Wide-Sense Nonblocking Multicast in a Class of Regular Optical WDM Networks

Chunling Zhou and Yuanyuan Yang, *Senior Member, IEEE*

***Abstract*—Multicast communication involves transmitting information from a single source node to multiple destination nodes, and is becoming an important requirement in high-performance networks. In this paper, we study multicast communication in a class of optical WDM networks with regular topologies such as linear arrays, rings, meshes, tori and hypercubes. For each type of network, we derive the necessary and sufficient conditions on the minimum number of wavelengths required for a WDM network to be wide-sense nonblocking for multicast communication under some commonly used routing algorithms.**

***Index Terms*—Multicast communication, optical networks, regular networks, routing algorithm, wavelength assignment algorithm, wavelength division multiplexing (WDM), wide-sense nonblocking.**

I. Introduction

OPTICAL networks offer the potential of interconnecting hundreds to thousands of users and providing capacities of the order of gigabits per second to each user. Advances in electro-optic technologies have made optical communication a promising networking choice to meet the increasing demands for higher channel bandwidth and lower communication latency of high-performance computing and communication applications.

Wavelength-division multiplexing (WDM) is an approach that can exploit the huge opto-electronic bandwidth mismatch by requiring that each end-user's equipment operates only at peak electronic rate, but multiple WDM channels from different end-users may be multiplexed on the same fiber. Under WDM, the spectrum is divided into multiple wavelengths, with each wavelength supporting a single communication channel operating at electronic rate. At any time, each lightpath between a pair of nodes is carried on a certain wavelength. If any two lightpaths do not share a physical fiber, they can be on the same wavelength. WDM networks have attracted many researchers over the past few years [1]–[21]. It is anticipated that the next generation of the Internet will employ WDM-based optical backbones [1].

Paper approved by W. C. Kwong, the Editor for Optical Communications of the IEEE Communications Society. Manuscript received February 2, 2001; revised May 18, 2001. This work was supported in part by the Army Research Office under Grant DAAH04-96-1-0234 and by the National Science Foundation under Grant CCR-0073085.

C. Zhou is with the Artemis International Corporation, Exton, PA 19341 USA (e-mail: Cindy_Zhou@artemispm.com).

Y. Yang is with the Department of Electrical and Computer Engineering, State University of New York at Stony Brook, Stony Brook, NY 11794 USA (e-mail: yang@ece.sunysb.edu).

Publisher Item Identifier S 0090-6778(02)00517-2.

Multicast communication involves transmitting information from a single source node to multiple destination nodes, and is becoming an important requirement in high-performance networks. Multicast communication has been extensively studied in parallel processing and electronic networking community; see, for example, [22]–[26], and has received much attention in the optical networking community recently [8]–[19]. Typical multicast applications include tele/video-conferencing and distributive services such as video channels and HDTV program distribution. Many such applications are of real-time nature and require not only high-bandwidth but also some quality-of-service (QoS) guarantees. It is this type of applications that motivate us to consider nonblocking multicast in WDM networks in this paper. This is because that with WDM, the high-bandwidth needs of such applications could be easily met, and with nonblocking multicast capability, short multicast latency (the time from the source to all destinations) could be guaranteed.

In general, an optical WDM network consists of routing nodes interconnected by point-to-point fiber links, which can support a certain number of wavelengths. We assume each link in the network is bidirectional and actually consists of a pair of unidirectional links with one link in each direction. A *connection* or a *lightpath* in a WDM network is an ordered pair of nodes (x, y) corresponding to transmission of a packet from source x to destination y. We assume that no wavelength converter facility is available in the network. Thus, a connection must use the same wavelength throughout its path. In this case, the lightpath satisfies the *wavelength-continuity constraint*. We also assume that no *light splitters* are equipped at each routing nodes.

The number of wavelengths required for a collection of connections in a WDM network under the wavelength-continuity constraint is determined using a graph $G = (V, E)$, the *conflict graph*, in which each connection in the network is represented by a vertex in G. An undirected edge connecting two vertices appears in G if and only if the corresponding connections share a physical fiber link. Color the vertices of G such that no two adjacent vertices have the same color (a *proper coloring*). Then the minimum number of colors in a proper coloring of G (i.e., the *chromatic number* of G) is the minimum number of wavelengths required for the corresponding connections in the original network. This problem in general is NP-complete. However, as seen later, for conflict graphs for multicast communication in most of the networks we consider, one can efficiently determine the chromatic number.

Due to the nonuniform nature of multicast communication, the study of multicast communication is in general much more difficult and complex than that of unicast communication. To fa-

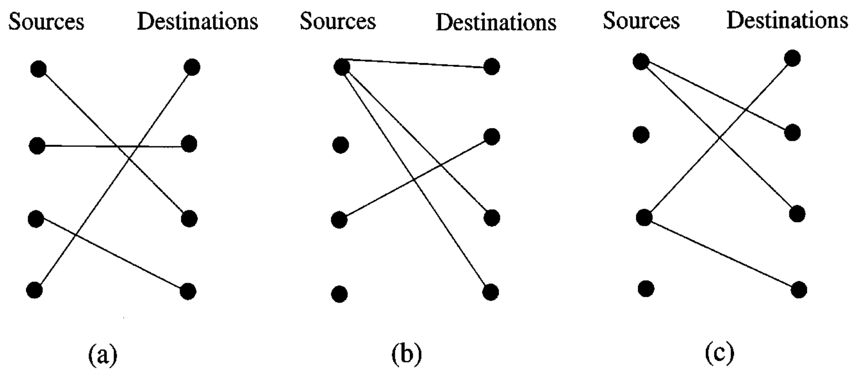

Fig. 1. Examples of multicast assignments in a 4-node network.

cilitate our analysis, in this paper we study a well-defined type of multicast communication pattern referred to as *multicast assignment*. A multicast assignment is a mapping from a set of network nodes (the source nodes) to a maximum set of network nodes (the destination nodes) with no overlapping allowed among the destination nodes of different source nodes. Fig. 1 gives several examples of multicast assignments in a 4-node network. Multicast assignments have been considered in the literature by many researchers, see, for example, [17], [22]–[25]. In a multicast assignment, the number of destination nodes from the same source node is referred to as the *fanout* of the source node. Two extreme cases of a multicast assignment are one-to-all broadcast and permutation. Clearly, an arbitrary multicast communication pattern can be decomposed into several multicast assignments.

