

Nonblocking WDM Multicast Switching Networks

Yuanyuan Yang, *Senior Member, IEEE*,
Jianchao Wang, *Member, IEEE Computer Society*, and Chunming Qiao

Abstract—With ever increasing demands on bandwidth from emerging bandwidth-intensive applications, such as video conferencing, E-commerce, and video-on-demand services, there has been an acute need for very high bandwidth transport network facilities. Optical networks are a promising candidate for this type of applications. At the same time, many bandwidth-intensive applications require multicast services for efficiency purposes. Multicast has been extensively studied in the parallel processing and electronic networking community and has started to receive attention in the optical network community recently. In particular, as WDM (wavelength division multiplexing) networks emerge, supporting WDM multicast becomes increasingly attractive. In this paper, we consider efficient designs of multicast-capable WDM switching networks, which are significantly different and, hence, require nontrivial extensions from their electronic counterparts. We first discuss various multicast models in WDM networks and analyze the nonblocking multicast capacity and network cost under these models. We then propose two methods to construct nonblocking multistage WDM networks to reduce the network cost.

Index Terms—Wavelength division multiplexing (WDM), optical networks, multicast, nonblocking, switching networks.

1 INTRODUCTION

WITH ever increasing demands on bandwidth from emerging bandwidth-intensive applications, such as video conferencing, E-commerce, and video-on-demand services, there has been an acute need for very high bandwidth transport network facilities whose capabilities are far beyond what current high-speed (ATM) networks can provide. Optical networks are a promising candidate for this type of application as the potential bandwidth of a single optical fiber is nearly 50 THz, which is about four orders of magnitude higher than that of a copper wire. *Wavelength division multiplexing (WDM)* is a technique that exploits the huge bandwidth of optical fibers by dividing the bandwidth into multiple channels (or wavelengths) that can operate concurrently, each at the highest data rate achievable in electronics. WDM optical networks have attracted many researchers over the past few years [1], [2], [3], [4], [5], [6], [7], [8], [9] and the next generation of the Internet is expected to employ WDM-based optical backbones [9].

Multicast is the ability to transmit information from a single source node to multiple destination nodes. Many bandwidth-intensive applications, such as those mentioned earlier, require multicast services for efficiency purposes. Multicast has been extensively studied in the parallel

processing and electronic networking community (for example, see [11], [12], [13], [14], [15], [16], [17], [18], [19]). Recently, it has also received attention in the optical network community [5], [6], [7], [8]. In particular, as WDM networks emerge, supporting WDM multicast becomes attractive due to the following reasons: First, multicast can be supported more efficiently by utilizing the inherent light splitting capability of optical switches than by copying data in electronics. Second, in a traditional (electronic) multicast switching network which consists of one or more stages of switches, each source node can send the same message to multiple destination nodes concurrently, but each destination node can receive at most one message at a time. To deal with the problem of multiple multicast connections with overlapped destinations, a complex scheduling algorithm is necessary to avoid conflicts among multiple multicast connections. Adopting WDM provides a way to enable each source node to send different messages to multiple sets of destination nodes and each destination node to receive different messages through multiple connections at the same time.

In this paper, we will consider efficient designs of nonblocking multicast-capable WDM switching networks (or simply switches or networks). Such designs require nontrivial extensions from the existing designs of electronic multicast switching networks. This is because, in addition to the differences between a WDM switch and an electronic switch in terms of how multicast is supported as mentioned above, a major challenge in designing a WDM switch is how to keep the data in the optical domain, eliminating the need for costly conversions between optical and electronic signals (or so-called O/E/O conversions). To meet the challenge, it is required that either the wavelength on which the multicast data is sent and received has to be the same or all-optical wavelength converters, which are also expensive,

- Y. Yang is with the Department of Electrical and Computer Engineering, State University of New York at Stony Brook, Stony Brook, NY 11794. E-mail: yang@ece.sunysb.edu.
- J. Wang is with DataTreasury Corporation, 201 Old Country Rd., Melville, NY 11747. E-mail: Jianchao.Wang@datatreasury.com.
- C. Qiao is with the Department of Computer Science and Engineering, State University of New York at Buffalo, Buffalo, NY 14260. E-mail: qiao@computer.org.

Manuscript received 31 May 2000; accepted 22 Aug. 2000.
For information on obtaining reprints of this article, please send e-mail to: tpd@computer.org, and reference IEEECS Log Number 112222.

need to be used to convert (the signals on) an input wavelength to an output wavelength. This implies potential trade-offs between the multicasting performance of a WDM switch and the number of wavelength converters needed, along with other design parameters. Note that, in order to avoid O/E/O conversions at intermediate nodes, it is also desirable for a multicast WDM switch to be nonblocking as blocked multicast data will be dropped (lost) due to the lack of optical RAM (or buffer).

In this paper, we will first describe various multicast models in WDM networks which specify the wavelengths that can be used by the multicast source and destinations. Then, under these models, we analyze the nonblocking multicast capacity (to be defined later) and network cost (in terms of the number of crosspoints and number of wavelength converters required in a nonblocking crossbar-based design). Finally, we propose nonblocking multi-stage networks to reduce the network cost.

2 NONBLOCKING MULTICAST IN WDM NETWORKS

First, we will discuss some concepts related to multicast in electronic networks which will be extended to multicast in WDM networks later on. In a *multicast connection*, one source node sends a message simultaneously to multiple destination nodes (a unicast connection is a special case of a multicast connection). A set of multicast connections is called a *multicast assignment* if they do not involve same source node (at the input side of a network) and same destination node (at the output side of a network). In a *nonblocking network*, the set of multicast connections in a multicast assignment can be simultaneously realized without any conflict. A multicast assignment is called a *full-multicast-assignment* if no new multicast connection can be added to this multicast assignment to form a new multicast assignment; otherwise, it is called a *partial-multicast-assignment*. In general, a multicast assignment, full or partial, is called an *any-multicast-assignment*.

Notice that a full-multicast-assignment is in fact the maximal set of multicast connections which can be established simultaneously in a network without conflict. In an $N \times N$ network, this means that each destination node needs to be connected to exactly one of the N source nodes and, hence, the number of full-multicast-assignments is N^N . Since, in an any-multicast-assignment, each destination can choose not to be connected, in addition to being connected to any one of the N source nodes, the number of any-multicast-assignments is $(N + 1)^N$.

In the following, we will extend these concepts and observations to multicast in WDM networks.

2.1 Multicast Models in WDM Networks

In a WDM network, each node at the input (output) side of a network is connected to one input (output) port of the network via a fiber link that has k wavelengths, $\lambda_1, \lambda_2, \dots, \lambda_k$. Fig. 1 illustrates an $N \times N$ k -wavelength WDM network. To fully utilize the bandwidth of a fiber, each node at the input (output) side is normally equipped with a multiwavelength transmitter (receiver) array consisting of k fixed-tuned transmitters (receivers). A multicast




Fig. 1. An $N \times N$ WDM network with k wavelengths.

connection in such a network uses a wavelength at an input port and one or more wavelengths at a set of output ports. Accordingly, a node at the input (output) side can be involved in up to k multicast connections simultaneously, which, as mentioned earlier, is one feature that makes a WDM switch different from its electronic counterpart. However, the restrictions are that no two wavelengths at the same output port can be used in the same multicast connection and, in addition, a wavelength at an output port cannot be used in more than one multicast connections simultaneously. Note that, because of the difference and, in particular, the restrictions described above, the multicast capacity of the WDM network shown in Fig. 1 is not equivalent to that of an $Nk \times Nk$ electronic network when $k > 1$.

Another feature that makes WDM networks different from their electronic counterparts and from each other has to do with the input and output wavelengths used in a multicast connection. More specifically, when realizing a multicast connection in a WDM network, there are different ways to assign wavelengths to source and destinations, which we refer to as *multicast models*. The first way is to assign the same wavelength to the source node and all the destination nodes of a multicast connection and is referred to as the *Multicast with Same Wavelength (MSW)* model. The second way is to assign the same wavelength to all destination nodes of a multicast connection, but the source node may use a different wavelength and is referred to as the *Multicast with Same Destination Wavelength (MSDW)* model. The third way is that the source node and each of the destination nodes may use a different wavelength and is referred to as the *Multicast with Any Wavelength (MAW)* model.

Fig. 2 illustrates these three multicast models. Clearly, based on the above description, a multicast connection under the MSW model is always allowed under the MSDW model and a multicast connection under the MSDW is always allowed under the MAW model, but not vice versa. In other words, MAW is a stronger model than MSDW, which in turn is stronger than MSW. Since a traditional electronic switching network can be viewed as a 1-wavelength WDM network, multicast in such a network is a special case of multicast under the MSW model.

