Multicasting in WDM networks with heterogeneous group weights

Jingyi He and Danny H. K. Tsang

Department of Electrical and Electronic Engineering,
Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong

eehjy@ust.hk; eetsang@ee.ust.hk

S.-H. Gary Chan

Department of Computer Science, Hong Kong University of Science and Technology,
Clear Water Bay, Kowloon, Hong Kong

gchan@cs.ust.hk

RECEIVED 10 FEBRUARY 2004; REVISED 5 MAY 2004;
ACCEPTED 5 MAY 2004; PUBLISHED 16 JUNE 2004

We address multicasting in a wavelength-routed WDM network in which each destination node has different, as opposed to the traditional uniform, weights for different multicast groups. The weight, in practice, may reflect the popularity of that multicast group at the node. The objective is therefore to serve the multicast groups so as to maximize the total weights of the multicast trees (or equivalently to minimize the weighted overall blocking rate). We propose for this purpose a routing and wavelength assignment (RWA) heuristic that serves the group of the largest weight first (LWF). Since LWF penalizes groups of small weights, we also propose a fairness improvement (FI) heuristic that runs on top of LWF to achieve better fairness among the multicast groups in terms of their respective blocking rates. We show that LWF significantly reduces the weighted overall blocking rate as compared with traditional schemes that do not take group weights into consideration, and FI is effective in improving the blocking fairness. Moreover, we show that FI does not trade off the overall blocking rate for fairness when the group weight heterogeneity is small, the group sizes are small, or the network is densely connected, because of more efficient use of the links.

OCIS codes: 060.2330, 060.4250.

1. Introduction

Wavelength-division multiplexing (WDM) is an effective technique for making use of the large amount of bandwidth in optical fibers to meet the bandwidth requirements of applications. It is generally believed that the next-generation Internet will be largely based on a WDM backbone [1]. Multicasting is the transmission of information from one source to multiple destinations simultaneously and is particularly important for many broadband services, such as videoconferencing, distance learning, and webcasting. Therefore multicasting support in future WDM networks is essential. In a wavelength-routed WDM network two end nodes communicate by setting up an all-optical lightpath that can span a number of physical links. When wavelength conversion is not possible in the network, wavelength continuity must be satisfied where all the links of a lightpath or multicast tree (i.e., light-tree) must use the same wavelength. In this paper we study multicasting in such networks.

A typical problem in multicasting is serving a number of multicast groups with a limited number of wavelengths with the objective of minimizing the overall blocking rate. The blocking rate can be either a session blocking rate or a destination blocking rate [2]. A session blocking rate is an appropriate metric for applications, such as distributed computing and videoconferencing, that require all the destinations to be reached for the communication to take place. The session is hence blocked without a complete multicast tree. For many other applications it is acceptable to serve some of the destination nodes instead of blocking the whole group; in this case a destination blocking rate is an appropriate measure. Traditionally, all the destination nodes are regarded as the same, and hence blocking one or the other does not make any difference in terms of the overall blocking rate.

In this paper we generalize the WDM multicast model by assigning a weight to each multicast group at each destination node. One may think of the weight as an indication of the popularity of a group at a certain node. We study the routing and wavelength assignment (RWA) problem in such a system: We determine the multicast tree and the wavelength to serve each group. We are interested in minimizing the weighted overall blocking rate, defined as the ratio of the total weight lost due to unserved nodes to the total weight of all the groups, and achieving fairness in the blocking rates among the groups. For example, if the weight of a node for a group is the number of group members at the node, then the weighted overall blocking rate is simply the overall end-user blocking rate. We show in Fig. 1 an example of such a WDM network with two multicast groups sourced at s_1 and s_2. The 2-tuple at each node represents the weight assigned to the two groups. If a group is assigned zero weight at a node, that means the node is not a member (or destination node) of the group. In this example, nodes *A*, *C*, and *D* are members of group 1 and nodes *A*, *B*, and *D* are members of group 2. In the following, for ease of exposition, we use the overall blocking rate to refer to the weighted overall blocking rate.

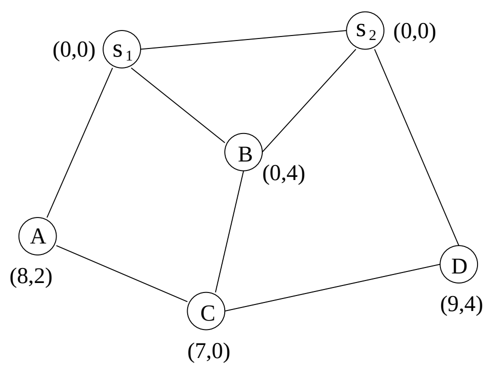

Fig. 1. Sample WDM multicast network of two groups. s_1 and s_2 denote the source nodes. Tuple (i, j) at each node represents the weight assigned to the two groups.

To achieve a low overall blocking rate, we propose a RWA heuristic that serves the group of the largest weight first (LWF). We show that LWF significantly reduces the overall blocking rate as compared with traditional schemes that do not take group weights into consideration. LWF, however, penalizes groups of small weights, making their blocking rate relatively high. To achieve more-uniform blocking rates and hence better fairness, we propose a fairness improvement (FI) algorithm that runs on top of LWF. The FI algorithm reallocates resources (through rerouting) from lightly blocked groups to serve more heavily blocked groups. We show that FI is effective in improving the blocking fairness. We also study the effect of several factors on the performance of these two heuristics, such as group weight heterogeneity, multicast group size, and network density in terms of the average nodal degree. We show that FI does not trade off the overall blocking rate for fairness when the group weight heterogeneity is small, the group sizes are small, or the network is densely connected, because of its more efficient use of the links.

