Journal of Interconnection Networks
Vol. 6, No. 4 (2005) 383–405

World Scientific
www.worldscientific.com

WAVELENGTH ASSIGNMENT IN MULTIFIBER OPTICAL STAR NETWORKS UNDER THE MULTICASTING COMMUNICATION MODE

RUSSELL BRANDT*

Department of Computer Science
University of California Santa Barbara, CA, 93106-5110 USA.
rbrandt@cs.ucsb.edu

TEOFILO F. GONZALEZ

Department of Computer Science
University of California Santa Barbara, CA, 93106-5110 USA.
teo@cs.ucsb.edu
http://www.cs.ucsb.edu/~teo

Received 23 January 2004
Revised 23 July 2005

This paper examines the wavelength assignment problem for single, dual, and multimessage multicasting over a star network with optical switching between fibers along the same wavelength. The specific problem we consider is given any star network, a predetermined number of fibers that connect its nodes, and a set of multicasts (or multidestination messages) to be delivered in one communication round, find a conflict free message transmission schedule that uses the least number of wavelengths per fiber. When the least number of wavelengths, λ_{min}, exceeds the number available, λ_{avl}, one may transform the schedule into one with $\lceil \lambda_{min}/\lambda_{avl} \rceil$ communication phases or rounds over the same network, but restricted to λ_{avl} wavelengths per fiber.

Keywords: Optical Networks; Wavelength Assignment; Multicasting.

1. Introduction

The ever increasing need for faster data transmission has led to extended research into the area of Optical Networks. Optical Networks allow much greater data transmission speeds than Electrical Networks, and optical switching has allowed us to retain these transmission speeds even when direct links between nodes are not available.[1,2] Furthermore, Wavelength Division Multiplexing (WDM) allows messages to be transmitted on different wavelengths (or channels) over the same fiber. In our star network the center node is a passive star coupler that joins the nodes in the network, making transmission between all nodes along the same wavelength

*Current affiliation: Lockheed Martin Corporation, russell.brandt@lmco.com.

completely optical.[2] Our networks can be single, dual, or multifiber networks depending on the number of fibers that connect adjacent nodes in the system. In dual fiber and multifiber networks, messages can be switched from one fiber to another along the same wavelength.[2,3]

High transmission speeds are needed for applications such as video conferencing, distributed data processing, scientific visualization, high speed supercomputing, and real-time medical imaging to name a few.[4] The need for these systems is growing and it is likely that future communication networks will include a large amount of multi-destination traffic.[1,4] Furthermore, since one of the biggest costs when building an optical network is the actual physical laying of the optical fibers, often many fibers may be installed at the same time, for about the same cost, resulting in multifiber networks.[5] Wide area testbeds are currently being developed, employing WDM technology to pass data over various wavelengths in real-time.[6] The problem we consider in this paper is given an $(n+1)$-node star network, a predetermined number of fibers, and a set of multidestination messages (or multicasts) to be exchanged between nodes in one communication round, find a conflict free transmission schedule that uses the least number of wavelengths per fiber. Hereafter we abbreviate wavelengths per fiber by λ/f, since λ normally identifies the number of wavelengths and f signifies fibers. Note that when the least number of wavelengths, λ_{min}, exceed the number available, λ_{avl}, one may transform our solution into one with $\lceil \lambda_{min}/\lambda_{avl} \rceil$ communication phases or rounds over the same network, but restricted to λ_{avl} wavelengths per fiber. Because of this simple equivalence we restrict our work to just finding a conflict free transmission schedule that uses the least number of wavelengths per fiber.

When we talk about our communication model we use *multicast* to indicate that a node sends a message to one or more nodes in the system (also called a *unicast* when sending to only one other node and a *broadcast* when sending to all other nodes). A star network is a set of nodes that communicate with each other by sending messages through an internal routing node (the passive star coupler in our case). The messages are then routed to all of the appropriate receiving nodes. The benefit of multifiber networks is that even though nodes must receive messages on the same wavelength from which they were sent, the receiving nodes are able to receive the message on *any* fiber. We call this optical rerouting of messages onto a *different fiber*, "switching" of the message. This is the central process that allows for better utilization of individual fibers in dual and multifiber networks.

Our problem falls under the general category of the wavelength assignment problem (WAP) which is to determine the wavelengths on which to send the required messages. A related problem is the scheduling and wavelength assignment (SWA) problem. The goal of SWA is to schedule the required messages on the available wavelengths in order to minimize the finish time. As we said before, for our problem one can easily transform a solution to the WAP to a solution to the SWA problem when the number of wavelengths per fiber is bounded by some fixed constant. A closely related problem is the wavelength and routing assignment problem (WRAP), in which both the routes and the wavelengths that each message uses must be determined. In our problem, the routes are fixed.

1.1. *Problem Definition*

We are given an $(n+1)$-node star network with each node i, other than the center node or star coupler, sending s_i multicasts. Every multicast requires the same amount of time to reach any subset of its destinations. The j^{th} multicast sent from node i, for $1 \leq j \leq s_i$, is denoted by (i, j) and has as destinations the set of nodes $d_{i,j}$. Since it does not make sense to send a multicast to the same node more than once, we represent $d_{i,j}$ by a set. Also, it does not make sense for a node to send a message to itself, so $\forall$ i and j, it must be that $i \notin d_{i,j}$. Consider Example 1 with $n = 4$, and three fibers. The routing requests s_i and $d_{i,j}$ are given in Table 1. The first two entries in Table 1 have the following meaning. Node 1 needs to send three messages. The first one is to be sent to node 2 and the other two messages must be sent to nodes 2 and 3. Node 2 does not send messages. As we show later on, under our communication rules all these messages can be transmitted in one communication round using only 2 wavelengths per fiber.

Table 1. Values s_i and $d_{i,j}$ for Example 1.

# Messages	Message Destinations		
$s_1 = 3$	$d_{1,1} = \{2\}$	$d_{1,2} = \{2,3\}$	$d_{1,3} = \{2,3\}$
$s_2 = 0$			
$s_3 = 3$	$d_{3,1} = \{1,2,4\}$	$d_{3,2} = \{2\}$	$d_{3,3} = \{4\}$
$s_4 = 1$	$d_{4,1} = \{1,2,3\}$		

We now specify the communication constraints in our model. Nodes can receive at most one message on each fiber-wavelength pair, and send at most one multicast on each fiber-wavelength pair, but not both. I.e., nodes cannot both send and receive at the same time along the same fiber-wavelength pair. At the central node in the star network we can switch a message from one fiber to another on the same wavelength, but we cannot switch the wavelength on which a message is sent. Therefore, it does not make sense to send a message on multiple fibers using the same wavelength because switching the messages across fibers can achieve the same result. However, it might be advantageous to send the same message on different wavelengths so that different nodes can receive the message on different wavelengths to avoid conflicts with all the other messages being transmitted.

In Figure 1 we give a communication assignment for Example 1. The figure depicts an assignment for each of the two wavelengths (w_1 and w_2). The assignment is represented in the rectangular box and indicates the switching that occurs in the passive star coupler. The nodes are represented twice, on the left side of the rectangle to signify sending of a message and on the right side to signify receiving of a message.

On the left side of each of the rectangles, the node-fiber pairs are indicated by a line segment that originates at a node and ends at the switching rectangle. For each node, we use the top line segment to represent fiber 1 (f_1), the next one is fiber 2 (f_2), and the bottom one is called f_3 and represents the third fiber. When one of these line segments does not continue inside the switching rectangle it indicates that there was no message being sent on that node-fiber pair. If the line segment continues as one or more line segments inside the switching rectangle, then the end points on the right side of the rectangle indicate the node-fiber pair destinations for that message. For example the message labeled A is transmitted to node 2, and the one labeled C is transmitted to nodes 1, 2, and 4. The former message corresponds to multicast $(1, 1)$ and the latter one corresponds to multicast $(3, 1)$. In Table 2 the multicasting messages (i, j) that need to be sent from each node are mapped to the actual communications labeled A through H in the schedule given in Figure 1. The only multicast that is transmitted along both wavelengths is $(4, 1)$ which corresponds to the message labeled D and H in Figure 1.

