Uniform vs. Non-uniform Band Switching in WDM Networks

Li-Wei Chen, Poompat Saengudomlert, Eytan Modiano

***Abstract*— We compare the effectiveness of uniform versus non-uniform waveband switching under the dual cost metrics of switching requirements and fiber capacity. We consider a star topology and begin by characterizing the optimal performance frontier achievable under no restrictions on waveband sizing, and provide algorithms employing non-uniform waveband sizing that approach or achieve this optimum. We then consider the special case of uniform waveband sizing, and show that the performance compares very favorably. We also extend our results to general topologies.**

I. Introduction

THE MAJORITY of the routing and switching literature for WDM networks has focused on how to minimize the total number of wavelengths [1], [2], [3], [4], [5], [6], [7], [8], since the number of wavelengths used specifies the amount of capacity required on each fiber. However, this single-resource approach does not take into account the switching necessary at each node which routes each call to its destination. Switching costs can easily dominate bandwidth costs in systems with a large number of calls. In a conventional WDM system, each input fiber is typically demultiplexed into wavelengths, and each wavelength relies on a $N \times N$ switch to route it to the appropriate output fiber, where N is the nodal degree. The number of switches required is equal to the number of wavelengths. For large networks with many wavelengths, this approach can require many, often expensive switches.

Waveband switching, also known as band switching, attempts to address this problem. The approach is based on the observation that if the number of input wavelengths per fiber is large relative to the number of output fibers, many of the wavelengths will need to be switched between the same fiber pairs. Waveband switching tries to group the wavelengths into wavebands such that all wavelengths in the same waveband can be switched together, allowing the processing to be performed at this coarser waveband level and reducing the number of switches required. With waveband switching each fiber would only be demultiplexed into wavebands, and the number of switches required would equal the number of wavebands. Since the number of wavebands required is typically much smaller than the number of wavelengths, this can greatly reduce the processing and switching costs.

In this paper, we consider the resources of interest to be the number of wavelengths and wavebands required by a given banding algorithm. Reducing the requirement for either quantity reduces the costs in the network. Conceptually, every banding algorithm can be represented by a point in a two-dimensional performance space, as illustrated in Figure 1, indicating the number of wavelengths and wavebands required by the algorithm. The shaded area in the figure represents the achievable region of performance over all possible algorithms. The goal is to characterize the optimal frontier of achievable performance. This frontier would give the optimal tradeoff between wavelengths and wavebands achievable.

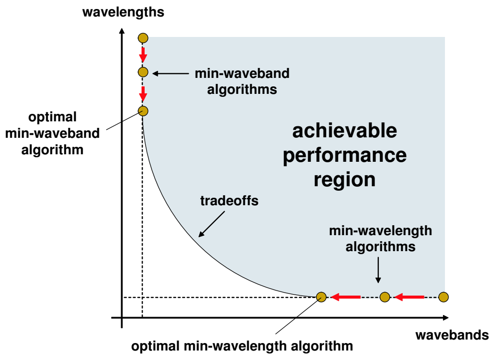

Fig. 1. The region of achievable wavelength-waveband tradeoffs.

There has been some work in the literature addressing the waveband switching problem. Many papers consider the problem of waveband allocation for static traffic. In [9], [10], [11], integer linear programming formulations are given for a variety of topologies, and the problem of optimal waveband allocation is shown to be NP-complete. In [12], efficient algorithms for dynamic traffic are considered under a simplified traffic model that limits traffic to a single source node and does not allow wavelength overprovisioning, even if it results in fewer wavebands.

In this paper, we consider dynamic traffic under the more general problem of determining the optimal tradeoff between wavelengths and wavebands in band switching. Furthermore, we allow a more general traffic model where every node is permitted to send traffic into the network. We first characterize the optimal tradeoff with no restrictions on waveband sizing, and derive efficient algorithms for achieving it. We then consider imposing the restriction that all wavebands must be uniformly sized, and show that the performance of uniform waveband algorithms compare very favorably with the optimum.

The authors are with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139, USA (email: lwchen@mit.edu, tengo@mit.edu, modiano@mit.edu)

This work was supported by the National Science Foundation (NSF) under Grant ANI-0073730.

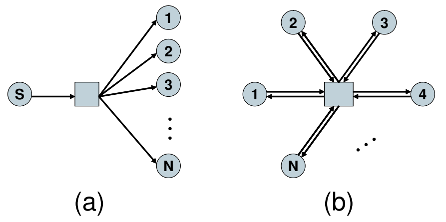

Fig. 2. (a) Single-source case, where a single source node sends a total of at most P calls to up to N destinations. (b) Multi-source case, where each of N nodes sends and receives a total of P calls.

A. System Model

In this paper, we adopt the P-port traffic model from [4], which assumes that P transmitters and receivers are available at each network node. This allows each node to send and receive a total of at most P calls at any given time. If the instantaneous traffic is represented by a matrix where each entry (i, j) consists of the number of calls sent from node i to node j, the P-port model constrains each row and column sum to be at most P. Any traffic set with a matrix obeying this constraint is termed *admissible*, and no calls in an admissible set may be blocked. Under this model, sufficient resources must be provisioned to support any admissible set. Call arrivals and departures may occur in arbitrary fashion, as long as the resultant traffic set remains admissible; these dynamic arrivals and departures are represented by transitions between different admissible sets. This model is attractive because it limits traffic in a realistic fashion based on hardware constraints, and also allows dynamic aspects of the traffic to be captured without making assumptions about the call statistics.

We primarily consider the star topology in this paper, with Section V describing extensions of our results to other topologies. This topology is representative of a hub or switch node in a network. All nodes are connected via bidirectional fibers to a central hub, which performs the switching. We assume no wavelength conversion, so calls must use the same wavelength on all hops. To avoid collision, no two calls may use the same wavelength in the same direction on the same fiber. We consider both the single-source case, where only a single node transmits, and the more general multi-source case, where every node may transmit. These situations are illustrated in Figure 2.

The problem of band switching under this model may be formulated as a matrix decomposition. Under the banding problem, we are given a traffic matrix C where each entry $[C]_{i,j}$ represents the number of calls transmitted from source node i to destination node j. In the single-source case, the traffic matrix is a vector of size $1 \times N$; for multi-source traffic, the traffic matrix is a square $N \times N$ matrix. Note that C may change over time due to call arrivals and departures. For a fixed C, the goal is to group the calls into bands such that calls within the same band that have the same source node go to the same destination. This can be expressed mathematically by

$$C \leq b_1 T_1 + b_2 T_2 + \ldots + b_B T_B \tag{1}$$

where each b_i is an integer representing the size of waveband i, and each T_i represents the corresponding switch configuration. T_i is therefore either a unit vector (single-source case) or a permutation matrix (multi-source).