To realize a multicast connection in a WDM network, a *light splitter* and some *wavelength converters* may be needed. A light splitter splits a signal on a wavelength to a set of signals on the same wavelength and a wavelength converter converts a signal on one wavelength to another wavelength. While splitters are made of glass and, hence, inexpensive,




Fig. 2. Three multicast models in a WDM multicast network. (a) MSW. (b) MSDW. (c) MAW.




Fig. 3. Placement of wavelength converters and splitters. (a) MSDW. (b) MAW.

wavelength converters are not. Intuitively, the number of wavelength converters needed may vary from one multicast model to another. For example, under the MSW model, no wavelength converter is needed. Under the MSDW model, one wavelength converter is needed for each multicast connection, which can be placed just in front of the splitter, as shown in Fig. 3a. Under the MAW model, the number of wavelength converters needed for each multicast connection is no less than the fan-outs of the multicast connection since at least one wavelength converter is needed at each output of the splitter, as shown in Fig. 3b. Note that, as will be shown in Figs. 5, 6, and 7, *light combiners* are also needed near the output ports of a WDM network. These (light) combiners combine multiple inputs into one output signal. They differ from WDM multiplexers in that, with combiners, only one of the inputs (e.g., fibers) to a combiner may carry a signal at any given time and that each input may use any wavelength to carry a signal. Like splitters, combiners are also passive devices which are not expensive.

The above discussion suggests that the stronger a multicast model is, the more the number of wavelength converters it may require, which implies cost-performance trade-offs. At the end of this section, we will analyze the cost of WDM multicast networks, which includes the number of wavelength converters required under different multicast models. As will be shown, the MSDW and MAW model require the same number of wavelength converters.

2.2 Multicast Capacity under Different Models

To quantify the performance of a WDM multicast network, we define the *multicast capacity* of a WDM network under a given multicast model to be the number of multicast assignments that can be realized in the network.

In the following, we analyze the multicast capacities of an $N \times N$ k -wavelength WDM multicast network under three different models, respectively. Clearly, the larger the multicast capacity, the better the performance (or, in other words, the stronger the multicast model). We will

determine the number of full-multicast-assignments, as well as that of any-multicast-assignments, and consider the three models in the order of increasing complexity involved in the analysis.

We start with the MSW model, which is the simplest to analyze.

Lemma 1. *For an $N \times N$ k -wavelength WDM multicast network under the MSW model, the multicast capacity is N^{Nk} for full-multicast-assignments and $(N+1)^{Nk}$ for any-multicast-assignments.*

Proof. Note that, under the MSW model, the same wavelength has to be used by a multicast connection at both the input and output sides. This is in addition to the two restrictions on how wavelengths can be used by multicast connections, as mentioned earlier in Section 2.1.

First, we consider the case of full-multicast-assignments. Since a wavelength at an output port can pair with the same wavelength at any one of the N input ports to form a multicast connection under the MSW model, it can be involved in N different multicast assignments. Given that each of the Nk wavelengths at the output side can be involved in N multicast assignments independently from any other wavelength at the output side under the MSW model, there are N^{Nk} different full-multicast-assignments.

Next, let's consider the case of any-multicast-assignments. For any given wavelength at the output side, there is one more possibility, which is not to pair with the any wavelength at the input side, than in the case of full-multicast-assignments. Therefore, there are $(N+1)^{Nk}$ possible any-multicast-assignments under the MSW model. \square

Based on the above results, it is obvious that an $N \times N$ k -wavelength WDM network under the MSW model is not the same as an $Nk \times Nk$ electronic network when $k > 1$ (whose multicast capability is $(Nk)^{Nk}$ for full-multicast

assignments and $(Nk + 1)^{Nk}$ for any-multicast assignments). As will be shown, this is true for the other two models as well. The following lemma states the results for the strongest model (MAW), which happens to be easier to analyze than the MSDW model.

Lemma 2. For an $N \times N$ k -wavelength WDM multicast network under the MAW model, the multicast capacity is

$$[P(Nk, k)]^N$$

for full-multicast-assignments and

$$\left[\sum_{j=0}^k P(Nk, k-j) \binom{k}{j} \right]^N$$

for any-multicast-assignments, where $P(x, i)$ is defined as

$$P(x, i) = x(x-1) \cdots (x-i+1)$$

and $\binom{k}{j}$ is a binomial coefficient.

Proof. Under the MAW model, no additional restrictions on how wavelengths are assigned to a multicast connection exist.

Consider full-multicast-assignments first. For the k wavelengths at an output port, λ_1 can pair with any one of the Nk wavelengths at the input side to form a multicast connection, λ_2 can pair with any one of the remaining $Nk - 1$ wavelengths at the input side to form another multicast connection, etc. In general, λ_i , where $1 \leq i \leq k$, can pair with any one of the remaining $Nk - i + 1$ wavelengths at the input side to form a multicast connection. Hence, the k wavelengths at the output port can be involved in $Nk(Nk - 1) \cdots (Nk - i + 1) \cdots (Nk - k + 1) = P(Nk, k)$ different multicast-assignments. Given that the wavelengths at each of the N output ports can be involved in different multicast assignments independently from each other, there are $[P(Nk, k)]^N$ possible full-multicast-assignments under the MAW model.

For any-multicast-assignments, j wavelengths at an output port, where $(0 \leq j \leq k)$, may choose not to pair with any wavelength at the input side, while the remaining $k - j$ wavelengths at the output port can be involved in different multicast assignments as described above. This will result in $\sum_{j=0}^k P(Nk, k-j) \binom{k}{j}$ different multicast assignments per output port and, hence, the total number of any-multicast-assignments under the MAW model is $\left[\sum_{j=0}^k P(Nk, k-j) \binom{k}{j} \right]^N$. \square

Finally, the MSDW model is analyzed and we have

Lemma 3. For an $N \times N$ k -wavelength WDM multicast network under the MSDW model, the multicast capacity is

$$\sum_{1 \leq j_1, \dots, j_k \leq N} P(Nk, \sum_{i=1}^k j_i) \prod_{i=1}^k S(N, j_i)$$

for full-multicast-assignments and

$$\sum_{\substack{1 \leq j_1 \leq N-l_1, \dots, 1 \leq j_k \leq N-l_k \\ 0 \leq l_1, \dots, l_k \leq N}} P\left(Nk, \sum_{i=1}^k j_i\right) \prod_{i=1}^k \binom{N}{l_i} S(N-l_i, j_i)$$

for any-multicast-assignments, where $S(N, j)$ is the Stirling number of the second kind [20], which is the number of ways to divide N elements to j disjoint groups.

The proof of this lemma is provided in the Appendix.

As a sanity check, when $k = 1$, a multicast network under either the MSW, MSDW, or MAW will be reduced to a traditional multicast network and its multicast capacity becomes N^N and $(N + 1)^N$ for full-multicast-assignments and any-multicast-assignments, respectively. We can verify that this is true from the above three lemmas. Verifications from Lemma 1 or 2 are trivial. For Lemma 3, when $k = 1$, we have

$$\sum_{1 \leq j \leq N} P(N, j) S(N, j) = N^N$$

and

$$\begin{aligned} & \sum_{1 \leq j \leq N-l} \sum_{0 \leq l \leq N} P(N, j) \binom{N}{l} S(N-l, j) \\ &= \sum_{0 \leq l \leq N} \binom{N}{l} \sum_{1 \leq j \leq N-l} P(N, j) S(N-l, j) \\ &= \sum_{0 \leq l \leq N} \binom{N}{l} N^{N-l} = (N+1)^N. \end{aligned}$$

Also, it can be verified that the multicast capacity of a WDM network of the same size, but under different models, increases in the order of MSW, MSDW, and MAW, which matches our intuition.

2.3 Network Cost under Different Models

In this paper, we will characterize the cost of a WDM multicast network by the number of crosspoints in addition to the number of wavelength converters mentioned earlier. The number of crosspoints is used as a representative measure of the hardware complexity of various switching circuits, such as the number of *semiconductor optical amplifier* (SOA) gates used to turn on and off output light beams in a WDM switch or the number of mirrors used to steer light beams in optical Micro-Electro-Mechanical Systems (MEMS) [10]. Though not a direct measure, the number of crosspoints may also be used to project the crosstalk and power loss inside a WDM switch.

In the rest of this section, we analyze the network cost of a nonblocking WDM network under the three multicast models mentioned earlier, assuming that the network is a crossbar-like switching fabric which may be implemented using light splitters and combiners as well as SOA gates. These SOA gates are active devices which, when turned on, permit light signals to go through and, when turned off, block light signals.