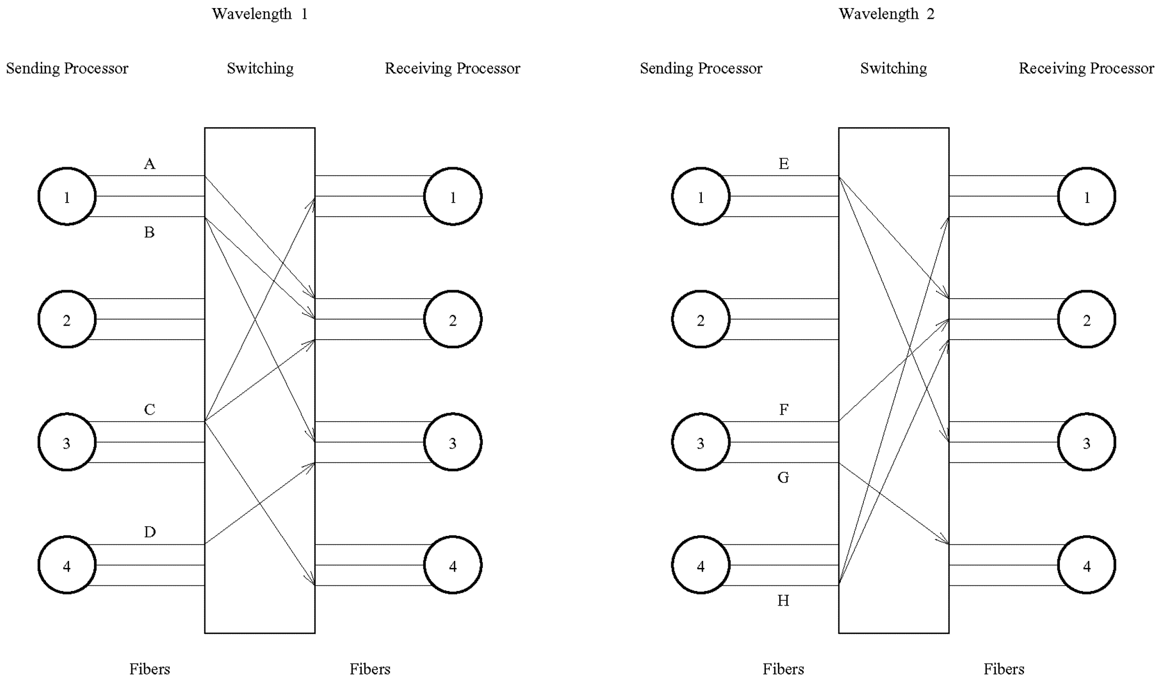

Fig. 1. Communication Assignment for Example 1.

For completeness, we now formally define assignments and communication schedules. Readers may skip the remaining part of this subsection if they are comfortable with our informal definitions. Given an $n+1$ node system with g fibers, and the sets of multicasts $d_{i,j}$, the assignment $S^{w}_{i,l}$ specifies all the node-fiber destination pairs for a message originating at the i^{th} node and being initially transmitted through

Table 2. Communication Schedule for Example 1.

Multicasts in Example 1	Transmission of Message in Figure 1
$(1,1)$	A
$(1,2)$	B
$(1,3)$	E
$(3,1)$	C
$(3,2)$	F
$(3,3)$	G
$(4,1)$	D and H

the l^{th} fiber over wavelength w. Figure 1 depicts all the assignments that are sufficient to transmit all the messages given in Example 1. For each wavelength w every node-fiber destination pair may be in at most one assignment, i.e., one of the sets $S^w_{i,l}$. The set S^w_0 denotes the node-fiber pairs that will not receive a message on wavelength w. In other words, an assignment is a partition of the nodes-fiber pairs for each wavelength w into sets as specified in (1.1).

$$\begin{array}{l} \forall w \ partition \ \{1,2,...,n\} \ X \ \{1,2,...,g\} \\ into \ S^w_0, \ S^w_{1,l}, \ S^w_{2,l}, \ ..., \ S^w_{n,l} \ for \ 1 \leq l \leq g \end{array} \quad (1.1)$$

In Figure 1 we give an assignment for the two wavelengths (w_1 and w_2). On wavelength w_1, the message labeled "A" is sent from node 1 on f_1 and received by node 2 on f_1. So, $S^1_{1,1} = \{(2,1)\}$. No messages are sent from node 1 on f_2, so $S^1_{1,2} = \emptyset$. The multicast labeled "B" is sent from node 1 on f_3 and received by both node 2 on f_2 and by node 3 on f_2, so $S^1_{1,3} = \{(2,2),(3,2)\}$. Remember that since nodes cannot send and receive on the same fiber-wavelength pair, once node 1 sends on $f_1 w_1$ it cannot receive on $f_1 w_1$. The nonempty sets $S^w_{i,l}$ for the assignment defined in Figure 1 are given in Table 3.

Table 3. Assignments given in Figure 1.

Wavelength w_1	Wavelength w_2
$S^1_0 = \{(1,1),(1,3),(3,1),(4,1),(4,2)\}$	$S^2_0 = \{(1,1),(1,2),(3,1),(3,3),(4,2),(4,3)\}$
$S^1_{1,1} = \{(2,1)\}$	$S^2_{1,1} = \{(2,1),(3,2)\}$
$S^1_{1,3} = \{(2,2),(3,2)\}$	$S^2_{3,1} = \{(2,2)\}$
$S^1_{3,1} = \{(1,2),(2,3),(4,3)\}$	$S^2_{3,3} = \{(4,1)\}$
$S^1_{4,1} = \{(3,3)\}$	$S^2_{4,3} = \{(1,3),(2,3)\}$

Note that if $S^w_{i,l} \neq \emptyset$, then $(i,l) \in S^w_0$, since for each wavelength w a node-fiber pair used to send a message cannot be used to receive a message. Additionally, for every fiber l the node-fiber pair (i,l) should not be in any set $S^w_{i,l}$ since nodes do not need to send a message to themselves.

A *communication schedule* for a problem instance is a partition of the assignments such that each multicast can be realized by a distinct set of assignments in the partition. In other words, a partition of the assignments $S_{k,l}^{w}$ into the sets $t_{i,j}$ is a communication schedule if for every multicast (i, j) in the problem instance we have that $d_{i,j}$ is exactly equal to all the nodes in the node-fiber destination pairs $S_{k,l}^{w} \in t_{i,j}$. Table 4 shows the nonempty sets $t_{i,j}$ for the assignments in Table 1 for the problem instance given in Example 1.

Table 4. Nonempty $t_{i,j}$ sets for Example 1..

Set $t_{i,j}$	Partition of $S_{k,l}^{w}$	Message Destinations
$t_{1,1}$	$\{S_{1,1}^{1}\}$	{2}
$t_{1,2}$	$\{S_{1,3}^{1}\}$	{2, 3}
$t_{1,3}$	$\{S_{1,3}^{2}\}$	{2, 3}
$t_{3,1}$	$\{S_{3,1}^{1}\}$	{1, 2, 4}
$t_{3,2}$	$\{S_{3,1}^{2}\}$	{2}
$t_{3,3}$	$\{S_{3,3}^{2}\}$	{4}
$t_{4,1}$	$\{S_{4,1}^{1}, S_{4,3}^{2}\}$	{1, 2, 3}

1.2. *Related Work*

The SWA problem was shown to be NP-Hard for both preemptive (operations can be stopped or preempted, and resumed at a later time) and non-preemptive (operations cannot be preempted) cases.[7] Bampis and Rouskas[7] developed efficient approximation algorithms for both cases. Li and Simha[3] consider the offline WAP over multiple fibers in a unicast only environment. The main result in Ref. 3 is that in a multifiber network, switching messages between the fibers increases wavelength utilization. For star networks, WAP is known to be NP-Complete over a single fiber[5], but in Ref. 3, optimal polynomial time algorithms for the cases of dual and multifiber networks are developed. For ring networks, the dual and multifiber cases are known to be NP-Complete and upper bounds are established for both cases.[3] Several papers consider multicasting environments over all optical single fiber WDM networks.[1,4] Thaker and Rouskas[2] survey multicast scheduling algorithms (MSAs) in single fiber star networks.