Any set of wavebands specifying a valid decomposition is sufficient to support the particular traffic matrix C to which it applies. We impose two additional constraints: we require that B and $\{b_i\}$ must be fixed over all admissible traffic sets. Fixing B is essential since B corresponds to the number of switches required, a hardware requirement that should not depend on random changes in the traffic. Fixing the band sizes $\{b_i\}$ removes the need for dynamically tunable filters, reducing costs. Under these two constraints, we require that banding algorithms be characterized by fixed values of B and $\{b_i\}$ such that for each admissible traffic set C, the algorithm is able to specify a decomposition according to (1) with specific switch configurations T_i for each waveband i. Recall that the performance of each banding algorithm can be judged by the number of wavelengths and wavebands it requires. The number of wavebands is given directly by B, while the total number of wavelengths can be calculated as $\sum_{i=1}^{B} b_i$.

As illustrated in Figure 1, our goal is to find the optimal achievable frontier. The most general formulation of this problem is to allow the waveband sizes $\{b_i\}$ to be non-uniform. Any uniform-waveband algorithm will then be a valid special case. Under the general formulation, we can divide this problem into three parts. Note that Figure 1 shows two asymptotes to the achievable region: one corresponding to the minimum possible number of wavelengths required (the *minimum-wavelength asymptote*, shown on the bottom of the achievable region), and the other corresponding to the minimum number of wavebands (the *minimum-waveband asymptote*, shown to the left of the achievable region). The first two parts of the problem focus on these asymptotes, and attempt to determine the optimal points on these lines. Specifically, we denote any algorithm with performance achieving the minimum possible number of wavelengths to be a *minimum-wavelength algorithm*. The problem of finding the best minimum-wavelength algorithm is known as the *minimum-wavelength problem*. Similarly, any algorithm using the minimum possible number of wavebands is a *minimum-waveband algorithm*; finding the best minimum-waveband algorithm is the *minimum-waveband problem*.

Once solutions to the minimum-wavelength and minimum-waveband problems are obtained, two points on the optimal frontier are known. It then remains only to find the best tradeoff between wavelengths and wavebands achievable between these two points. Ideally, such a tradeoff should present a curve which adheres as closely as possible to the asymptotes, presenting the best possible tradeoff. Obtaining the best possible such tradeoff is the subject of the final component of the banding problem.

In Section II, the banding problem is investigated for the single-source traffic case. In Sections III and IV, banding for general multi-source traffic is considered. We will compare the special case of uniformly-sized wavebands to the more general non-uniform case, and show that uniform waveband sizing compares very favorably. Finally, Section V describes extending the results of the paper to general network topologies.

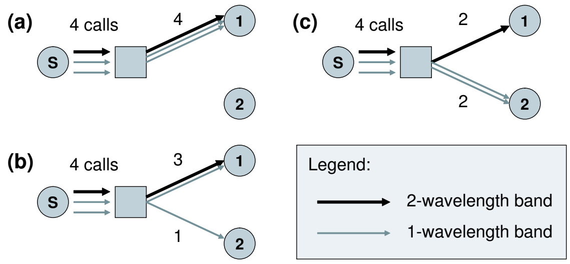

Fig. 3. The switching configurations for the 3 unique maximal traffic sets of Example 1. The traffic sets shown are: (a) [4,0] (b) [3,1] (c) [2,2].

II. Waveband Switching for Single-Source Traffic

In this section, we consider the banding problem for the case of single-source traffic. Under this scenario, a single source node sends up to P units of traffic to be switched to N possible destination nodes. The primary purpose of investigating the single-source model is to derive intuition for the design of good banding algorithms that will be beneficial in addressing the multi-source traffic case in the next section. The single-source model also has some merit in cases of a one-to-many traffic scenario.

Example 1: Consider the case of $P = 4$, $N = 2$. In this example, the source sends 4 calls, distributed among 2 destinations. There are only 5 possible maximal traffic sets, which (expressed in vector form) are $[4,0]$, $[3,1]$, $[2,2]$, $[1,3]$, and $[0,4]$. Clearly at least 4 wavelengths are required to support the traffic, since there are 4 calls. We can show that if we restrict ourselves to using only 4 wavelengths (i.e. we consider the minimum-wavelength problem), the minimum number of wavebands required is 3: one band of size 2, and two bands of size 1. By exhaustive verification we can prove that this waveband sizing is sufficient for all possible traffic sets:

$$
\begin{aligned}
[4,0] &= 2\cdot[1,0] + 1\cdot[1,0] + 1\cdot[1,0] \\
&= 2e_1 + e_1 + e_1 \\
[3,1] &= 2e_1 + e_1 + e_2 \\
[2,2] &= 2e_1 + e_2 + e_2
\end{aligned}
$$

where e_i is a unit vector with the i^{th} entry equal to 1.

Note that, as required, the sizes of each band and total number of bands are fixed, and only the accompanying unit vectors (which correspond to the switch configurations for each band) change between traffic sets. The switching of each waveband for each scenario is illustrated in Figure 3. In this example, the savings in switching is not large because the number of calls is not very large relative to the number of destinations. As the number of calls increases, the savings will increase as well.

The number of wavebands can be further reduced if the use of additional wavelengths is permitted. One possibility is to have one band of size 3, and one band of size 2. This reduces the number of wavebands to two, and these two wavebands can still support all 5 possible traffic sets. However, the total number of wavelengths used has increased to 5. Efficient methods of making these sorts of tradeoffs will be discussed.

We consider the two special cases of the minimum-wavelength and minimum-waveband problems. Recall that the solutions to these two problems will provide two points on the optimal achievable performance frontier. We will defer the discussion of obtaining good tradeoffs between these two points until the multi-source case.

A. The Minimum-Wavelength Problem

Recall that for the minimum-wavelength problem, we consider only banding algorithms that use the minimum possible number of wavelengths. Under the P-port model, a minimum of P wavelengths are clearly necessary and sufficient: up to P calls can be sent, and if each wavelength is individually switched, P wavelengths can support all the calls. The goal is to find the optimal algorithm that uses only P wavelengths.

Since the number of wavelengths equals the maximum possible number of calls, given a maximal admissible traffic set, each wavelength must be used to support a call. Therefore, the minimum-waveband problem is equivalent to finding a method of partitioning the P wavelengths into wavebands such that, for any admissible traffic set, there exists a method for assigning calls to wavebands such that every wavelength is assigned a call. Furthermore, the optimal minimum-wavelength algorithm should accomplish this while at the same time minimizing the total number of wavebands. We show in this section that the optimal minimum-wavelength banding algorithm is a greedy algorithm. Specifically, the greedy algorithm chooses waveband sizes recursively, where at each step a waveband is chosen to be as large as possible subject to the constraint that every wavelength in that band can always be assigned a call under any maximal admissible traffic set. (We say that in this case every wavelength can be *fully utilized*.)

Define $b_{max}(N,P)$ to be the maximum waveband size that we can guarantee will be fully utilized by any traffic set sending P calls to N destinations. Since all calls in the waveband must go to the same destination, this is equivalent to providing a guarantee that a destination node can always be found (under *any* admissible P-port traffic set) which receives at least $b_{max}(N,P)$ calls. To illustrate, when considering the example in Figure 3, we note that over all admissible traffic sets, a destination can always be found which receives at least 2 calls, leading to the conclusion that $b_{max}(N,P) = 2$ in that case. In general, we can guarantee that at least one of the destinations receives $\lceil P/N \rceil$ calls. Furthermore, this is the largest number for which we can make this guarantee; this follows from the fact that one admissible traffic set is where the traffic is divided evenly (up to a difference of one wavelength due to integer constraints) among all destinations, and no destination receives more than $\lceil P/N \rceil$ calls under this traffic set. Therefore $b_{max}(N,P) = \lceil P/N \rceil$.