We briefly review previous work on WDM multicasting as follows. There has been much study focusing on the optimal RWA for a multicasting request, possibly under various constraints such as sparse light splitting and limited light power [3–14]. On the other hand, other studies address the RWA problem for satisfying multiple multicasting requests simultaneously [15–20]. The requests can be either a batch request that needs to be served simultaneously or a long-term traffic matrix that is used to facilitate the network (logical topology) design and configuration. Our study falls into the second group but differs in two ways. First, our study uses a more general model in which different weights are assigned to the groups at the nodes to reflect their importance, which has not been considered before. Second, our study focuses more on RWA with a limited number of wavelengths, rather than on minimizing the number of wavelengths needed to serve all the groups. Therefore our objectives are not only to minimize the blocking rate but also to serve the groups in a fair manner.

The rest of this paper is organized as follows: In Section 2, we present the system model and problem formulation. The LWF heuristic is presented in Section 3, and the FI heuristic is presented in Section 4. Illustrative numerical results for the performance of the two heuristics are given in Section 5. We conclude the paper in Section 6.

2. System Model

The WDM network is modeled as an undirected graph $G = (V,E)$, where V is the set of nodes with $|V| = N$ and E is the set of links with $|E| = L$. The nodes are labeled $v_1, v_2, \ldots, v_N$. All of them have light-splitting capability, and hence the multicast tree built on the graph is always realizable. The set of wavelengths available on each link is $\Lambda = \{\lambda_1, \lambda_2, \ldots, \lambda_W\}$, where W is the total number of wavelengths. There is no wavelength conversion capability at the network nodes, and hence the wavelength continuity must be met.

There are M multicast groups. Each multicast group is represented as $g_i = (s_i, D_i)$, where $s_i \in V$ is the source node and $D_i \subset V$ is the set of destination nodes of the multicast group. We define group spread α_i for group i as the fraction of destination nodes of group i in the network. Clearly, $\alpha_i = |D_i|/(N-1)$. w_{ij} is the weight assigned to group i at node $v_j \in V$, where

$$w_{ij} = \begin{cases} \mathbf{R}^+ & \text{if } v_j \in D_i \\ 0 & \text{otherwise} \end{cases}, \tag{1}$$

for $1 \le i \le M$ and $1 \le j \le N$.

The RWA problem is to determine the route and wavelength assignment for each group such that the overall blocking rate can be minimized given the group distributions $g_i = (s_i, D_i), i = 1, \ldots, M$. Clearly, we only need to consider M to be larger than W (otherwise, all the multicast groups could be fully served). We denote the RWA for group i as a tuple (T_i, λ'_i), where $T_i \subset E$ is the set of links used to serve the group and $\lambda'_i \in \Lambda$ is the wavelength used on these links (all these links use the same wavelength due to wavelength continuity). Note that T_i may or may not be a complete multicast tree spanning all the destination nodes in the group. Let V_i denote the set of nodes on T_i and η denote the overall blocking rate. Then we have

$$\eta = \frac{\sum_{i=1}^{M} \sum_{v_j \notin V_i} w_{ij}}{\sum_{i=1}^{M} \sum_{j=1}^{N} w_{ij}}. \tag{2}$$

The RWA problem can hence be formally stated as follows:

$$\begin{array}{lll} \text{Determine:} & (T_i, \lambda'_i), & \text{for } i = 1, \ldots, M; \\ \text{To minimize:} & \eta; & \\ \text{Subject to:} & \lambda'_i \neq \lambda'_j, & \text{if } T_i \bigcap T_j \neq \emptyset. \end{array}$$

The constraint simply means that those multicast (sub)trees sharing some links must have different wavelengths.

To evaluate the fairness among the blocking rates of the groups, we first define the blocking rate of group i, η_i, as the ratio of the weight lost due to the unserved nodes to the weight of all its destination nodes:

$$\eta_i = \frac{\sum_{v_j \notin V_i} w_{ij}}{\sum_{j=1}^{N} w_{ij}}. \tag{3}$$

The fairness index (according to Jain [21]) is then given by

$$f = \frac{\left[\sum_{i=1}^{M} (1-\eta_i)\right]^2}{M \sum_{i=1}^{M} (1-\eta_i)^2}. \tag{4}$$

Note that $f \in (0, 1]$, with the larger value meaning better fairness and $f = 1$ meaning perfect fairness.

3. Minimizing the Overall Blocking Rate

Because the multicast RWA problem is nondeterministic polynomial-time- (NP-) complete [16], we propose a heuristic to minimize the overall blocking rate in this section. Note that the aforementioned RWA problem consists of two basic subproblems, namely, the routing problem (determining T_i) and the wavelength assignment problem (determining λ_i'). There are in general two ways to solve the two subproblems: In static RWA the multicast tree of a group is first constructed irrespective of the wavelength availability on the links and is then assigned a certain wavelength from the available wavelength pool. In dynamic RWA a multicast tree is built on each wavelength graph (a subgraph of G by removing the links on which the wavelength is not available) and the best one is chosen. Obviously, dynamic RWA performs better than static RWA for better use of the wavelengths [2]. We hence focus on dynamic RWA in this paper.

With dynamic RWA, after a group is served, the corresponding wavelength graph is reduced due to the removal of the used links. In this process some nodes would be disconnected in the wavelength graph because all the links incident on them have been removed. If this particular wavelength is used to serve another group later, only a subset of the destination nodes can be served. In other words, the multicast groups served later are more likely to be blocked. To minimize the overall blocking rate (or equivalently, to maximize the overall weight of the served groups), we serve the group of the LWF.