In this paper we extend the above research by considering the offline version of the WAP in multifiber multicast networks. We use an optical star network as our model and develop bounds for single, dual, and multifiber networks using WDM. As pointed out before, our results can be easily extended to the SWA version of our problem.

1.3. *Conventions and Outline*

We introduce the notation $(\alpha \mid \beta \mid \gamma)$ to refer to subsets of problem instances as follows. The first and second terms, α and β, specify that every node in the system can receive at most α messages and send at most β multicasts. Additionally, there is at least one node in the system that receives α messages and sends β multicasts. We consider the cases where the values of α and β are 1, 2, or n. Corresponding to the values of β, the cases are called *single, dual* and *multimessage multicasting*, respectively. The number of fibers is γ. We assume that all fibers have the same number of wavelengths.

For every system, ($\alpha \mid \beta \mid \gamma$), we exhibit a *lower bound* and an *upper bound* on the number of wavelengths required per fiber. By the lower bound x we mean that *there exists* at least one problem instance, in this specific system, that requires at least x wavelengths per fiber in order to achieve conflict free transmission of all its messages. By the upper bound y, we mean that *every* problems instance, in this specific system, can achieve conflict free transmission of all its messages with at most y wavelengths per fiber.

Throughout the paper we use the terms *wavelength* and *color* interchangeably. Additionally, although the internal routing node is always present, in this paper we simplify our descriptions and figures by ignoring it and showing all messages going from one node in the system directly to another node in the system. We call this representation the *message directed multigraph.* The message directed graph for Example 1 is given in Figure 2. Strictly speaking it is not a multigraph because there are thin lines joining (bundling together) directed edges (messages) in order to represent multidestination messages. We corrupt existing notation and refer to the graphs as multigraphs. The message directed multigraph is called the *message multigraph* when all messages have a single destination (unicasting), and we ignore the direction of the edges. A solution is an assignment of fiber-wavelength pairs to both ends of each edge in such a way that the restrictions discussed above are satisfied. E.g., the wavelength must be the same on both ends of an edge; no two different messages originating or ending at the same vertex can be assigned to the same fiber-wavelength pair; etc. When there is just one fiber and all messages are unicasts, our problem corresponds to the chromatic index problem which is defined below.

The *chromatic index* (or edge coloring problem) of a multigraph is the minimum number of colors required to color the edges of a multigraph so that no two edges emanating from the same vertex have the same color. The corresponding decision problem is known to be NP-Complete[8] and there are several well known approximation techniques for the edge coloring problem. The best-known approximation algorithm appears in Ref. 9. Our problem is equivalent to the edge coloring problem when all messages are unicasts and the number of fibers is equal to one. It is not known whether or not our problem is NP-complete when there are two or more fibers.

For every problem instance a trivial lower bound is simply $\left\lceil \frac{\alpha+\beta}{\gamma} \right\rceil$ wavelengths per fiber. About half of the lower bounds we have obtained are equal to this trivial bound and the remaining lower bounds are obtained using examples that range

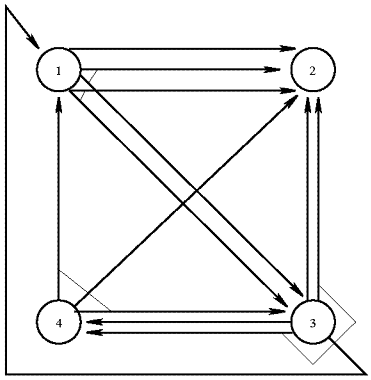

Fig. 2. Message Directed Graph for Example 1.

from simple to very complex problem instances. Almost all of our results are given in Table 5. LB stands for Lower Bound, and UB stands for Upper Bound in Table 5. When there is just one value it means we have tight bounds, i.e. LB=UB. It is assumed that $n > 2$ and $g > 2$, except for problem $(2 \mid n \mid g)$ where the lower bound is valid for $g \geq 4$; and, the UB valid for $g \geq 3$. Detailed proofs of all the results appear in.[10]

This paper includes the most interesting of the results in Table 5. The remaining results listed in the table appear in Ref. 10 and were obtain by using similar arguments to the ones used in this paper. The paper is organized as follows. In Section 2 we present our (simple) bounds for the $(1 \mid 1 \mid g)$ problem. In Section 3 we present our lower bound of four for the $(2 \mid 1 \mid 1)$ problem and then we show that every such problem instance never needs more than 5 colors (wavelengths). Furthermore, such coloring can be easily constructed in $O(n)$ time. In Section 4 we present a complex lower bound and a simple upper bound for the $(n \mid 1 \mid g)$ problem. A lower bound for the $(2 \mid 2 \mid 2)$ problem is established in Section 5. Section 6 has an interesting upper bound for the $(n \mid n \mid g)$ problem. We present our conclusions in Section 7.

2. Bounds for the $(1 \mid 1 \mid g)$ Problem

2.1. *One Fiber (LB = UB = 3)*

Let us consider single fiber networks where every node can receive up to one message and send up to one multicast. The lower bound and the upper bound are both equal to 3 wavelengths when there is one fiber. To establish the lower bound, consider the simple problem instance with three nodes forming a loop (Figure 3). It requires three colors because each of the three edges must be colored differently than the other two.

Next, we show that the upper bound is also three. Every problem instance with nodes sending up to one multicast and receiving up to one message can be viewed as a collection of disjoint subgraphs. Each subgraph is a tree, plus there could be

Table 5. Lower and Upper Bounds for the Number of Wavelengths required per Fiber.

Fibers	send ≤ 1 rec. ≤ 1	send ≤ 1 rec. ≤ 2	send ≤ 1 rec. $\leq n$
1	3	LB = 4 UB = 5	LB = 2n UB = $2n + \left\lceil \frac{n}{2} \right\rceil$
2	1	2	n
g	1	1	$\left\lceil \frac{n}{g-1} \right\rceil$

Fibers	send ≤ 2 rec. ≤ 1	send ≤ 2 rec. ≤ 2	send ≤ 2 rec. $\leq n$
1	3	6	3n
2	2	3	LB = n UB = $\left\lceil \frac{3n}{2} \right\rceil$
g	1	1	$\left\lceil \frac{n}{g-2} \right\rceil$

Fibers	send $\leq n$ rec. ≤ 1	send $\leq n$ rec. ≤ 2	send $\leq n$ rec. $\leq n$
1	n+1	2n+2	n^2+n
2	$\left\lceil \frac{n+1}{2} \right\rceil$	LB = $\left\lceil \frac{n+2}{2} \right\rceil$ UB = n+1	LB = $\left\lceil \frac{\frac{n^2}{2}+n}{2} \right\rceil$ UB = $\left\lceil \frac{n^2+n}{2} \right\rceil$
g	$\left\lceil \frac{n+1}{g} \right\rceil$	LB = $\left\lceil \frac{n+2}{g} \right\rceil$ UB = $\left\lceil \frac{2n+2}{g} \right\rceil$	LB = $\left\lceil \frac{\frac{n^2}{g}+n}{g} \right\rceil$ UB = $\left\lceil \frac{2n^2}{g^2-g} \right\rceil$

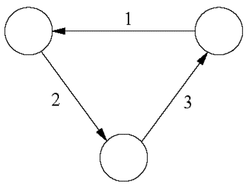

Fig. 3. Three Colorable Multigraph.

a single additional edge in the network from a node to the root. We call the edges in each tree, *tree edges*; and the single additional edge is called the *back edge*. We color each of the disjoint trees with three colors as follows. Color the tree edges from the first level to the second level with color 1. Color the tree edges from the next level to the following level with color 2. Repeat this process through all of the levels, alternating between color 1 and color 2. Finally, color the back edge with color 3. Clearly, this coloring is a valid one.

