightwave Networks," *IEEE Journal of Lightwave Technology*, vol. 9, pp. 1386–1397, Oct 1991.

[31] T. S. Wailes and D. G. Meyer, "Multiple Channel Architecture: A New Optical Interconnection Strategy for Massively Parallel Computers," *IEEE Journal of Lightwave Technology*, vol. 9, pp. 1702–1716, Dec 1991.

[32] G. R. Hill, "Wavelength Domain Optical Network Techniques," *Proceedings of the IEEE*, vol. 77, pp. 121–132, Jan 1989.

[33] G. Shtengel, H. Temkin, P. Brunsenbach, T. Uchida, M. Kim, C. Parsons, W. E. Quinn, and S. E. Swirhun, "High-Speed Vertical-Cavity Surface Emitting Laser," *IEEE Photonics Technology Letters*, vol. 5, pp. 1359–1362, 1993.

[34] R. D. Smith and S. D. Personick, "Receiver Design for Optical Fiber Communication Systems," in *Semiconductor Devices for Optical Communications* (H. Kressel, ed.), ch. 4, Springer-Verlag, Heidelberg, 1987.

## List of Figures


1 An example of the optical multi-mesh hypercube network: a $(4,4,3)$-OMMH (128 nodes) interconnection is shown. Two links with the same labels are connected for the wraparound connections of the torus. Only a few addresses are shown in the parenthesis for clarity. Solid lines represent hypercube connections and dashed lines torus connections. . . . . 23

2 A $(4,4,3)$-OMMH interconnection network, another isomorphic view. Wraparound connections of the torus are omitted and only a few addresses are shown in the parenthesis for clarity. Solid lines represent hypercube connections and dashed lines torus connections. 24

3 Comparison of (a) diameter, (b) link complexity, (c) total number of links, and (d) normalized average message distance of the hypercube and the OMMH when the two networks have the same number of nodes. . . . . 25

4 Space-invariant optical implementation of a 5-cube network using a binary phase grating. 26

5 Wavelength assignment for a 5-cube. Wavelengths are assigned such that no two nodes in the same row or column have an identical wavelength. . . . . 27

6 A 2-D side view of a 5-cube module to interface with torus links for the construction of the $(l,m,5)$-OMMH network. . . . . 28

7 A 2-D top view of a 5-cube module to interface with torus links for the construction of the $(l,m,5)$-OMMH network. . . . . 29

8 (a) A quadrant beam splitter using GRIN lenses. (b) Geometry of the quadrant beam splitter. . . . . 30