In order to avoid the need for costly conversions between optical and electronic signals (so called O/E/O conversions) at immediate nodes, it is desirable for a multicast WDM network to be *nonblocking* as blocked multicast data will be dropped (lost) due to the lack of optical RAM (or buffer). A network is said to be nonblocking for multicast assignments if for any legitimate multicast connection request from a source node to a set of destination nodes, it is always possible to provide a connection path through the network to satisfy the connection request without any disturbance to existing connections. If the path selection must follow a routing algorithm to maintain the nonblocking connecting capability, the network is said to be *wide-sense nonblocking*. In this paper, we will focus on determining the minimum number of wavelengths required for a WDM network to be wide-sense nonblocking for arbitrary multicast assignments (denoted as w_w). In other words, we will determine the condition on which any multicast assignments can be embedded in a WDM network on-line under the routing algorithm. We will consider several typical regular networks, such as linear arrays, rings, meshes, tori, and hypercubes under some commonly used routing algorithms. This class of network has been extensively studied in the literature and many properties of them have been well-understood. Therefore, we choose these networks as a starting place to study the complex multicast issue in optical WDM networks. We expect the results obtained for it to be representative of the performance obtainable from practical long-haul networks, which also tend to have small degree and several paths of similar distance between each source-destination pair. In addition to long-haul networks, these networks themselves can be a good candidate for a multicast optical cross-connect (OXC) architecture due to their nonblocking multicast capability. The conditions for embedding multicast assignments in these WDM networks off-line (i.e., *rearrangeable nonblocking*, a weaker type of nonblocking capability than wide-sense) were studied by Zhou and Yang in [18], [19], and the conditions for embedding permutation assignments were derived by Qiao and Mei in [9], [10].

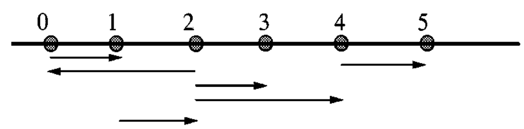

Fig. 2. A random sequence of connections in a WDM linear array.

The rest of the paper is organized as follows. Sections II–V derive the necessary and sufficient conditions on the number of wavelengths for linear arrays, rings, meshes and tori, and hypercubes, respectively. Finally, Section VI summarizes the results obtained in this paper and points out some future work.

II. LINEAR ARRAYS

Fig. 2 shows a six-node WDM linear array and a random set of connections in it. Clearly, there are only two possible directions for any connection in a linear array and the routing algorithm is unique. The following theorem gives the wide-sense nonblocking condition for a WDM linear array.

Theorem 1: The necessary and sufficient condition for a WDM linear array with N nodes to be wide-sense nonblocking for any multicast assignment is the number of wavelengths $w_w = N - 1$.

Proof: Sufficiency: In a linear array, each connection between a source and destination is either rightward or leftward. In a WDM network, the connections on the same link in different directions may use the same wavelength. To determine the sufficient number of wavelengths for a multicast assignment, we assume the connections in the same direction use different wavelengths. At any network state, if an existing connection is released, the wavelength it used can be reused by a new connection in either direction. In the case of a new connection is requested, notice that there are at most $N-1$ connections in the same direction in a multicast assignment, and there exist at most $N-2$ rightward or leftward connections. Without loss of generality, suppose the new connection is a rightward connection. Then we can assign it a wavelength that is different from the

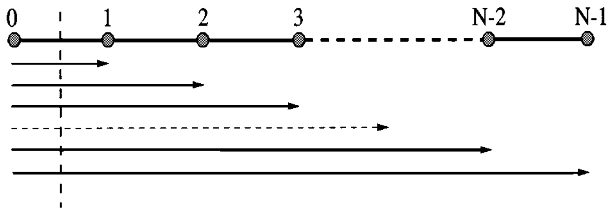

Fig. 3. A worst case multicast assignment in a linear array.

wavelengths the rightward connections are using. Thus, $N-1$ wavelengths are sufficient for all rightward connections. At the same time, a new rightward connection may use the same wavelength as a leftward connection since a rightward connection and a leftward connection can never interfere with each other even if they share the same fiber link. Similarly, a new leftward connection may use the same wavelength as a rightward connection. Hence, $N-1$ wavelengths are sufficient for all connections in any multicast assignment. Thus, we have $w_w \leq N-1$.

Necessity: Consider a multicast assignment as shown in Fig. 3, in which node 0 is the only source of all $N-1$ rightward connections. Regardless of the order of the connection requests, all these $N-1$ connections share link $0 \rightarrow 1$. Thus, at least $N-1$ wavelengths are required for this multicast assignment. That is, $w_w \geq N-1$.

Therefore, $N-1$ wavelengths are sufficient and necessary for a WDM linear array with N node to be wide-sense nonblocking for any multicast assignment, that is, we have $w_w = N-1$. ■

Next, we will present the wavelength assignment algorithm for a linear array. Whenever a new connection is requested, the following algorithm is used to assign a wavelength for it. In this algorithm, two tables for recording the wavelengths used are maintained, where T_R is for rightward connections, and T_L is for leftward connections. Table T_W is used for keeping the available wavelengths.

Wavelength Assignment Algorithm in a Linear Array

Step 1) Initially, let both T_R and T_L be empty and T_W contain $N-1$ available wavelengths.

Step 2) When a new connection is requested, if it is rightward, assign a wavelength that is in T_W but not in T_R to this new connection, and add this wavelength to T_R; if it is leftward, assign a wavelength that is in T_W but not in T_L to this new connection, and add this wavelength to T_L.

Step 3) When an existing connection is released, if it is rightward, delete the wavelength it used from T_R; if it is leftward, delete the wavelength it used from T_L.

End.

III. RINGS

In this section, we analyze the wide-sense nonblocking conditions for WDM rings. We will first consider unidirectional rings and then bidirectional rings.

A. Unidirectional Rings

In a unidirectional ring, the routing algorithm is unique. Without loss of generality, we assume the direction of a ring is counter-clockwise.

Theorem 2: The necessary and sufficient condition for a unidirectional WDM ring with N nodes to be wide-sense nonblocking for any multicast assignment is the number of wavelengths $w_w = N$.

Proof: Sufficiency: Based on the definition of a multicast assignment, there are a total of N connections in any multicast assignment. If a new connection is requested, assign it a wavelength that is not currently used. If an existing connection is released, its wavelength can be reused. Therefore, N wavelengths are sufficient for any multicast assignment. Hence we have $w_w \leq N$.