The basic idea of LWF is to build multicast trees on all the wavelength graphs for unserved groups, calculate the tree weights, find the group–wavelength combination of the largest weight, and serve that group with the corresponding wavelength. The next group to be served is the one with the largest weight by use of the remaining wavelengths. Let a_{ij} denote the weight of the served nodes in group g_i with wavelength λ_j. The above process is simply finding the largest element a_{lm} in matrix $\{a_{ij}\}$, removing row l and column m, and finding the next-largest element. Initially, all the groups are unserved and $\{a_{ij}\}$ is an $M \times W$ matrix. After the first W groups are served with different wavelengths, the algorithm proceeds to the second round of RWA, in which $\{a_{ij}\}$ is an $(M-W) \times W$ matrix. Note that each time a wavelength is used to serve a group the corresponding wavelength graph is updated by removing the links on the multicast tree of that group. The algorithm stops either when all the multicast groups have been served (maybe partially) or when some groups are completely blocked and hence $\{a_{ij}\}$ is an all-zero matrix. The details of the LWF algorithm are shown in Algorithm 1.

We illustrate in Fig. 2 the last two rounds of RWA of a LWF operation, in which four groups are still to be served by two wavelengths. Note that after serving g_2 and g_3 with λ_1

Algorithm 1 LWF

```
/* Initialization */
T_i ← ∅; λ'_i ← 0
U ← {g_i | i = 1, ..., M}        /* The set of unserved groups */
G_j ← G                          /* The graphs on the W wavelengths */
while U ≠ ∅ do
  for all g_i ∈ U do
    for j = 1 to W do
      Build the multicast tree for g_i on G_j: T_ij
      Calculate the weight of the nodes served by T_ij: a_ij
    end for
  end for
  if {a_ij} is all-zero then
    STOP                         /* No more groups can be served */
  else
    A ← {g_i | a_ij > 0}; B ← {λ_j | a_ij > 0}
    while A ≠ ∅ && B ≠ ∅ do
      a_lm ← max_{g_i∈A, λ_j∈B} {a_ij}        /* The largest weight */
      T_l ← T_lm; λ'_l ← λ_m                  /* g_l is served by (T_l, λ'_l) */
      Update G_m by removing links on T_lm
      U ← U − g_l; A ← A − g_l; B ← B − λ_m
    end while
  end if
end while
```

and λ_2, respectively, new multicast trees are built for g_1 and g_4 on the updated graphs on the two wavelengths, and hence the weights are different from the previous round.

The algorithm used to build the multicast trees can be any of the existing ones, such as the Kou–Markowsky–Berman (KMB) algorithm, the minimum spanning tree (MST) heuristic, or the shortest-path tree (SPT) heuristic [22]. In this study we chose the SPT heuristic because of its simplicity and low source-to-destination delay. The SPT heuristic simply calculates the shortest paths from the source to all the destinations and aggregates them by removing duplicated links. It has a complexity of $O(L+N\log N)$[22]. The complexity of LWF is primarily due to the building of the SPTs and the process of searching for the largest element of and then updating (i.e., removing the corresponding row and column of) matrix $\{a_{ij}\}$. In each round the SPT heuristic is run for $|U|$ groups on W wavelength graphs, leading to a complexity of $O[|U|W(L+N\log N)]$. If we use an Adelson-Velskii and Landis (AVL) tree as the data structure for the process of searching and updating matrix $\{a_{ij}\}$, the complexity is $O[|U|W\log(|U|W)]$. In general, W groups can be served in each round, and hence $|U|$ decreases from M to 0 in a step size of W. Therefore the LWF algorithm has the complexity of $O\left[M^2(L+N\log N) + M^2\log(MW)\right]$ or $O\left[M^2(L+N\log N + \log M + \log W)\right]$.

4. Improving the Fairness among Multicast Groups

Obviously, LWF achieves a low blocking rate by giving priority to larger-weight groups, leading to unfairness in the blocking rate. In this section we propose a FI heuristic that runs on top of LWF to address this.

The basic idea of FI is to serve more nodes of the most severely blocked multicast group by making use of the links previously assigned to other lightly blocked groups, so that the maximum blocking rate among the multicast groups can be reduced. FI takes the results of LWF (i.e., $\{T_i, \lambda'_i, \eta_i\}$) as the input and outputs a modified RWA with better fairness among

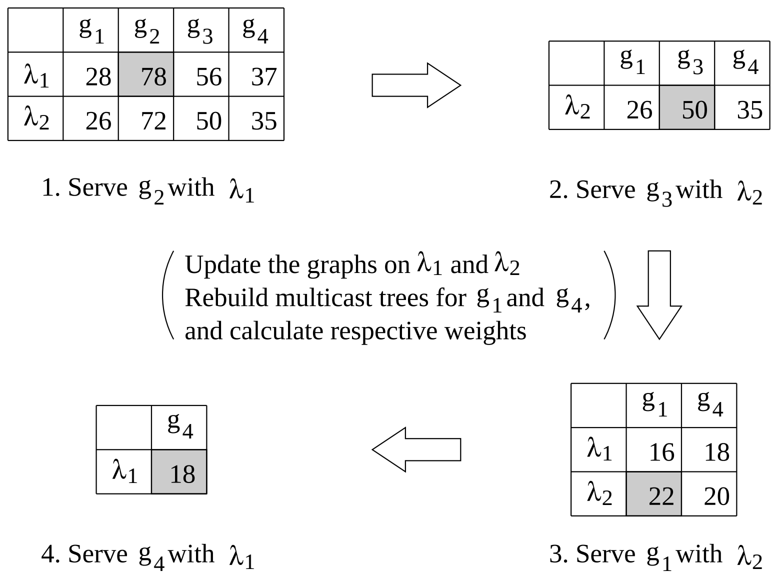

Fig. 2. Illustration of the LWF operation.

the multicast groups. FI starts by determining the set of wavelengths that are assigned to more than one group and then does the rearrangement for each of them. Note that the wavelength assignments are not changed, and the algorithm rearranges (reroutes) only the multicast trees served by the same wavelength. If a wavelength is used to serve only one multicast group, that group would be completely served and no rearrangement is needed.

