Optimal Wavelength Assignment for Uniform All-to-All Traffic in WDM Tree Networks

Poompat Saengudomlert, Eytan H. Modiano, and Robert G. Gallager
(*tengo@mit.edu, modiano@mit.edu, gallager@mit.edu*)
Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
77 Massachusetts Ave, Cambridge, MA 02139

Abstract

We consider the problem of wavelength assignment (WA) for l-uniform traffic in an arbitrary tree network, where each end node transmits l wavelengths to every other end node. We provide a lower bound on the number of wavelengths needed to support l-uniform traffic and a WA algorithm that meets this bound. We also show that in a tree network wavelength converters cannot be used to reduce the required number of wavelengths.

1 Introduction

In a wavelength division multiplexed (WDM) network, the fiber bandwidth is divided into multiple frequency bands often called wavelengths. Using optical switches at the network nodes, some wavelengths can be selected at each node for termination and electronic processing, and others selected for optical bypassing. In an *all-optical network* architecture, each traffic session optically bypasses electronic processing at each node on its path other than the source node and the destination node. One important benefit of this architecture is a significant cost saving from the use of fewer and/or smaller electronic switches in the network. We consider all-optical networks in this paper.

Without optical wavelength conversion, routing of traffic sessions is subjected to the *wavelength continuity constraint*, which dictates that the lightpath corresponding to a given session must travel on the same wavelength on all links from the source node to the destination node. Using wavelength converters potentially allows the network to support a larger set of traffic. However, such converters are likely to be expensive. Hence, several researchers have focused on the problem of routing and wavelength assignment (RWA) assuming no wavelength conversion. We focus on this same problem.

A large body of literature investigates the RWA problem under the wavelength continuity constraint. We can categorize existing results into two groups based on whether static or dynamic provisioning of routes and wavelengths is performed. We consider static provisioning in this paper. For static provisioning, the traffic to be supported is assumed known and fixed over time. The goal is often to minimize the number of wavelengths used in the network [1, 2], or to maximize the number of supported traffic sessions for a fixed number of wavelengths [3, 4, 5]. For general traffic in arbitrary network topologies, these static RWA problems are known to be NP-complete [3]. Consequently, bounds on the optimal costs have been derived [4, 6], and several RWA heuristics have been developed [1, 4, 5, 7, 8].

For some specific traffic and network topologies, analytical solutions to static RWA problems can be obtained. In [9], an RWA algorithm using the minimum number of wavelengths to support uniform all-to-all traffic in a bidirectional ring was developed. In this paper, we solve the static RWA problem for uniform all-to-all traffic in an arbitrary tree topology. Since there is no routing problem in a tree topology, our RWA algorithm only has to perform wavelength assignment (WA) and will be referred to as a WA algorithm.

In section 2, we define l-uniform traffic and formulate the WA problem for l-uniform traffic in an arbitrary tree topology. In section 3, we present our WA algorithm which uses the minimum number of wavelengths. Finally, we summarize the results in section 4.

2 Problem Formulation

Consider an all-optical WDM network with no wavelength conversion. Adjacent nodes are connected by two unidirectional fibers, one in each direction. In addition, all fibers contain the same number of wavelengths. In an arbitrary tree topology, we assume there are $N > 2$ end nodes,[1] which are the leaf nodes of the tree.[2] For the purpose of WA, we can assume that each non-leaf node has degree at least 3.[3] If a non-leaf node has degree less than 3, it can be removed from the tree without changing the WA problem. Assume that each traffic session has a rate of one wavelength. At a given time, only one session can use a specific wavelength in a fiber, but multiple sessions on a specific wavelength can use the same node.

Define *l-uniform traffic* to be static traffic in which each end node transmits l wavelengths (i.e. sessions) to, and receives l wavelengths from, each of the other end nodes.[4] Note that l-uniform traffic requires $l(N-1)$ transmitters and $l(N-1)$ receivers at each end node. Since the traffic is static, these transmitters and receivers need not be tunable. Moreover, we can use fixed optical switches at each switching node.

For an arbitrary tree topology with N end nodes, let W_l denote the minimum number of wavelengths which, if provided in each fiber, can support l-uniform traffic with no wavelength conversion. Our goal is to find the value of W_l and design a WA algorithm which uses W_l wavelengths to support l-uniform traffic. In addition, let L_l denote the minimum number of wavelengths which, if provided in each fiber, can support l-uniform traffic given full wavelength conversion at all nodes. It is clear that $L_l \leq W_l$.

3 Optimal WA Algorithm

Consider the WA problem for 1-uniform traffic. The results are later extended, in a straightforward manner, to l-uniform traffic. For a given tree, let $\mathcal{T}$ denote a set of bidirectional links. Each link e in $\mathcal{T}$ corresponds to a cut which separates the N end nodes into two sets, denoted by $\mathcal{N}_{e,1}$ and $\mathcal{N}_{e,2}$. The amount of traffic (in wavelengths) on a fiber across link e is equal to $|\mathcal{N}_{e,1}||\mathcal{N}_{e,2}|$. Let w^* denote the maximum traffic over all the fibers. Clearly, L_1 is equal to w^*, as given below.

$$L_1 \quad = \quad w^* \quad = \quad \max_{e \in \mathcal{T}} |\mathcal{N}_{e,1}||\mathcal{N}_{e,2}| \tag{1}$$

[1] The WA problem for a tree with two end nodes is trivial.

[2] With slight modification, our WA algorithm can also be applied to a tree with non-leaf end nodes [10].