9 A model for the optical power budget calculation of the torus subnetwork in the OMMH network. . . . . 31

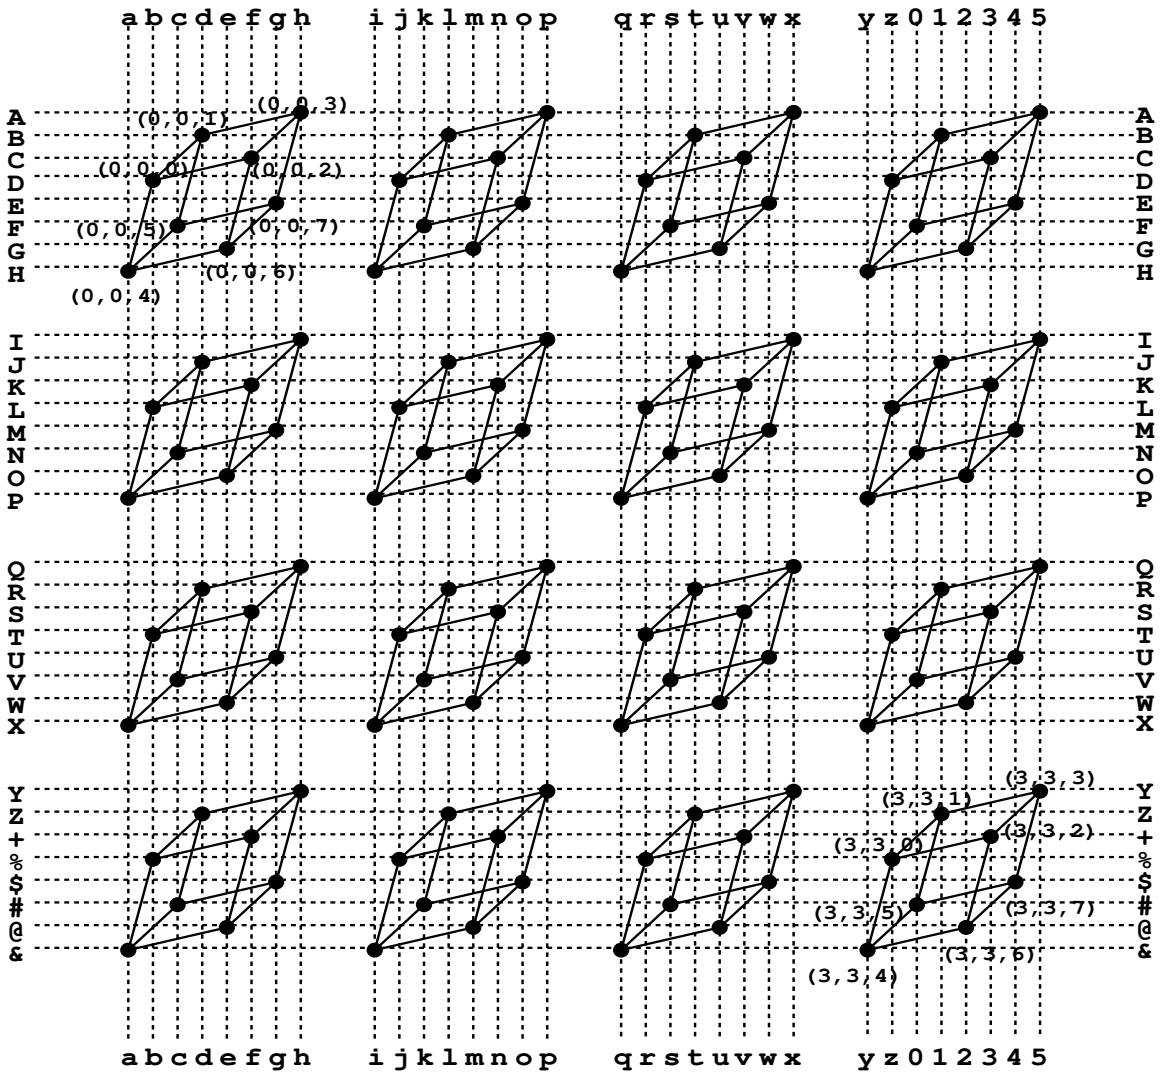


Figure 1: An example of the optical multi-mesh hypercube network: a $(4,4,3)$-OMMH (128 nodes) interconnection is shown. Two links with the same labels are connected for the wraparound connections of the torus. Only a few addresses are shown in the parenthesis for clarity. Solid lines represent hypercube connections and dashed lines torus connections.