Necessity: Consider the multicast assignment $\Pi_N = \{i \rightarrow (i+N-1) \bmod N | 0 \leq i \leq N-1\}$ (Fig. 4(a) shows an example of such assignment for $N=6$). Since each node in the network is the source of one connection and each connection covers $N-1$ nodes that are the sources of other connections, each connection shares links with other $N-1$ connections. This implies that each connection shares links with all other connections. By the definition of a conflict graph, in this case we know that the conflict graph of this multicast assignment is an N-node complete graph as shown in Fig. 4(b). Clearly, the chromatic number of the conflict graph is N. Thus, at least N wavelengths are required for this multicast assignment. That is, $w_w \geq N$.

Therefore, N wavelengths are sufficient and necessary for a unidirectional WDM ring with N nodes to be wide-sense nonblocking for any multicast assignment, that is, $w_w = N$. ■

Now, we describe the wavelength assignment algorithm in a unidirectional ring. In this algorithm two tables are maintained: one is for the currently used wavelengths T_u, and another is for the available wavelengths T_n.

Wavelength Assignment Algorithm in a Unidirectional Ring

Step 1) Initially, let T_u be empty, and assign N wavelengths to T_n.

Step 2) When a new connection is requested, assign to it a wavelength that is in T_n, but not in T_u, add the new wavelength to T_u, and delete it from T_n.

Step 3) When an existing connection is released, add the wavelength it used to T_n, and delete it from T_u.

End.

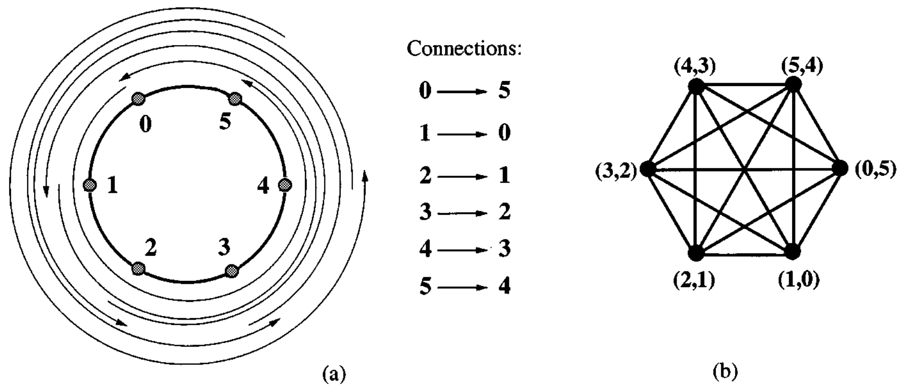

Fig. 4. (a) A worst-case multicast assignment in a unidirectional ring. (b) The associated conflict graph.

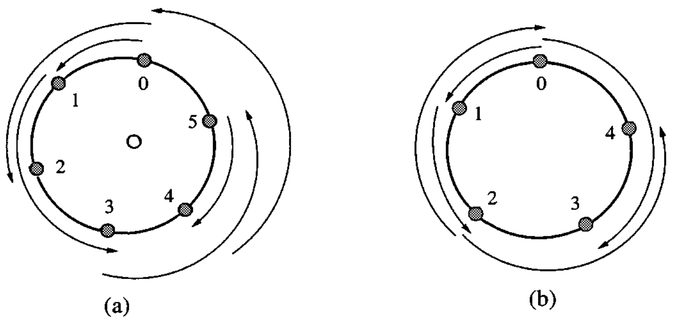

Fig. 5. A random sequence of connections in a bidirectional ring with N nodes. (a) N is even. (b) N is odd.

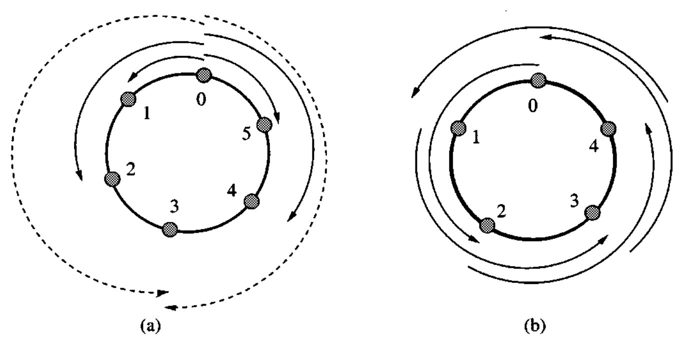

Fig. 7. A worst-cases multicast assignment in a bidirectional ring. (a) N is even. (b) N is odd.

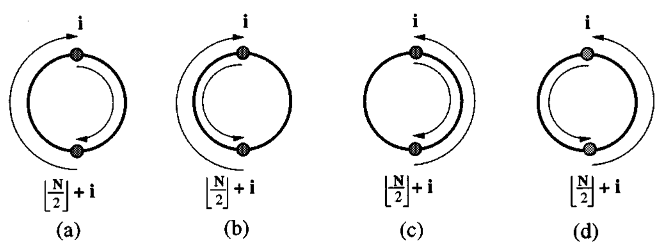

Fig. 6. Four possible relationships between the connection destined to node i and the connection destined to note $(i + \lfloor N/2 \rfloor) \bmod N$.

B. Bidirectional Rings

In a bidirectional ring, there are two possible paths for a connection between any two nodes: clockwise or counter-clockwise. In our analysis of wide-sense nonblocking conditions, we adopt the *shortest path routing* algorithm. Fig. 5 shows a random sequence of connections in a bidirectional ring under shortest path routing.

Theorem 3: The necessary and sufficient condition for a bidirectional WDM ring with N nodes to be wide-sense nonblocking for any multicast assignment under shortest path routing is the number of wavelengths $w_w = \lceil N/2 \rceil$.

Proof: Sufficiency: In a multicast assignment, there are at most N connections. Consider two nodes with $\lfloor N/2 \rfloor$ nodes apart on the ring: node i and node $i + \lfloor N/2 \rfloor \bmod N$, where $1 \le i \le (N/2) - 1$. As shown in Fig. 6, under the shortest path routing there are four possible relationships between the connection destined to node i and the connection destined to node $(i + \lfloor N/2 \rfloor) \bmod N$. When these two connections are in different directions, they cannot interfere with each other. When they are in the same direction, the path of each connection covers no more than $\lfloor N/2 \rfloor$ links. Thus, the paths of these two connections cannot be overlapped, which means they can use the same wavelength. Therefore, we can divide the N connections in a multicast assignment into $\lceil N/2 \rceil$ pairs with the connections in each pair using the same wavelength. In particular, when N is even, pair i $(0 \le i \le (N/2) - 1)$ contains the connection destined to node i and the connection destined to node $(i + (N/2)) \bmod N$, and when N is odd, pair 0 contains connection destined to node 0, and pair i $(1 \le i \le \lceil N/2 \rceil - 1)$ contains connections destined to nodes i and $(i + \lfloor N/2 \rfloor) \bmod N$. When a new connection is requested, we first determine which pair it is in, and then assign to it the wavelength corresponding to that pair.