[3] Since we assume that each link consists of two fibers, one in each direction, the indegree and the outdegree of any given network node are the same. We simply refer to their value as the node degree.

[4] We reserve the terms *transmit* and *receive* for the end nodes which source and sink traffic sessions. Intermediate nodes which only switch traffic but neither source nor sink traffic are not considered transmitting or receiving traffic.

In this paper, we show that $W_1 \le w^*$, which implies $W_1 = L_1 = w^*$. We do so by constructing a WA algorithm. Figure 1 illustrates an example in which a greedy WA algorithm fails to support 1-uniform traffic using w^* wavelengths. In this example, inspection shows that $w^* = 2$. Note that the same wavelength is assigned to the oppositely directed sessions between the same pair of nodes, e.g. sessions (1,2) and (2,1) on wavelength λ_1. After assigning wavelength λ_1 to sessions (1,2) and (2,1) and wavelength λ_2 to sessions (1,3) and (3,1), neither λ_1 nor λ_2 can be assigned to support session (2,3). It follows that more than $w^* = 2$ wavelengths are required. Therefore, this example tells us that the design of a WA algorithm using w^* wavelengths is not trivial. Figure 1 also demonstrates that, in order to use w^* wavelengths, we may need to support the oppositely directed sessions between the same pair of nodes on different wavelengths.

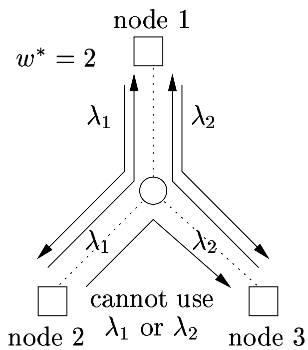

sequence of sessions for WA	corresponding sequence of WA steps
(1,2)	(1,2) on λ_1
(2,1)	(2,1) on λ_1
(1,3)	(1,3) on λ_2
(3,1)	(3,1) on λ_2
(2,3)	(2,3) not on λ_1 or λ_2

(i, j) denotes a session from node i to node j.

Figure 1: An example in which a greedy approach requires more than w^* wavelengths.

We now derive a few useful properties related to w^*. Let e^* denote the link associated with w^*. There may be multiple choices for e^*. The exact choice does not matter in the following discussion. We refer to e^* as the *bottleneck link* since it is the link with the maximum traffic on a fiber. Link e^* separates the leaf nodes into two sets $\mathcal{N}_{e^*,1}$ and $\mathcal{N}_{e^*,2}$. Without loss of generality, choose $\mathcal{N}_{e^*,1}$ such that $|\mathcal{N}_{e^*,1}| \le |\mathcal{N}_{e^*,2}|$. Since we assume there are more than two leaf nodes, $\mathcal{N}_{e^*,2}$ must contain multiple leaf nodes. Define the *bottleneck node* v^* to be the end point of e^* opposite to $\mathcal{N}_{e^*,1}$, i.e. the subtree connected to v^* by e^* contains all the leaf nodes in $\mathcal{N}_{e^*,1}$, as illustrated in figure 2a.

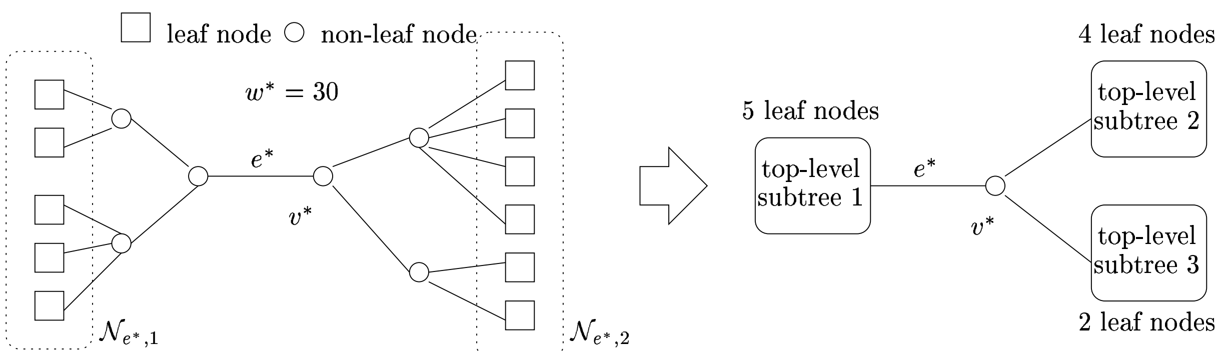

(a) (b)

Figure 2: The bottleneck link e^* and the bottleneck node v^*.

We refer to each subtree connected to v^* as a *top-level subtree*. Note that a top-level subtree can be a single node. Figure 2b shows the top-level subtrees associated with the tree in figure 2a. Let d^* be the degree of v^*. Since v^* is a non-leaf node, $d^* \ge 3$. It follows that there are $d^* \geq 3$ top-level subtrees. Let $\mathcal{S}_i$, $1 \leq i \leq d^*$, denote the set of all the leaf nodes in top-level subtree i, and $x_i = |\mathcal{S}_i|$. The following lemma provides useful properties of the top-level subtrees connected to v^* as well as bounds on w^*. Throughout the paper, proofs of the lemmas and the theorems are omitted. They are available in [10].