Single-Source Greedy Algorithm:

1) Let $P_1 = P$ be the number of calls remaining and N be the number of nodes. Let $i = 1$.
2) Let waveband i be of size $b_i = b_{max}(N, P_i) = \lceil P_i/N \rceil$.
3) Locate a destination receiving at least $b_{max}(N, P_i)$ calls. Route waveband i to this destination, and assign $b_{max}(N, P_i)$ calls to it. The number of calls remaining becomes $P_{i+1} = P_i - b_{max}$.
4) If $P_{i+1} > 0$, let $i \leftarrow i + 1$ and go to Step 2.

Example 2: We revisit Example 1 and show how the greedy algorithm is used to obtain the minimum-wavelength waveband sizes used in Figure 3. In that example, $P = P_1 = 4$ and $N = 2$. In the first iteration, the greedy algorithm chooses the first waveband to be of size $\lceil P_1/N \rceil = \lceil 4/2 \rceil = 2$. As a corollary, we are guaranteed that 2 calls can be assigned to this waveband, leaving $P_2 = 2$ calls unassigned. The next two iterations partition the remaining wavelengths into bands of a single wavelength each, for a final partition of $\{2, 1, 1\}$.

We must now show that choosing the waveband sizes using the greedy algorithm is optimal for the minimum-wavelength problem. The full proof, omitted here for brevity, is based on establishing that the minimum number of wavebands required to support a given number of calls is non-decreasing in the number of calls. It therefore follows that the optimal approach for choosing each waveband size is to choose the band as large as possible, thereby minimizing the amount of traffic which remains (and therefore the subsequent number of wavebands).

The greedy algorithm provides a method for optimally determining waveband sizes for the minimum-wavelength problem. This also implicitly provides a way of determining the minimum number of wavebands required (i.e. by running the algorithm and counting the number of wavebands produced). We can also derive an explicit upper bound on the minimum number of wavebands required in the minimum-wavelength scenario. We proceed by relaxing the integer constraints on $b_{max}(N, P)$. Let P_k be the number of calls remaining after running the k^{th} iteration of the greedy algorithm. The series progresses as follows:

$$P_k = \left(1 - \frac{1}{N}\right)^k \cdot P \tag{2}$$

If $P \leq N$, then the number of bands B is simply equal to P since each band is composed of only a single wavelength. Therefore consider $P > N$ and determine the number of bands k required to reduce the number of unassigned wavelengths to N. It is straightforward to show that $k = \frac{\log(N/P)}{\log(1-1/N)}$. Then the total number of wavebands is simply $k + N$. Since relaxing the ceiling constraints underestimates the size of each waveband, this gives an upper bound on the number of wavebands B, namely:

$$B \leq \begin{cases} N + \frac{\log\left(\frac{N}{P}\right)}{\log\left(1-\frac{1}{N}\right)} & , P > N \\ P & , P \leq N \end{cases} \tag{3}$$

From (3), we can also make the additional observation that if $P \leq N$, the number of bands B equals the number of wavelengths P, and there is no savings from banding in the minimum-wavelength case as each wavelength continues to be switched individually.

B. *The Minimum-Waveband Problem*

The optimal minimum-waveband algorithm is the one that requires the fewest wavelengths subject to using only the minimum number of wavebands. In addition to providing a second point on the optimal frontier, this will establish the minimum cost in wavelengths required to obtain the maximum possible reduction in switching.

Since all wavelengths in the same waveband must be switched to the same destination, and there are N possible destination, a minimum of N wavebands are are necessary. One (inefficient) approach that requires only N wavebands is to statically switch one waveband to each destination, and provision P wavelengths per waveband; since there are a total of only P calls, this is sufficient to support any admissible traffic set. Our goal is to find a better, optimal algorithm using only N wavebands that minimizes the number of wavelengths used. We first obtain a lower bound on the number of wavelengths required using the following lemma.

Lemma 1: Consider a banding algorithm that uses N wavebands, and order the wavebands from smallest to largest. Let b_i be the size of the i^{th} waveband. If the source sends up to P calls to the N destination nodes, b_i is bounded by

$$b_i \geq \left\lfloor \frac{P - N + i}{i} \right\rfloor, \quad i = 1, \ldots, N \tag{4}$$

Corollary: The total number of wavelengths W required is bounded by the sum of the bounds on the individual waveband sizes, namely

$$W \geq \sum_{i=1}^{N} \left\lfloor \frac{P - N + i}{i} \right\rfloor \tag{5}$$

This summation can be shown to increase as $O(P \log N)$.

Proof: The proof proceeds by constructing an admissible traffic set which requires b_i to have at least $\left\lfloor \frac{P-N+i}{i} \right\rfloor$ wavelengths. Consider the traffic set where the source S sends $\left\lfloor \frac{P-N+i}{i} \right\rfloor$ calls each to the first i nodes, and a single call to each remaining node. The total traffic in this construction is

$$i \cdot \left\lfloor \frac{P - N + i}{i} \right\rfloor + (N - i) \leq P$$

and therefore it is admissible. Since each destination receives at least one call, each of the N wavebands goes to a different destination.

Without loss of generality, we assign the largest i wavebands to nodes 1 through i. Each of these wavebands must support $\left\lfloor \frac{P-N+i}{i} \right\rfloor$ calls. Therefore $b_i \geq \left\lfloor \frac{P-N+i}{i} \right\rfloor$. ■

Since Lemma 1 is a lower bound, any minimum-waveband algorithm which achieves the bound is optimal. We next present an algorithm that can support any admissible traffic set using wavebands of the minimal sizes specified by (4). Since this minimum-waveband algorithm would use no more wavelengths than the lower bound, it is therefore optimal.

Min-Band Algorithm:

1) Index the N waveband in order of decreasing size, so that waveband i has size $b_i = \left\lfloor \frac{P-N+i}{i} \right\rfloor$, where $i = 1, \ldots, N$. Note that b_1 is the largest waveband, and b_N is the smallest.
2) Let $i = 1$.
3) Locate the destination node with the greatest number of remaining calls. Switch waveband i to that node. Assign up to b_i calls to waveband i, and remove these calls from the traffic set.
4) If no calls remain, the algorithm terminates. Otherwise, increment i and return to Step 2.

By design, the algorithm uses only wavebands of the minimum size, and therefore meets the lower bound. It remains only to show that it is able to support any admissible set. First, suppose that each destination receives at least a single call. In this case, we can rank each destination in decreasing order of number of calls received, so that the first destination receives the most calls. Then the min-band algorithm allocates the i^{th} waveband to the i^{th} destination. Since each destination receives at least one call, the first i destinations receive at most $P-(N-i) = P-N+i$ calls, and the i^{th} destination receives at most $\lfloor (P-N+i)/i \rfloor$ calls. Since the i^{th} waveband has size $\lfloor (P-N+i)/i \rfloor$, it suffices to accommodate the calls. The proof in the case where some destinations do not receive calls is more cumbersome but follows the same approach.