The rerouting procedure on each wavelength graph is as follows: First, from the set of groups served by the same wavelength, find the group (g_l) with the largest blocking rate ($\eta_{\max}$) and the group (g_m) with the lowest blocking rate (may be zero). Then for each of the blocked nodes in g_l try to reconnect it to its current multicast tree T_l through the shortest path P, assuming the links on tree T_m can also be used by P in addition to the free links. In serving this node with path P, some nodes on tree T_m can become unreachable from the root. However, they can be reattached to T_m through a set of alternative shortest paths Q. When both T_l and T_m are updated, the new blocking rate of the two groups is calculated. If the larger blocking rate is lower than the previous $\eta_{\max}$, it means that the largest blocking rate on this wavelength has been reduced and the fairness is improved. Then we re-sort the blocking rate of the groups and try to make another round of improvement. If the rerouting does not offer any improvement, we discard the changes and try to reconnect the next blocked node in g_l. Note that the blocked nodes are tried in random order in this study. If all the blocked nodes in g_l have been tried without success, it means that rerouting group g_m cannot help, and we try the group with the next-lowest blocking rate. The rerouting procedure on this wavelength stops until the worst-case blocking rate can no longer be improved. The details of the FI algorithm are shown in Algorithm 2.

We show in Fig. 3 an example of the FI operation based on Fig. 1. Two multicast groups sourced at s_1 and s_2 are served by the same wavelength. Because group 1 has a larger weight (24) than group 2 (10), LWF serves group 1 first. The output is shown in Fig. 3(a), where group 1 is completely served, whereas only node B in group 2 can be served with a blocking rate of 0.6. With this result, FI first tries to reconnect blocked node D in group 2 to its existing multicast tree, using links originally used by group 1. As a result, the same node as a member of group 1 must be blocked. FI then tries to reconnect it to the

Algorithm 2 FI

```
/* The set of wavelengths serving multiple groups */
A ← {λ_k | ∃i ≠ j, λ'_i = λ'_j = λ_k}
/* Rerouting multicast trees on each such wavelength graph */
for all λ_k ∈ A do
  NEXT: B ← {g_i | λ'_i = λ_k}          /* Groups served by λ_k */
  η_l, η_max ← max_{g_i∈B}{η_i}          /* The largest η_i */
  B ← B − g_l
  while B ≠ ∅ do
    η_m ← min_{g_i∈B}{η_i}               /* The smallest η_i */
    C ← {blocked nodes of g_l}
    for all v ∈ C do
      Reconnect v to T_l through shortest path P, assuming links on T_m also usable
      Reconnect the blocked nodes of g_m, if possible, to T_m through shortest paths
      Update T_l and T_m and calculate the new η_l and η_m
      η'_max ← max{η_l, η_m}
      if η'_max < η_max then
        GOTO NEXT                  /* Improved, new g_l and g_m */
      else
        Restore T_l, T_m, η_l, η_m    /* No improvement, next v */
      end if
    end for
    B ← B − g_m                 /* g_m cannot help g_l, try next */
  end while
end for
```

partial tree of group 1 through an alternative path and succeeds through the link between (9,4) and (7,0). After this rerouting, group 1 is still completely served, and the blocking rate of group 2 is reduced to 0.2. Then FI tries to serve another blocked node A of group 2. However, reconnecting it to the current partial tree, e.g., through node C to D, will block node D as a member in group 1, resulting in a blocking rate of 0.375 for group 1. This value is larger than 0.2, and hence this change will not be applied. The output of FI is therefore group 1 with a blocking rate of 0 and group 2 with a blocking rate of 0.2, which is a better fairness than the output of LWF.

The complexity of the FI algorithm can be analyzed as follows: The core of the algorithm reroutes two multicast trees served by the same wavelength. The rerouting procedure involves reconnecting nodes to an existing tree through the shortest path, which has the worst-case complexity of $O(L+N\log N)$. The number of nodes to be reconnected is $O(N)$. The complexity of the rerouting procedure is hence $O[N(L+N\log N)]$. The rerouting procedure is run on each wavelength graph between the most-blocked group and all the other groups. The number of times the rerouting procedure is run is $O(M)$. Therefore the worst-case complexity of the FI algorithm is $O[MN(L+N\log N)]$. Note that, if the multicast groups represent the long-term traffic load, both LWF and FI can be computed offline in advance to solve the RWA problem.

5. Illustrative Numerical Results

In this section we present illustrative simulation results for the performance of LWF and FI in terms of the overall blocking rate and fairness among the multicast groups. For comparison, we consider two other RWA schemes that do not take the group weights into consideration. One is that all the groups are treated the same and are served in a random order (RO). The other is that the groups are served in decreasing order of their number of destination nodes or largest group first (LGF). Obviously, if all the groups have the same weight, LWF

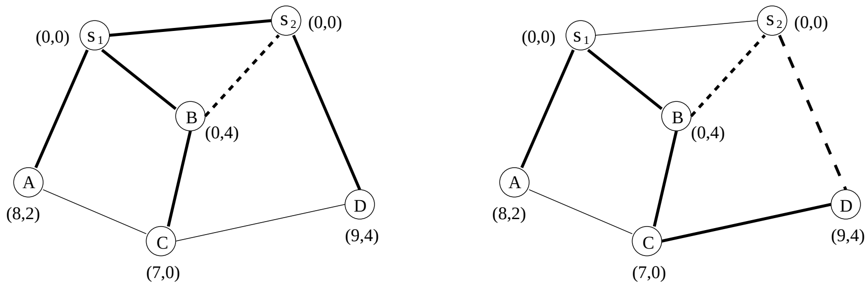

(a) Multicast trees by LWF

(b) Multicast trees by LWF+FI

Fig. 3. Illustration of the FI operation. The bold solid lines represent the multicast tree for group 1, and the dashed lines represent the multicast tree for group 2. (a) Output of LWF, where the blocking rates of the two groups are $\eta_1 = 0$ and $\eta_2 = 0.6$. (b) Output after applying FI, where the blocking rates of the two groups are $\eta_1 = 0$ and $\eta_2 = 0.2$.

is simply LGF.