Lemma 1 *Number the top-level subtree connected to the bottleneck node v^* by the bottleneck link e^* as top-level subtree 1, and the rest of the top-level subtrees from 2 to d^*, where d^* is the degree of v^*. Then, (1) $x_i \leq x_1 \leq N/2$ for all $1 \leq i \leq d^*$, and (2) $\frac{1}{d^*}\left(1 - \frac{1}{d^*}\right) N^2 \leq w^* \leq \frac{N^2}{4}$.*

Before describing our WA algorithm, we describe some of the ideas behind it. Define a *local session* to be a traffic session whose source and destination are in the same top-level subtree. Accordingly, a *non-local session* has its source and its destination in different top-level subtrees. Note that a non-local session has to travel through the bottleneck node v^*, whereas a local session does not have to travel all the way to v^* and back to its destination, i.e. each session never uses the same link twice in opposite directions.

From lemma 1, we can always number the d^* top-level subtrees such that $x_1 \geq x_2 \geq ... \geq x_{d^*}$. It follows that $w^* = x_1(N - x_1)$. Our WA algorithm first assigns wavelengths to all of the non-local sessions. It then assigns wavelengths to all the local sessions in each top-level subtree. Consider top-level subtree 1. Since there are in total $x_1(N-1)$ local and non-local sessions transmitted from nodes in this subtree while there are only $x_1(N - x_1)$ wavelengths available, it is clear that we need to reuse some wavelengths previously assigned to non-local sessions to support local sessions. Such wavelength reuse is the cause of the main complexity in the design of an efficient WA algorithm.

Let $n_{i,j}$ denote leaf node j in top-level subtree i, where $1 \leq i \leq d^*$ and $1 \leq j \leq x_i$. With respect to $n_{i,j}$, define a *reusable wavelength* to be a wavelength used by $n_{i,j}$ to receive a non-local session (from a node in a different top-level subtree), but not used by $n_{i,j}$ to transmit a non-local session (to a node in a different top-level subtree). Figure 3 shows two examples in which λ_1 is a reusable wavelength with respect to $n_{i,j}$.

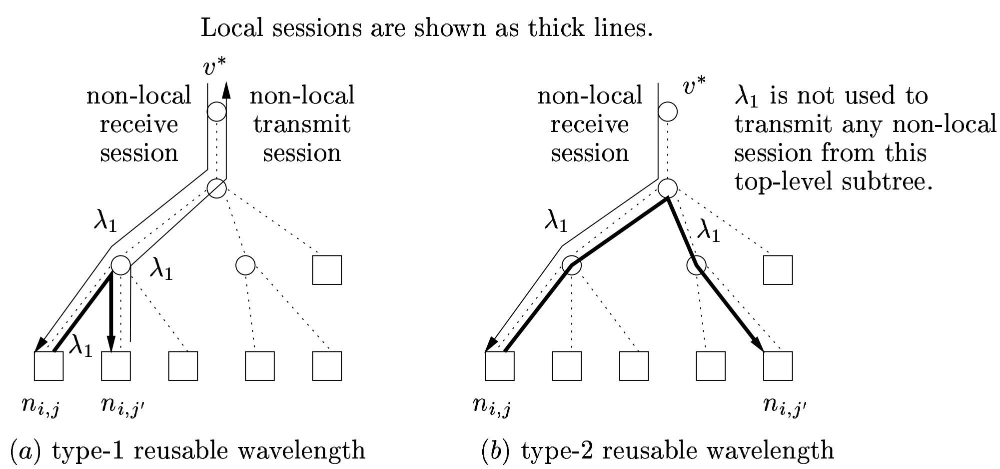

Figure 3: Reusable wavelength λ_1 with respect to node $n_{i,j}$.

With respect to node $n_{i,j}$, define a *type-1* reusable wavelength to be a reusable wavelength which is also used by a different node in the same top-level subtree (i.e. top-level subtree i) to transmit a non-local session. For example, in figure 3a, with respect to $n_{i,j}$, λ_1 is a type-1 reusable wavelength. In addition, with respect to $n_{i,j}$, define a *type-1 local node* to be a different node in the same top-level subtree which transmits a non-local session on a reusable wavelength (with respect to $n_{i,j}$). For example, in figure 3a, with respect to $n_{i,j}$, $n_{i,j'}$ is a type-1 local node.

With respect to $n_{i,j}$, define a *type-2* reusable wavelength to be a reusable wavelength which is not type-1, i.e. it is not used by any other node in the same top-level subtree to transmit a non-local session. For example, in figure 3b, with respect to $n_{i,j}$, λ_1 is a type-2 reusable wavelength. In addition, with respect to $n_{i,j}$, define a *type-2 local node* to be a different node in the same top-level subtree which is not type-1, i.e. it does not transmit a non-local session on any reusable wavelength (with respect to $n_{i,j}$). For example, in figure 3b, with respect to $n_{i,j}$, if $n_{i,j'}$ does not use any reusable wavelength (with respect to $n_{i,j}$) to transmit a non-local session, then $n_{i,j'}$ is a type-2 local node.