Therefore, for $\lceil N/2 \rceil$ pairs of connections, $\lceil N/2 \rceil$ wavelengths are sufficient. Hence, we have $w_w \le \lceil N/2 \rceil$.

Necessity: Consider N is even first. Look at a worst case multicast assignment as shown in Fig. 7. In this multicast assignment, node 0 is the only source of $N - 1$ connections to all the other nodes. Regardless of the order the $N - 1$ connection requests are satisfied, the $(N/2) - 1$ connections destined to nodes $1, \ldots, (N/2) - 1$, respectively, share link $0 \to 1$. Similarly, the $(N/2) - 1$ connections destined to nodes $N - 1, \ldots, (N/2) + 1$, respectively, share link $0 \to N - 1$. The connection destined to node $N/2$ has two possible directions: if it is clockwise, it uses link $0 \to N - 1$; if it is counter-clockwise, it uses link $0 \to 1$. In either case, there are $N/2$ connections sharing link $0 \to 1$

or link $0 \rightarrow N-1$. Thus, at least $N/2$ wavelengths are required for all connections in this multicast assignment. We have $w_w \geq N/2$ when N is even.

Next, let's consider N is odd. The worst case multicast assignment in this case is $\Pi_N = \{i \rightarrow (i+(N-1)/2) \bmod N | 0 \leq i \leq N-1\}$. This multicast assignment is actually a permutation assignment in which each node is a source of exactly one connection and a destination of exactly one connection. We can see that the path of each connection in this multicast assignment covers $(N-1)/2$ links. Thus, each connection shares links with other $(N-1)/2$ connections, which implies that the number of wavelengths needed for this multicast assignment is at least $(N-1)/2+1$. That is, $w_w \geq (N-1)/2+1 = \lceil N/2 \rceil$ when N is odd.

Now we can see that in both cases, $\lceil N/2 \rceil$ wavelengths are sufficient and necessary for a bidirectional WDM ring with N nodes to be wide-sense nonblocking for any multicast assignment. Hence, $w_w = \lceil N/2 \rceil$. ■

The wavelength assignment algorithm for a bidirectional ring is given bellow. In this algorithm, we number the $\lceil N/2 \rceil$ available wavelengths as $w_0, \ldots, w_{\lceil N/2 \rceil - 1}$.

```
Wavelength Assignment Algorithm in a
Bidirectional Ring
  Case 1) N is even. When a new con-
        nection (i,j) is requested, if
        0 ≤ j ≤ (N/2) − 1, assign wave-
        length w_j to this connection. If
        (N/2) ≤ j ≤ N − 1, assign wavelength
        w_{j−(N/2)}, to this connection.
  Case 2) N is odd. When a new con-
        nection (i,j) is requested, if
        0 ≤ j ≤ (N−1)/2, assign wave-
        length w_j to this connection. If
        (N+1)/2 ≤ j ≤ N − 1, assign wave-
        length w_{j−(N−1)/2} to this connec-
        tion.
  End.
```

IV. Meshes and Tori

In this section, we consider WDM mesh and torus networks under *row-major shortest path routing* algorithm.

Definition 1: Under the row-major shortest path routing, for a connection request $((x_0, y_0), (x_1, y_1))$ in a mesh or a torus, the path is deterministically from node (x_0, y_0) to node (x_0, y_1) in row x_0 along the shortest path first, then to node (x_1, y_1) in column y_1 along the shortest path.

We will first discuss the wide-sense nonblocking condition for WDM meshes.

A. Meshes

Definition 2: For a connection in a mesh under row-major shortest path routing, if the connection goes right at the first step from the source, we refer to it as a rightward connection. If the connection goes left at the first step, we refer to it as a leftward connection. Otherwise, we refer to it as a straight connection.

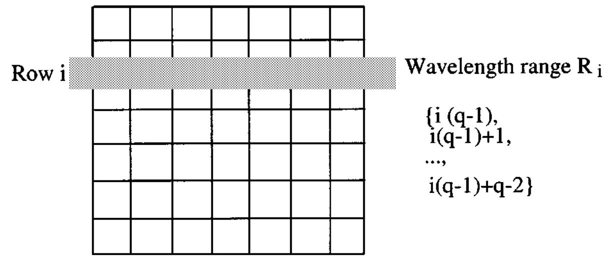

Fig. 8. Wavelength range for row i in a $p*q$ WDM mesh.

The following theorem gives the wide-sense nonblocking condition for a mesh.

Theorem 4: The necessary and sufficient condition for a WDM mesh with p rows and q columns to be wide-sense nonblocking for any multicast assignment under row-major shortest path routing is the number of wavelengths $w_w = p*(q-1)$.

Proof: Sufficiency: Let the wavelengths be numbered from 0. We divide all wavelengths into p equal length ranges, $R_0, R_1, \ldots, R_{p-1}$, and let the connections destined to row i use the wavelengths within range R_i for $0 \leq i \leq p-1$ as shown in Fig. 8. Consequently, the connections destined to the same column will use different wavelengths. Thus, the connections destined to the same column never interfere with each other even if they share a link. Now we consider the connections destined to the same row i to determine how many wavelengths are sufficient for one row. Among all the connections to the same row, if the sources of two connections are in different rows, these two connections cannot interfere with each other in any row, which means they can use the same wavelength. Thus, it will suffice to consider the connections originated from the same row i. Now the only connections need to be considered are those originated from the same row and destined to the same row. In this case, each row can be considered as a linear array with q nodes. By Theorem 1, we know that for any multicast assignment, $q-1$ wavelengths are sufficient for a linear array with q nodes to be wide-sense nonblocking. Hence, $p*(q-1)$ wavelengths are sufficient for a p-row mesh.

We can assign a wavelength to a new connection to row i as follows. If the new connection is a rightward connection, we assign to it a wavelength in range R_i but is not currently used by any rightward connections. If it is a leftward connection, we assign to it a wavelength in range R_i but is not currently used by any leftward connections. If it is a straight connection, we treat it as a leftward connection if the number of leftward connections is less than that of rightward connections in row i. Otherwise, treat it as a rightward connection. When an existing connection is released, we can simply delete it. Hence, for the $p*q$ connections in a mesh, $p*(q-1)$ wavelengths are sufficient, which means $w_w \leq p*(q-1)$.