In studying their performance we consider the following factors: The first is the weight heterogeneity of the multicast groups. We use the following model for the group weights. All the destination nodes of a group g_i have the same weight denoted by w_i.We do not further differentiate the nodes in the same group for simplicity, and this does not affect the performance of the proposed heuristics because the affecting factors are the total weight and the blocking rate of the group instead of the intragroup weight distribution. The weights of the M groups follow the geometric relationship, i.e., $w_{i+1}/w_i = 1 - p$ for $i = 1, \ldots, M-1$, where $p \in [0,1)$ is the heterogeneity factor. Clearly, larger p means higher heterogeneity, and, when $p = 0$, all the groups have the same weight.

The second factor is the group size. Serving a group with a larger group spread tends to use more links in the network and hence leaves less resources for serving other groups and for rerouting. For simplicity, we consider that all the groups statistically have the same group spread (denoted by α); i.e., for each group the source node is randomly chosen in the network and any other node is a destination with probability α.

The third factor is the network density, which can be measured by average nodal degree d. With the same number of wavelengths and network nodes, a larger d means more links and hence wavelength channels for serving the multicast groups.

We use the National Science Foundation Network (NSFNET) backbone as shown in Fig. 4 to study the effect of the weight heterogeneity and group size. It consists of 14 nodes ($N = 14$) and 21 links ($L = 21$), with $d = 3$. The label on each link represents the corresponding link delay. When studying the effect of network density, we generate 20-node networks with different average nodal degrees by use of the Georgia Tech Internetwork Topology Models (GT-ITM) topology generator (with the Waxman 2 method) developed by the Georgia Institute of Technology [23].

The baseline parameters in our simulations are $M = 8$, $W = 5$, $p = 0.2$, $\alpha = 0.7$, and $d = 3$. Each data point is an average of 100 simulation runs (i.e., 100 group distribution scenarios).

5.A. Effect of Weight Heterogeneity

We show in Fig. 5 the overall blocking rate achieved by the four RWA schemes as weight heterogeneity p increases. With RO and LGF the overall blocking rate is rather independent of p since the two schemes do not consider the group weights. In contrast, LWF significantly decreases the blocking rate as p increases because a larger p means that the weight of the blocked groups accounts for a smaller portion of the total weight. When FI is applied, we observe that the overall blocking rate is increased at large p values but decreased at small p values. The reason is as follows. With FI, to serve a blocked node of a small-weight group, the blocked nodes of the larger-weight group can either be served through rerouting or be blocked. In the former case the fairness can be improved without sacrificing the overall blocking rate, as the example in Fig. 3 shows. In the latter case the overall blocking rate would likely be increased. When p is small, the weight among the groups is rather uniform, and hence the increase in the overall blocking rate in the latter case is minor and the overall effect of the two cases is a reduced η. On the other hand, when p is large, blocking a large-weight node may lead to a substantial increase in the overall blocking rate. Therefore the overall effect is an increased η.

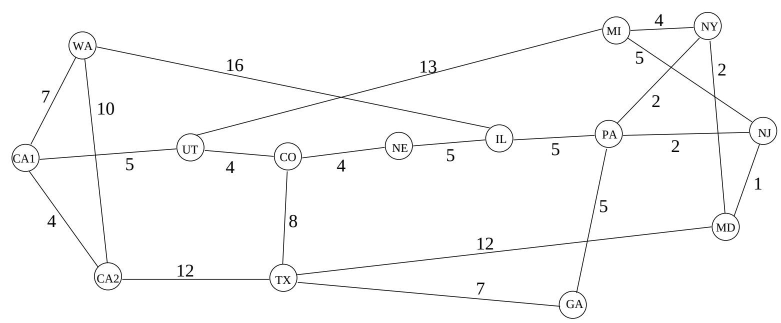

Fig. 4. NSFNET backbone used in the simulation.

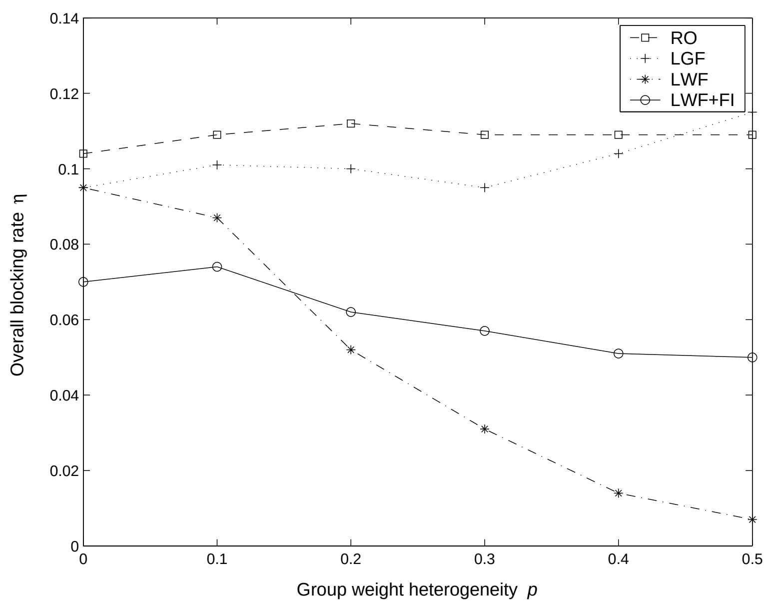

Fig. 5. Overall blocking rate η versus group weight heterogeneity p.