Notice that, by the above definitions, with respect to any given node $n_{i,j}$, each node $n_{i,j'}$, $j' \neq j$, is either a type-1 or type-2 local node. The following lemma indicates our strategy of assigning wavelengths to the local sessions transmitted from $n_{i,j}$ using reusable wavelengths with respect to $n_{i,j}$

Lemma 2 *With respect to node $n_{i,j}$, we have the following properties.*

1. *Node $n_{i,j}$ can transmit a local session to type-1 local node $n_{i,j'}$ on a type-1 reusable wavelength (with respect to $n_{i,j}$) which is used by $n_{i,j'}$ to transmit a non-local session.*
2. *Node $n_{i,j}$ can transmit a local session to type-2 local node $n_{i,j'}$ on any type-2 reusable wavelength (with respect to $n_{i,j}$).*

Lemma 2 suggests the following method of assigning wavelengths to the local sessions. Consider the local sessions transmitted from node $n_{i,j}$ in top-level subtree $\mathcal{S}_i$. There are $x_i - 1$ such sessions. Let $\mathcal{P}_{i,j}^{(1)}$ and $\mathcal{P}_{i,j}^{(2)}$ be the sets of type-1 and type-2 local nodes with respect to $n_{i,j}$ respectively. Notice that type-1 local nodes (with respect to $n_{i,j}$) have associated with them distinct reusable wavelengths (with respect to $n_{i,j}$). From statement 1 of lemma 2, $n_{i,j}$ can use a distinct type-1 reusable wavelength (with respect to $n_{i,j}$) to transmit a local session to each type-1 local node in $\mathcal{P}_{i,j}^{(1)}$. It remains to provide wavelengths for the local sessions to type-2 local nodes (with respect to $n_{i,j}$).

We later state in our WA algorithm that it is always possible to assign wavelengths to the non-local sessions so that there are at least $|\mathcal{P}_{i,j}^{(2)}|$ type-2 reusable wavelengths with respect to each node $n_{i,j}$ in $\mathcal{S}_i$. Given $|\mathcal{P}_{i,j}^{(2)}|$ type-2 reusable wavelengths with respect to $n_{i,j}$, statement 2 of lemma 2 implies that $n_{i,j}$ can use a distinct type-2 reusable wavelength (with respect to $n_{i,j}$) to transmit a local session to each type-2 local node in $\mathcal{P}_{i,j}^{(2)}$.

We repeat the same process for all the leaf nodes. Since all the reusable wavelengths (with respect to the same node or with respect to different nodes) in each top-level subtree are distinct,[5] different local sessions (transmitted from the same node or from different nodes) never use the same wavelength.

In conclusion, the condition that there are at least $|\mathcal{P}_{i,j}^{(2)}|$ type-2 reusable wavelengths with respect to each node $n_{i,j}$ is a sufficient condition for the WA of all the local sessions to exist. We state this conclusion formally in the following lemma, which is later used to develop our WA algorithm.

Lemma 3 *If there are at least $|\mathcal{P}_{i,j}^{(2)}|$ type-2 reusable wavelengths with respect to node $n_{i,j}$ for all $1 \leq i \leq d^*$ and $1 \leq j \leq x_i$, then we can assign wavelengths to all the local sessions as follows. Consider the local sessions transmitted from $n_{i,j}$ in $\mathcal{S}_i$.*

[5] Otherwise, there would be multiple sessions using the same wavelength on the fiber from v^* to some top-level subtree.

1. *To transmit a local session to a node in* $\mathcal{P}_{i,j}^{(1)}$, $n_{i,j}$ *uses a type-1 reusable wavelength (with respect to* $n_{i,j}$*) which is used by that node to transmit a non-local session.*
2. *To transmit a local session to a node in* $\mathcal{P}_{i,j}^{(2)}$, $n_{i,j}$ *uses a distinct type-2 reusable wavelength (with respect to* $n_{i,j}$*).*

Our WA algorithm operates in three phases. In phase 1, we assign wavelength bands each of which is used by the non-local sessions from one top-level subtree to another. In phase 2, we perform WA for individual non-local sessions based on the wavelength bands obtained from phase 1. The goal of phase 2 is to assign wavelengths in such a way that enough type-1 and type-2 reusable wavelengths exist to support all local traffic. Finally, in phase 3, we perform WA for local sessions independently in each top-level subtree. The following is our WA algorithm which uses w^* wavelengths to support 1-uniform traffic in an arbitrary tree topology. We refer to this algorithm as the *tree WA algorithm.*

Algorithm 1 (Tree WA Algorithm) (Use w^* wavelengths.) Number the top-level subtrees so that the numbers of leaf nodes, denoted by $x_1, ..., x_{d^*}$, satisfy $x_1 \geq x_2 \geq ... \geq x_{d^*}$. Note that $w^* = x_1(N - x_1)$.

Phase 1: Assign the wavelength band for the non-local sessions from one top-level subtree to another as follows. For convenience, let $\Lambda_{(i,i')}$ denote the wavelength band for the non-local sessions from $\mathcal{S}_i$ to $\mathcal{S}_{i'}$. Note that $\Lambda_{(i,i')}$ contains $x_i x_{i'}$ wavelengths.