Necessity: Consider the multicast assignment in which node (0,0) is the source of $p*(q-1)$ rightward connections. Under row-major shortest path routing, regardless of the order these $p*(q-1)$ connections are satisfied, they must share link $(0,0) \rightarrow (0,1)$ as shown in Fig. 9. This indicates that at least $p*(q-1)$ wavelengths are required for all connections in the multicast assignment, which implies $w_w \geq p*(q-1)$.

Hence, we have $w_w = p*(q-1)$. ■

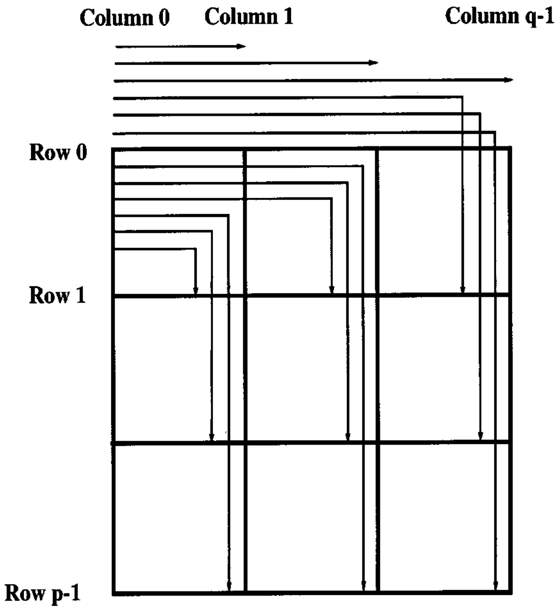

Fig. 9. A worst-case multicast assignment in a $p \times q$ WDM mesh under row-major routing.

We now give the wavelength assignment algorithm for a $p \times q$ mesh. Let the available wavelengths be numbered as $w_0, w_1, \ldots, w_{p*(q-1)-1}$, and the wavelengths be divided into p ranges, where range i has $q-1$ wavelengths $(w_{i*(q-1)}, \ldots, w_{i*(q-1)+q-2})$, which corresponds to row i. Two tables T_R^i and T_L^i are maintained for each row i to record the used wavelengths for rightward connections and leftward connections, respectively.

Wavelength Assignment Algorithm in a Mesh

Step 1) Initially, let tables T_R^i and T_L^i be empty.

Step 2) When a new connection to row i is requested, if it is a rightward connection, assign to it a wavelength that is in range i, but is not in T_R^i, and add this wavelength to T_R^i. If it is a leftward connection, assign to it a wavelength that is in range i, but is not in T_L^i, and add this wavelength to T_L^i. If it is a straight connection, compare the sizes of T_R^i and T_L^i. If $|T_R^i| \geq |T_L^i|$, treat the new connection as a leftward connection. Otherwise, treat the new connection as a rightward connection.

Step 3) When an existing connection to row i is released, if it is rightward connection, delete it from T_R^i; if it is leftward connection, delete it from T_L^i.

End.

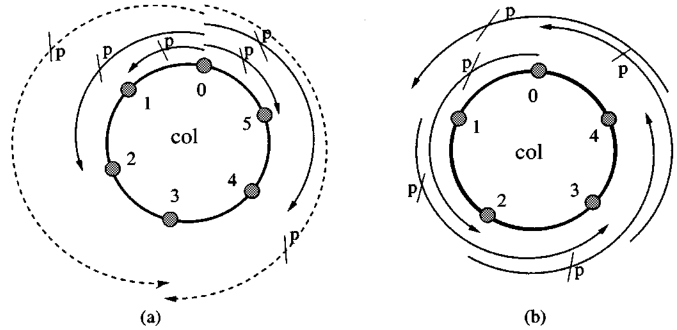

Fig. 10. Top view of a worst-case multicast assignment in a $p \times q$ WDM torus. (a) q is even. (b) q is odd.

B. Tori

The following theorem gives the wide-sense nonblocking condition for a WDM torus.

Theorem 5: The necessary and sufficient condition for a WDM torus with p rows and q columns to be wide-sense nonblocking for any multicast assignment under row-major shortest path routing is the number of wavelengths $w_w = p * \lceil q/2 \rceil$.

Proof: Sufficiency: Similar to a mesh, we can let the connections destined to different rows use the wavelengths in a different range so that the connections destined to the same column cannot interfere with each other even when they use the same link. Also, it is clear that under row-major shortest path routing, if the sources of two connections are in different rows, these two connections can never meet in the same row. Therefore, we only need to consider the case that the sources are in the same row. If the connections don't interfere in a row, they cannot interfere in the entire torus. Thus, if x wavelengths are sufficient for a row, $p * x$ wavelengths are sufficient for the entire torus. In fact, we can consider the torus as p layers of bidirectional rings with q nodes in each ring. When a new connection destined to row i is requested, assign a wavelength to it from the wavelength range of row i. Since there are at most q connections destined to row i, by Theorem 3 we know that $x = \lceil q/2 \rceil$. Then in a torus, $p * x = p * \lceil q/2 \rceil$ wavelengths are sufficient. That is, $w_w \leq p * \lceil q/2 \rceil$.

Necessity: We first assume q is even. Consider the multicast assignment in Fig. 10(a), in which node (0,0) is the only source of all connections. Clearly, the connections destined to columns $1, \ldots, (q/2) - 1$, respectively, share link $(0,0) \rightarrow (0,1)$ regardless the order these connections are satisfied. Similarly, the connections destined to columns $(q/2)+1, \ldots, q-1$, respectively, share link $(0,0) \rightarrow (0, q-1)$. Also, the connections destined to column $q/2$ may use either link $(0,0) \rightarrow (0,1)$ or link $(0,0) \rightarrow (0, q-1)$ because both links are on a shortest path. If the connections destined to column $q/2$ use link $(0,0) \rightarrow (0,1)$, $p*(q/2)$ connections share link $(0,0) \rightarrow (0,1)$. Similarly, if the connections destined to column $q/2$ use link $(0,0) \rightarrow (0, q-1)$, $p*(q/2)$ connections share the link $(0,0) \rightarrow (0, q-1)$. Hence, at least $p*(q/2)$ wavelengths are required. That is, $w_w \geq p*(q/2)$ when q is even.