We next show in Fig. 6 the fairness index versus p. Clearly, FI effectively improves the fairness. Note that the fairness is measured in terms of the blocking rate (or served portion) of the groups. With the same blocking rate, a large-weight group would lead to a larger overall blocking rate than a small-weight group, but their effect on the overall fairness is the same. Therefore LWF achieves fairness similar to that of RO and LGF. Moreover, the fairness of all four schemes is independent of p.

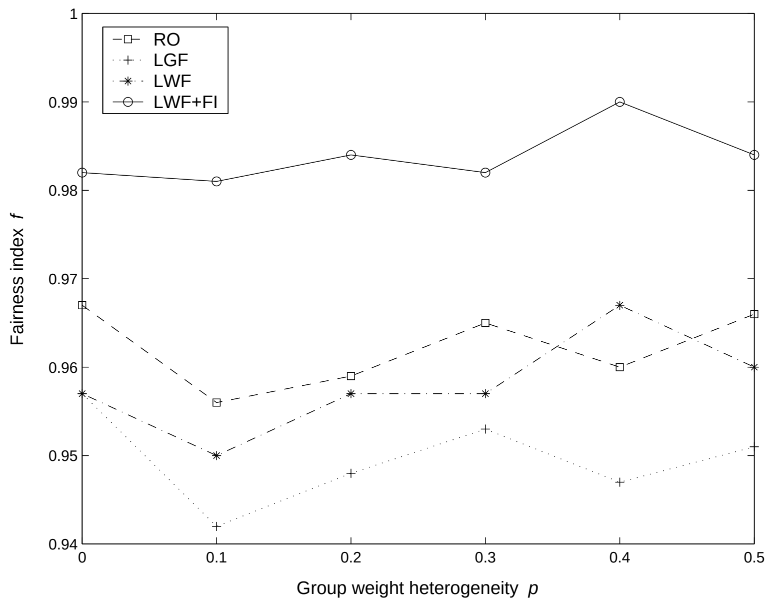

Fig. 6. Fairness index f versus group weight heterogeneity p.

5.B. Effect of Group Size

We show in Fig. 7 overall blocking rate η versus group spread α. The blocking rate increases with α because serving more destination nodes requires more resources (wavelength channels on the links). We observe that LWF substantially reduces the blocking rate as compared with RO and LGF (by approximately 50% in this example). We also observe that FI increases the blocking rate of LWF at only large values of α. At small α more links are available for rerouting, and hence the blocking rate with FI can even be reduced compared with the case without FI.

We show in Fig. 8 the fairness index versus α. We see that the fairness index decreases as α increases. This is because more links are to be used to serve a group. Consequently, after serving the first W groups of the largest weights, fewer links are left for serving the rest of the groups, which hence have a larger blocking rate. Since the first W groups can always be completely served (with a blocking rate of zero), the blocking rate difference among the groups becomes larger as α increases. Therefore the fairness index decreases. We observe that the fairness index of RO, LGF, and LWF decreases fast, whereas FI achieves much better fairness.

5.C. Effect of Network Density

We show in Fig. 9 overall blocking rate η versus average nodal degree d. As d increases, η decreases because more links are available to route the traffic. When d is large enough, all the groups will be completely served irrespective of the RWA scheme used. However, when there are limited link resources, LWF achieves significant reduction in the overall blocking rate compared with RO and LGF. Similar to the case for small group spread, FI does not trade off the overall blocking rate for fairness when d is large, due to its more effective rerouting with more available links.

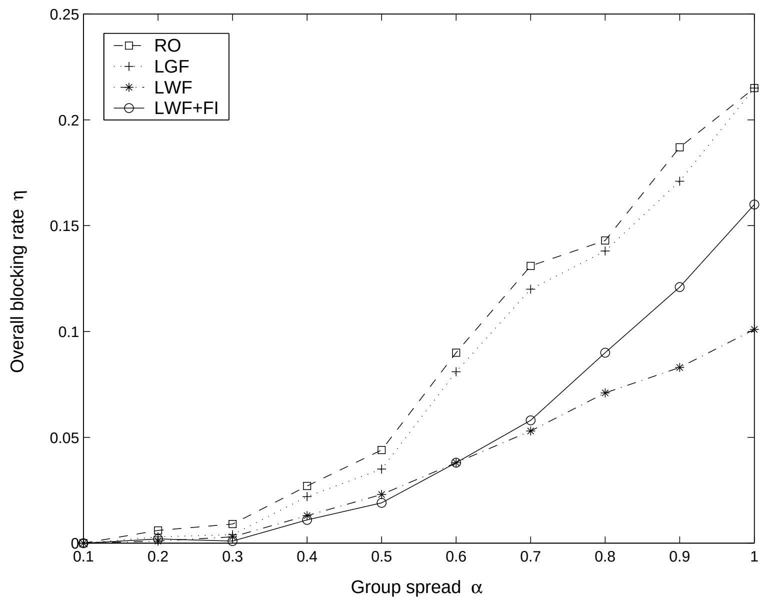

Fig. 7. Overall blocking rate η versus group spread α.