Example 3: Consider the case where $P = 22$ calls are distributed among $N = 4$ destinations. The optimum minimum-waveband algorithm requires 4 wavebands. According to Lemma 1, the first waveband is of size $b_1 = \lfloor P-N+1 \rfloor = \lfloor 22-4+1 \rfloor = 19$. Similarly, $b_2 = 10$, $b_3 = 7$, and $b_4 = 5$.

These wavebands can support any P-port admissible traffic set, $P = 22$. For example, consider the traffic set $C = [5, 8, 7, 2]$. The first waveband is assigned to node 2, the destination with the most traffic, and carries all 8 calls. Similarly, b_2 is assigned to node 3, b_3 is assigned to node 1, and b_4 is assigned to node 4. Using the matrix decomposition notation of (1), this can be written as

$$\begin{aligned} [5,8,7,2] &\leq 19e_2 + 10e_3 + 7e_1 + 5e_4 \\ &= [7, 19, 10, 5] \end{aligned}$$

Note that here, the total number of wavelengths available to each destination (represented by the vector on the right-hand side of the equation) is greater than the number of calls: more wavelengths were provisioned than absolutely required for this particular traffic set. This over-provisioning in wavelengths is necessary in order to guarantee that *all* admissible traffic sets can be accommodated.

It is instructive to compare this to the decomposition obtained by the greedy algorithm of Section II-A. The first waveband for the greedy algorithm consists of $\lceil P/N \rceil = \lceil 22/4 \rceil = 6$ wavelengths. The remaining waveband sizes can be shown to be $\{4, 3, 3, 2, 1, 1, 1, 1\}$. One possible switching configuration for these waveband sizes is:

$$\begin{aligned} [5,8,7,2] &= 6e_2 + 4e_3 + 3e_1 + 3e_3 + 2e_1 + e_2 \\ &\quad + e_2 + e_4 + e_4 \\ &= [5, 8, 7, 2] \end{aligned}$$

The greedy algorithm, since it is a minimum-wavelength algorithm, did not over-provision any wavelengths; however, more wavebands were required. This example also illustrates an important point. In the decomposition given by (1), equality is guaranteed to hold for all C if and only if the wavebands were allocated using a minimum-wavelength algorithm (such as the greedy algorithm). Algorithms such as the min-band algorithm allow some overprovisioning of wavelengths in order to further decrease the number of wavebands required.

III. Waveband Switching for multi-source Traffic

We now consider the more general case of multi-source traffic. In this scenario, N nodes are connected to a central hub. Each node is assumed to have a hardware limitation of P transmitters and receivers, and can therefore send and receive up to P calls. The hub must switch the calls, at a band level, from the appropriate source to destination nodes. In our discussion, we will assume that self-traffic is allowed; the case without self-traffic is similar and leads to comparable results. We will show that many of the concepts in this scenario parallel those in the single-source case. (The primary differences are that the traffic set now consists of a traffic matrix rather than a vector, and the switching configuration for each waveband will now consist of a permutation matrix rather than a unit vector.) We will again begin by considering two special cases, the minimum-wavelength and minimum-waveband problems, followed by investigating algorithms that provide a tradeoff between these two cases.

A. The Minimum-Wavelength Problem

Recall that for the minimum-wavelength problem, we constrain ourselves to the domain of banding algorithms which use only the minimum possible number of wavelengths. Since each node may send up to P calls, it is clear that at least P wavelengths are necessary. In [13] it is shown that this is also sufficient. The challenge is therefore to first partition the P wavelengths into wavebands, and second, to develop an algorithm that will provide a valid wavelength assignment for any admissible traffic set using these wavebands.

We first address the partitioning of the wavebands. We consider the cases of interest to be maximal traffic sets, since we can add fictitious calls to any non-maximal set to construct a maximal one. We define a wavelength to be *fully utilized* if it is used to carry a call on every link. Mathematically, this is equivalent to stating that the matrix of calls supported by that wavelength forms a permutation matrix. We say that a waveband is fully utilized if every wavelength in that waveband is fully utilized.

Every minimum-wavelength algorithm must be able to fully utilize every waveband under any admissible maximal traffic set. We have already seen in Section II-A that a greedy approach is optimal for this type of problem. Recall that the greedy algorithm worked recursively by partitioning, at each step, the largest waveband size possible subject to the constraint that it could be fully utilized by any admissible traffic set. Intuitively, this was because the number of wavebands required turned out to be non-increasing in the amount of traffic; therefore at each step the optimal solution was to minimize the residual amount of traffic. The only remaining problem is that largest fully utilizable waveband size, $b_{max}(N, P)$, must be derived for the multi-source case.

We can view $b_{max}(N, P)$ as the largest number of identical permutation matrices that we are guaranteed to be able to find within *any* $N \times N$ matrix with row and column sums equal to P. This is equivalent to stating that sufficiently many calls must exist within any admissible traffic set with parameters N and P to fully utilize the waveband.

Example 4: Consider the case where $P = 9$ and $N = 3$. One admissible maximal traffic set satisfying these constraints is given by

$$C_1 = \begin{bmatrix} 1 & 5 & 3 \\ 3 & 2 & 4 \\ 5 & 2 & 2 \end{bmatrix}$$

The largest waveband that could be fully utilized by this particular traffic set is 4 wavelengths, as shown below:

$$C_1 = 4 \cdot \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 1 & 1 & 3 \\ 3 & 2 & 0 \\ 4 & 2 & 2 \end{bmatrix}$$

The first term consists of 4 identical permutation matrices representing calls that can be used to fully utilize a waveband of size 4 or less; the second term forms the remaining calls. Note that this shows only that b_{max} must be *at most* 4; it may be possible that some other traffic set may require an even smaller waveband for full utilization. The challenge will be to obtain in closed form an equation for b_{max} without performing this sort of exhaustive examination of all possible maximal admissible traffic sets. This is given by the following lemma.

Lemma 2: For any $N \times N$ matrix with row and column sums at most P, there exist at least b_{max} identical permutation matrices within the matrix, with b_{max} given by

$$b_{max} = \begin{cases} \left\lceil \frac{4P}{N(N+2)} \right\rceil & , N \text{ even} \\ \left\lceil \frac{4P}{(N+1)^2} \right\rceil & , N \text{ odd} \end{cases} \tag{6}$$

Furthermore, there exists at least one $N \times N$ matrix where no more than b_{max} identical permutations can be found.

Proof: See Appendix A. ■

The relation in (6) can be used in conjunction with the greedy algorithm to optimally partition the wavebands. Since the waveband sizing is independent of the particular traffic set and depends only on the network parameters N and P, wavebands of constant size can be used for any admissible traffic set. By using the expression $b_{max}(N, P)$ to size each waveband, we have guaranteed that a permutation of the appropriate size can be found to fully utilize each waveband.

Example 5: We continue with the 3-node network from Example 4 where $P = 9$. Using the greedy algorithm, we would determine that the largest waveband should be $\left\lceil \frac{4P}{(N+1)^2} \right\rceil = \left\lceil \frac{(4)(9)}{(4)^2} \right\rceil = 3$.