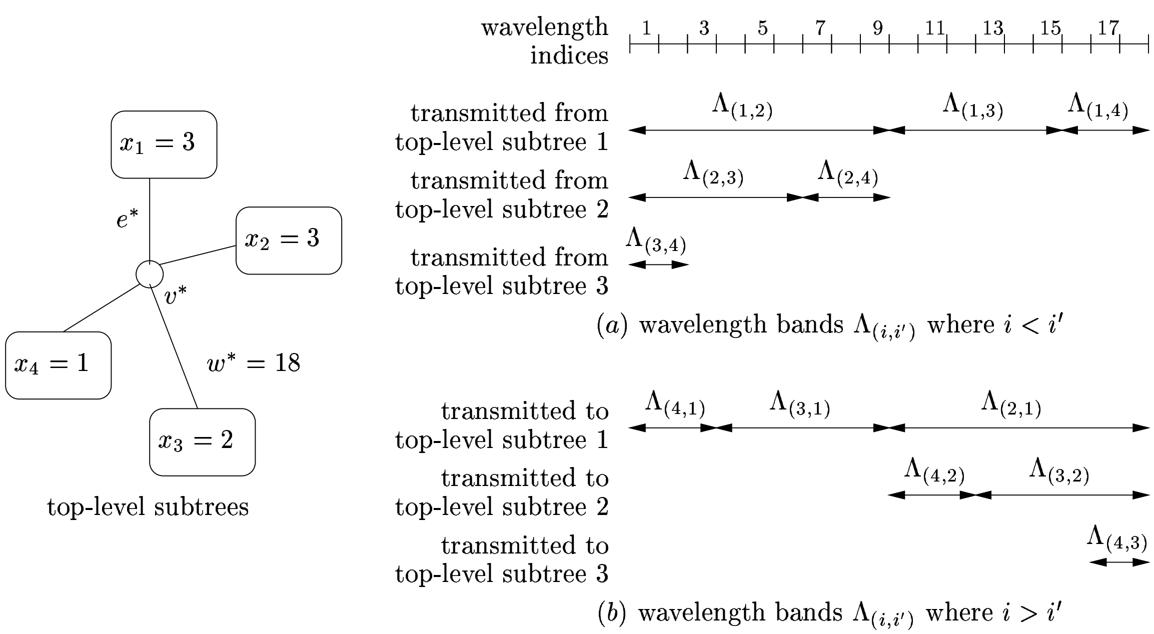

Figure 4: Phase 1 of the tree WA algorithm.

Figure 4 specifies the wavelength bands between all pairs of top-level subtrees. To obtain wavelength band $\Lambda_{(i,i')}$, where $i < i'$, follow the diagram in figure 4a. There are $d^* - 1$ rows of wavelength bands. In row i, $1 \leq i \leq d^* - 1$, we assign consecutive wavelengths starting from wavelength 1 (from left to right) to wavelength bands $\Lambda_{(i,i+1)}$, ..., $\Lambda_{(i,d*)}$. For example, the wavelength band $\Lambda_{(1,3)}$ contains 6 wavelengths with indices 10 to 15. On the other hand, to obtain wavelength band $\Lambda_{(i,i')}$, where $i' < i$, follow the diagram in figure 4b. There are $d^* - 1$ rows of wavelength bands. In row i', $1 \leq i' \leq d^* - 1$, we assign consecutive wavelengths starting from wavelength w^* (from right to left) to wavelength bands $\Lambda_{(i'+1,i')}$, ..., $\Lambda_{(d^*,i')}$. For example, the wavelength band $\Lambda_{(4,2)}$ contains

3 wavelengths with indices 10 to 12. Although a specific example is illustrated, the general scheme should be clear.

It can be shown that, in each top-level subtree, the assigned receive wavelength bands do not overlap [10], i.e. there is no wavelength collision between two non-local receive sessions in two different bands. In addition, the assigned transmit wavelength bands do not overlap. As a result, there is no wavelength collision among the non-local transmit sessions and among the non-local receive sessions in each top-level subtree.

As an example to show how the scheme works, consider two wavelength bands $\Lambda_{(1,4)}$ and $\Lambda_{(2,4)}$ for non-local receive sessions in top-level subtree 4. The highest wavelength index in $\Lambda_{(2,4)}$, denoted by $\lambda^+_{(2,4)}$, is $x_2x_3 + x_2x_4$. The lowest wavelength index in $\Lambda_{(1,4)}$, denoted by $\lambda^-_{(1,4)}$, is $x_1x_2 + x_1x_3 + 1$. Since $x_1 \geq x_2 \geq ... \geq x_d$, it follows that $x_1x_2 \geq x_2x_3$ and $x_1x_3 \geq x_2x_4$. Thus,

$$\lambda^-_{(1,4)} \quad = \quad x_1x_2 + x_1x_3 + 1 \quad > \quad x_2x_3 + x_2x_4 \quad = \quad \lambda^+_{(2,4)}.$$

It follows that a non-local session in wavelength band $\Lambda_{(1,4)}$ and a non-local session in wavelength band $\Lambda_{(2,4)}$ never share the same wavelength and therefore do not collide.

Phase 2: In this phase, we assign wavelengths to individual non-local sessions based on the wavelength bands obtained from phase 1. Our goal is to assign wavelengths so that there are at least $|\mathcal{P}^{(2)}_{i,j}|$ type-2 reusable wavelengths with respect to node $n_{i,j}$ for all $1 \leq i \leq d^*$ and $1 \leq j \leq x_i$, as suggested by lemma 3.

We first perform partial WA as follows. For each wavelength band $\Lambda_{(i,j)}$ containing x_ix_j wavelengths (used for the non-local sessions from $\mathcal{S}_i$ to $\mathcal{S}_j$), we break the band up into x_j subbands of x_i contiguous wavelengths. The first subband is assigned to be receive wavelengths for node $n_{j,1}$. The second subband is assigned to be receive wavelengths for $n_{j,2}$, and so on. For example, based on the example in figure 4, in top-level subtree 1, node $n_{1,1}$ receives three non-local sessions from top-level subtree 2 on the subband of $\Lambda_{(2,1)}$ containing wavelengths 10, 11, and 12. Notice that we have not specified which node in top-level subtree 2 uses a specific wavelength (10, 11, or 12) to transmit to $n_{1,1}$.