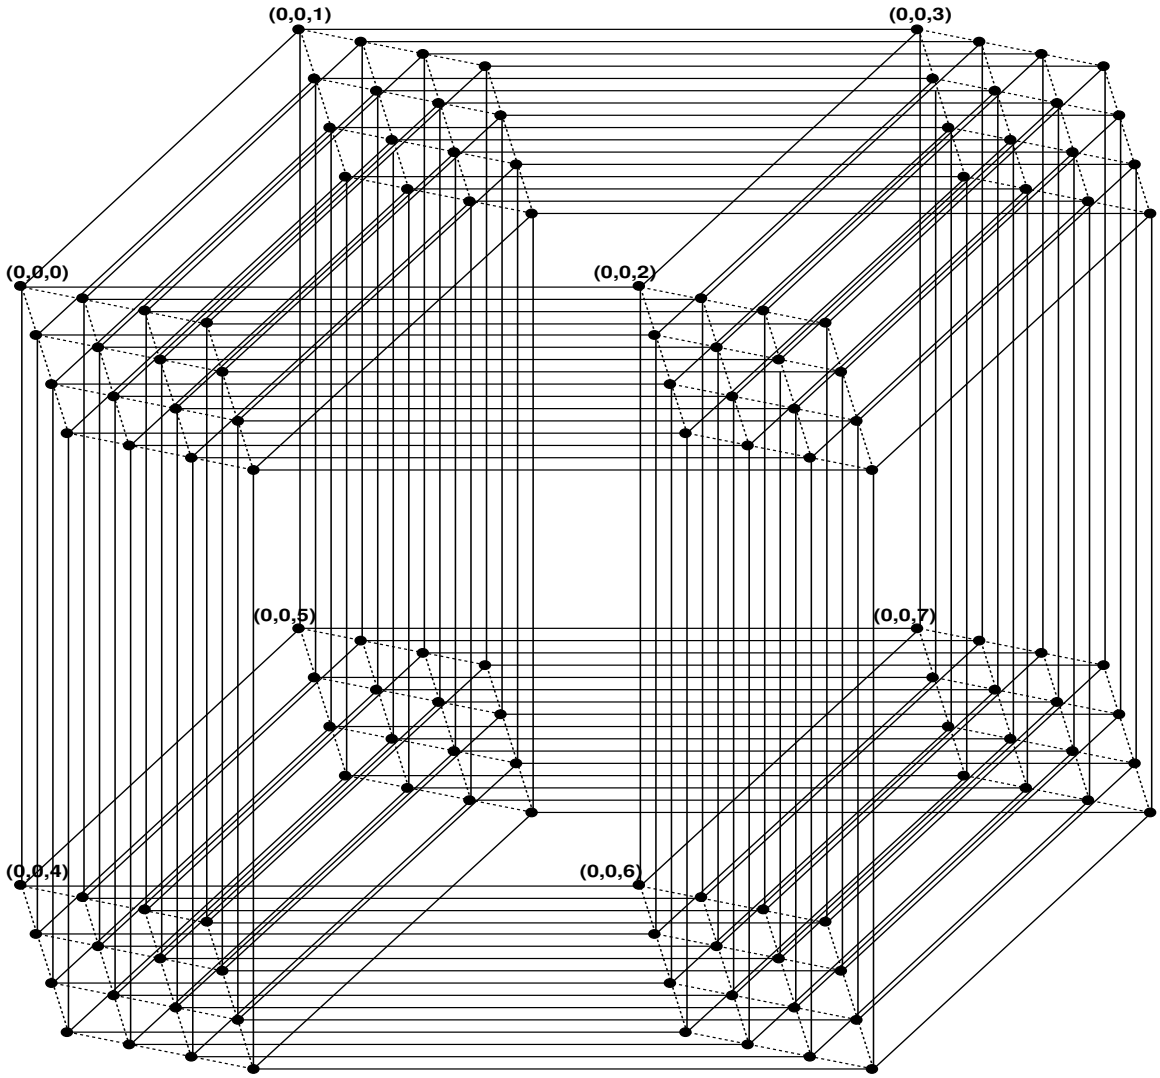


Figure 2: A $(4, 4, 3)$-OMMH interconnection network, another isomorphic view. Wraparound connections of the torus are omitted and only a few addresses are shown in the parenthesis for clarity. Solid lines represent hypercube connections and dashed lines torus connections.