After this step, $9 - 3 = 6$ wavelengths remain to be partitioned. We repeat this process until all wavelengths have been assigned to bands. The final waveband partition is $\{3, 2, 1, 1, 1, 1\}$. By choice of b_{max}, this partitioning can be fully utilized by any admissible traffic set. For example, consider the traffic set C_1 from Example 4. One possible decomposition is

$$\begin{aligned} \begin{bmatrix} 1 & 5 & 3 \\ 3 & 2 & 4 \\ 5 & 2 & 2 \end{bmatrix} &= 3 \cdot \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} + 2 \cdot \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \\ &+ 1 \cdot \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix} + 1 \cdot \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \\ &+ 1 \cdot \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} + 1 \cdot \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} \end{aligned}$$

Note that the equality in this decomposition indicates that all wavelengths are fully utilized. Furthermore, by the optimality of the greedy algorithm, we are guaranteed that this is the minimum possible number of wavebands subject to the minimum-wavelength constraint.

In principle, with the greedy algorithm, we can obtain the exact minimum number of wavebands required under the minimum-wavelength constraint simply by iterating through the algorithm and counting the number of wavebands produced. We can also obtain in closed form an upper bound on the number of wavebands by relaxing the integer constraints. Using arguments analogous to the single-source case, we can show that an upper bound on the number of wavebands B is given by

$$B \leq \begin{cases} \frac{(N+1)^2}{4} + \frac{\log\left[\frac{(N+1)^2}{4P}\right]}{\log\left[1 - \frac{4}{(N+1)^2}\right]} & , P > \frac{(N+1)^2}{4} \\ P & , P \leq \frac{(N+1)^2}{4} \end{cases} \tag{7}$$

It can be show from (7) that the number of wavebands required grows as $O(N^2 \log(P/N^2))$. Since the number of wavebands required by the greedy algorithm is the minimum possible for any minimum-wavelength algorithm, this allows us to quantify the maximum switching reduction possible without wavelength inefficiency.

B. The Minimum-Waveband Problem

Recall that a minimum-waveband algorithm is defined to be a banding algorithm that uses the minimum possible number of wavebands. Since all wavelengths in the same waveband must go to the same destination, and there are N possible different destinations, a minimum of N wavebands are required. One way to achieve this is to statically provision P wavelengths between each source-destination pair, using a total of PN wavelengths. The minimum-waveband problem is therefore to find a better, dynamic algorithm that uses fewer wavelengths. Since the single-source traffic model is a special case of the multi-source model, we can also use (5) to provide a lower bound on the number of wavelengths required:

$$W \geq \sum_{i=1}^{N} \left\lfloor \frac{P - N + i}{i} \right\rfloor$$

However, in this case it is possible to show that the bound is not tight. We therefore do not know the achievable minimum number of wavelengths, only that it cannot be less than that specified by (5).

We next propose a wavelength-efficient minimum-banding algorithm which requires $O(P\sqrt{N})$ total wavelengths, which improves on the $O(PN)$ worst case. The algorithm operates by decomposing the traffic set into N sub-matrices, and attempts to group entries with heavy weights and light weights into separate sub-matrices. This will allow some wavebands to use less than the worst case of P. The algorithm relies on the following lemma:

Lemma 3: Consider a P-port traffic set on an N-node star. For any value of k such that $1 \leq k \leq N$, there exists a decomposition satisfying (1) where

$$\begin{aligned} b_1 &= \ldots = b_k = P \\ b_{k+1} &= \ldots = b_N = \left\lceil \frac{P}{k+1} \right\rceil \end{aligned}$$

Proof: See Appendix B. ∎

Corollary: Any P-port N-node traffic set can be routed using k bands of size P and $N-k$ bands of size $\left\lceil \frac{P}{k+1} \right\rceil$.

The proof of Lemma 3 forms the basis of the SQRT(N) algorithm. Since it holds for any value of k, it is logical to use the value of k which results in the fewest total number of wavelengths used. To determine this, we write down the expression for the number of wavelengths required:

$$W_k = kP + (N-k)\left(\frac{P}{k+1}\right) \quad (8)$$

It can be shown that this expression is minimized at $k = \sqrt{N+1}-1$. If we relax the integer constraint on k and substitute this back into the equation, we obtain:

$$W = 2P\left(\sqrt{N+1}-1\right)$$

leading to the observation that the SQRT(N) algorithm requires $O(P\sqrt{N})$ wavelengths. The results of this section show that the maximum amount of switching reduction can be achieved by, at worst, a factor of $\sqrt{N}$ increase in the number of wavelengths.

Example 6: We examine the case of Example 5 under the minimum-waveband restriction. Using the SQRT(N) algorithm, we see that for $N = 3$ and $P = 9$, we should choose $k = \sqrt{N+1}-1 = 1$. We therefore require only 1 waveband of size 9 and $N-1 = 2$ bands of size $\lceil P/(k+1) \rceil = 5$, producing a final waveband sizing of $\{9,5,5\}$, for a total of 19 wavelengths and 3 bands. We can compare this with the optimal minimum-wavelength solution of $\{3,2,1,1,1,1\}$, which uses the minimum of 9 wavelengths but requires 6 wavebands.

Shown below is one possible call assignment using these waveband sizes for the specific traffic set C_1.

$$\begin{bmatrix} 1 & 5 & 3 \\ 3 & 2 & 4 \\ 5 & 2 & 2 \end{bmatrix} \leq 9 \cdot \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} + 5 \cdot \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} + 5 \cdot \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 5 & 9 & 5 \\ 5 & 5 & 9 \\ 9 & 5 & 5 \end{bmatrix}$$

Note that wavelengths were overprovisioned in this case for some source-destination pairs. However, the additional wavelengths are required to guarantee that all admissible traffic sets can be accommodated by the 3 wavebands.

C. Hybridization: Wavelength-Waveband Tradeoffs

To obtain a tradeoff in performance between the greedy algorithm (which has minimal wavelength usage) and the SQRT(N) algorithm (which has minimal waveband usage), we can use hybrid algorithms that size some wavebands using the greedy algorithm and the remaining using SQRT(N). By varying how many wavebands are sized using each algorithm, a family of hybrid algorithms can be obtained that can be shown to have the following wavelength-waveband tradeoff:

$$W_{dual} = \left\lceil 1 + 2\sqrt{N}\left(1 - \frac{4}{(N+1)^2}\right)^{B_{dual}-N} \right\rceil \cdot P$$

Figure 4 shows the performance of this approach for a star with $N = 10$, $P = 1000$. The two asymptotes represent the SQRT(N) and greedy algorithms, which operate in the minimum-waveband and minimum-wavelength regimes, respectively. The dual-algorithm approach essentially interpolates between the minimum-waveband and minimum-wavelength cases to produce algorithms with intermediate waveband and wavelength requirements.