Theorem 2.1. The lower and upper bound for the $(1 \mid 1 \mid 1)$ problem is three. Furthermore the upper bound can be generated for any problem instance in linear time with respect to the input length.

Proof. The proof of the lower and upper bound follows from the above discussion. It is simple to show that the above wavelength assignment procedure can be easily constructed in linear time with respect to the number of nodes and edges in the graph. □

2.2. *Two or More Fibers (LB = UB = 1)*

In this case, the lower bound and the upper bound are both equal to $1\lambda/f$. The lower bound is obvious since there is a node that receives a message and sends a multicast. An algorithm to "color" the messages is a simple one. All messages are sent out on f_1w_1 (fiber one and wavelength one). Next, we will utilize the dual fiber network and switch all the messages to fiber two, i.e., f_2w_1. Since every message is sent out on f_1w_1, received on f_2w_1, and every node can receive at most one message, there will not be any conflicts. Therefore, every problem instance can be colored with one color per fiber. Note that when there are more than two fibers available the upper bound does not decrease.

Theorem 2.2. The lower and upper bound for the $(1 \mid 1 \mid g)$ problem, for $g \geq 2$, is 1 color per fiber. Furthermore the upper bound can be generated for in linear time with respect to the input length.

Proof. The proof of this theorem follows from the above discussion. □

3. Bounds for the (2 | 1 | 1) Problem

For single message multicasting where every node can receive up to two messages over a single fiber, we show a lower bound of 4 λ/f and an upper bound of 5 λ/f. It is important to note that all the figures in this section do not have the line that bundles multidestination messages. The reason for this was to make the figures easier to follow. This does not create an ambiguity because every node sends at most one multicast.

Let us now establish the lower bound with the problem instance given in Figure 4. We now show that this network cannot be colored with 3 colors (which correspond to f_1w_1, f_1w_2, and f_1w_3).

Theorem 3.1. The lower bound for the $(2 \mid 1 \mid 1)$ problem is 4 colors per fiber.

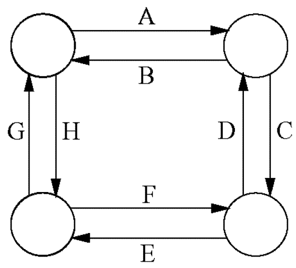

Fig. 4. Multigraph requires four colors.

Proof. The proof is by contradiction. Suppose that the above multigraph can be colored with three colors. Without loss of generality we can assign color 1 to edge A, color 2 to edge C, and color 3 to edge D. Since there are only three colors available, it must be that at each vertex, the edges emanating from that vertex are colored identically. Therefore, edge B must be colored with color 2, because edge C has been assigned color 2. Similarly, edge H must be colored with color 1, and edge E must be colored with color 3. Now edges F and G cannot be colored with any of the three colors because it creates a conflict (Figure 5). □

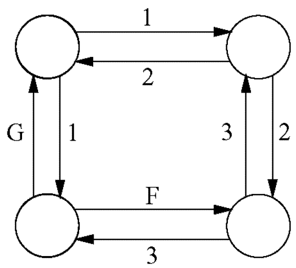

Fig. 5. Edges F and G cannot be colored with colors 1, 2, or 3.

Figure 6 shows a color assignment using 4 colors for the multigraph in Figure 4.

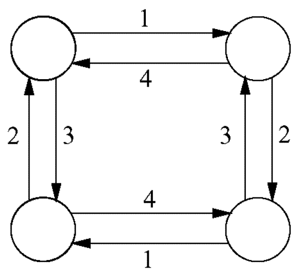

Fig. 6. Multigraph is four colorable.

To show the upper bound of 5 λ/f we first consider the subproblem where every node has exactly two outgoing edges (i.e., every multicast has two destinations) and two incoming edges. We present a constructive proof that shows that every such problem instance is colorable with 5 colors. Then we show how to use this result to color all problem instances with 5 colors. The resulting algorithm takes $O(n)$ time.

Lemma 3.1. *Every problem instance where every node has exactly 2 outgoing edges (i.e., each multicast has two destinations) and exactly 2 incoming edges can be colored with 5 colors (which correspond to f_1w_1 through f_1w_5). Furthermore, this process can be easily implemented to take $O(n)$ time, where n is the number of vertices in the multigraph.*

Proof. Our proof is constructive. We consider each node one at a time. When considering a node x we color its incoming edges and in some special cases, we must recolor some previously colored edges. We will refer to the nodes where the two incoming messages to node x originate the "parents" of node x and label them $P1$ and $P2$. If any of these edges have not yet been colored then assume (without actual assigning) they have been colored with any one of the five colors without creating a conflict with the other incoming edge to $P1$ or $P2$. Since both messages leaving a node belong to the same multicast, they can be colored with the same or with different colors. This allows us to ignore how the messages that nodes $P1$ and $P2$ send to nodes other than x are colored. There are three cases depending on the number of different colors assigned to the incoming edges to $P1$ and $P2$. The proofs for the first two cases are simple and similar to each other, while the proof of the third case is more complex.

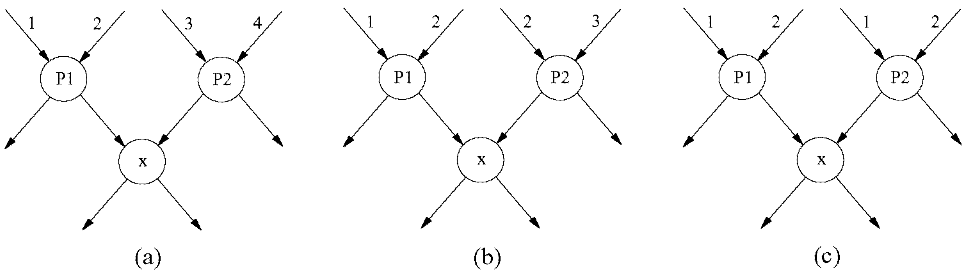

Fig. 7. The three different cases.

Case 1: The incoming messages to $P1$ have different colors than those incoming to $P2$.

Without loss of generality assume that the incoming messages to $P1$ are colored 1 and 2; and the incoming messages to $P2$ are colored 3 and 4 (Figure 7(a)). Now, without considering the edges emanating from node x, the edge from $P1$ to x could be colored 3, 4, or 5; and the edge from $P2$ to x could be colored 1, 2, or 5. The number of these available colors would be the least if the edges leaving node x have been assigned two different colors. This could eliminate at most two color choices from both of node $x's$ incoming edges. But even so, node $x's$ incoming edges could still be colored with the remaining colors, because in all cases there will be at least one color available in one of the incoming edges to node x that is not available in the other incoming edge. Given this we know it is possible to color node $x's$ incoming edges with the remaining colors.

Case 2: The incoming messages to $P1$ have one color in common with those incoming to $P2$.

Without loss of generality, we may assume that the incoming messages to $P1$ are assigned colors 1 and 2; and the incoming messages of $P2$ to colors 2 and 3 (Figure 7(b)). Now, without considering the edges emanating from node x, the edge from $P1$ to x could be colored 3, 4, or 5; and the edge from $P2$ to x could be colored 1, 4, or 5. The number of these available colors would be the least if the edges leaving node x have been assigned two different colors. This could eliminate at most two color choices from both of node $x's$ incoming edges; but in all cases either there will be at least one color available in one of the incoming edges to node x that is not available in the other and vice versa, or one incoming edge has at least one available color and the other one has two. Given this we know node $x's$ incoming edges can be colored with the remaining colors.

Case 3: The incoming messages to $P1$ have both colors in common with those incoming to $P2$.