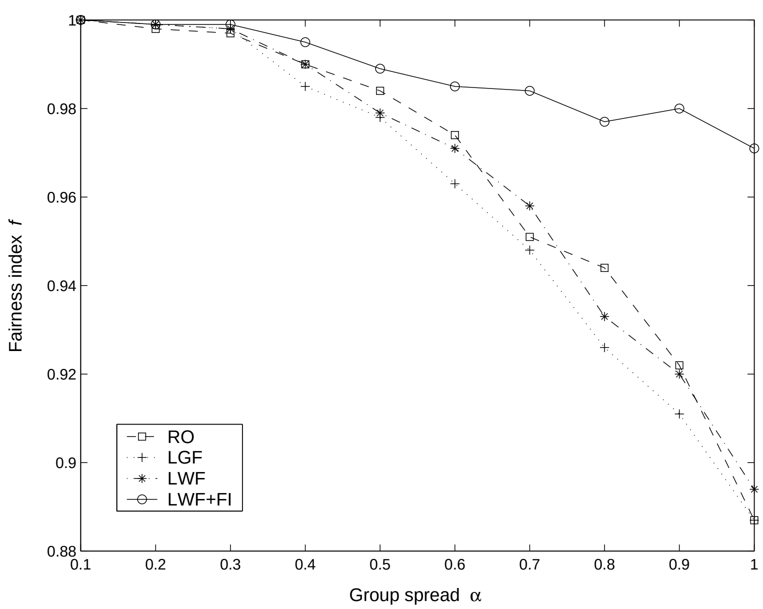

Fig. 8. Fairness index f versus group spread α.

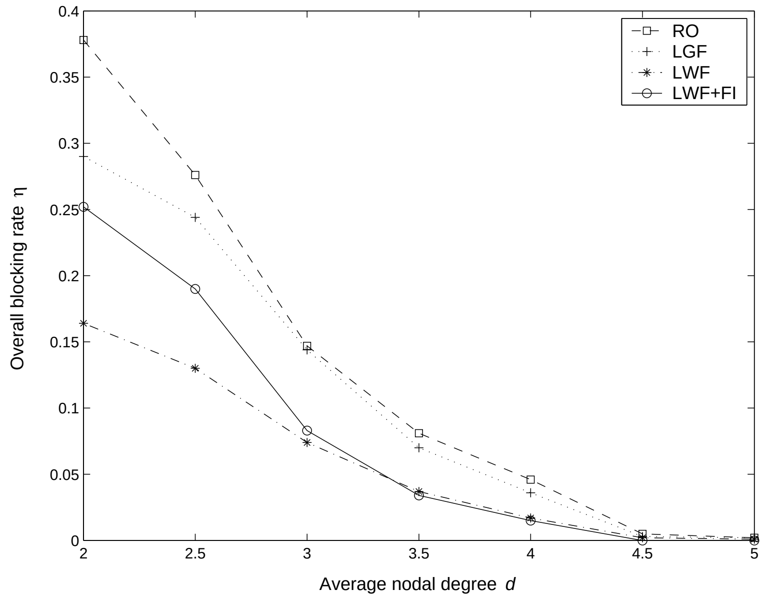

Fig. 9. Overall blocking rate η versus average nodal degree d.

We show in Fig. 10 the fairness index versus d. As d increases, more groups enjoy low blocking rate, and hence better fairness can be achieved. The FI is dependent on d. When d is small, the number of links available for rerouting in the FI stage is small, and hence the improvement is limited. When d is large, LWF can already achieve low blocking rates for the groups. Consequently, the improvement achieved by FI is also minor. In conclusion, FI is most effective in moderately connected networks (i.e., with a d of approximately 2.5).

6. Conclusions

In this paper we have considered a generalized WDM multicast network model in which different multicast groups can have different weights at each node, so as to reflect different popularity of the multicast groups at each node. We have addressed two important issues in such networks: minimizing the weighted overall blocking rate and serving the groups fairly in terms of their respective blocking rates. We have proposed two heuristics, LWF and FI, to minimize the overall blocking rate and improve the fairness, respectively. We show that LWF significantly reduces the weighted overall blocking rate as compared with traditional schemes that do not take group weights into consideration. FI is also shown to be effective in improving the blocking fairness. Moreover, it does not trade off the overall blocking rate for fairness when the group weight heterogeneity is small, the group sizes are small, or the network is densely connected, because of more efficient use of the links.

As a final comment, we note that this study, as well as most previous studies on multicast RWA, assumes that a multicasting session can monopolize the bandwidth of the wavelength assigned to it and hence does not take traffic grooming into consideration [24]. However, in many cases the multicast sessions may demand much lower bandwidth. Therefore multicast RWA with traffic grooming could be an interesting future research topic.

Acknowledgments

This work was supported, in part, by the Area of Excellence Scheme in Information Technology of the University Grant Council (AoE/E-01/99) and by the Competitive Earmarked

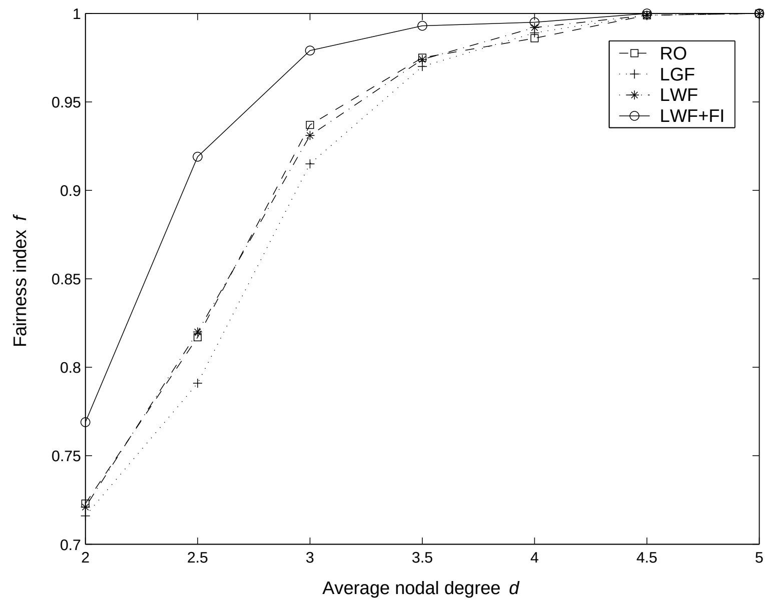

Fig. 10. Fairness index f versus average nodal degree d.