Note that a WDM network under a given multicast model is nonblocking if every any-multicast-assignment under the multicast model can be realized. Since different multicast models have a different number of any-multicast-assignments, one should not use the cost only to determine

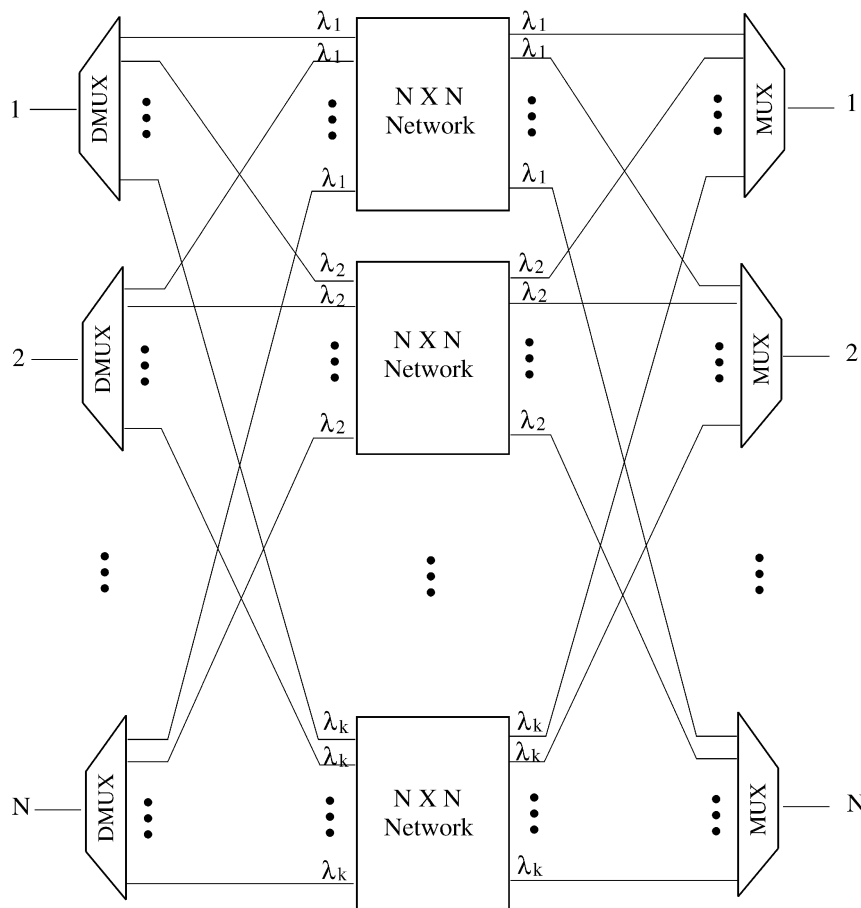


Fig. 4. An $N \times N$ k -wavelength WDM network consisting of k parallel $N \times N$ 1-wavelength networks.

which multicast model is better. In fact, it is reasonable to expect a nonblocking network under a weaker multicast model, such as the MSW model, to have a lower cost than a network under a stronger multicast model, such as the MAW model. However, as we will see in the next section, such cost analysis will lead us to more efficient nonblocking *multistage* networks where the number of crosspoints can be greatly reduced. The results from such cost analysis are also useful for making cost-performance trade-offs in designs (when combined with the results on the multicast capacity obtained so far) as will be discussed later.

2.3.1 The Number of Crosspoints

For an $N \times N$ k -wavelength WDM multicast network under the MSW model, the number of crosspoints is kN^2 (excluding those which might be involved in the wavelength multiplexers and demultiplexers at the input and output sides, respectively). In fact, since, in any multicast connection, the source and destinations must use the same wavelength, the network is equivalent to k parallel multicast-capable $N \times N$ 1-wavelength networks (i.e., space switches) coupled together as shown in Fig. 4. Each of these $N \times N$ 1-wavelength networks may be implemented as shown in Fig. 5.

In an $N \times N$ k -wavelength WDM multicast network under the MSDW model, the number of crosspoints is k^2N^2 since any of the Nk wavelengths at the input side may be connected to any of the Nk wavelengths at the output side.

An example when $N = 3$ and $k = 2$ is shown in Fig. 6. For the same reason, the number of crosspoints in a WDM multicast network under the MAW model is also k^2N^2 and a corresponding example is shown in Fig. 7.

2.3.2 The Number of Wavelength Converters

Clearly, for an $N \times N$ k -wavelength WDM multicast network under the MSW model, no converter is needed. However, for a WDM multicast network of the same size under either the MSDW or MAW model, Nk wavelength converters are required. As shown in Fig. 3a, under the MSDW model, a converter can be placed for each of the Nk wavelengths at the network input side to convert a source wavelength to a possibly different wavelength, which is then split to reach multiple output ports. Under the MAW model shown in Fig. 3b, a converter needs to be placed for each of the Nk wavelengths at the output side such that, after a source wavelength is split into multiple outputs, the wavelength can be converted to a different wavelength at each output port independently of other output ports (this placement also works for the MSDW model).

2.4 Comparison of Different Models

We summarize the results obtained in this section in Table 1. We list the multicast capacities, the numbers of crosspoints, and the numbers of wavelength converters required for nonblocking WDM multicast networks under different

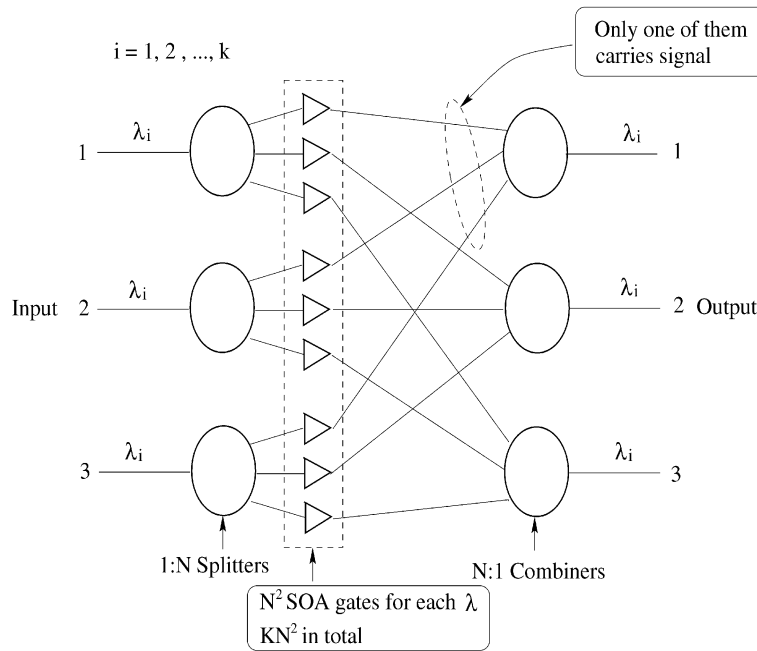


Fig. 5. A possible implementation of an $N \times N$ 1-wavelength multicast network (when $N = 3$).

models. Note that, while cost-performance trade-offs exist between the MSW and the MAW models, it is obvious that the MSDW model is not desirable because its cost is the same as that of the MAW model, but its performance is inferior to that of the MAW model.

3 MULTISTAGE WDM SWITCHING NETWORKS

In this section, we will investigate how to use a multistage network to reduce the number of crosspoints. We will first consider an $N \times N$ three-stage network, which has r switching modules of size $n \times m$ in the input stage, m switching modules of size $r \times r$ in the middle stage, and r switching modules of size $m \times n$ in the output stage with $N = nr$ and $m \geq n$, as shown in Fig. 8. There is exactly one link between every two switching modules in two consecutive stages and each switching module is assumed to be multicast-capable. In general, a network can have any odd number of stages and be built in a recursive fashion from these switching modules, which are in fact regarded as networks of a smaller size.

Extensive research has been done for such multistage networks [11], [12], [13], [14], [15], [16] in the electronic domain. A critical design issue is how to ensure that the network is nonblocking and, at the same time, minimize the number of crosspoints in the network, which requires that the number of middle stage switches m be minimized. Yang and Masson [14] gave a sufficient condition for a nonblocking multicast network, $m \geq 3(n - 1) \frac{\log r}{\log \log r}$, which yields the best available design for this type of multicast networks. In [16], this condition was proven to be necessary under several commonly used routing strategies.

In the rest of the section, we will extend the results to nonblocking WDM multicast networks under the MSW and MAW models, as well as the MSDW model for the sake of completeness. It is worth mentioning that every link in the

multistage network is now a WDM link with k wavelengths. In addition, the network cost for a WDM switching network also includes the number of wavelength converters. Most importantly, while the switching modules in an electronic multistage network are homogeneous (i.e., all based on crossbars), those in a WDM multistage network can be heterogeneous (i.e., under different models such as MSW, MSDW, and MAW), which presents a major challenge in the analysis.