Without loss of generality we can assign the incoming messages of $P1$ and the incoming messages of $P2$ each to colors 1 and 2 (Figure 7(c)). If an edge leaving node x is an incoming edge to $P1$ or $P2$, then the incoming edges to node x can be colored with the remaining colors, as in Case 2. In all other cases, we proceed as follows. Without considering the edges emanating from node x, the edge from $P1$ to x and the edge from $P2$ to x could each be colored 3, 4, or 5. If the outgoing edges from node x are colored, then remove such coloring. Those edges, leaving node x, may each be colored with two colors in such a way that they will not conflict with the coloring of the other edges at the nodes where they end. Let S_1 be the set of colors of which the first edge leaving node x can be colored and let S_2 be the set of colors of which the second edge leaving node x can be colored.

If S_1 and S_2 have a color in common, then these edges (the edges leaving node x) can be assigned one such color and there will be at least two possible colors that can be assigned to node $x's$ incoming edges. Therefore, a valid coloring is possible in this case.

On the other hand, let us consider the case when S_1 and S_2 do not have a color in common. Since S_1 and S_2 have two colors each, there is at least one color in S_1 or S_2 that is not color 3, 4, or 5. Assume, without loss of generality, that such color is $s \in S_1$. Now, assign color s to the first edge leaving node x, and assign one of the colors from S_2, let us call it t, to the second edge leaving node x. The incoming edges to node x may then be assigned to colors $\{3, 4, 5\} - t$. Therefore, a valid coloring is also possible in this case.

It is simple to show that the above process can be implemented to take $O(n)$ time, where n is the number of vertices in the multigraph. □

Let us now return to the more general problem where nodes can send up to one multicast to any number of destinations, and receive up to two messages. We now show that any such instance can be colored using 5 colors.

Theorem 3.2. Every instance of the (2 | 1 | 1) problem can be colored with 5 colors. Furthermore, this process can be easily implemented to take $O(n)$ time, where n is the number of vertices in the multigraph.

Proof. Let G_1 be any message directed multigraph. Every node in G_1 has an in-degree of at most two and an out-degree of at most n. Consider any node x of out-degree zero or one (see Figure 8). Given any 5-coloring of $G_1 - \{incoming\ edges\ to\ x\}$ we now show that the incoming edges to node x in G_1 can be colored with the colors available. For any node x of out-degree zero (Figure 8(a)) or out-degree one (Figure 8(b)) with in-degree two we can color its incoming edges without conflicts as follows. The edge from node p to node x cannot be of the same color as the two edges being received by node p (Figure 8). Similarly for the edge from node q to node x, two color choices are not possible. This leaves at least 3 colors available to color 2 edges. The number of available colors decreases at most by one when we color the outgoing edge to node x. In the worst case this leaves at least 2 colors available to color 2 edges. Therefore, there is always a valid coloring for the two incoming edges to x in G_1. It is simple to see that the same argument holds for any node x of out-degree zero or out-degree one with in-degree less than two, because there are fewer restrictions.

We now transform G_1 into G_2 such that every node in G_2 is such that either the out-degree is one and the in-degree is zero, or the out-degree is at least two and the in-degree is at most two. The message directed multigraph G_2 is constructed by using the above transformation as follows: Initially we define multigraph G_2 to be a copy of G_1. Now, we delete from G_2 the incoming edges to every node with

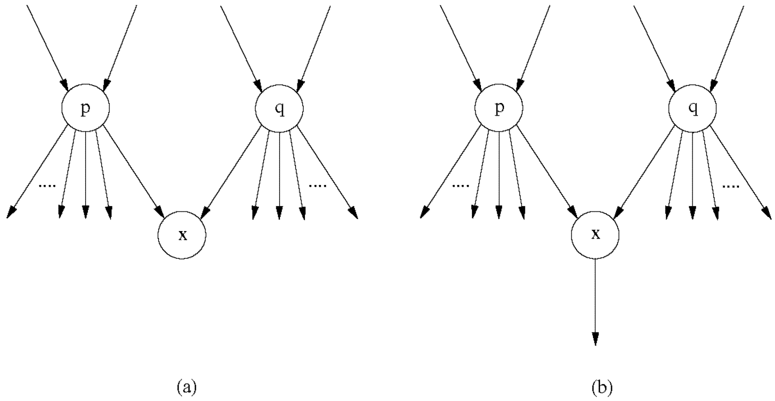

Fig. 8. Coloring a Multigraph with five colors.

out-degree zero or one. We repeat this process until no such vertex exists. Then we delete the isolated vertices (nodes without any incoming or outgoing arcs), if any. The resulting multigraph is such that for every node either the out-degree is one and the in-degree is zero *or* the out-degree is at least two and the in-degree is at most two.

Let id_i be the in-degree of vertex i in G_2 and let od_i be the out-degree of vertex i in G_2. Clearly, $\Sigma_{G_2} id_i = \Sigma_{G_2} od_i$. By construction, we know that for each i either $id_i = 0$ and $od_i = 1$, or $id_i \leq 2$ and $od_i \geq 2$. In any multigraph with n nodes, let y be the number of nodes with $id_i = 0$ and $od_i = 1$, and let $n - y$ be the number of nodes with $od_i \geq 2$. Therefore, $\Sigma_{G_2} id_i \leq (0 * y) + 2(n - y)$ and $\Sigma_{G_2} od_i \geq (1 * y) + 2(n - y)$.

Since $\Sigma_{G_2} id_i = \Sigma_{G_2} od_i$, it must then be that $2(n-y) \geq y + 2(n-y)$ and therefore $y \leq 0$. So, $\Sigma_{G_2} id_i \leq 2n$ and $\Sigma_{G_2} od_i \geq 2n$. Now, since $\Sigma_{G_2} id_i = \Sigma_{G_2} od_i$, it must be that every node i is of in-degree 2 and out-degree 2. By Lemma 3.1, such a problem is 5-colorable. Therefore, every problem instance where all nodes can send at most one multicast and receive at most two messages can be colored with 5 colors.

It is simple to see that this process can be implemented to take $O(n)$ time, where n is the number of nodes in the multigraph. □

To demonstrate the constructive proof of Theorem 3.2, consider the problem instance in Figure 9. Initially, ignore the numbers shown on the edges. The first step is to delete the edges incident to nodes I and J in Figure 9(a) because these nodes each have zero outgoing edges. We obtain the multigraph in Figure 9(b). Now delete from the multigraph in Figure 9(b), the incoming edges to nodes G and H because these nodes have an out-degree of one and zero, respectively. The resulting multigraph is given in Figure 9(c). Next, from Figure 9(c), delete the incoming edges to node F and obtain Figure 9(d). We have reduced the problem to one in which all nodes have an in-degree of two and an out-degree of two, when we ignore the isolated vertices. Now, as illustrated in Figure 9(d), we color the edges using the constructive proof of Lemma 3.1. We choose nodes one at a time in an arbitrary

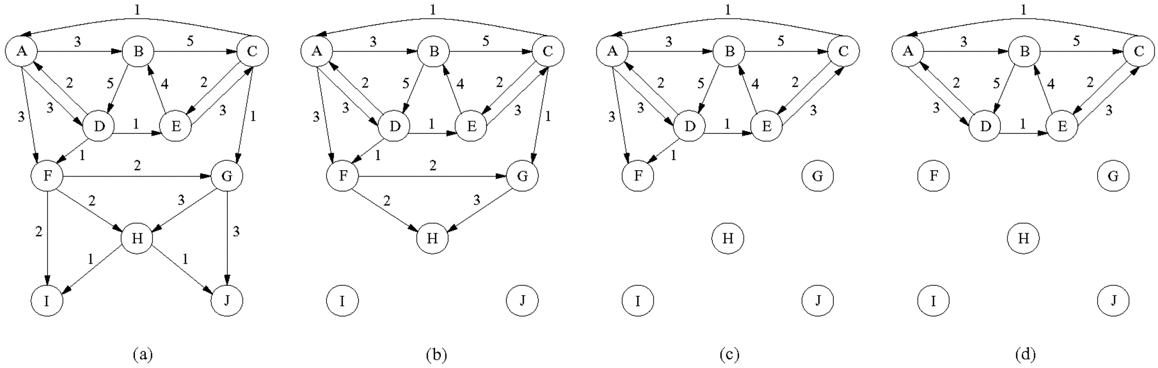

Fig. 9. (a) Example for constructive proof of Theorem 3.2. The steps of the constructive proof are illustrated in (b), (c) and (d).

order and color the incoming edges with the smallest available color. Specifically, we start with node A and color its incoming edges 1 and 2. Next, choose node E and color its incoming edges 1 and 2. Node B is next, with its incoming edges being colored 3 and 4. Node C follows, with its incoming edges being colored 3 and 5, as does node D with its incoming edges also being colored 3 and 5. The process continues with the coloring of the edges deleted during the first three reduction steps. We color the edges incoming to nodes F, G, H, I, and J with the resulting coloring shown in Figure 9(a).