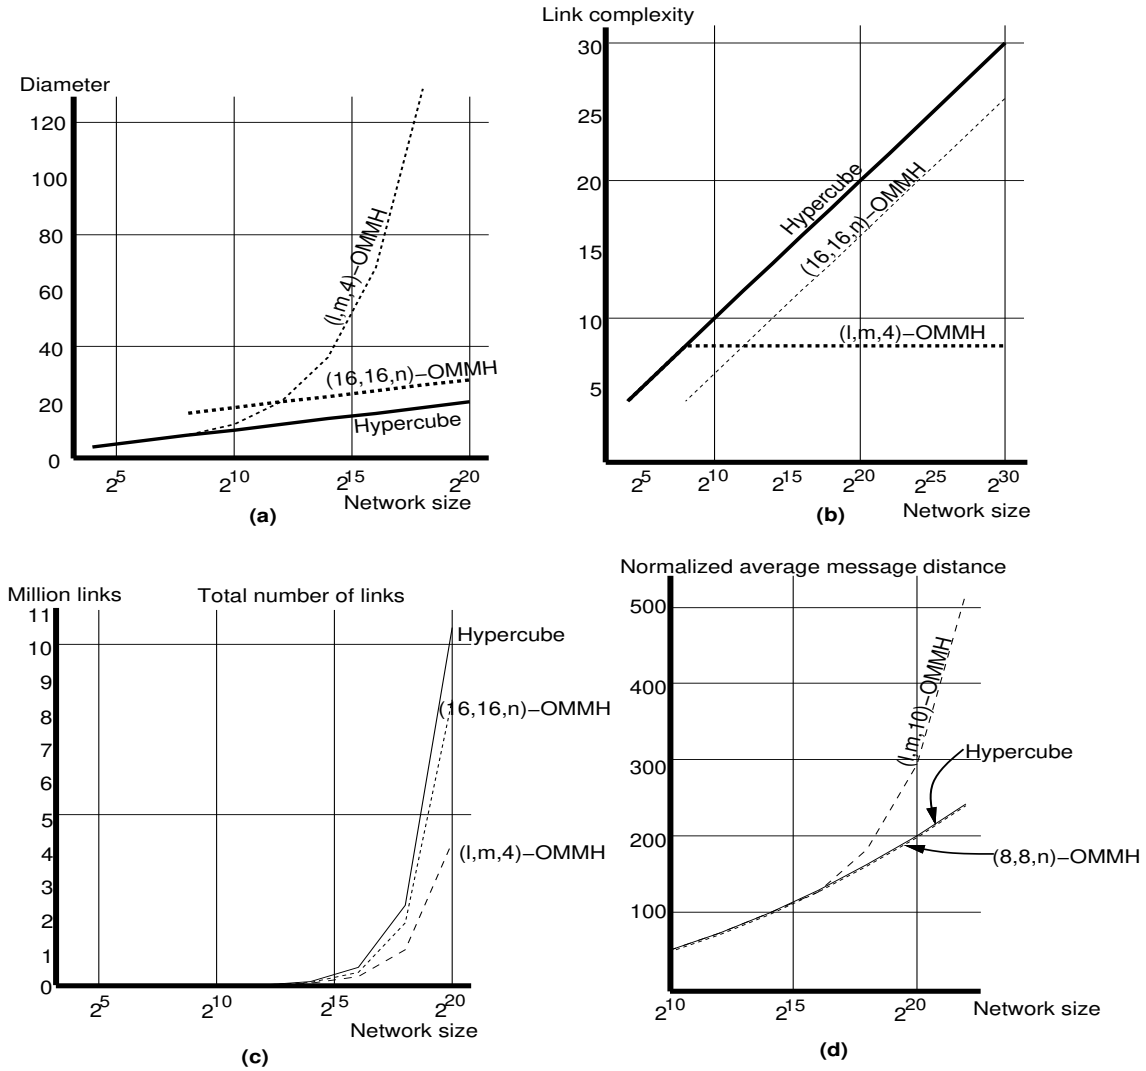


Figure 3: Comparison of (a) diameter, (b) link complexity, (c) total number of links, and (d) normalized average message distance of the hypercube and the OMMH when the two networks have the same number of nodes.

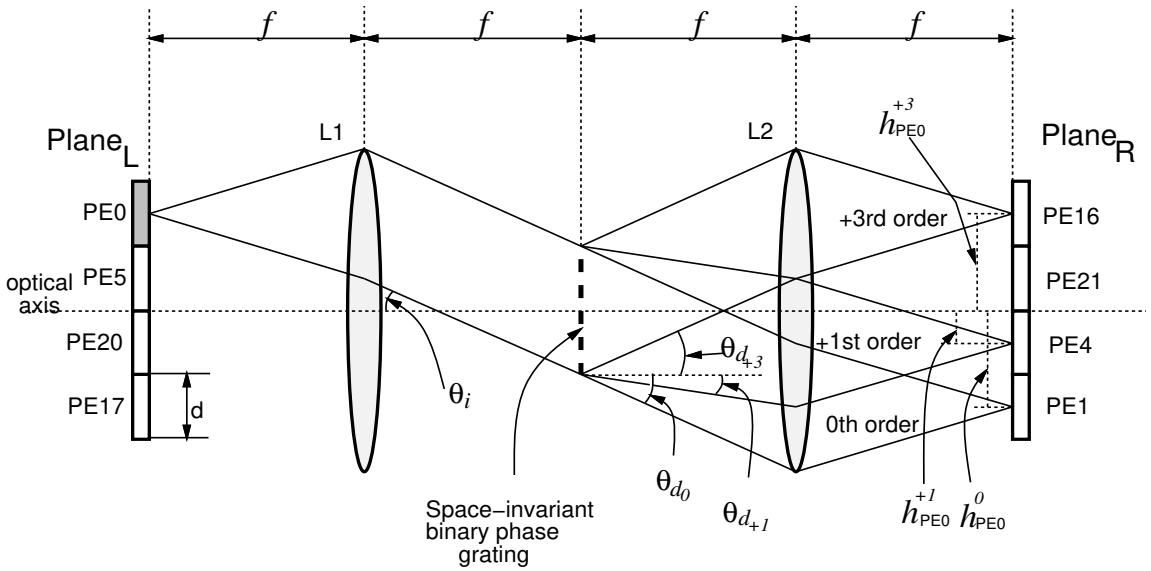


Figure 4: Space-invariant optical implementation of a 5-cube network using a binary phase grating.