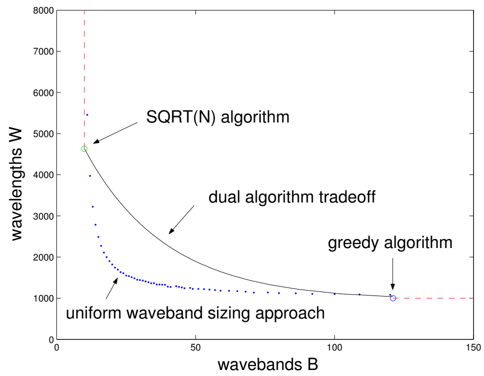

Fig. 4. Performance of the non-uniform banding algorithms of Section III compared to the uniform banding algorithm of Section IV in a star with $N = 10$, $P = 1000$.

IV. The Uniform Waveband Approach

Thus far in discussing multi-source traffic, we have allowed wavebands to be non-uniformly sized. Here we consider fixing all wavebands to a constant size b and derive the number of wavelengths and wavebands required. By varying b, a family of banding algorithms with varying numbers of wavebands and wavelengths can be obtained. Somewhat surprisingly, we will see that in the multi-source case, by using uniform waveband sizes, most of the maximum banding gain can be obtained at a very small cost in additional wavelengths.

We begin by first deriving the minimum number of wavebands required for a fixed waveband size b.

Theorem 1: Given a fixed band size b, the necessary and sufficient minimum number of wavebands required to support P-port traffic in an N-node star is

$$B_{uniform} = N + \left\lfloor \frac{P-N}{b} \right\rfloor \quad (9)$$

Corollary: The corresponding necessary and sufficient minimum number of wavelengths required is

$$W_{uniform} = bN + b\left\lfloor \frac{P-N}{b} \right\rfloor \quad (10)$$

We first prove necessity of (9) by providing an example which requires at least this number of wavelengths. Sufficiency will be shown by construction. (10) then follows directly from the fact that each band is of size b.

Consider the traffic set where node 1 sends a single call to nodes 1 to $N-1$, and $P-(N-1)$ calls to node N. In this case, $N-1$ bands are required to support traffic to the first $N-1$ nodes, while $\left\lceil \frac{P-(N-1)}{b} \right\rceil$ bands are required to support traffic to node N. This gives a lower bound on the number of wavebands of

$$\begin{aligned} B_{uniform} &\geq (N-1) + \left\lceil \frac{P-(N-1)}{b} \right\rceil \\ &= N + \left\lfloor \frac{P-N}{b} \right\rfloor \end{aligned}$$

where the last step follows from the observation that $\left\lceil \frac{P-(N-1)}{b} \right\rceil = \left\lfloor \frac{P-N}{b} \right\rfloor + 1$.

We will show using a bipartite matching approach that these quantities are sufficient as well. We will first construct a bipartite multigraph. A multigraph differs from a graph in that it allows multiple edges between the same two nodes. We consider two sets of nodes $V_1 = \{s_1, \ldots, s_N\}$ and $V_2 = \{d_1, \ldots, d_N\}$. For a given admissible traffic set, define the number of calls from node i to node j to be $c_{i,j}$. Then create $\lceil \frac{c_{i,j}}{b} \rceil$ edges connecting node s_i to d_j. The complete set of edges $\mathcal{E}$ represents the traffic to be carried, now in units of wavebands (each of size b) instead of wavelengths.

The number of edges adjacent to each source node s_i can be obtained by summing $\lceil \frac{c_{i,j}}{b} \rceil$ over j. Let the number of destinations receiving non-zero traffic from node i be N_i. Without loss of generality, assume these are nodes 1 through N_i. We decompose the traffic to these destinations by

$$c_{i,j} = \omega_{i,j} + r_{i,j} \tag{11}$$

where $r_{i,j}$ is chosen such that $\omega_{i,j}$ is a nonnegative integer multiple of b, and $1 \le r_{i,j} \le b$. We can then express the summation of interest as

$$\begin{aligned} \sum_{i=1}^{N} \left\lceil \frac{c_{i,j}}{b} \right\rceil &= \sum_{i=1}^{N} \left\lceil \frac{\omega_{i,j} + r_{i,j}}{b} \right\rceil \\ &= \sum_{i=1}^{N_i} \left\lceil \frac{r_{i,j}}{b} \right\rceil + \sum_{i=1}^{N_i} \frac{\omega_{i,j}}{b} \\ &= N_i + \frac{1}{b} \sum_{i=1}^{N_i} \omega_{i,j} \end{aligned} \tag{12}$$

where the second step relies on $\frac{\omega_i}{b}$ being integer, and the third on the fact that $r_i \le b$.

By summing (11) over i and noting that $\sum_{i=1}^{N_i} c_{i,j} = P$, we can obtain the following useful relation:

$$\begin{aligned} P &= \sum_{i=1}^{N_i} \omega_{i,j} + \sum_{i=1}^{N_i} r_{i,j} \\ \Rightarrow \sum_{i=1}^{N_i} \omega_{i,j} &= P - \sum_{i=1}^{N_i} r_{i,j} \\ &\le P - N_i \end{aligned}$$

where the last line results from observing that $r_{i,j} \ge 1$. Using this result, (12) becomes

$$\begin{aligned} \sum_{i=1}^{N} \left\lceil \frac{c_{i,j}}{b} \right\rceil &\le N_i + \frac{P - N_i}{b} \\ &= \left(1 - \frac{1}{b}\right) N_i + \frac{P}{b} \\ &\le \left(1 - \frac{1}{b}\right) N + \frac{P}{b} \\ &= N + \frac{P - N}{b} \end{aligned}$$

Since both the summation on the left and N on the right are integers, by taking the floor of both sides we can conclude

$$\sum_{i=1}^{N} \left\lceil \frac{c_{i,j}}{b} \right\rceil \le N + \left\lfloor \frac{P - N}{b} \right\rfloor$$

Similar arguments can be used to show that the number of edges adjacent to each destination node d_j is given by

$$\sum_{j=1}^{N} \left\lceil \frac{c_{i,j}}{b} \right\rceil \le N + \left\lfloor \frac{P - N}{b} \right\rfloor$$

Therefore each node has at most $N + \lfloor \frac{P-N}{b} \rfloor$ edges adjacent to it. The following lemma will now prove useful.

Lemma 4: In a bipartite multigraph where each node is adjacent to at most k edges, a partitioning exists that divides the edges into at most k matchings.

Proof: See [14]. ■

By Lemma 4, at most an equal number of matchings are required. Since calls in each matching can share the same waveband, at most $N + \lfloor \frac{P-N}{b} \rfloor$ wavebands are required. The number of wavelengths follows directly from the fact that each waveband is of size b.

Example 7: Consider a star with $N = 10$ and $P = 1000$, and consider uniform band sizes of $b = 40$. By Theorem 1, $N + \lfloor \frac{P-N}{b} \rfloor = 10 + \lfloor \frac{1000-10}{50} \rfloor = 29$ wavebands are required. Since each waveband consists of 50 wavelengths, a total of $50 \cdot 29 = 1711$ wavelengths are used as well.