4. Bounds for the (n | 1 | g) Problem

For the ($n \mid 1 \mid g$) problem the lower bound and the upper bound are both equal to $\lceil \frac{n}{g-1} \rceil \lambda/f$, where $g \geq 2$ is the number of fibers. To show the lower bound we give a problem instance that cannot be colored using $(\lceil \frac{n}{g-1} \rceil -1)\lambda/f$. The problem instance we construct is such that no matter what solution we provide using $\lceil \frac{n}{g-1} \rceil \lambda/f$, there is a sequence of n-node sets that satisfy the following properties. Each of these n-node sets has the same colors available for their multicasts. The last of these n-node sets has less than $\frac{n}{g}$ of such colors available. Furthermore, all of these nodes in this last set will send a message to the same node x. Therefore, we will have a conflict because there will not be enough color-fiber pairs available for the messages that node x receives. Therefore, there is no solution that uses $\lceil \frac{n}{g-1} \rceil \lambda/f$.

We now describe the structure of our problem instance. There are $k+1$ levels of nodes with ϕ_i nodes in level i, for $0 \leq i \leq k$. There are n nodes at level zero ($\phi_0 = n$) and each of these nodes sends a message to all the nodes at level 1. The arguments that we use require that there be at least n nodes in level 1, all receiving their messages on the same n element subset of fiber-wavelength (fiber-color) pairs no matter how the edges from level zero to level one are colored. Since each node receives its messages in an n-element subset of the fiber-color pairs, then we will be able to guarantee that at least n nodes at level 1 receive all their messages in the same n-element subset of fiber-color pairs by letting $\phi_1 = \binom{(\lceil \frac{n}{g-1} \rceil -1)g}{n}(n-1)+1$.

Inductively, our arguments will require that a set of n nodes at level i, for $1 < i < k$, that receive their messages on the same n-element subset of fiber-color pairs will all send their multicast to a set of nodes at level $i+1$ and that an n element subset of those nodes receive their messages on the same fiber-color pair. Since we do not know which subset of n nodes receives their messages on the same n-element subset of fiber-color pairs, our construction will have every n-element subset of n nodes at level i send its message to a set of ϕ_1 distinct nodes at the next level. Therefore, $\phi_i = \binom{\phi_{i-1}}{n}\phi_1$ for $i > 1$. Note that this multigraph has a huge number of nodes; however, this number of nodes is finite. The number of nodes could be decreased because not all the colors will be available for the send operation of the special n nodes identified at level i, but doing this would complicate our proof.

Suppose that there is a solution to the problem instance defined above using $(\lceil \frac{n}{g-1} \rceil - 1)\ \lambda/f$ colors on g fibers. Let us now identify a set of n nodes, S_i, at each level i, for $i \geq 1$. The first set S_1 are n nodes in level 1 that receive all their messages in the same n-element subset of fiber-color pairs, which we know exist. The set of n nodes at level i, referred to by S_i, for $i \geq 2$, is an n-element subset of nodes at level i all of which receive their messages on the same set of fiber-color pairs at level i and whose n incoming messages originate at the nodes in S_{i-1}. By using arguments similar to the ones above we can establish that such set exists. Let β_i be the colors on which all the messages of a node in S_i are received. Clearly, all the nodes in S_i receive their messages on the same n-element subset of fiber-color pairs, therefore all the nodes in S_i receive their messages on the same colors and therefore β_i is the same for all of them. We will show that $\beta_1 > \beta_2 > \ldots > \beta_k$. Then by making k large enough we will reach a contradiction since there will g fibers and fewer than n/g colors which means that there will not be enough fiber-color pairs for the n messages arriving at any node in S_k.

Before we prove this fact it is important to consider a problem instance of the form given above with $n = 21$, $g = 3$ and $k = 4$. We claim that the set of multicasts cannot be delivered with $(\lceil \frac{n}{g-1} \rceil - 1) = 10$ wavelengths per fiber. Suppose there is a solution. Identify the sets S_1, S_2, S_3 by the procedure defined above. Now let us consider the 21 nodes in S_1 that all receive their messages on the same fiber-color pairs. Assume without loss of generality that they receive them on all the $\beta_1 = 10$ colors. Since there are three fibers, at least one color must be used on the three fibers. Therefore the set of nodes S_1 must send all their messages on at most 9 colors and all the nodes in S_2 must receive their messages on the same 9 colors, i.e, $\beta_2 = 9$. Since there are three fibers, at least three color must be used on the three fibers. Therefore the set of nodes S_2 must send all their messages on at most $9 - 3$ colors plus the unused color (since only nine colors were used to receive the messages). Hence, all the nodes in S_3 must receive their messages on these 7 colors, i.e, $\beta_3 = 7$. Since there are three fibers all of these three colors must be used on the three fibers. So the set of nodes S_3 must send all their messages using at most 3 colors (the unused colors since only seven colors were used to receive all the messages). But then there will be only 9 fiber-color pairs for every node receiving messages from all the n nodes in S_3. So there is a conflict which contradicts our assumption that all the messages can be transmitted using 10 colors. In what follows we generalize these arguments to show that it is impossible to use $\beta \leq \lceil \frac{n}{g-1} \rceil - 1$ colors on problem instances with n nodes and g fibers. It is important to note that our argument will not hold on the above problem instance when there are 11 colors. The reason for this is that all of these colors will be available to send messages at every level because it is possible to use each color in at most 2 fibers for the messages incoming to each set S_i.

Let us generalize our argument to apply to problem instances with any number of nodes and fibers. Let β be the number of colors in the g fibers in which all the messages that node x in S_i (at any level $i > 1$) receives. Clearly, $\beta \leq \lceil \frac{n}{g-1} \rceil - 1$. It must be that $\beta g \geq n$, as otherwise there is no feasible coloring of the incoming messages at the node. Therefore, $\lceil \frac{n}{g} \rceil \leq \beta \leq \lceil \frac{n}{g-1} \rceil - 1$. The number of colors that are available in the system and were not used in the input messages to node x is $(\lceil \frac{n}{g-1} \rceil - 1 - \beta)$. Clearly all of these colors can be used for the multicast operation of node x. In addition to these colors available for the messages to be sent out, we also have some of the β colors used for receiving the messages. In particular, the colors that were not used in all the fibers. We will bound the number of such colors.