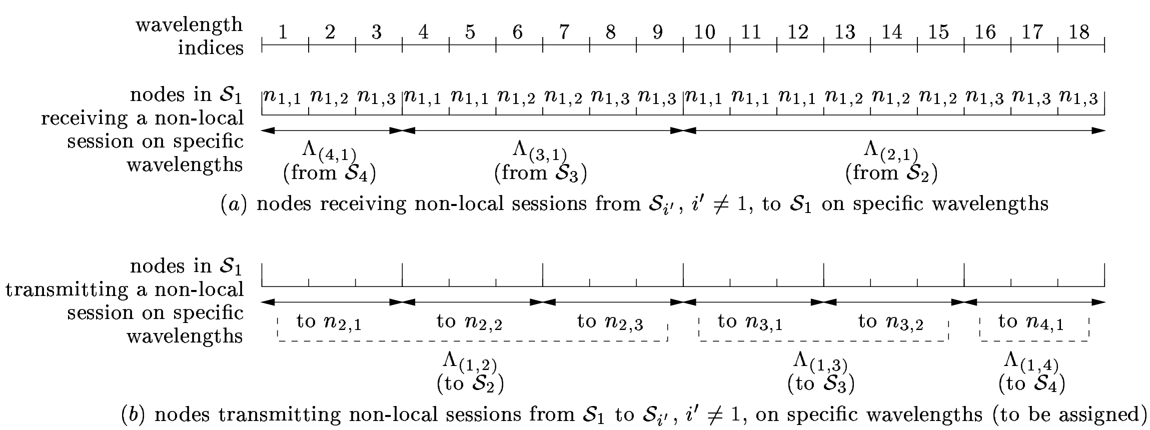

Figure 5: The result of the partial WA in phase 2 of the tree WA algorithm for top-level subtree 1 in figure 4.

Figure 5 illustrates the result of the partial WA in top-level subtree 1. Note that the partial WA also specifies the subbands used by the nodes in $\mathcal{S}_1$ to transmit to each node in $\mathcal{S}_{i'}$, $i' \neq 1$, as shown in figure 5b. For example, in $\Lambda_{(1,2)}$, wavelengths 1, 2, and 3 are used for the non-local sessions from $\mathcal{S}_1$ to $n_{2,1}$.

It remains to specify the source nodes for specific wavelengths in each subband, i.e.

filling the empty slots in each subband in figure 5b with $n_{1,1}$, $n_{1,2}$, and $n_{1,3}$. Such specifications in top-level subtree 1 can be done independently from the similar specifications in all the other top-level subtrees since the lightpaths corresponding to each subband traverse the same set of fibers outside top-level subtree 1. In other words, the WA outside top-level subtree 1 looks the same regardless of how we fill the empty slots in figure 5b. Furthermore, such specifications yield, for each node, the corresponding type-1 and type-2 reusable wavelengths together with type-1 and type-2 local nodes.

Assume for now that the set of wavelengths used to receive and to transmit non-local sessions are the same in a given top-level subtree i. (This is the case for top-level subtree 1 and any other subtree i with $x_i = x_1$. However, the assumption does not always hold, e.g. top-level subtree 3 in figure 4.) We now describe how to assign the source nodes in each wavelength subband so that $|P_{i,j}^{(2)}| = 0$, i.e. no type-2 local node with respect to $n_{i,j}$, for each node $n_{i,j}$ in $\mathcal{S}_i$. Note that the condition $|P_{i,j}^{(2)}| = 0$ yields the sufficient condition in lemma 3 for the WA of all the local sessions in top-level subtree i to exist, i.e. there are at least $|P_{i,j}^{(2)}|$ type-2 reusable wavelengths with respect to each node $n_{i,j}$ in $\mathcal{S}_i$. The goal $|P_{i,j}^{(2)}| = 0$ is equivalent to $|P_{i,j}^{(1)}| = x_i - 1$. That is, we must ensure that, each node $n_{i,j'}$, $j' \neq j$, in $\mathcal{S}_i$ transmits at least one non-local session on one of the wavelengths used by $n_{i,j}$ to receive non-local sessions.

We can visualize the problem of assigning the source nodes in each subband using a bipartite graph. We consider each top-level subtree separately. For top-level subtree i, construct a *partial WA bipartite graph* denoted by $(\mathcal{V}_1, \mathcal{V}_2, \mathcal{E})$ as follows. The set $\mathcal{V}_1$ contains the $N - x_i$ leaf nodes outside top-level subtree i, i.e. $\{n_{i',j'} : i' \neq i, 1 \leq j' \leq x_{i'}\}$. The set $\mathcal{V}_2$ is equal to $\mathcal{S}_i$, i.e. $\{n_{i,j} : 1 \leq j \leq x_i\}$. In the set of edges $\mathcal{E}$, an edge joins $n_{i',j'}$ in $\mathcal{V}_1$ and $n_{i,j}$ in $\mathcal{V}_2$ for each wavelength that is used to receive a non-local session from a node in $\mathcal{S}_i$ by $n_{i',j'}$, and is used to receive a non-local session by $n_{i,j}$. There may be multiple edges between the same pair of nodes.