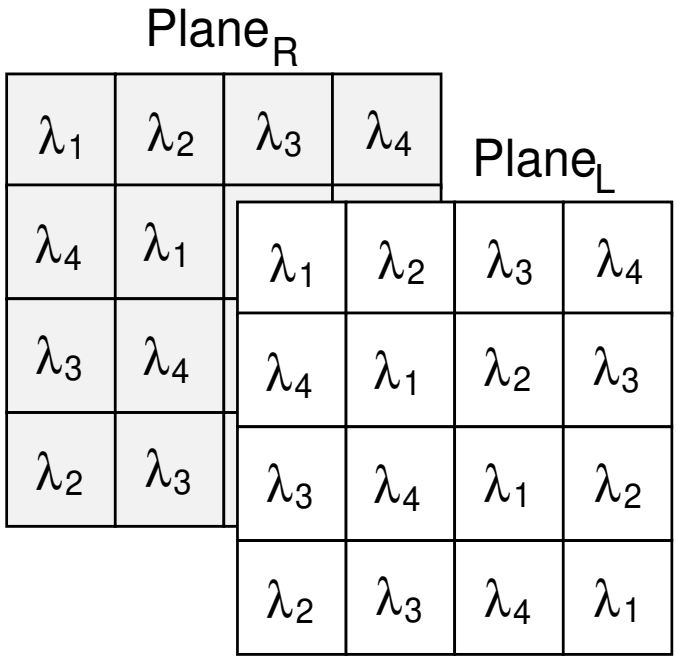


Figure 5: Wavelength assignment for a 5-cube. Wavelengths are assigned such that no two nodes in the same row or column have an identical wavelength.