3.1 Terminologies

There are many ways to construct a WDM multicast network under the MSW, MSDW, or MAW model using switching modules which can also be under different models. However, based on the results from the last section, the MSW model has the lowest network cost and the smallest multicast capacity, while the MAW model has the highest cost and the largest multicast capacity. Hence, it becomes natural for us to consider the following two extreme methods to construct a three-stage WDM multicast network under either the MSW, MSDW, or MAW model. Using what we call the *MSW-dominant* construction method, the switching modules in both the input stage and middle stage adopt the MSW model and the switching modules in the output stage adopt the MSW, MSDW, or MAW model, respectively, which determines the model the network as a whole will be under (see Fig. 9a). On the other hand, using what we call the *MAW-dominant* construction method, the switching modules in both the input stage and middle stage adopt the MAW model and the switching modules in the output stage adopt the MSW, MSDW, or MAW models, respectively (see Fig. 6b). The reason for considering the MAW-dominant construction is that, as shown in Fig. 10 (where wavelength converters are not shown), a multicast connection may be blocked at a middle-stage MSW switch due to its restricted wavelength

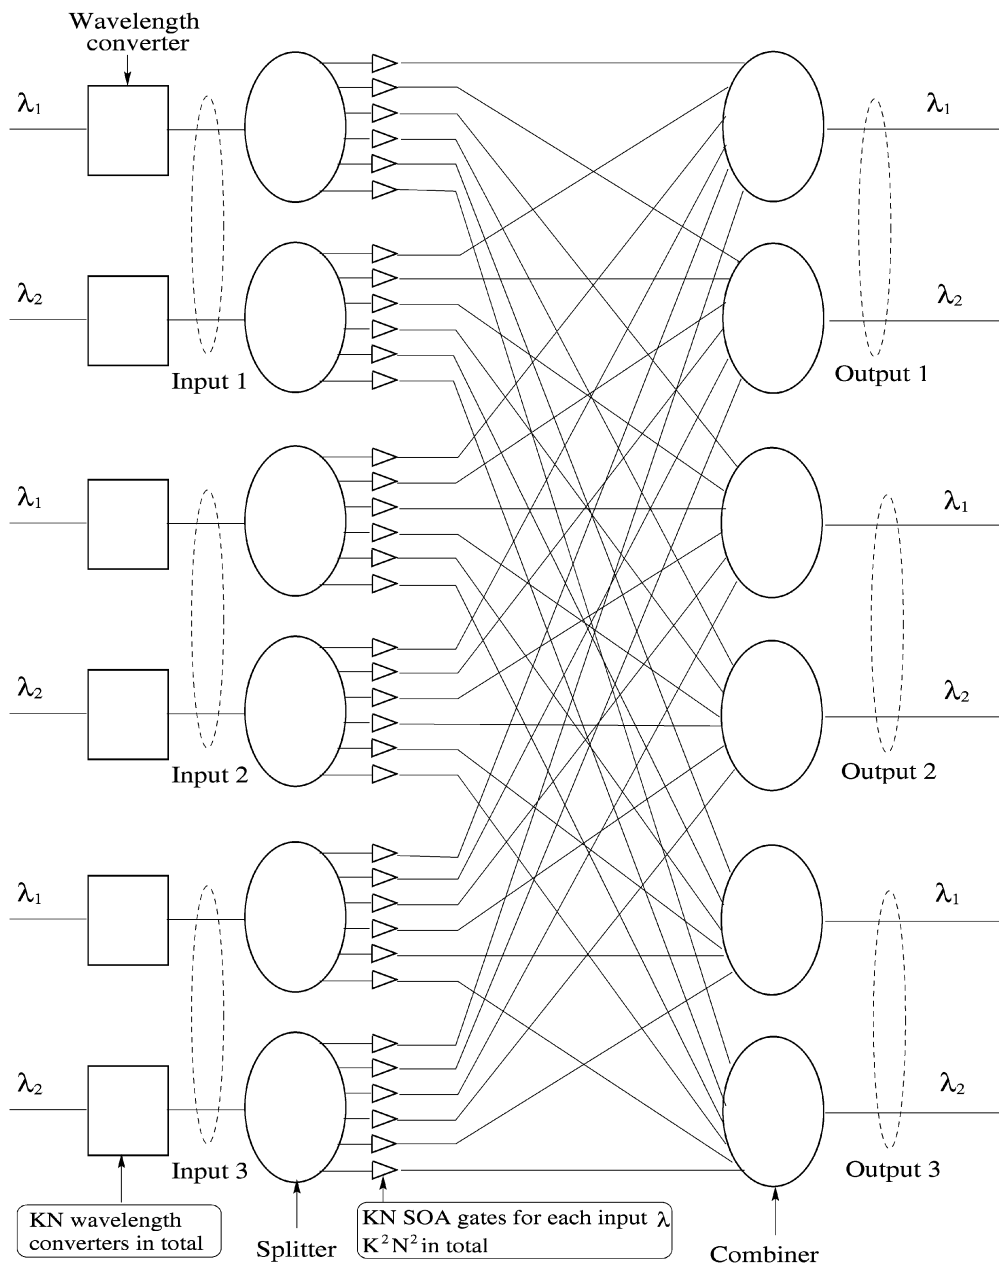


Fig. 6. An example of an $N \times N$ k -wavelength WDM multicast network under the MSDW model (when $N = 3$ and $k = 2$).

assignment requirement, while if MAW switches are used in the first two stages, as in the MAW-dominant construction, such a blocking will not occur.

We can also extend these two methods when constructing a WDM multicast network with an arbitrary number of stages as follows: In the MSW-dominant construction, all switching modules except those in the last stage adopt the MSW model and, in the MAW-dominant construction, all switching modules except those in the last stage adopt the MAW model.

Note that an $N \times N$ k -wavelength nonblocking multistage WDM network under a given model will have the same multicast capacity as a crossbar-based $N \times N$ k -wavelength WDM network under the same model and, hence, only the nonblocking conditions and network cost for three-stage WDM multicast networks under the MSW, MSDW, and

MAW models based on these two constructions will be analyzed in the next section.

The following notations will be used in the analysis, which will be useful to represent a multicast connection in terms of the set of output stage switches reachable from an input wavelength. Let $O = \{1, 2, \dots, r\}$ denote the set of all output stage switches numbered from one to r . Since there can be up to k multicast connections from a middle stage switch $j \in \{1, 2, \dots, m\}$ to an output stage switch $p \in O$ (one on each wavelength), we shall use M_j to represent the *destination multiset* (whose base set is O), where p may appear more than one time if more than one multicast connection goes from j to p . The number of times p appears in M_j or the number of multicast connections from j to p is called the *multiplicity* of p in multiset M_j .