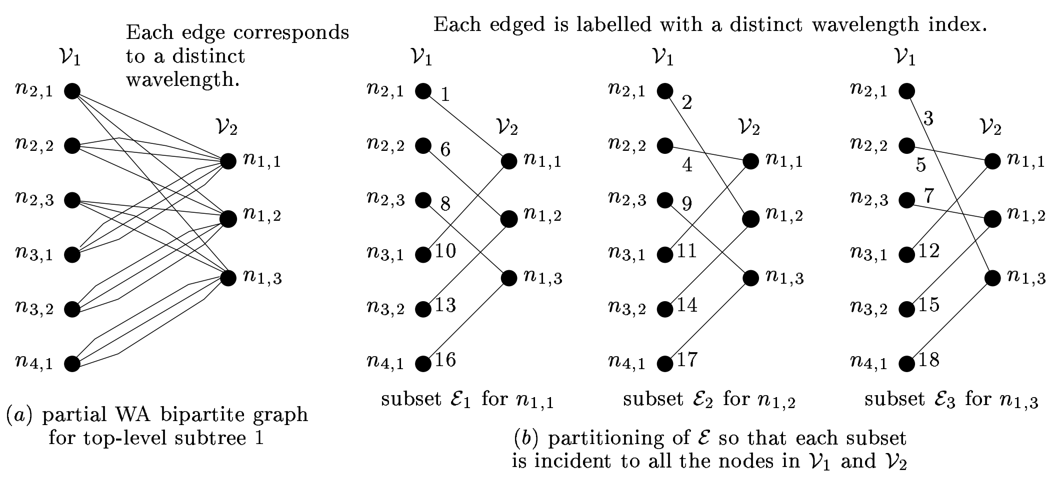

Figure 6: Partial WA bipartite graph for top-level subtree 1 in figure 4.

For example, figure 6a shows the partial WA bipartite graph specified by the partial WA in top-level subtree 1 in figure 5. In particular, the edge between $n_{2,1}$ in $\mathcal{V}_1$ and $n_{1,1}$ in $\mathcal{V}_2$ corresponds to wavelength 1 which is used both to receive a non-local session from $\mathcal{S}_1$ by $n_{2,1}$ and to receive a non-local session by $n_{1,1}$. Two edges between $n_{2,3}$ in $\mathcal{V}_1$ and $n_{1,3}$ in $\mathcal{V}_2$ correspond to wavelengths 8 and 9 which are used both to receive a non-local session from $\mathcal{S}_1$ by $n_{2,3}$ and to receive a non-local session by $n_{1,3}$.

From the assumption that, in top-level subtree i, the set of non-local transmit wavelengths is equal to the set of non-local receive wavelengths, it follows that every non-local transmit wavelength corresponds one-to-one to an edge in the partial WA bipartite graph. Since each node $n_{i',j'}$ in $\mathcal{V}_1$ receives x_i non-local sessions from $\mathcal{S}_i$, each node $n_{i',j'}$ has degree x_i. Since each node $n_{i,j}$ in $\mathcal{V}_2$ receives $N - x_i$ non-local sessions, each node $n_{i,j}$ in $\mathcal{V}_2$ has degree $N - x_i$. In addition, there are in total $x_i(N - x_i)$ edges in the partial WA bipartite graph.

We next partition the set of edges, or equivalently the set of non-local transmit wavelengths from $\mathcal{S}_i$, into x_i subsets each with $N - x_i$ edges. Each subset of wavelengths are then used by some node $n_{i,j}$ in $\mathcal{S}_i$ (or equivalently $\mathcal{V}_2$) to transmit its $N - x_i$ non-local sessions to the $N - x_i$ nodes in $\mathcal{V}_1$. Thus, it is necessary that each subset of edges contains $N - x_i$ edges and is incident on all the nodes in $\mathcal{V}_1$, or else there would be a node in $\mathcal{V}_1$ not reachable from $\mathcal{S}_i$ in some subset of wavelengths. To achieve the goal of having $|\mathcal{P}^{(2)}_{i,j}| = 0$ for each $n_{i,j}$ in $\mathcal{S}_i$, we require in addition that each subset of edges is incident to all the nodes in $\mathcal{V}_2$. To see why this additional requirement is a sufficient condition for our goal, consider a given node $n_{i,j}$ in $\mathcal{S}_i$. Since every subset of edges is incident on $n_{i,j}$, it follows that each of the other nodes in $\mathcal{S}_i$ transmits a non-local session on a wavelength used by $n_{i,j}$ to receive a non-local session, and is thus a type-1 local node with respect to $n_{i,j}$. Therefore, with respect to each $n_{i,j}$ in $\mathcal{S}_i$, there is no type-2 local node, i.e. $|\mathcal{P}^{(2)}_{i,j}| = 0$.

Therefore, we want to partition $\mathcal{E}$ into x_i subsets of $N - x_i$ edges so that each subset is incident to all the nodes in $\mathcal{V}_1$ and $\mathcal{V}_2$. We showed in [10] that this partitioning problem can be solved by reducing it to a bipartite matching problem. For example, for the partial WA bipartite graph in figure 6a, figure 6b shows one possible partitioning of $\mathcal{E}$ such that each subset of edges is incident to all the nodes in $\mathcal{V}_1$ and $\mathcal{V}_2$.

As mentioned above, after the partition of $\mathcal{E}$, we assign the wavelengths corresponding to each subset of $\mathcal{E}$ to each $n_{i,j}$ in $\mathcal{S}_i$ to transmit its non-local sessions. For example, according to figure 6b, we assign subsets $\mathcal{E}_1$, $\mathcal{E}_2$, and $\mathcal{E}_3$ to $n_{1,1}$, $n_{1,2}$, and $n_{1,3}$ respectively. Node $n_{1,1}$ transmits its non-local sessions on wavelengths 1, 6, 8, 10, 13, and 16 to $n_{2,1}$, $n_{2,2}$, $n_{2,3}$, $n_{3,1}$, $n_{3,2}$, and $n_{4,1}$ respectively.