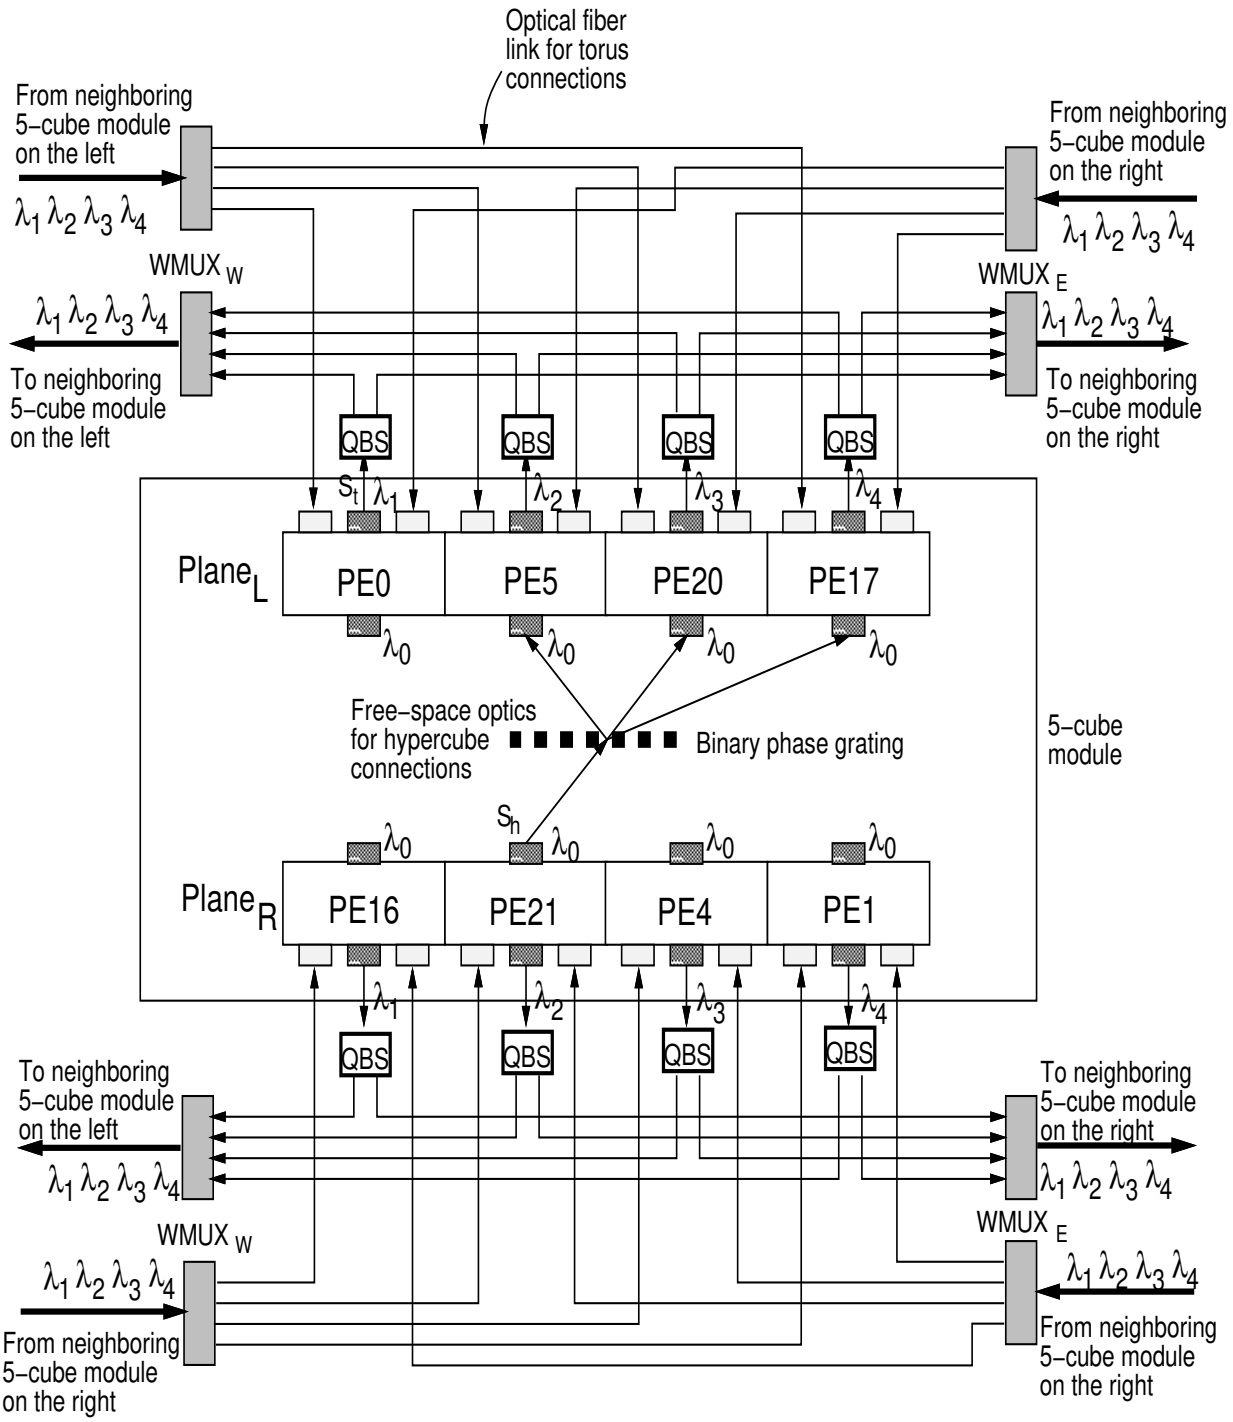


Figure 6: A 2-D side view of a 5-cube module to interface with torus links for the construction of the $(l, m, 5)$-OMMH network.