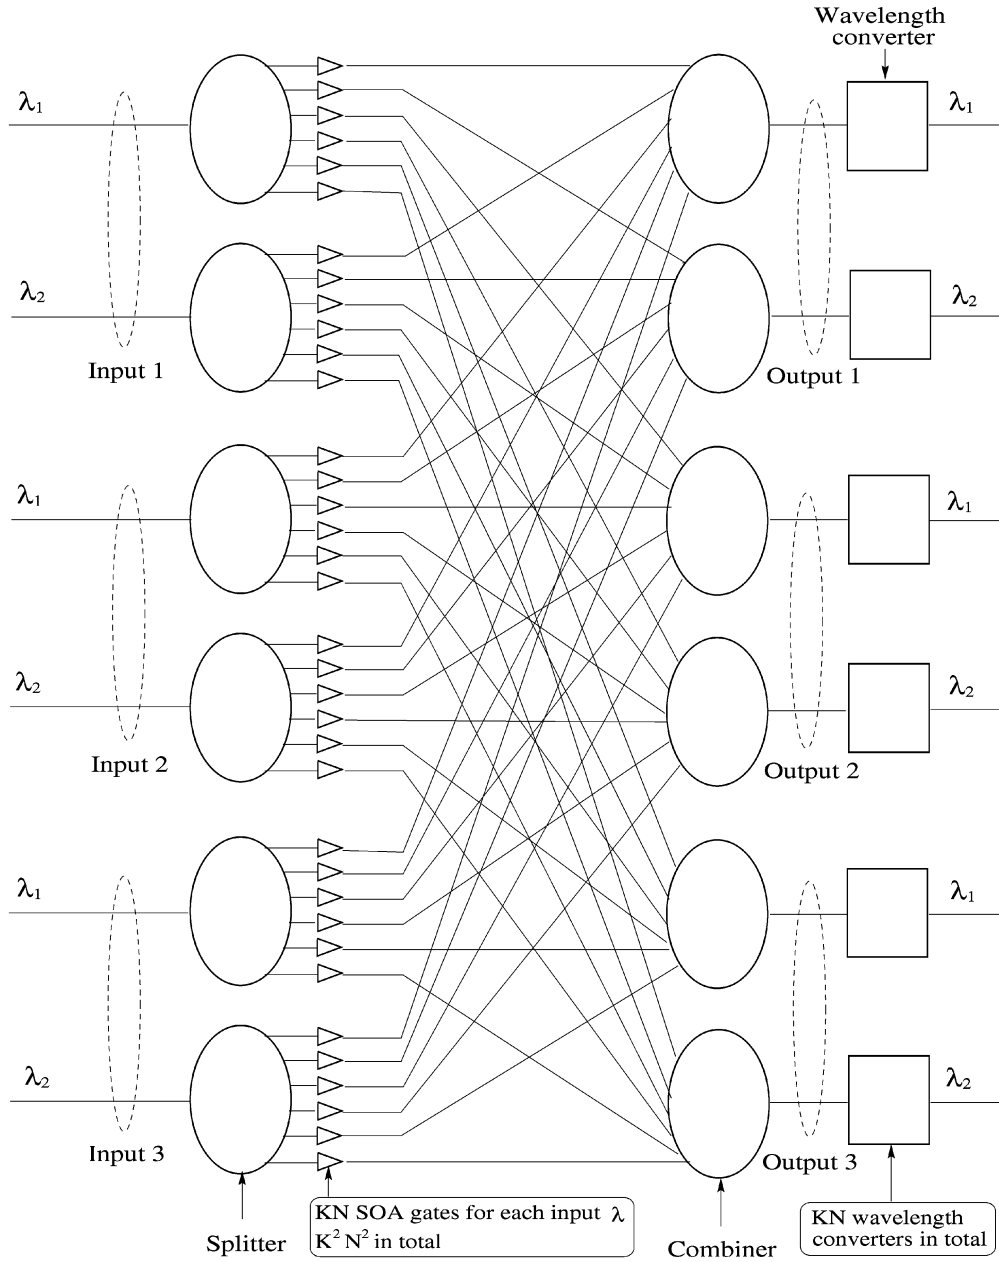


Fig. 7. An example of an $N \times N$ k -wavelength WDM multicast network under the MAW model (when $N = 3$ and $k = 2$).

In addition, denote an input wavelength λ_{l_i} at an input port i by (i, λ_{l_i}) , where $i \in \{1, 2, \dots, nr\}$ and $1 \leq l_i \leq k$. Assume that a WDM network is currently providing a set of multicast connections from some input wavelengths to some output wavelengths. We shall refer to the set of middle stage switches that can be reached by the signals carried on the input wavelength (i, λ_{l_i}) as the *available middle switches* for that input wavelength.

3.2 Nonblocking Condition for the MSW-Dominant Construction

For a WDM multicast switching network under either the MSW, MSDW, or MAW model and adopting the MSW-dominant construction method (i.e., by implementing the switching modules in the first two stages under MSW), the nonblocking condition can be reduced to the case of an electronic multicast switching network where only one

wavelength is used. In fact, for a multicast connection with source input wavelength (i, λ_{l_i}) , it can be realized using only wavelength λ_{l_i} in the first two stages and then realized in the third stage under either the MSW, MSDW, or MAW model. Therefore, we can simply ignore other wavelengths and consider multicast routing using only wavelength λ_{l_i} . Consequently, the results of nonblocking condition of a traditional multicast network can be directly applied to a WDM multicast network under the MSW model.

When we limit the network to using only one wavelength, the destination multiset M_j is no longer a multiset but becomes an ordinary set (i.e., each element can only appear once). In this case, we simply call it a *destination set* and can directly use the results in [14]. More specifically, by applying the routing strategy that limits each multicast connection to using no more than x middle stage switches,

TABLE 1
Comparison of WDM Multicast Networks Under Different Models

Model	Capacity (full-multicast-assignments)
MSW	N^{Nk}
MSDW	$\sum_{1 \leq j_1, \dots, j_k \leq N} P(Nk, \sum_{i=1}^k j_i) \prod_{i=1}^k S(N, j_i)$
MAW	$[P(Nk, k)]^N$
Model	Capacity (any-multicast-assignments)
MSW	$(N+1)^{Nk}$
MSDW	$\sum_{\substack{1 \leq j_1 \leq N-l_1, \dots, 1 \leq j_k \leq N-l_k \\ 0 \leq l_1, \dots, l_k \leq N}} P(Nk, \sum_{i=1}^k j_i) \prod_{i=1}^k \binom{N}{l_i} S(N-l_i, j_i)$
MAW	$\left[\sum_{j=0}^k P(Nk, k-j) \binom{k}{j} \right]^N$
Model	# Crosspoints
MSW	kN^2
MSDW	k^2N^2
MAW	k^2N^2
Model	# Converters
MSW	0
MSDW	kN
MAW	kN

where x is a number which we can use to optimize the network cost, as in [14], we can obtain the following lemma.

Lemma 4. *A new multicast connection request with fanout r can be satisfied using some x ($x \geq 1$) middle stage switches, say, i_1, \dots, i_x , from among the available middle switches if and only if the intersection of the destination sets of these x middle stage switches is empty, i.e., $\bigcap_{j=1}^x M_{i_j} = \phi$.*

Furthermore, by considering all possible multicast connections from the same input stage switch, we can establish the following theorem.

Theorem 1. *A WDM multicast network adopting the MSW-dominant construction is nonblocking if*

$$m > \min_{1 \leq x \leq \min\{n-1, r\}} \left\{ (n-1) \left(x + r^{\frac{1}{x}} \right) \right\}. \quad (1)$$

The proofs for the above lemma and theorem are similar to those in [14] and, hence, are omitted. Note that the minimum value of m can be obtained from (1) by minimizing the righthand side expression over all possible values of x .

3.3 Nonblocking Condition for the MAW-Dominant Construction

In this subsection, we consider WDM multicast networks under either the MSW, MSDW, or MAW model and adopting the MAW-dominant construction method (i.e., switching modules in the first two stages are implemented as MAW). Since this construction method provides more choices for establishing multicast connections in the first two stages, it is interesting to know if it can lead to a better nonblocking condition (i.e., a smaller m). In this construction, we must consider destination multiset M_j

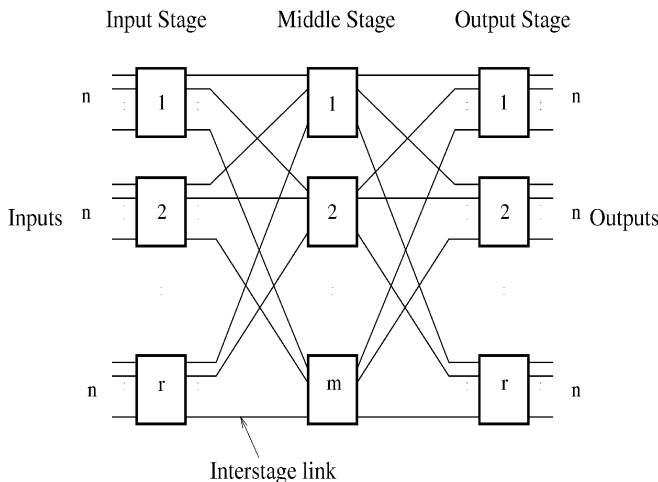


Fig. 8. A three-stage switching network.

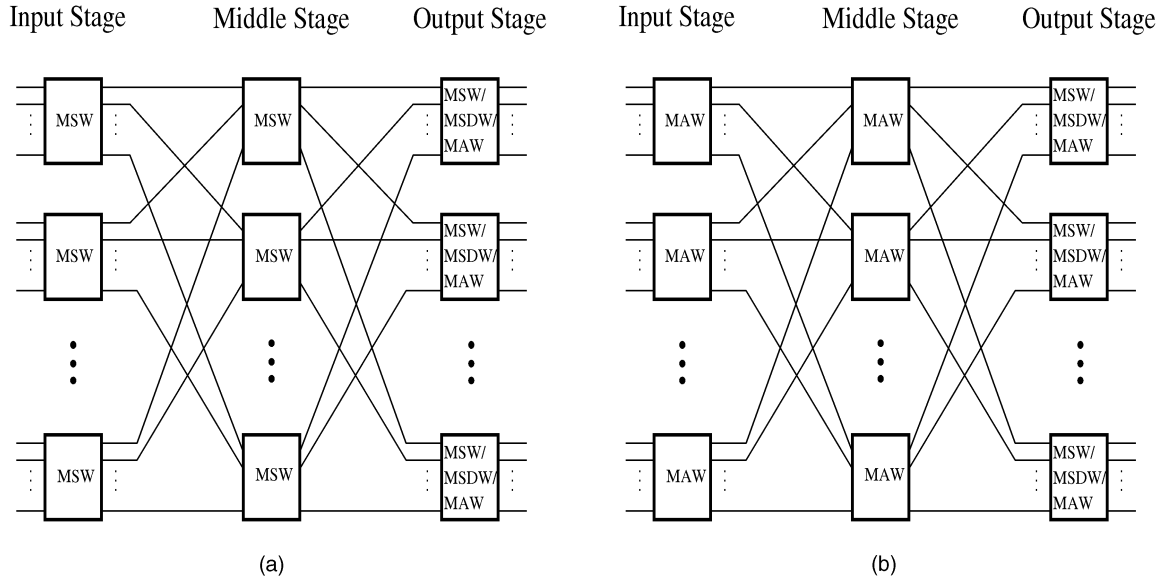


Fig. 9. Two constructions of multistage networks. (a) MSW-dominant construction. (b) MAW-dominant construction.