In the case of q is odd consider the multicast assignment $\Pi_{p*q} = \{(0,j) \rightarrow (i, (j + (q-1)/2) \bmod q) | 0 \leq i \leq p -$

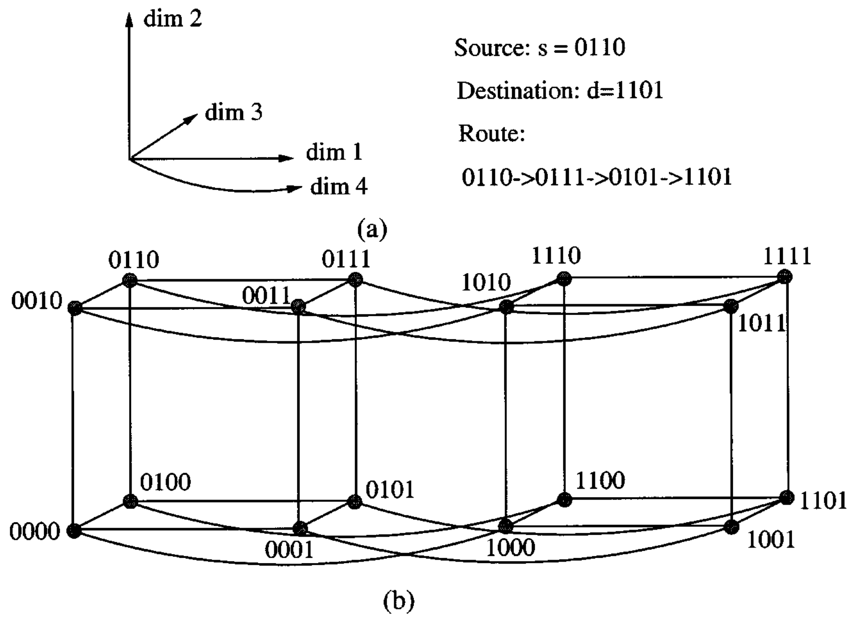

Fig. 11. E-cube routing in a hypercube with 16 nodes.

$1,\ 0 \leq j \leq q-1\}$ as shown in Fig. 10(b). In this multicast assignment, each node in row 0 (a ring with q nodes) is the source of p connections destined to the same column. Clearly, these p connections go through the same path in row 0 under row-major shortest path routing. Thus, they must be assigned different wavelengths. By Theorem 3, we know that a ring with q nodes needs $(q-1)/2+1$ wavelengths. Hence, for p layers of rings, $p*((q-1)/2+1)$ wavelengths are required. That is, $w_w \geq p*((q-1)/2+1) = p*\lceil q/2 \rceil$ when q is odd.

Therefore, we have $w_w = p*\lceil q/2 \rceil$. ■

The wavelength assignment algorithm for a torus is described as follows. In this algorithm we number the available wavelengths as $w_0, w_1, \ldots, w_{p*\lceil q/2 \rceil}$, and divide the wavelengths into p ranges, where range i $(0 \leq i \leq p-1)$ has $\lceil q/2 \rceil$ wavelengths, $w_{i*\lceil q/2 \rceil}, \ldots, w_{i*\lceil q/2 \rceil + \lceil q/2 \rceil -1}$. For each connection destined to row i, a wavelength within range i is assigned.

```
Wavelength Assignment Algorithm in a Torus
```

Case 1) q is even. When a new connection to node (i,j) is requested, if $0 \leq j \leq (q/2)-1$, assign to it wavelength w_k, where $k = i*(q/2)+j$; if $(q/2) \leq j \leq q-1$, assign to it wavelength w_k, where $k = i*(q/2)+(j-(q/2))$.

Case 2) q is odd. When a new connection to node (i,j) is requested, if $0 \leq j \leq (q-1)/2$, assign to it wavelength w_k, where $k = i*(q+1/2)+j$; if $(q+1)/2 \leq j \leq q-1$, assign to it wavelength w_k, where $k = i*(q+1/2)+(j-(q-1)/2)$.

End.

V. HYPERCUBES

In our analysis of wide-sense nonblocking conditions for WDM hypercubes, *e-cube routing* algorithm is used, which is a deterministic shortest path routing algorithm with dimensions crossed from the lowest to the highest. Specifically, e-cube routing is defined as follows.

Definition 3: In an n-cube with $N = 2^n$ nodes, let each node b be binary-coded as $b = b_{n-1}b_{n-2}\ldots b_1 b_0$, and the n dimensions of a hypercube be numbered as $i = 1, 2, \ldots, n$, where the ith dimension corresponds to the $(i-1)$st bit in the node address. Suppose the source node $s = s_{n-1}s_{n-2}\ldots s_1 s_0$ and the destination node $d = d_{n-1}d_{n-2}\ldots d_1 d_0$. Let $v = v_{n-1}\ldots v_1 v_0$ be any node along the route from s to d. In e-cube routing, a route from s to d with a minimum number of steps is uniquely determined as follows:

Step 1) Compute the direction bit $r_i = s_{i-1} \oplus d_{i-1}$ for all n dimensions $(i = 1, \ldots, n)$. Initially, dimension $i = 1$ and $v = s$.

Step 2) Route from the current node v to the next node $v \oplus 2^{i-1}$ if $r_i = 1$. Skip this step if $r_i = 0$.

Step 3) Increment the dimension to $i+1$ (i.e., $i \leftarrow i+1$). If $i \leq n$, go to step 2, else exit.

Fig. 11 illustrates an example of e-cube routing in a four-dimensional hypercube. For $s = 0110$ and $d = 1101$, we have $r = r_4 r_3 r_2 r_1 = 1011$. We first route from s to $s \oplus 2^0 = 0111$ since $r_1 = 0 \oplus 1 = 1$. Then route from $v = 0111$ to $v \oplus 2^1 = 0101$ since $r_2 = 1 \oplus 0 = 1$. We skip dimension $i = 3$ because $r_3 = 1 \oplus 1 = 0$. Finally, we route from $v = 0101$ to $v \oplus 2^3 = 1101 = d$ since $r_4 = 1$. Note that the route is determined from dimension 1 to dimension 4 in order. If the ith bit of s and d agree, no routing is needed along dimension i. Otherwise, move from the current node to the next node along the same dimension. The procedure is repeated until the destination node is reached.

To present our results for WDM hypercubes, we need more definitions.