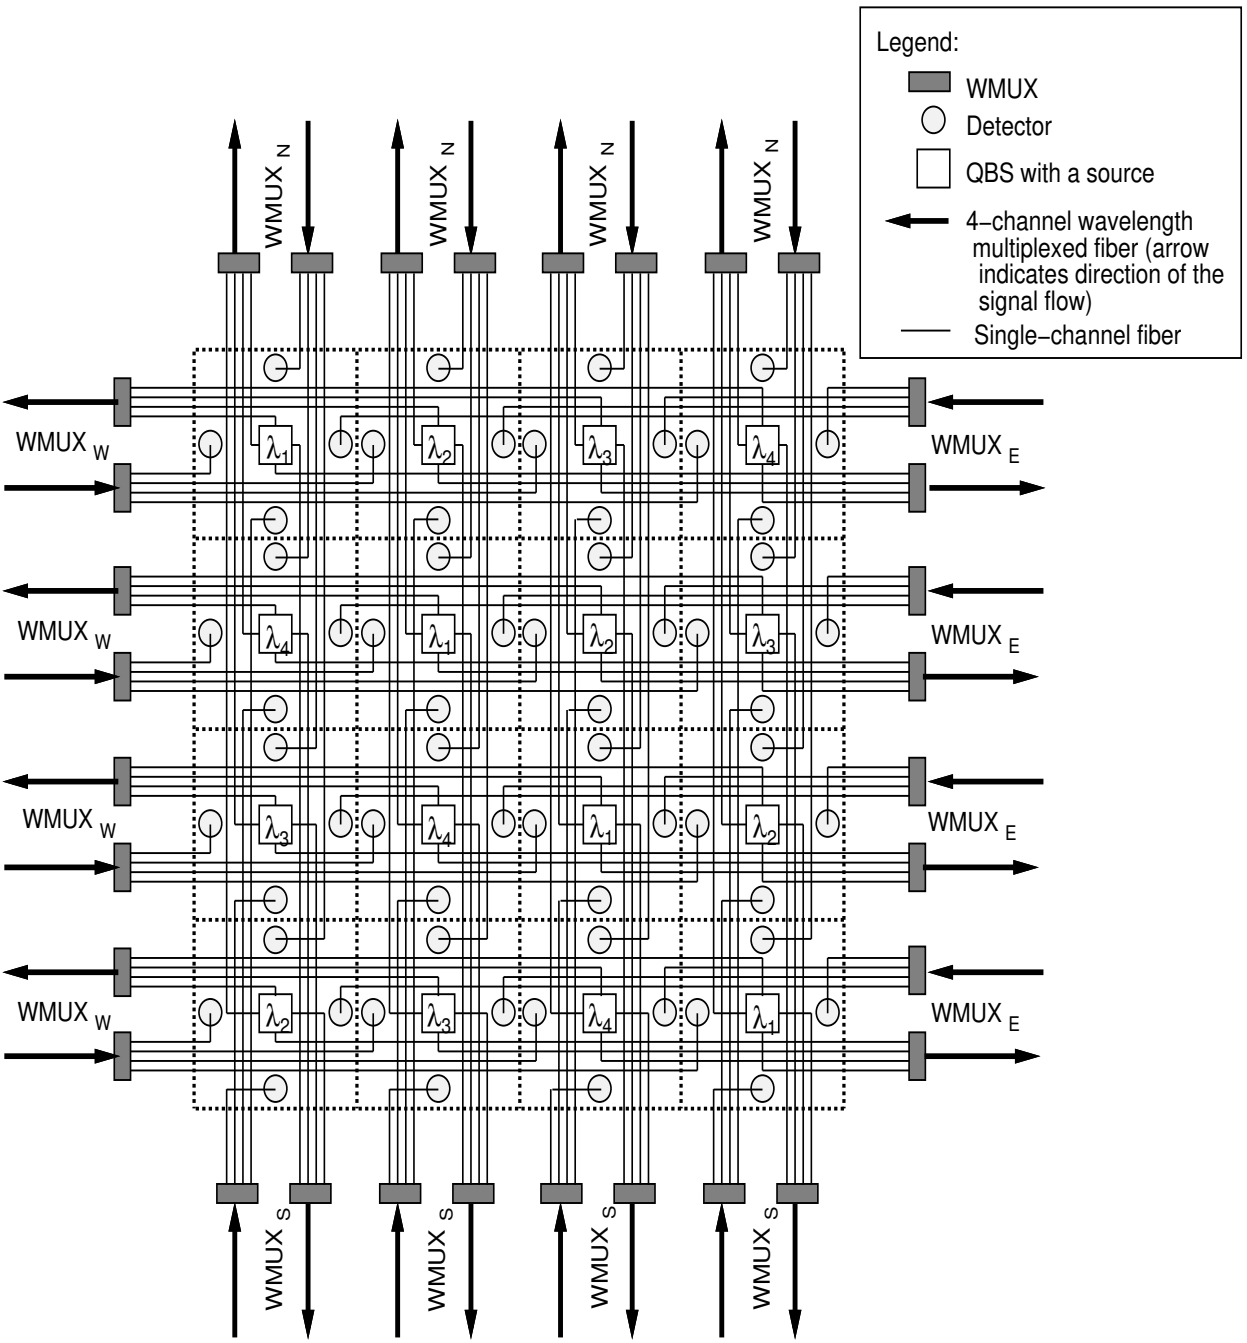


Figure 7: A 2-D top view of a 5-cube module to interface with torus links for the construction of the $(l, m, 5)$-OMMH network.

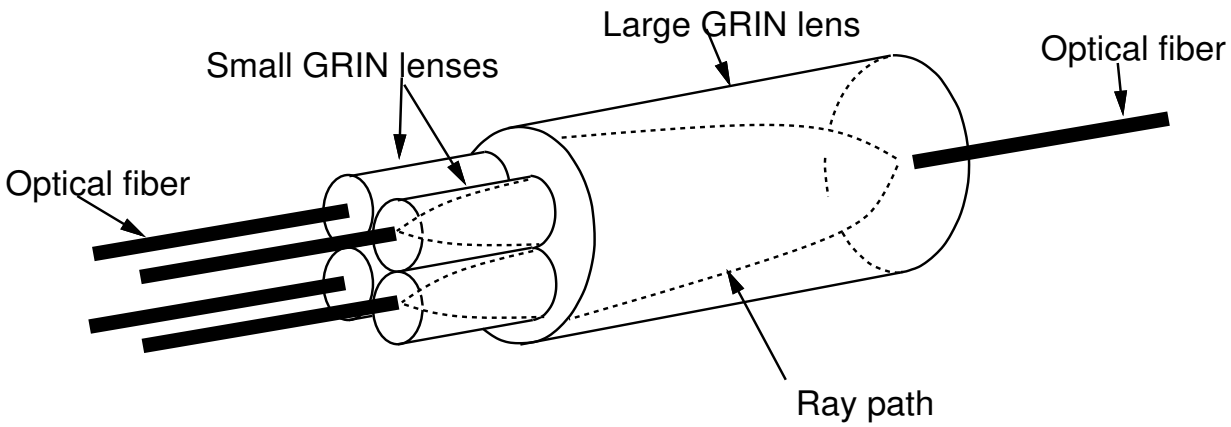


(a)