We now have a method for obtaining necessary and sufficient conditions on the number of wavebands and wavelengths are necessary and sufficient. By ignoring the integrality constraints, we can solve for the number of wavelengths as a function of the number of wavebands and obtain an approximate characterization of the waveband-wavelength tradeoff using uniformly-sized wavebands:

$$W_{uniform} = P + \left(\frac{P - N}{B_{uniform} - N} - 1 \right) N \tag{13}$$

Figure 4 illustrates the performance of the uniform waveband algorithm for a star with $N = 10$, $P = 1000$. The two asymptotes represent the SQRT(N) and greedy algorithms, which operate in the minimum-waveband and minimum-wavelength regimes, respectively. Note that although the uniform waveband algorithm performs poorly in the minimum-waveband regime (where it requires many more wavelengths than the SQRT(N) algorithm), the performance improves dramatically once a few additional wavebands are introduced. By around 40 wavebands, it requires only slightly more wavelengths than the greedy algorithm, which uses 121 wavebands. We observe that by allowing slightly more wavelengths than the minimum-wavelength case, the fixed-waveband algorithm can greatly reduce the number of wavebands required, approaching the minimum-waveband bound significantly.

From this graph, two observations can be made:

1) As the number of wavebands increases, the performance of the uniform-waveband algorithm appears to approach the optimal performance of the greedy algorithm. In particular, at the right endpoint, it appears to be almost wavelength-efficient.
2) Because of the slow increase in the number of wavelengths required as the number of wavebands decreases, it appears that the majority of the reduction in the number of wavebands can be achieved at very little cost in wavelength inefficiency (as compared to the greedy algorithm).

The first observation can be verified by comparing the number of wavelengths used by the uniform-waveband algorithm to the greedy algorithm. For P large we can approximate B_{greedy} by $\frac{N^2}{4}\log\left(\frac{4P}{N^2}\right)$. For this number of bands, the number of wavelengths used by the uniform-waveband algorithm is approximately

$$\begin{aligned} W_{uniform} &\approx P + \frac{PN}{B} \\ &= \left[1 + \frac{1}{\frac{N^2}{4}\log\left(\frac{4P}{N^2}\right)}\right] \cdot P \\ &= (1+\alpha)\cdot P \end{aligned}$$

where α is a term that goes to zero as P increases. Recall that the greedy algorithm, which was wavelength-efficient, uses the minimum of P wavelengths. Therefore the performance of the uniform-waveband algorithm approaches the optimum asymptotically in the minimum-wavelength regime.

It is also possible to show analytically by slope analysis that $W_{uniform}$ approaches its final value very quickly; this gives rise to the second observation, which is extremely significant from a practical perspective. If we are interested in building an actual implementation, it indicates that a majority of the gain from using banding can be achieved with very little wavelength inefficiency. For example, in the graph of Figure 4, the processing granularity can be reduced from 1000 wavelengths (without banding) to 30 wavebands, a reduction of 97%, at a cost of only a 50% increase in the number of wavelengths.

V. Banding on General Topologies

Thus far all our results have been for the star topology. In this section we extend the preceding banding results to general topologies for which routing algorithms for P-port traffic are known.

Recall that we have shown that banding can be considered as a matrix decomposition problem, where for a given admissible traffic set C, our goal is to decompose it into the sum of a fixed number B of weighted permutation matrices:

$$C \le b_1 T_1 + b_2 T_2 + \ldots + b_B T_B$$

where the band sizes $\{b_i\}$ and the total number of wavebands B are constant for all traffic sets. The goal was to minimize, over all possible admissible traffic sets, the two cost parameters corresponding to the number of wavebands B and the number of wavelengths $\sum_{i=1}^{B} b_i$.

In the star, each permutation T_i could be accommodated using a single waveband consisting of b_i wavelengths. This approach can be extended to other topologies in a straightforward manner, with the main difference being that each permutation T_i may now require multiple wavebands of size b_i to support it. In general, if the RWA algorithm requires $\phi(N)$ wavelengths for permutation traffic on the topology, and a banding algorithm is considered which uses B wavebands and W wavelengths on a star, then the extension of that banding algorithm to the new topology requires:

$$\begin{aligned} B_{total} &= B \cdot \phi(N) \\ W_{total} &= W \cdot \phi(N) \end{aligned}$$

Routing algorithms for permutation traffic exist in the literature for rings with [8] and without [4] conversion, trees [13], and torus networks [7].

For example, in the case of a ring, [4] provides an optimal RWA algorithm for the bidirectional ring topology without conversion using the minimum number of wavelengths. Specifically, it shows that $\lceil N/3 \rceil$ wavelengths are necessary and sufficient to support any single-port traffic set. Since T_i is a permutation matrix, it can be supported using $\lceil N/3 \rceil$ wavelengths. Each set of calls $b_i T_i$ from the decomposition of (14), can therefore be supported using $\lceil N/3 \rceil$ wavebands, each consisting of b_i wavelengths.

Using this approach, the entire traffic set T can be supported using B sets of wavebands, where each such set i consists of $\lceil N/3 \rceil$ wavebands of b_i wavelengths. The total numbers of wavebands and wavelengths are

$$\begin{aligned} B_{total} &= B\left\lceil \frac{N}{3} \right\rceil \\ W_{total} &= \left\lceil \frac{N}{3} \right\rceil \sum_{i=1}^{B} b_i = W\left\lceil \frac{N}{3} \right\rceil \end{aligned}$$

VI. Conclusion

In this paper, we considered waveband switching as a method for reducing overall network costs. We provide wavelength-efficient algorithms that use the minimum possible number of wavebands, and show that the optimal approach is to use a greedy algorithm. We provided minimum-waveband algorithms that allow for small wavelength inefficiencies in return for reducing the number of wavebands down to just the nodal degree. We use these results to help characterize the optimal performance frontier. We also provided a uniform waveband approach that compares very favorably to the optimal performance frontier and achieves large reductions in switching requirements at very little cost in wavelength inefficiency. Finally, we extend our results to general topologies where permutation traffic routing algorithms are known.

Appendix A

The proof of Lemma 2 relies on bipartite graphs. Any $N \times N$ matrix C can be represented by a bipartite graph consisting of two sets of N nodes, $S = \{s_1, \ldots, s_N\}$ and $D = \{d_1, \ldots, d_N\}$. An edge exists between nodes s_i and d_j if the corresponding matrix entry $[C]_{i,j}$ is non-zero, and the edge is labelled with the value $[C]_{i,j}$.

Define a *maximal matching* to be a set of N edges such that exactly one edge is incident on each node. It is straightforward to show that any maximal matching corresponds to a permutation matrix. Furthermore, the number of such permutation matrices that exist within C is equal to the weight of the smallest edge in the matching. Therefore, to prove the lemma, it is sufficient to show that (1) for any admissible traffic set, there always exists a matching with minimum edge weight at least b_{max}, and (2) for any value of $b > b_{max}$, there exists some traffic set for which no maximal matching with minimum edge weight b exists.

We make use of Hall's Theorem in the proof. Hall's Theorem is stated as follows:

Hall's Theorem [14]*:* In a bipartite graph $(S, D, \mathcal{E})$, where $\mathcal{E}$ is the set of edges, define the neighborhood of a subset $v \in S$ to be those nodes in D which are connected via some edge in $\mathcal{E}$ to some node in v. Then there exists a maximal matching if and only if, for every subset $v \in S$, its neighborhood $N(v)$ has size $|N(v)| \geq |v|$.