We know that $\beta \leq \lceil \frac{n}{g-1} \rceil - 1$ (or equivalently $\beta < \frac{n}{g-1}$), which reduces to $(\beta g - \beta) < n$ and therefore $(\beta g - n) < \beta$. Additionally, we know $(\beta g - n) \geq 0$, since $\beta \geq \lceil \frac{n}{g} \rceil$. Therefore, the number of fiber-wavelength (fiber color) pairs that are not used in the receiving end of node x is $(\beta g - n)$. Since $(\beta g - n) < \beta$, we know that at most $(\beta g - n)$ of these unused fiber-wavelength pairs have a unique color and are available for the multicast operation (sending messages) of node x. Therefore, the maximum number of colors available for the multicasting send operation of node x is:

$$(\beta g - n) + \left(\left\lceil \frac{n}{g-1} \right\rceil - 1 - \beta \right) \tag{4.1}$$

Lemma 4.1. *If node x receives all its messages in β different colors, then, when the number of fibers g is at least two, the maximum number of colors available for its multicasting send operation is less than β, for $\lceil \frac{n}{g} \rceil \leq \beta \leq \lceil \frac{n}{g-1} \rceil - 1$. In other words,*

$$(\beta g - n) + (\lceil \frac{n}{g-1} \rceil - 1 - \beta) < \beta.$$

Proof. As established above, the maximum number of colors available for any node's multicasting send operation is $(\beta g - n) + (\lceil \frac{n}{g-1} \rceil - 1 - \beta)$. Now, we need to show that this expression is less than β; or equivalently, $\beta(g-1) - n + \lceil \frac{n}{g-1} \rceil - 1 - \beta < 0$. Let $\beta = \lceil \frac{n}{g-1} \rceil - 1 - h$ for some $0 \leq h \leq (\lceil \frac{n}{g-1} \rceil - \lceil \frac{n}{g} \rceil - 1)$. Substituting in the above expression, we have $(\lceil \frac{n}{g-1} \rceil - 1 - h)(g-1) - n + h < 0$. Next, we let $n = c(g-1) + k$

for some positive integers c and $0 \leq k < g-1$; so, the expression becomes:

$$\left(\left\lceil \frac{c\,(g-1)+k}{g-1} \right\rceil - 1 - h\right)(g-1) - n + h < 0 \tag{4.3}$$

If $k > 0$, then Equation 4.3 becomes $(c+1-1-h)(g-1)-n+h < 0$. This expression reduces to $-h(g-2)-k < 0$, and this always holds because $g \geq 2, h \geq 0$ and $k > 0$. On the other hand, if $k = 0$, then Equation 4.3 becomes $(c-1-h)(g-1)-n+h < 0$. This expression reduces to $-(g-1)-h(g-2) < 0$, and this always holds because $g \geq 2$ and $h \geq 0$.

Therefore, if a node receives all messages in β different colors, then the maximum number of colors available for its multicasting send operation is less than β. □

It is important to note that when $\lceil \frac{n}{g} \rceil \leq \beta \leq \lceil \frac{n}{g-1} \rceil$ a lemma similar to the previous one cannot be proved. We now establish our lower bound for the problem instance defined above.

Theorem 4.1. *For $g \geq 2$ the lower bound for $(\ n \mid 1 \mid g\ )$ is $\lceil \frac{n}{g-1} \rceil\ \lambda/f$.*

Proof. By proving Lemma 4.1, we have established that, in our problem instance when the number of fibers is greater than or equal to two ($g \geq 2$), $\beta_i > \beta_{i+1}$. Since the number of colors which is used as input to the nodes in sets S_i is always decreasing as we increase i, by having at most $(\lceil \frac{n}{g-1} \rceil - \lceil \frac{n}{g} \rceil)$ levels of nodes, the number of colors which are used for a subset of n nodes will fall below $\frac{n}{g}$. Therefore, there will not be enough fiber-wavelength pairs to color the incoming message for these nodes. This contradicts that our problem instance could be colored with $\lceil \frac{n}{g-1} \rceil - 1\ \lambda/f$ colors. □

Theorem 4.2. *For $g \geq 2$ the upper bound for $(\ n \mid 1 \mid g\ )$ is $\lceil \frac{n}{g-1} \rceil\ \lambda/f$. Furthermore, a color assignment can be constructed in linear time with respect to the input length.*

Proof. Now let us establish an upper bound for the maximum number of colors per fiber which can be used to color all problem instances. To obtain the upper bound of $\lceil \frac{n}{g-1} \rceil\ \lambda/f$, color the incoming messages at every node (in any order) with $f_1 w_1$ through $f_1 w_{\lceil \frac{n}{g-1} \rceil}$, $f_2 w_1$ through $f_2 w_{\lceil \frac{n}{g-1} \rceil}$, and so on up to $f_{g-1} w_1$ through $f_{g-1} w_{\lceil \frac{n}{g-1} \rceil}$. (Note: when $\frac{n}{g-1}$ is not an integer, not all of these fiber-wavelength pairs will be necessary.) Next, we will utilize fiber g for the multicasts and switch messages to their appropriate fibers. I.e., for every message to be received on $f_y w_z$ (for $y = 1, 2, \ldots, g-1$ and $z = 1, 2, \ldots, \lceil \frac{n}{g-1} \rceil$), send it out on $f_g w_z$ and then switch it to $f_y w_z$. Clearly, our method establishes that every problem instance can be colored with $\lceil \frac{n}{g-1} \rceil\ \lambda/f$. It is simple to show that this coloring can be constructed in linear time with respect to the input length. □

5. Lower Bound for (2 | 2 | 2)

For this problem the lower bound is equal to 3 λ/f. To establish our lower bound, we give a problem instance with six nodes that requires at least 3 λ/f (see Figure 10). We now show that it is not possible to color the problem instance in Figure 10 with 2 λ/f. Figure 10 consists of all unicasts except for edges C and D, which are part of the same multicast. Note that edges C and D must be sent on the same fiber-wavelength pair in order for node $x's$ edges to be colorable in 2 λ/f since there are two fibers. Edges can be assigned to f_1w_1, f_2w_1, f_1w_2, or f_2w_2. Since we can switch messages on the same wavelength from one fiber to another, for this example we will ignore the fiber number and concentrate on the fact that every node can have at most two edges on w_1 and at most two edges on w_2. Since there are two fibers and two wavelengths, it has to be that the messages C and D are sent on the same wavelength in a feasible solution. To see the conflicts that arise, let us focus our attention on edges A and C.

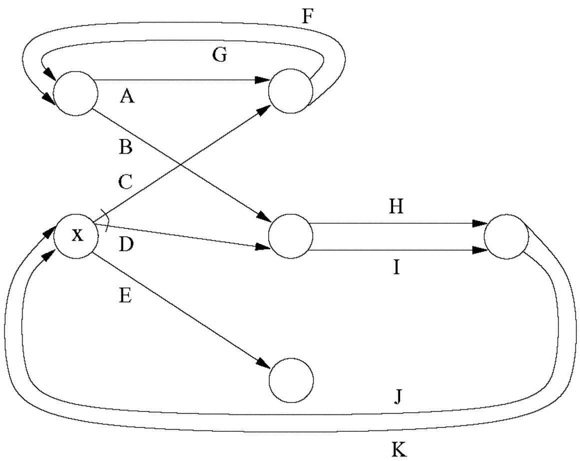

Fig. 10. Requires 3 λ/f.

Edges A and C must either be assigned to the same wavelength or to different ones. First, we consider assigning edges A and C the same wavelength (without loss of generality we can choose w_1). This forces edges F and G to be assigned to w_2, which in turn forces edge B to be assigned to w_1. Furthermore, since edge D is on w_1 (because edge C and D must be assigned to the same wavelength, see comment above), edges H and I must be assigned to w_2. This forces edges J and K to be on w_1 which creates a conflict at node x. Therefore, assigning edges A and C the same wavelength does not allow for a feasible solution.

Next, we consider assigning edges A and C to different wavelengths (without loss of generality we can let edge A be w_1 and edge C be w_2). This forces edges F and G to be on different wavelengths, which in turn forces edge B to be assigned to w_2. Since edge D is on w_2 (because there are two wavelengths and two fibers only), edges

H and I must be assigned to w_1. This forces edges J and K to be assigned to w_2 which creates a conflict at node x. Therefore, assigning edges A and C to different wavelengths does not allow for a feasible solution.