having a multiplicity larger than one. More specifically, denote the multiset M_j as

$$M_j = \{1^{i_{j1}}, 2^{i_{j2}}, \dots, r^{i_{jr}}\}, \quad (2)$$

where $0 \leq i_{j1}, i_{j2}, \dots, i_{jr} \leq k$ are the multiplicities of elements $1, 2, \dots, r$, respectively. Notice that if every multiplicity is less than k , i.e., $0 \leq i_{j1}, i_{j2}, \dots, i_{jr} \leq k - 1$, then the maximal multicast connection that can be realized through middle stage switch j without interfering with any existing connections is $\{1, 2, \dots, r\}$ (in terms of the set of output stage switches reachable from middle stage switch j). In general, the maximal multicast connection that can go through middle stage switch j is $\{p | 0 \leq i_{jp} < k \text{ for } 1 \leq p \leq r\}$. Now, the question is what is the maximal multicast connection that can go through two middle stage switches j and h ? We thus define the intersection of multisets as follows:

$$M_j \cap M_h = \{1^{\min\{i_{j1}, i_{h1}\}}, 2^{\min\{i_{j2}, i_{h2}\}}, \dots, r^{\min\{i_{jr}, i_{hr}\}}\}. \quad (3)$$

Then, the maximal multicast connection that can go through two middle stage switches j and h is the maximal multicast connection that can go through a middle stage switch with its destination multiset equal to multiset $M_j \cap M_h$.

From the point of realizing a multicast connection, which is characterized as an ordinary set, we can see that those

elements in M_j with multiplicity k cannot be used. Accordingly, we define the cardinality of M_j as

$$|M_j| = |\{i_{jp} | i_{jp} = k \text{ for } 1 \leq p \leq r\}| \quad (4)$$

and the null of M_j as

$$M_j = \phi \text{ iff } |M_j| = 0. \quad (5)$$

It can be easily verified that Lemma 4 still holds when we consider M_j as a multiset and use the definitions of intersection, cardinality, and null in this section.

Now, we are in a position to extend the result of Theorem 1 to the case of the destination multiset being a multiset with multiplicity no more than k , as defined in (2), and the operations of such multisets are defined in (3), (4), and (5).

Lemma 5. For any x , $1 \leq x \leq \min\{n - 1, r\}$, let m' be the maximum number of middle stage switches whose destination multisets satisfy that their multiplicities are no more than k , that there are at most $(nk - 1) 1s, (nk - 1) 2s, \dots, (nk - 1) rs$ distributed among the m' destination multisets, and that the intersection of any x of the destination multisets is not empty. Then, we have

$$m' \leq (n - 1)r^{\frac{1}{x}}.$$




Fig. 10. An example showing the difference between having an MSW switch and an MAW switch at the middle stage: A multicast connection involving λ_1 of input 1 and two outputs (λ_1 of output 1 and λ_2 of output 2) will be blocked if the switch is MSW, but not if the switch is MAW.

The proof of this lemma is an extension of the results in [14] and is provided in the Appendix.

The immediate corollary of the above lemma is:

Corollary 1. *In a WDM multicast network adopting the MAW-dominant construction, for a new multicast connection request with fanout r' , $1 \leq r' \leq r$, if there exist more than $(n-1)r^{\frac{1}{x}}$, $1 \leq x \leq \min\{n-1, r'\}$, available middle switches for this connection request, then there will always exist x middle stage switches through which this new connection request can be satisfied.*

Accordingly, we can establish:

Theorem 2. *A WDM multicast network adopting the MAW-dominant construction is nonblocking for any multicast assignments if*

$$m > \min_{1 \leq x \leq \min\{n-1, r\}} \left\{ \left\lfloor \left(n - \frac{1}{k} \right) x \right\rfloor + (n-1)r^{\frac{1}{x}} \right\}. \quad (6)$$

Proof. Recall that the routing strategy for realizing multicast connections is to realize each multicast connection by using no more than x middle stage switches. By Corollary 1, if we have more than $(n-1)r^{\frac{1}{x}}$ available middle switches for a new multicast connection request, we can always choose x middle stage switches to realize this connection request. Now, there may be at most $nk-1$ other input wavelengths, each of which is used for a multicast connection. By the routing strategy, each of them is connected to no more than x wavelengths on different outputs of this input stage switch and then connected to no more than x middle stage switches. Notice that, unlike a traditional network, in a WDM network, two wavelengths at different input ports can be connected to two wavelengths of the same output of an input stage switch and then connected to the same middle stage switch. Now, the question is what is the number of middle stage switches which are not available for a new multicast connection in the worst case? Since each fiber link has k wavelengths and if all the k wavelengths of the fiber link connecting the current input stage switch to a middle stage switch are used by k existing multicast connections, this middle stage switch is not available. Therefore, we can have at most

$$\left\lfloor \frac{(nk-1)x}{k} \right\rfloor = \left\lfloor \left(n - \frac{1}{k} \right) x \right\rfloor$$

middle stage switches which are not available for a new multicast connection request. Thus, the total number of middle stage switches required, m , is greater than the number of unavailable middle switches in the worst case plus the maximum number of available middle switches needed to realize a multicast connection. The minimum value for m is obtained from (6) by minimizing the right hand side expression for all possible values of x . \square

Although the minimum values of m derived from Theorems 1 and 2 are only sufficient for nonblocking in the MSW-dominant and MAW-dominant constructions, respectively, by using an approach similar to that used in [16] one can also obtain matching values of m that are necessary.

3.4 Network Cost and Comparison

From Theorems 1 and 2, one may deduce that number of middle stage switches required for nonblocking in the MAW-dominant construction would be slightly larger than that in the MSW-dominant construction. In addition, since a switching module under the MAW model has more crosspoints than a switching module under the MSW model, a multistage WDM network adopting the MAW-dominant construction under any one of the three multicast models will have more crosspoints than that adopting the MSW-dominant construction under the same multicast model (especially in the first two stages). Similarly, the former requires more wavelength converters. Meanwhile, given a multicast model, the multicast capacity does not change whether the construction is MAW-dominant or MSW-dominant. Therefore, we conclude that, for a multi-stage WDM multicast network under either the MSW, MSDW, or MAW model, the MSW-dominant construction is a better choice.

In the following, we will discuss multistage WDM networks adopting the MSW-dominant construction. Notice that the condition shown in Theorem 1 can be reduced to $m \geq 3(n-1) \frac{\log r}{\log \log r}$ by letting $x = \frac{2 \log r}{\log \log r}$. Now, we are in position to calculate the network cost under different models. More specifically, given that the switching modules in the first two stages are under the MSW model and the switching modules in the last stage are under either the MSW, MSDW, or MAW, based on the results from the last section and letting $n = r = N^{\frac{1}{2}}$, we know that the number of crosspoints for an $N \times N$ network under the MSW model is

$$\begin{aligned} r \cdot knm + m \cdot kr^2 + r \cdot kmn &= kmr(2n+r) \\ &= O\left(kN^{\frac{1}{2}} \frac{\log N}{\log \log N}\right) \end{aligned}$$

and that the number of crosspoints under the MSDW or MAW model is

$$\begin{aligned} r \cdot knm + m \cdot kr^2 + r \cdot k^2mn &= kmr[(k+1)n+r] \\ &= O\left(k^2N^{\frac{1}{2}} \frac{\log N}{\log \log N}\right). \end{aligned}$$

For the number of wavelength converters, we only need to consider the switching modules in the output stage. Therefore, the numbers of wavelength converters are 0, $r \cdot mk = O\left(kN \frac{\log N}{\log \log N}\right)$, and $r \cdot nk = kN$ for a network under the MSW, MSDW, and MAW model, respectively.

Note that the MSDW model results in more wavelength converters than the MAW model. This is because, under the MSDW model, wavelength converters are assumed to be placed in front of an output stage switch (with m links), while, under the MAW model, they are placed on the other side of the output stage switch (with only n links, where $n < m$). Though one may reduce the number of wavelength converters under the MSDW model by using a better wavelength converter placement strategy (e.g., by placing the wavelength converters in the middle of the $m \times n$ switching module, it will still require the same number of wavelength converters as that required by the MAW model.