TABLE I
NUMBER OF WAVELENGTHS REQUIRED FOR A WDM NETWORK TO BE WIDE-SENSE NONBLOCKING FOR ANY MULTICAST ASSIGNMENT

Network Topology	Routing Algorithm	No. Wavelengths Required
N node linear array	—	$N-1$
N node unidirectional ring	Shortest path	N
N node bidirectional ring	Shortest path	$\lceil \frac{N}{2} \rceil$
$p \times q$ mesh	Row-major	$p*(q-1)$
$p \times q$ torus	Row-major	$p*\lceil \frac{q}{2} \rceil$
n-dimensional hypercube	e-cube	2^{n-1}

Definition 4: Based on the binary address of each node, a hypercube can be divided into two subcubes. In the 0-subcube, the first (i.e., the lowest) bit of the binary address of each node is 0. In the 1-subcube, the first bit of the binary address of each node is 1.

Definition 5: For a connection in a hypercube, if the destination of this connection is in the 0-subcube, it is referred to as a 0-connection; if the destination of the connection is in the 1-subcube, it is referred to as a 1-connection.

Now, we are in the position to give the wide-sense nonblocking condition for WDM hypercubes.

Theorem 6: The necessary and sufficient condition for a WDM hypercube with $N = 2^n$ nodes to be wide-sense nonblocking for any multicast assignment under e-cube routing is the number of wavelengths $w_w = N/2 = 2^{n-1}$.

Proof: Sufficiency: For a 0-connection in any multicast assignment, if the source is in the 1-subcube, the first step of the routing is to correct the first bit from 1 to 0. Then the first link on the path is from the 1-subcube to the 0-subcube. For any 0-connection, if the source is in the 0-subcube, its path can never get into the 1-subcube. Similarly for 1-connections. Therefore, at most $N/2$ connections go from the 1-subcube to the 0-subcube. Similarly, at most $N/2$ connections go from the 0-subcube to the 1-subcube. Clearly, any 0-connection and any 1-connection cannot interfere with each other even if they use the same link, since they are in different directions. Thus, any 0-connection and any 1-connection can use the same wavelength. Therefore, for the links between these two subcubes, $N/2$ wavelengths are sufficient. In the meanwhile, there are at most $N/2$ connections within the 0-subcube and at most $N/2$ connections within the 1-subcube. Suppose all 0-connections use different wavelengths and 1-connections use different wavelengths. Then there is not interference between any two connections, which means $N/2$ wavelengths are sufficient in any subcube. When a new connection is requested, if it is a 0-connection, we assign to it a wavelength that is not used by any 0-connection. As we know, before the new connection request is satisfied, there are at most $(N/2)-1$ 0-connections, which means at most $(N/2)-1$ wavelengths have been used by 0-connections. If it is a 1-connection, similarly, at most $(N/2)-1$ wavelengths have been used by 1-connections. Therefore, $N/2$ wavelengths are sufficient, that is, $w_r \leq N/2$.

Necessity: Consider the multicast assignment in which node 0 is the only source of all 1-connections and node $N-1$ is the only source of all 0-connections, as shown in Fig. 12. Since node 0 is in the 0-subcube and all destinations of the 1-connections are in the 1-subcube, the path starts from the 0-subcube then goes to the 1-subcube. Thus all 1-connections share link $00\ldots0 \rightarrow 00\ldots1$ regardless of the order they were satisfied. Similarly, all 0-connections share link $11\ldots1 \rightarrow 11\ldots0$. Therefore, at least $N/2$ wavelengths are required for link $00\ldots0 \rightarrow 00\ldots1$ and link $11\ldots1 \rightarrow 11\ldots0$, which indicates $w_r \geq N/2 = 2^{n-1}$.

Hence, $N/2$ wavelengths are sufficient and necessary for a WDM hypercube with 2^n nodes to be wide-sense nonblocking for any multicast assignment under e-cube routing. That is, $w_w = N/2 = 2^{n-1}$. ■

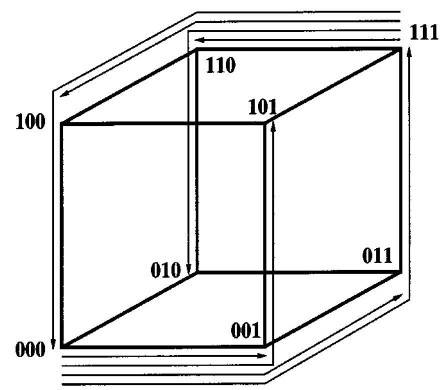

Fig. 12. A worst-case multicast assignment in a WDM hypercube under e-cube routing.

We now present the wavelength assignment algorithm for a hypercube. In this algorithm, two tables for used wavelengths T_0 and T_1 are maintained for 0-connections and 1-connections respectively. The available wavelengths are stored in table T_W.

```
Wavelength Assignment Algorithm in a
Hypercube
  Step 1) Initially, set T_0 and T_1 to
        empty, and assign wavelengths
        (w_0, w_1, ..., w_{N/2}) to T_W.
  Step 2) When a new connection is re-
        quested, if it is a 0-connection,
        assign to it a wavelength that is
        in T_W, but is not used in T_0, then
        add this wavelength to T_0. If it
        is a 1-connection, assign to it
        a wavelength that is in T_W, but
```

```
      is not used in T1, then add this
      wavelength to T1.
Step 3) When an existing connection is
      released, if it is a 0-connec-
      tion, delete the wavelength it
      used from T0. If it is a 1-con-
      nection, delete the wavelength it
      used from T1.
End.
```

VI. Conclusions

In this paper, we have studied multicast communication in a class of WDM optical networks with regular topologies. The networks studied include: linear arrays, rings, meshes, tori, and hypercubes. We have obtained the necessary and sufficient conditions on the minimum number of wavelengths required for a WDM network to be wide-sense nonblocking for arbitrary multicast assignment for each type of these networks under commonly used routing algorithms. Our results are summarized in Table I. An interesting future work is to generalize the approach developed in this paper to other routing algorithms and other network topologies including irregular networks. Another interesting issue is to determine nonblocking conditions for multicast assignments in WDM networks with wavelength converters and/or light splitting switches.

References

[1] B. Mukherjee, "WDM optical communication networks: Progress and challenges," *IEEE J. Select. Areas Commun.*, vol. 18, Oct. 2000.

[2] ——, *Optical Communication Networks*. New York: McGraw-Hill, 1997.

[3] "Special issue on multiwavelength optical networks," *J. Lightwave Technol.*, vol. 14, June 1996.

[4] "Special issue on optical networks," *IEEE J. Select. Areas Commun.*, vol. 14, June 1996.