Theorem 5.1. The lower bound for (2 | 2 | 2) is to 3 λ/f.

Proof. By the above discussion. □

We should note that edge E, although not yet mentioned directly, is necessary. Without edge E, the edges leaving node x could be viewed as unicasts, meaning edges C and D could be sent on different fiber-wavelength pairs. We should point out that, a multigraph that consists of only unicasts can always be colored with 2 λ/f. One such coloring can be generated by an elegant linear time algorithm. For completeness, the following theorem establishes this result.

Theorem 5.2. The upper bound for the (2 | 2 | 2) problem when all the multicasts are unicasts is 2 λ/f. Furthermore, a color assignment can be constructed in linear time with respect to the input length.

Proof. Simply add edges to the system so that every node has 2 incoming and 2 outgoing edges. Then, we may view the edges as a set of node disjoint Euler circuits. Note that a circuit that goes through n nodes in the multigraph has $2n$ edges, so in every Euler circuit every node is visited twice. We consider each Euler circuit one at a time and traverse it. We assign wavelength w_1 to the first edge and w_2 to the second one and continue alternating until we have assigned wavelengths to all the edges. This will form a valid color assignment for the nodes in the circuit. Repeating this operation on all the circuits generates a coloring for the message multigraph. □

6. Upper Bound for the (n | n | g) Problem

In this section, we establish an upper bound of $m = \frac{n^2}{i(g-1)}$ λ/f, where i is any positive integer such that $1 \leq i \leq g-1$, for the (n | n | g) problem. The value for m is minimum when $i = \frac{g}{2}$ in which case $m = \frac{2n^2}{g^2-g}$ which is less than two times the lower bound given in Table 5.

The following descriptions assume that the values of i, m, and n are such that when used to assign wavelengths always result in integer values for the expressions below. When this is not the case similar results have been obtained, so we refer the reader to Ref. 10. The messages to be *received* by all nodes are assigned as follows.

The idea is to use i fibers (f_1, f_2, ..., f_i) to receive the messages at every node. We will use each fiber for $\frac{n}{i}$ of the messages. Since each fiber has m wavelengths, we will allow the use of $\frac{m}{\frac{n}{i}}$ wavelengths for each message. So the first incoming message at every node can be assigned to fiber one (f_1) and any of the wavelengths from set

S_1 (see Figure 11). The second incoming message at every node can be assigned to f_1 and any of the wavelengths from set S_2; and so on up to the $(\frac{n}{i})^{th}$ message which can be assigned to f_1 and any of the wavelengths from set $S_{\frac{n}{i}}$. The next $\frac{n}{i}$ messages are assigned similarly, but using f_2. This coloring process continues until the last set of $\frac{n}{i}$ messages which use f_i and the same sets of wavelengths. At this point, using this technique, all of the incoming messages at every node can be assigned a unique fiber-wavelength pair. The appropriate wavelength to use for every incoming

1	2	$\cdots$	$\frac{im}{n}$	S_1
$\frac{im}{n}+1$	$\frac{im}{n}+2$	$\cdots$	$\frac{2im}{n}$	S_2
$\vdots$	$\vdots$		$\vdots$	$\vdots$
$(\frac{n}{i}-1)\frac{im}{n}+1$	$(\frac{n}{i}-1)\frac{im}{n}+2$	$\cdots$	m	$S_{\frac{n}{i}}$
T_1	T_2	$\cdots$	$T_{\frac{im}{n}}$	

Fig. 11. Sets of wavelengths S_i and T_i.

message will be determined based on the multicast's wavelength assignments, which are described below. Note that we have used a total of i fibers and m λ/f to assign the messages *received* at every node.

Next we discuss the multicasts sent from every node. In order to avoid conflicts every message in every multicast must be sent using a wavelength in each of the sets $S_1, S_2, ..., S_{\frac{n}{i}}$ on some fiber; and no other multicast emanating from this node can use these fiber-wavelengths pairs. To accomplish this we define a set T_j as the set of all of the j^{th} elements in each of the sets $S_1, S_2, ..., S_{\frac{n}{i}}$ (when viewing these sets as order sets). Figure 11 gives a possible definition of the sets S and T. Clearly, there are $\frac{im}{n}$ different T sets. So, each fiber can be used for $\frac{im}{n}$ different multicasts. Therefore, the total number of fibers needed to send the n multicasts at every node is $\frac{n^2}{im}$ (fibers f_{i+1}, f_{i+2}, ..., $f_{i+\frac{n^2}{im}}$). Therefore $g = i + \frac{n^2}{im}$ which implies $m = \frac{n^2}{i(g-1)}$.

Theorem 6.1. An upper bound for the ($n \mid n \mid g$) problem is $m = \frac{n^2}{i(g-1)}$ λ/f, for any positive integer i such that $1 \leq i \leq g-1$. The number of wavelengths per fiber is minimum and equal to $\frac{2n^2}{g^2-g}$ λ/f when $i = \frac{g}{2}$. Furthermore, a color assignment can be constructed in linear time with respect to the input length.

Proof. By the above discussion. □

7. Conclusion

We determined lower and upper bounds on the number of wavelengths required per fiber for star networks where the messages are routed using Wavelength Division Multiplexing. Single, dual, and multimessage multicasting were considered along with single, dual, and multifiber optical networks. If, as networks develop, the available fibers increases beyond the amount of traffic in the network, (*i.e.* $g \gg n$), many of our results reduce to 1 λ/f; however, this seems unlikely to happen given current trends. Future work could include continued efforts to obtain tight bounds for all of the remaining systems within star networks along with finding bounds for ring networks and more general network topologies.

The most interesting open problem is to tighten the lower and upper bound for the ($2 \mid 1 \mid 1$) problem. We generated a huge number of problem instances all of which could be colored with four colors. However, it does not seem possible to make our algorithm that use five colors to only use four colors, even when every multicast has two destinations. Another interesting problem is to narrow the gap between the lower and upper bound for the ($n \mid n \mid g$) problem. It is not known whether our problems are NP-hard when there are multicasts, we conjecture they are NP-hard.

Acknowledgments

The authors wish to thank the referees for the comments and suggestions to improve our presentation.

References

1. G. Rouskas and M. Ammar, "Multi-Destination Communication Over Tunable-Receiver Single-Hop WDM Networks," *IEEE Journal on Selected Areas in Communications*, 15(3), April, 1997, pp. 501–511.
2. D. Thaker and G. Rouskas, "Multi-Destination Communication in Broadcast WDM Networks: A Survey," *Optical Networks*, 3(1), Jan./Feb. 2002, pp. 34–44.
3. G. Li and R. Simha, "On the Wavelength Assignment Problem in Multifiber WDM Star and Ring Networks," *IEEE/ACM Transactions On Networking*, 9(1), February 2001, pp. 60–68.
4. G. Rouskas and M. Ammar, "Multi-Destination Communication Over Single-Hop Lightwave WDM Networks," *Proceedings of INFOCOM '94*, June 1994, pp. 1520–1527.
5. A. Ferreira, S. Perennes, A. Richa, H. Rivano and N. Stier, "On the Design of Multifiber WDM Networks," *AlgoTel '02*, May 2002, pp. 25–32.
6. A. Chien, "OptIPuter Software: System Software for High Bandwidth Network Applications," *http://www-csag.ucsd.edu/projects/Optiputer.html.* September, 2002.
7. E. Bampis and G. Rouskas, "On Scheduling Problems with Applications to Packet-Switched Optical WDM Networks," *Proceedings of INFOCOM '01*, August 19-24, 2001, pp. 163–172.
8. I. Holyer, "The NP-Completeness of Edge-Coloring," *SIAM J. Comput.*, 10(4), November, 1981, pp. 718–720.
9. C. Berge, *Graphs and Hypergraphs*, Chapter 12, 1976, (translated by E. Minieka).
10. R. Brandt, "Multicasting using WDM in Multifiber Optical Star Networks," *Thesis*, UCSB, September 2003.