TABLE 2

Cost Comparison of Multistage and Crossbar WDM Multicast Networks under Different Models (CB-Crossbar, MS-Multistage)

Model	# Crosspoints	# Converters
MSW/CB	kN^2	0
MSW/MS	$O(kN^{\frac{1}{2}} \frac{\log N}{\log \log N})$	0
MSDW/CB	k^2N^2	kN
MSDW/MS	$O(k^2N^{\frac{1}{2}} \frac{\log N}{\log \log N})$	$O(kN \frac{\log N}{\log \log N})$
MAW/CB	k^2N^2	kN
MAW/MS	$O(k^2N^{\frac{1}{2}} \frac{\log N}{\log \log N})$	kN

We summarize the network costs of WDM networks under the three models based on either crossbar (CB) or multistage (MS) construction in Table 2. First, by comparing a multistage network with a crossbar network, one can see that, under any of the three models, a multistage network has significantly fewer crosspoints. The number of wavelength converters is the same except, under the MSDW model, a multistage network may have slightly more wavelength converters. Second, by comparing multistage networks under different models, it is clear that, since the MSDW model has at least the same network cost as the MAW model, it is less cost-effective than the MAW model in the multistage construction, just as in the case for crossbar-based networks. In addition, similarly to the case for crossbar networks, a multistage multicast network under the MSW has a lower network cost than a network under the MAW model, while the latter provides a larger multicast capacity and, hence, the two represent cost-performance trade-offs in the design.

4 CONCLUSIONS

In this paper, we have proposed efficient designs for WDM multicast switching networks. We have considered three different ways to assign wavelengths to a multicast connection (i.e., multicast models) in a WDM network, namely, Multicast with the Same Wavelength (MSW), Multicast with the Same Destination Wavelength (MSDW), and Multicast with Any Wavelength (MAW). We have analyzed the nonblocking performance (in terms of the multicast capacity) and network cost (in terms of the number of crosspoints and wavelength converters needed) for a crossbar-based WDM network under these multicast models. Finally, we have proposed two methods, namely, the MSW-dominant and MAW-dominant method, to construct a nonblocking multistage WDM network with reduced network cost and compared various designs of WDM multicast switching networks under the MSW, MSDW, and MAW models. Our results have indicated that the MSDW model is not cost-effective in either crossbars or multistage networks, while the MSW and MAW models

represent cost-performance trade-offs in the design of both crossbars and multistage networks. In addition, the MSW-dominant construction method is more suitable for designing multistage WDM multicast networks than the MAW-dominant construction method.

APPENDIX

In this appendix, we provide the proofs for Lemmas 3 and 5, respectively.

Proof for Lemma 3. Under the MSDW model, all destination nodes in a multicast connection use the same wavelength, which may not be the same as the source wavelength.

We consider full-multicast-assignments first. Among Nk wavelengths at the network output side, there are exactly N wavelengths with λ_i for $1 \leq i \leq k$, each of which belongs to one of N different output ports. For an integer j_i ($1 \leq j_i \leq N$), we divide N output wavelengths of λ_i to j_i groups, each of which corresponds to the set of destination wavelengths of a multicast connection. There are $S(N, j_i)$ ways for such a division, where $S(N, j_i)$ is the Stirling number of the second kind [20]. Notice that the divisions for different wavelengths λ_i are independent. Given $1 \leq j_1, j_2, \dots, j_k \leq N$, we have $\prod_{i=1}^k S(N, j_i)$ ways to divide N λ_1 s to j_1 groups, N λ_2 s to j_2 groups, \dots , N λ_k s to j_k groups to form $\sum_{i=1}^k j_i$ multicast connections. Also, we need to choose the source wavelengths from Nk input wavelengths for these multicast connections and there are $P(Nk, \sum_{i=1}^k j_i)$ ways to do it. Therefore, we have $P(Nk, \sum_{i=1}^k j_i) \prod_{i=1}^k S(N, j_i)$ different full-multicast-assignments for k given numbers j_1, j_2, \dots, j_k . After we consider all possible values of j_1, j_2, \dots, j_k , we can obtain the multicast capacity for the MSDW model as given in Lemma 3 for full-multicast-assignments.

Finally, we consider any-multicast-assignments for the MSDW model. For N wavelength λ_i s at N output ports, assume l_i of them are not in use and divide the rest of $N - l_i$ wavelengths to j_i groups, each of which corresponds to the set of destination wavelengths of a multicast connection, where $0 \leq l_i \leq N$ and $1 \leq j_i \leq N - l_i$. There are $\binom{N}{l_i} S(N - l_i, j_i)$ different ways for such assignments. Similarly, the assignments for k different wavelengths $\lambda_1, \dots, \lambda_k$ are independent. Noticing that there are $P(Nk, \sum_{i=1}^k j_i)$ ways to choose source wavelengths for $\sum_{i=1}^k j_i$ multicast connections, we obtain that the number of different any-multicast-assignments is $P(Nk, \sum_{i=1}^k j_i) \prod_{i=1}^k \binom{N}{l_i} S(N - l_i, j_i)$ for given l_1, \dots, l_k and j_1, \dots, j_k . Hence, we can obtain the multicast capability as given in this lemma by considering all possible values of l_i s and j_i s. \square

Proof for Lemma 5. Suppose these m' middle switches are $1, 2, \dots, m'$ with destination multisets $M_1, M_2, \dots, M_{m'}$, which are nonempty multisets by the assumptions. Clearly, by using (4) and (5), we have that

$$0 < |M_i| \leq r \text{ for } 1 \leq i \leq m'.$$

Notice that at most $(nk-1)$ 1s, $(nk-1)$ 2s, \dots , $(nk-1)$ rs are distributed among the m' multisets. Moreover, from (4), k multiple of the same element in M_i contributes a value 1 to $|M_i|$. Thus, for any j ($1 \leq j \leq r$), $(nk-1)$ multiple of element j contributes a value no more than $\lfloor \frac{nk-1}{k} \rfloor$ to $\sum_{i=1}^{m'} |M_i|$. We have

$$\sum_{i=1}^{m'} |M_i| \leq \left\lfloor \frac{nk-1}{k} \right\rfloor r = \left\lfloor n - \frac{1}{k} \right\rfloor r = (n-1)r.$$

Let M_{j_1} be the multiset such that

$$|M_{j_1}| = \min_{1 \leq i \leq m'} |M_i|.$$

Then, we obtain that

$$m'|M_{j_1}| \leq \sum_{i=1}^{m'} |M_i| \leq (n-1)r$$

and, thus,

$$m' \leq \frac{(n-1)r}{|M_{j_1}|}, \quad (7)$$

by noting that $|M_{j_1}| > 0$.

Without loss of generality, suppose that, in multiset M_{j_1} , the $|M_{j_1}|$ elements each with multiplicity k are $1, 2, \dots, |M_{j_1}|$, that is, only $i_{j_1 1}, i_{j_1 2}, \dots, i_{j_1 |M_{j_1}|}$ are equal to k .

Now, consider m' new multisets $M_{j_1} \cap M_i$ for $1 \leq i \leq m'$. From (3), (4), and the assumption that the intersection of any two multisets is nonempty, we have

$$0 < |M_{j_1} \cap M_i| \leq |M_{j_1}|$$

and only elements $1, 2, \dots, |M_{j_1}|$ in multiset $M_{j_1} \cap M_i$ may have multiplicity k and, thus, can make a contribution to the value of $|M_{j_1} \cap M_i|$. Notice that at most $(nk-1)$ 1s, $(nk-1)$ 2s, \dots , $(nk-1)$ $|M_{j_1}|$ s are distributed in the m' multisets $M_{j_1} \cap M_i$ for $1 \leq i \leq m'$. Again, by using a similar analysis as above, we obtain that

$$\sum_{i=1}^{m'} |M_{j_1} \cap M_i| \leq \left\lfloor \frac{nk-1}{k} \right\rfloor |M_{j_1}| = (n-1)|M_{j_1}|. \quad (8)$$

Let M_{j_2} be the multiset such that

$$|M_{j_1} \cap M_{j_2}| = \min_{1 \leq i \leq m'} |M_{j_1} \cap M_i|.$$

Then, we can similarly have

$$m' \leq \frac{(n-1)|M_{j_1}|}{|M_{j_1} \cap M_{j_2}|} \quad (9)$$

by noting that $|M_{j_1} \cap M_{j_2}| > 0$.

In general, for $2 \leq y < x$, we can have

$$m' \leq \frac{(n-1) |\bigcap_{l=1}^{y-1} M_{j_l}|}{|\bigcap_{l=1}^y M_{j_l}|} \quad (10)$$

and $|\bigcap_{l=1}^y M_{j_l}| > 0$ from the assumption that the intersection of no more than x original multisets is nonempty.