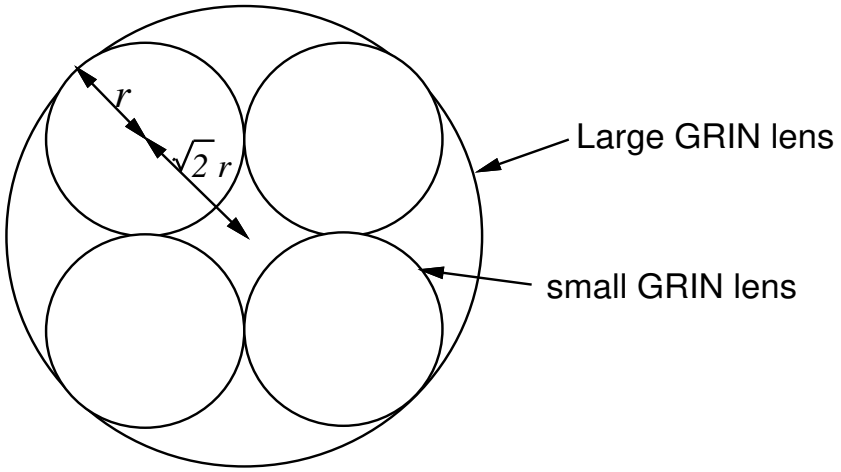


(b)

Figure 8: (a) A quadrant beam splitter using GRIN lenses. (b) Geometry of the quadrant beam splitter.

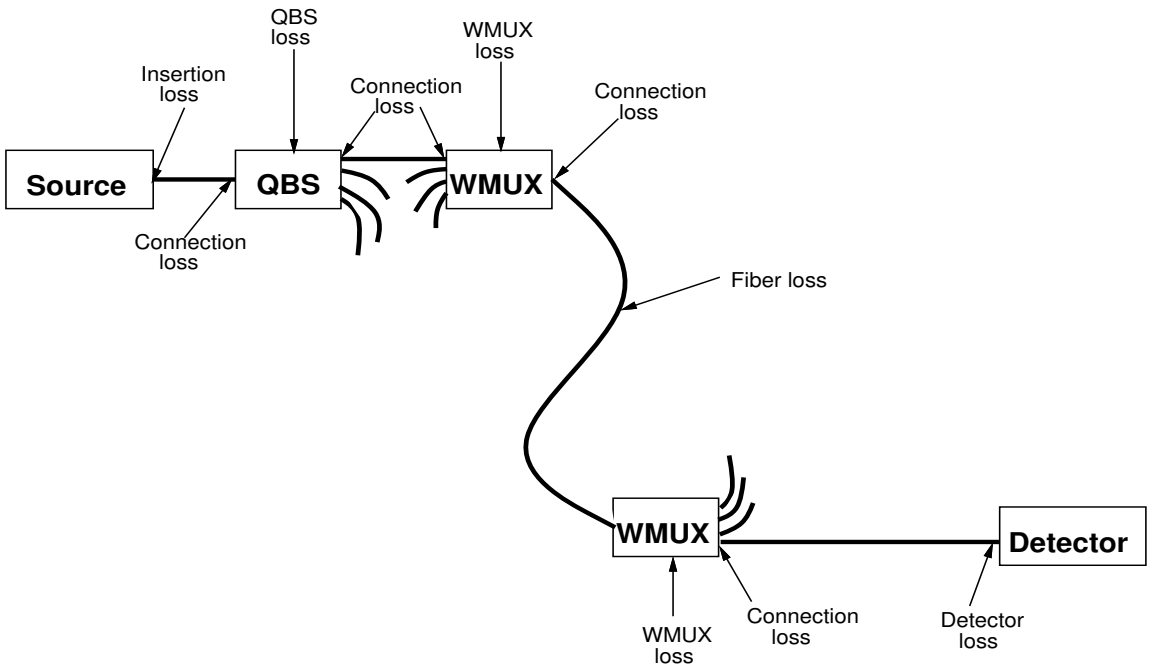


Figure 9: A model for the optical power budget calculation of the torus subnetwork in the OMMH network.