Similar techniques can be used to perform phase 2 for the top-level subtrees which do not satisfy the previous assumption that the set of non-local transmit wavelengths is equal to the set of non-local receive wavelengths (e.g. top-level subtree 3 in figure 4). We omit the details which can be found in [10].

Phase 3: In this phase, we assign wavelengths to local sessions in each top-level subtree. The assignment based on lemma 3 can be carried out independently in different top-level subtrees. From phase 2, in top-level subtree i, there are at least $|\mathcal{P}^{(2)}_{i,j}|$ type-2 reusable wavelengths with respect to node $n_{i,j}$ for all $1 \leq j \leq x_i$. Thus, we can assign wavelengths to all the local sessions as suggested by the statements in lemma 3.

The construction of the tree WA algorithm implies the following theorem.

Theorem 1 *In an arbitrary tree topology with 1-uniform traffic among leaf nodes, W_1 is given by*

$$W_1 \quad = \quad L_1 \quad = \quad w^* = \quad \max_{e \in \mathcal{T}} |\mathcal{N}_{e,1}||\mathcal{N}_{e,2}|.$$

Theorem 1 tells us that wavelength conversion cannot decrease the wavelength requirement for 1-uniform traffic in an arbitrary tree topology. In addition, from statement 2 of lemma 1, the minimum value of w^* is at least $\frac{1}{d^*}(1 - \frac{1}{d^*})N^2$. The tree topologies with w^*

close to this lower bound are the ones in which each top-level subtree has approximately N/d^* leaf nodes. Roughly speaking, it is desirable to have all the top-level subtrees support an equal amount of traffic.

It is a simple extension to establish that $W_l = lW_1$. First, we use the same argument as in the derivation of w^* in (1) to show that the bottleneck link e^* carries lw^* wavelengths in each fiber. Thus, $L_l \geq lw^*$. To show that $W_l \leq lw^*$, we apply the tree WA algorithm l times on l disjoint sets each of which contains w^* wavelengths. We state the result formally as a corollary to theorem 1.

Corollary 1 *For an arbitrary tree topology with l-uniform traffic among leaf nodes, W_l is given by*

$$W_l \quad = \quad L_l \quad = \quad lw^* = \quad l \max_{e \in \mathcal{T}} |\mathcal{N}_{e,1}||\mathcal{N}_{e,2}|.$$

4 Conclusion

We developed a WA algorithm which uses the minimum number of wavelengths to support l-uniform traffic in an arbitrary tree topology. Our algorithm involves finding the bottleneck node v^* in the tree, defining top-level subtrees connected to v^*, and assigning wavelengths to the non-local sessions such that these wavelengths can be reused to support local sessions in each top-level subtree. We observed that in a tree topology wavelength converters cannot be used to reduce the required number of wavelengths for l-uniform traffic. In addition, for fixed numbers of end nodes N and top-level subtrees d^*, the required number of wavelengths W_l is minimized when all the top-level subtrees contain approximately the same number of end nodes, i.e. N/d^* nodes.

References

[1] D. Banerjee and B. Mukherjee, "A practical approach for routing and wavelength assignment in large wavelength-routed optical networks," *IEEE JSAC*, vol. 14, no. 5, pp. 903-908, June 1996.

[2] R. Ramaswami and K.N. Sivarajan, "Design of logical topologies for wavelength-routed optical networks," *IEEE JSAC*, vol. 14, no. 5, pp. 840-851, June 1996.

[3] I. Chlamtac, A. Ganz, and G. Karmi, "Lightpath communications: an approach to high bandwidth optical WAN's," *IEEE Trans. on Comm.*, vol. 40, no. 7, pp. 1171-1182, July 1992.

[4] R. Ramaswami and K.N. Sivarajan, "Routing and wavelength assignment in all-optical networks," *IEEE/ACM Trans. on Networking*, vol. 3, no. 5, pp. 489-500, October 1995.

[5] Z. Zhang and A.S. Acampora, "A heuristic wavelength assignment algorithm for multihop WDM networks with wavelength routing and wavelength reuse," *IEEE/ACM Trans. on Networking*, vol. 3, no. 3, pp. 281-288, June 1995.

[6] S. Banerjee, J. Yoo, and C. Chen, "Design of wavelength-routed optical networks for packet switched traffic," *Jour. of Lightwave Tech.*, vol. 15, no. 9, pp. 1636-1646, September 1997.

[7] K.C. Lee and V.O.K. Li, "A wavelength rerouting algorithm in wide-area all-optical networks," *Jour. of Lightwave Tech.*, vol. 14, no. 6, pp. 1218-1229, June 1996.

[8] B. Mukherjee *et al.*, "Some principles for designing a wide-area WDM optical network,"*IEEE/ACM Trans. on Networking*, vol. 4, no. 5, pp. 684-696, October 1996.

[9] G. Wilfong, "Minimizing wavelengths in an all-optical ring network," in *Proceedings of International Symposium on Algorithms and Computation*, pp. 346-355, December 1996.

[10] P. Saengudomlert, *Architectural Study of High-Speed Networks with Optical Bypassing*, MIT Ph.D. Thesis, September 2002.