On the other hand, also from this assumption, we have m' multisets $(\bigcap_{l=1}^{x-1} M_{j_l}) \cap M_i$ which are all nonempty. This means that $|(\bigcap_{l=1}^{x-1} M_{j_l}) \cap M_i| \geq 1$ for $1 \leq i \leq m'$ and

$$m' \leq \sum_{i=1}^{m'} |(\bigcap_{l=1}^{x-1} M_{j_l}) \cap M_i|.$$

Then, by using a similar argument as that used in deriving (8), we can establish

$$m' \leq (n-1) |\bigcap_{l=1}^{x-1} M_{j_l}|. \quad (11)$$

Finally, m' must be no more than the geometric mean of the righthand sides of (7), (10), and (11) and, thus, we obtain

$$m' \leq (n-1)r^{\frac{1}{x}}.$$

□

ACKNOWLEDGMENTS

Yuanyuan Yang was supported in part by the US National Science Foundation (NSF) under grant number CCR-0073085 and US Army Research Office under grant number DAAH04-96-1-0234. The research of Jianchao Wang was performed while the author was working at GTE Laboratories, Waltham, Massachusetts. Chunming Qiao was supported in part by the NSF under grant number ANIR-9801778.

REFERENCES

- [1] "Special Issue on Multiwavelength Optical Networks," *IEEE/OSA J. Lightwave Technology*, vol. 14, June 1996.
- [2] "Special Issue on Optical Networks," *IEEE J. Selected Areas in Comm.*, vol. 14, June 1996.
- [3] "Special Issue on Dense Wavelength Division Multiplexing Techniques for High Capacity and Multiple Access Communication Systems," *IEEE J. Selected Areas in Comm.*, vol. 8, Aug. 1990.
- [4] "Special Issue on WDM Networks," *J. High-Speed Networks*, vol. 4, nos. 1-2, 1995.
- [5] L.H. Sahasrabudde and B. Mukherjee, "Light-Trees: Optical Multicasting for Improved Performance in Wavelength-Routed Networks," *IEEE Comm.*, vol. 37, no. 2, pp. 67-73, 1999.
- [6] R. Malli, X. Zhang, and C. Qiao, "Benefit of Multicasting in All-Optical WDM Networks," *Proc. Conf. All-Optical Networks*, pp. 209-220, Nov. 1998.
- [7] C. Qiao, M. Jeong, A. Guha, X. Zhang, and J. Wei, "WDM Multicasting in IP over WDM Networks," *Proc. Int'l Conf. Network Protocols (ICNP)*, pp. 89-96, Nov. 1999.
- [8] E. Iannone, M. Listanti, and R. Sabella, "Multicasting in Optical Transport Networks," *J. Optical Comm.*, vol. 19, no. 3, pp. 90-98, Mar. 1998.
- [9] B. Mukherjee, *Optical Communication Networks*. McGraw-Hill, 1997.

- [10] L. Lin, E. Goldstein, L. Lunardi, and R. Tkach, "Optical Cross-Connects for High-Capacity Lightwave Networks," *J. High Speed Networks*, vol. 8, no. 1, pp. 17-34, 1999.
- [11] G.M. Masson and B.W. Jordan, "Generalized Multi-Stage Connection Networks," *Networks*, vol. 2, pp. 191-209, 1972.
- [12] F.K. Hwang and A. Jajszczyk, "On Nonblocking Multiconnection Networks," *IEEE Trans. Comm.*, vol. 34, pp. 1,038-1,041, 1986.
- [13] P. Feldman, J. Friedman, and N. Pippenger, "Wide-Sense Nonblocking Networks," *SIAM J. Discrete Math.*, vol. 1, no. 2, pp. 158-173, May 1988.
- [14] Y. Yang and G.M. Masson, "Nonblocking Broadcast Switching Networks," *IEEE Trans. Computers*, vol. 40, no. 9, pp. 1,005-1,015, Sept. 1991.
- [15] Y. Yang, "A Class of Interconnection Networks for Multicasting," *IEEE Trans. Computers*, vol. 47, no. 8, pp. 899-906, Aug. 1998.
- [16] Y. Yang and G.M. Masson, "The Necessary Conditions for Clos-Type Nonblocking Multicast Networks," *IEEE Trans. Computers*, vol. 48, no. 11, pp. 1,214-1,227, Nov. 1999.
- [17] C. Lee and A.Y. Oruç, "Design of Efficient and Easily Routable Generalized Connectors," *IEEE Trans. Comm.*, vol. 43, nos. 2, 3, 4, pp. 646-650, 1995.
- [18] Y. Yang and J. Wang, "A New Self-Routing Multicast Network," *IEEE Trans. Parallel and Distributed Systems*, vol. 10, no. 12, pp. 1,299-1,316, Dec. 1999.
- [19] J. Duato, S. Yalamanchili, and L.M. Ni, *Interconnection Networks: An Engineering Approach*. Los Alamitos, Calif.: IEEE Press, 1997.
- [20] K.P. Bogart, *Introductory Combinatorics*, second ed. Harcourt Brace Jovanovich, 1990.



Yuanyuan Yang received the BEng and MS degrees in computer science and engineering from Tsinghua University, Beijing, China, in 1982 and 1984, respectively, and the MSE and PhD degrees in computer science from Johns Hopkins University, Baltimore, Maryland, in 1989 and 1992, respectively. Dr. Yang is a tenured associate professor of computer engineering, with a joint appointment in computer science, at the State University of New York at Stony Brook, where she directs the High-

Performance Computing and Networking Research Lab. Dr. Yang's 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 networks which can support efficient collective communications. She is the principal investigator of two on-going research projects funded by the US National Science Foundation (NSF) and US Army Research Office (ARO). Dr. Yang 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 the *Journal of Parallel and Distributed Computing*. Dr. Yang holds two US patents in the area of multicast communication networks, with five more patents pending. She has served on the program/organizing committees of a number of international conferences in her areas of research. Dr. Yang is an editor of the *IEEE Transactions on Parallel and Distributed Systems*. She is a senior member of the IEEE and a member of the ACM, IEEE Computer Society, and IEEE Communication Society.



Jianchao Wang received the BEng degree in computer science and engineering from Tsinghua University, Beijing, China, and the MS and PhD degrees in computer science from the Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China. Dr. Wang was a principal member of the technical staff at GTE Laboratories Incorporated in Waltham, Massachusetts, from 1995-2000. He recently joined DataTreasury corporation in Melville, NY.

Prior to joining GTE Labs, Dr. Wang worked at the Institute of Computing Technology, Chinese Academy of Sciences, Johns Hopkins University, and Legent Corporation in Marlboro, Massachusetts. Dr. Wang has received numerous excellence/achievement awards from GTE and Legent and has five US patents granted or pending. He was listed in the 1998 edition of *International Who's Who in Information Technology*. Dr. Wang's research interests include IP telephony, databases, programming languages, computer communication networks, computer algorithms, and fault-tolerant computing. Dr. Wang is a member of IEEE Computer Society and the ACM.



Chunming Qiao received the Andrew-Mellon Distinguished Doctoral Fellowship Award from the University of Pittsburgh and earned his PhD in computer science in 1993. Prior to that, he was admitted to the top-ranked University of Science and Technology (USTC) of China at age 15 and earned his BS degree in computer engineering. Dr. Qiao is currently a tenured associate professor in the Computer Science and Engineering Department of the State Uni-

versity of New York at Buffalo and the director of the LANDER (Lab for Advanced Networking Design, Evaluation, and Research), which conducts cutting edge research on optical networks, wireless/mobile communications, and the Internet. He pioneered research on the next generation optical internet (or IP over WDM networks) and, particularly, optical burst switching (OBS). He has published 30 papers in leading technical journals such as *IEEE/OSA Journal on Lightwave Technology*, *IEEE Journal on Selected Area in Communications (JSAC)*, *IEEE/ACM Transactions on Networking*, *IEEE Transactions on Communications*, *IEEE Transaction on Parallel and Distributed Systems*, as well as *IEEE Transactions on Computers*, and given several invited talks and keynote speeches.

Dr. Qiao's research has been sponsored by the US National Science Foundation, Alcatel Corporate Research Center, Nokia Research Center, and Telcordia (formerly Bellcore). He is serving as a cochair for the program section on IP/WDM and optical internet at APOC '2001 (a premier conference on optical communications in Asia-Pacific) and has been the chair for the program on optical layer and internetworking technology (a part of the SPIE's Photonic East 2000), a cochair for the annual All-Optical Networking Conference since 1997, a program vice cochair for the 1998 International Conference on Computer Communications and Networks (IC3N), expert panelist, technical program committee member, and session organizer/chair for several other conferences and workshops. He is also the chair of the recently established Technical Group on Optical Networks (TGON) sponsored by SPIE, the editor-at-large for the IEEE Communication Society in the area of optical networking, an editor of the *Journal of High-Speed Networks (JHSN)*, and *Optical Networks Magazine*.