Research Grant of the Research Grant Council (HKUST6014/01E and HKUST6199/02E) in Hong Kong.

References and Links

[1] B. Mukherjee, *Optical Communication Networks* (McGraw-Hill, New York, 1997).

[2] G. Sahin and M. Azizoglu, "Multicast routing and wavelength assignment in wide area networks," in *All-Optical Networking: Architecture, Control, and Management Issues,* J. M. Senior and C. Qiao, eds., Proc. SPIE **3531**, 196–208 (1998).

[3] J. Wang, B. Chen, and R. N. Uma, "Dynamic wavelength assignment for multicast in all-optical WDM networks to maximize the network capacity," IEEE J. Sel. Areas Commun. **21**, 1274–1284 (2003).

[4] S. Sankaranarayanan and S. Subramaniam, "Comprehensive performance modeling and analysis of multicasting in optical networks," IEEE J. Sel. Areas Commun. **21**, 1399–1413 (2003).

[5] G. Xue and R. Banka, "Bottom-up construction of dynamic multicast trees in WDM networks," in *Proceedings of the IEEE 2003 International Performance, Computing, and Communications Conference* (IEEE, New York, 2003), pp. 49–56.

[6] A. Zsigri, A. Guitton, and M. Molnar, "Construction of light-trees for WDM multicasting under splitting capability constraints," in *10th International Conference on Telecommunications* (IEEE, New York, 2003), pp. 171–175.

[7] R. Libeskind-Hadas and R. Melhem, "Multicast routing and wavelength assignment in multihop optical networks," IEEE/ACM Trans. Netw. **10**, 621–629 (2002).

[8] B. Chen and J. Wang, "Efficient routing and wavelength assignment for multicast in WDM networks," IEEE J. Sel. Areas Commun. **20**, 97–109 (2002).

[9] X. Zhang, J. Wei, and C. Qiao, "Constrained multicast routing in WDM networks with sparse light splitting," in *INFOCOM 2000: Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies* (IEEE, New York, 2000), pp. 1781–1790.

[10] S. Yan, M. Ali, and J. Deogun, "Route optimization of multicast sessions in sparse light-splitting optical networks," in *Global Telecommunications Conference (GLOBECOM '01)* (IEEE, New York, 2001), pp. 2134–2138.

[11] M. Ali, “Optimization of splitting node placement in wavelength-routed optical networks,” IEEE J. Sel. Areas Commun. **20**, 1571–1579 (2002).
[12] M. Ali and J. S. Deogun, “Cost-effective implementation of multicasting in wavelength-routed networks,” J. Lightwave Technol. **18**, 1628–1638 (2000).
[13] M. Ali and J. S. Deogun, “Power-efficient design of multicast wavelength-routed networks,” IEEE J. Sel. Areas Commun. **18**, 1852–1862 (2000).
[14] K.-D. Wu, J.-C. Wu, and C.-S. Yang, “Multicast routing with power consideration in sparse splitting WDM networks,” in *IEEE International Conference on Communications, 2001. ICC 2001* (Institute of Electrical and Electronics Engineers, New York, 2001), pp. 513–517.
[15] N. K. Singhal and B. Mukherjee, “Protecting multicast sessions in WDM optical mesh networks,” J. Lightwave Technol. **21**, 884–892 (2003).
[16] X.-H. Jia, D.-Z. Du, X.-D. Hu, M.-K. Lee, and J. Gu, “Optimization of wavelength assignment for QoS multicast in WDM networks,” IEEE Trans. Commun. **49**, 341–350 (2001).
[17] J. He, S.-H. G. Chan, and D. H. K. Tsang, “Routing and wavelength assignment for WDM multicast networks,” in *Global Telecommunications Conference (GLOBECOM ’01)* (IEEE, New York, 2001), pp. 1536–1540.
[18] R. K. Pankaj, “Wavelength requirements for multicasting in all-optical networks,” IEEE/ACM Trans. Netw. **7**, 414–424 (1999).
[19] L. H. Sahasrabuddhe and B. Mukherjee, “Light-trees: optical multicasting for improved performance in wavelength-routed networks,” IEEE Commun. Mag. **37**(2), 67–73 (1999).
[20] M. Mellia, A. Nucci, A. Grosso, E. Leonardi, and M. A. Marsan, “Optimal design of logical topologies in wavelength-routed optical networks with multicast traffic,” in *Global Telecommunications Conference (GLOBECOM ’01)* (IEEE, 2001), pp. 1520–1525.
[21] R. Jain, D. Chiu, and W. Hawe, “A quantitative measure of fairness and discrimination for resource allocation in shared computer systems,” DEC Research Report TR-301 (Digital Equipment Corporation, 1984).
[22] F. K. Hwang, D. S. Richards, and P. Winter, *The Steiner Tree Problem* (Elsevier, New York, 1992).
[23] E. W. Zegura, K. Calvert, and S. Bhattacharjee, “How to model an internetwork,” in *INFOCOM ’96: Fifteenth Annual Joint Conference of the IEEE Computer Societies—Networking the Next Generation* (IEEE, New York, 1996), pp. 594–602.
[24] R. Dutta and G. N. Rouskas, “Traffic grooming in WDM networks: past and future,” in IEEE Netw. **16**, 46–56 (2002).