[5] "Special issue on WDM networks," *J. High-Speed Networks*, vol. 4, no. 1/2, 1995.

[6] "Special issue on WDM networks," *Eur. Trans. Telecommun.*, vol. 11, no. 1, Jan./Feb. 2000.

[7] "Special issue on optical networks," *J. Lightwave Technol.*, vol. 18, Dec. 2000.

[8] L. H. Sahasrabuddhe and B. Mukherjee, "Light-trees: Optical multicasting for improved performance in wavelength-routed networks," *IEEE Communication*, vol. 37, no. 2, pp. 67–73, 1999.

[9] R. Malli, X. Zhang, and C. Qiao, "Benefit of multicasting in all-optical WDM networks," in *SPIE Conf. on All-Optical Networks*, vol. 3531, Nov. 1998, pp. 209–220.

[10] C. Qiao, M. Jeong, A. Guha, X. Zhang, and J. Wei, "WDM multicasting in IP over WDM networks," in *Int. Conf on Network Protocols (ICNP)*, Nov. 1999, pp. 89–96.

[11] E. Iannone, M. Listanti, and R. Sabella, "Multicasting in optical transport networks," *J. Opt. Commun.*, vol. 19, no. 3, pp. 90–98, Mar. 1998.

[12] M. Listanti, A. Cervelli, and R. Sabella, "Strategies and algorithms for routing both unicast and multicast paths in WDM networks," *Eur. Trans. Telecommun.*, vol. 11, no. 1, pp. 43–54, Jan./Feb. 2000.

[13] G. Sahin and M. Azizog-lu, "Routing and wavelength assignment in all-optical networks with multicast traffic," *Eur. Trans. Telecommun.*, vol. 11, no. 1, pp. 55–62, Jan./Feb. 2000.

[14] Z. Ortiz, G. N. Rouskas, and H. G. Perros, "Maximizing multicast throughput in WDM networks with tuning latencies using the virtual receiver concept," *Eur. Trans. Telecommun.*, vol. 11, no. 1, pp. 63–72, Jan./Feb. 2000.

[15] M. Ali and J. S. Deogun, "Cost-effective implementation of multicasting in wavelength-routed networks," *J. Lightwave Technol.*, vol. 18, pp. 1628–1638, Dec. 2000.

[16] X. Zhang, J. Y. Wei, and C. Qiao, "Constrained multicast routing in WDM networks with sparse light splitting," *J. Lightwave Technol.*, vol. 18, pp. 1917–1927, Dec. 2000.

[17] Y. Yang, J. Wang, and C. Qiao, "Nonblocking WDM multicast switching networks," *IEEE Trans. Parallel Distrib. Syst.*, vol. 11, pp. 1274–1287, Dec. 2000.

[18] C. Zhou and Y. Yang, "Embedding multicast in a class of optical WDM networks," in *Proc. Int. Conf. on Parallel and Distributed Processing Techniques and Applications*, vol. 2, Las Vegas, NV, July 1998, pp. 597–604.

[19] ——, "Multicast communication in a class of rearrangeable optical WDM networks," in *Proc. IEEE 5th Int. Conf. on Massively Parallel Processing Using Optical Interconnections (MPPOI '98)*, Las Vegas, NV, June 1998, pp. 68–75.

[20] C. Qiao and Y. Mei, "Off-line permutation embedding and scheduling in multiplexed optical networks with regular topologies," *IEEE/ACM Trans. Networking*, vol. 7, pp. 241–250, Apr. 1999.

[21] ——, "A comparative study of cost effective multiplexing approaches for online permutation embedding and scheduling in optical networks," *IEEE Trans. Parallel Distrib. Comput.*, vol. 61, pp. 1–17, Jan. 2001.

[22] P. Feldman, J. Friedman, and N. Pippenger, "Wide-sense nonblocking networks," *SIAM J. Discr. Math.*, vol. 1, no. 2, pp. 158–173, May 1988.

[23] Y. Yang and G. M. Masson, "Nonblocking broadcast switching networks," *IEEE Trans. Computers*, vol. 40, pp. 1005–1015, 1991.

[24] C. Lee and A. Y. Oruç, "Design of efficient and easily routable generalized connectors," *IEEE Trans. Commun.*, vol. 43, pp. 646–650, 1995.

[25] Y. Yang and J. Wang, "A new self-routing multicast network," *IEEE Trans. Parallel Distrib. Syst.*, vol. 10, pp. 1299–1316, Dec. 1999.

[26] J. Duato, S. Yalamanchili, and L. M. Ni, *Interconnection Networks: An Engineering Approach*. Los Alamitos, CA: IEEE Computer Society Press, 1997.

Chunling Zhou received the B.S. degree in Computer Science and Engineering from Tsinghua University, Beijing, China, and the M.S. degree in Computer Science from the University of Vermont, Burlington, Vermont. Her Master's thesis is on multicasting in optical WDM networks.

She is currently a senior software engineer in Artemis International Corporation, a leading vendor in Project Management software, in Exton, PA.

Yuanyuan Yang (S'91–M'92–SM'98) received the B.Eng. and M.S. degrees in Computer Science and Engineering from Tsinghua University, Beijing, China, in 1982 and 1984, respectively, and the M.S.E. and Ph.D. degrees in Computer Science from The Johns Hopkins University, Baltimore, MD, in 1989 and 1992, respectively.

She is an Associate Professor of Computer Engineering, with a joint appointment in Computer Science, at the State University of New York at Stony Brook, where she also directs the High-Performance Computing and Networking Research Laboratory. Her research interests include parallel and distributed computing and systems, high speed networks, optical networks, high performance computer architecture, and fault-tolerant computing. Her recent work focuses on the designs and performance analyses of the networks which can support efficient collective communications. She is the principal investigator of two on-going research projects funded by the National Science Foundation (NSF) and Army Research Office (ARO). She has published extensively in major journals, book chapters, and refereed conference proceedings related to these research areas, such as IEEE Transactions on Computers, IEEE Transactions on Parallel and Distributed Systems, IEEE Transactions on Communications, and *Journal of Parallel and Distributed Computing.* She holds three U.S. patents in the area of multicast communication networks, with four more patents pending. She has served as a chair or on the program/organizing committees of a number of international conferences and workshops in her areas of research. Dr. Yang is an Associate Editor for the IEEE Transactions on Parallel and Distributed Systems and a Guest Editor for *Optical Networks Magazine*. She is a member of Association for Computing Machinery, a member of IEEE Computer and Communications Societies, and a Distinguished Visitor of IEEE Computer Society.