Applying this to our problem, Hall's Theorem states that a maximal matching with minimum edge weight b exists if and only if no subset of source nodes v exists such that the b-neighborhood $N_b(v)$, consisting of destination nodes connected to some source node in v via an edge of at least weight b, has size less than $|v|$. Supposing for a moment that $|v| = n$ and $N_b(v) = m$, a maximal matching exists for a given traffic set iff $m \geq n$ for all possible v in that traffic set. Furthermore, since we require the waveband to be fully utilizable for *all* admissible traffic sets, we must have $m \geq n$ for all possible v for *all possible admissible traffic sets*.

We can consider, for any fixed n and m, when it is possible to find an admissible traffic set where $m < n$. Consider the maximal admissible traffic set where the maximum mP calls are sent from v to nodes in $N_b(v)$, leaving $(n-m)P$ calls being sent to the $N - m$ non-b-neighborhood nodes. If the $(n-m)P$ calls are divided evenly among the $n(N-m)$ edges going to non-b-neighborhood nodes, the largest edge weight is $\left\lceil \frac{(r-m)P}{r(N-m)} \right\rceil$. Since all edges to non-neighborhood nodes must be of weight less than b, this construction is valid whenever $b > \left\lceil \frac{(r-m)P}{r(N-m)} \right\rceil$. We have therefore shown that the condition $b \leq \left\lceil \frac{(r-m)P}{r(N-m)} \right\rceil$ is necessary.

It is possible to show that, for a fixed n and m, this is a necessary and sufficient constraint on b. We can therefore minimize this expression over all n and m and choose b_{max} to be equal to this value:

$$\begin{aligned} b_{max} &= \min_n \left\{ \min_m \left\lceil \frac{(n-m)P}{n(N-m)} \right\rceil \right\} \\ &= \min_n \left\lceil \frac{P}{n[N-(n-1)]} \right\rceil \\ &= \left\lceil \min_n \frac{P}{n[N-(n-1)]} \right\rceil \end{aligned}$$

where in the last line, the minimization can be brought inside due to the monotonicity of the floor function. Ignoring integrality constraints, the right-hand size is minimized at $n^* = \frac{N+1}{2}$. If N is odd, then $\frac{N+1}{2}$ is integer and a valid choice for n. If N is even, then since $\frac{P}{n[N-(n-1)]}$ is convex, the minimizing value must be either $\frac{N}{2}$ or $\frac{N}{2}+1$; both result in the same value of the minimizing function. This gives the final value of b_{max} from (6).

Appendix B

The proof of Lemma 3 will be by construction. We will first decompose the traffic matrix C into two sub-matrices: the "heavy" matrix C_H, containing all entries with weight greater than $\left\lceil \frac{P}{k+1} \right\rceil$, and the "light" matrix C_L, containing *no* entries greater than $\left\lceil \frac{P}{k+1} \right\rceil$.

We first assign any entry in C greater than $\left\lceil \frac{P}{k+1} \right\rceil$ to C_H. Note that at this point each row and column in C_H contains at most k entries. (If any row or column exceeds k entries, then that row or column in C must have had a sum greater than P, meaning C is not an admissible traffic set.) We next continue assigning entries in C to C_H until each row and column of C_H has exactly k entries.

Suppose there exists a row in C_H that contains fewer than k entries. Then there must also be a column that has fewer than k entries. (This follows from the fact that C_H is square; if all columns have k entries, and each row has no more than k entries, then all rows must also have k entries.) Locate the entry corresponding to that row and column in C, and assign it to C_H. Repeat until each row has k entries. By the same reasoning as before, all columns must now have k entries also. It is well known that any such matrix can be decomposed into at most k matrices with only one non-zero entry per row and column. Therefore, by performing this further decomposition and noting that all entries in C_H are at most P, we have shown that C_H can be supported by at most k wavebands of size P.

Assign all remaining entries in C to C_L; this gives C_L therefore has exactly $N - k$ entries per row and column. This can similarly be decomposed into $N - k$ matrices with only one entry per row and column; since each entry is at most $\left\lceil \frac{P}{k+1} \right\rceil$, we can support C_L using at most $N - k$ wavebands of size $\left\lceil \frac{P}{k+1} \right\rceil$.

References

[1] L. Li and A. K. Somani, "Dynamic wavelength routing using congestion and neighborhood information," *IEEE/ACM Trans. Networking*, vol. 7, pp. 779–786, October 1999.

[2] D. Banerjee and B. Mukherjee, "Wavelength-routed optical networks: linear formulation, resource budgeting tradeoffs, and a reconfiguration study," *IEEE/ACM Trans. Networking*, vol. 8, pp. 598–607, October 2000.

[3] R. Ramaswami and K. N. Sivarajan, "Routing and wavelength assignment in all-optical networks," *IEEE/ACM Trans. Networking*, vol. 3, pp. 489–500, October 1995.

[4] A. Narula-Tam, P. J. Lin, and E. Modiano, "Efficient routing and wavelength assignment for reconfigurable WDM networks," *IEEE J. Select. Areas Commun.*, vol. 20, pp. 75–88, January 2002.

[5] O. Gerstel, G. Sasaki, S. Kutten, and R. Ramaswami, "Worst-case analysis of dyanmic wavelength allocation in optical networks," *IEEE/ACM Trans. Networking*, vol. 7, pp. 833–845, December 1999.

[6] P. Saengudomlert, E. Modiano, and R. G. Gallager, "An on-line routing and wavelength assignment algorithm for dynamic traffic in a WDM bidirectional ring," in *Proc. JCIS*, March 2002.

[7] P. Saengudomlert, E. Modiano, and R. G. Gallager, "On-line routing and wavelength assignment for dynamic traffic in WDM ring and torus networks," in *Proc. IEEE INFOCOM*, April 2003.

[8] L.-W. Chen and E. Modiano, "Efficient routing and wavelength assignment for reconfigurable WDM networks with wavelength converters," in *Proc. IEEE INFOCOM*, April 2003.

[9] A. A. M. Saleh and J. M. Simmons, "Architectural principles of optical regional and metropolitan access networks," *J. Lightwave Technol.*, vol. 17, pp. 2431–2448, December 1999.

[10] M. Lee, J. Yu, Y. Kim, C. Kang, and J. Park, "Design of hierarchical crossconnect WDM networks employing a two-stage multiplexing scheme of waveband and wavelength," *IEEE J. Select. Areas Commun.*, vol. 20, pp. 166–171, January 2002.

[11] X. Cao, V. Anand, Y. Xiong, and C. Qiao, "A study of waveband switching with multilayer multigranular optical cross-connects," *IEEE J. Select. Areas Commun.*, vol. 21, pp. 1081–1095, September 2003.

[12] R. Izmailov, S. Ganguly, V. Kleptsyn, and A. C. Varsou, "Non-uniform waveband hierarchy in hybrid optical networks," in *Proc. IEEE INFOCOM*, April 2003.

[13] P. Saengudomlert, E. Modiano, and R. G. Gallager, "Dynamic wavelength assignment for WDM all-optical tree networks," in *Proc. Allerton*, September 2003.

[14] M. N. S. Swamy and K. Thulasiraman, *Graphs, Networks, and Algorithms*, Wiley, New York, 1981.