Optimal Waveband Switching in WDM Networks

Li-Wei Chen, Poompat Saengudomlert, Eytan Modiano

Abstract—Switching traffic together in bundles of wavelengths called *wavebands* can greatly reduce the switching costs of the network. We consider the problem of partitioning wavelengths into wavebands for star networks using the minimum number of total wavebands. We provide a greedy algorithm for waveband partitioning and show that it is optimal in that it requires the minimum number of wavebands subject to using the minimum possible number of wavelengths. We also give an algorithm for allocating calls from any admissible traffic set to the wavebands in a non-blocking manner. Finally, we show that the increase in the number of wavebands required is logarithmic in the number of calls and polynomial in the size of the network.
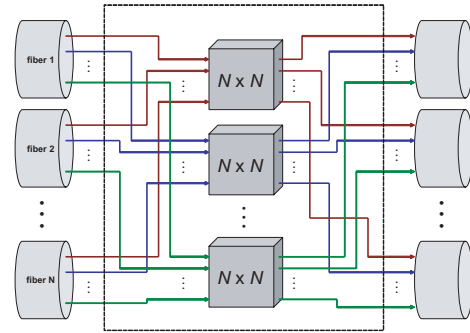
Fig. 1. Switching requirements for a node with degree N. The node has N input fibers and N output fibers, and needs to be able to switch any of the W wavelengths on each input to any output. This requires W $N \times N$ switches.

I. INTRODUCTION

With traffic demands continuing to increase rapidly each year, wavelength-division multiplexing (WDM) has emerged as an attractive solution for increasing capacity in optical networks. WDM allows multiple data streams to be carried using the same fiber link, as long as each data stream occupies a different wavelength [1].

Much of the work currently in the RWA literature considers routing at the wavelength level ([2], [3], [4], [5], [6], [7], [8], [9]). As the number of wavelengths increases, the cost of accessing and managing all these wavelengths at each node grows rapidly, making this approach difficult to scale. In particular, each node needs to be able to switch each wavelength independently, increasing the switching complexity rapidly as the size of the network grows: for a node with degree N and W incoming wavelengths per fiber, W $N \times N$ switches are required, as illustrated in Figure 1. This has led to the idea of banding, where wavelengths are grouped together into wavebands that are then routed together, allowing RWA to be performed at the waveband level [10]; ideally the number of wavebands is much smaller, and processing can be done at this more coarse level, reducing costs. If switching can be performed on only B wavebands, then only B $N \times N$ switches are required.

There has been some work addressing the question of how best to partition the available wavelengths into wavebands. In [11], the authors investigate the case where a single source node initiates calls to multiple destinations, and all calls share the same first link while the remainder of the path to each destination is disjoint. The problem is to partition the wavelengths on the first link such that switching at the end of that link need only be done at the waveband level. They provide a partitioning algorithm that optimally minimizes the number of wavebands needed subject to using the minimum number of wavelengths.

We consider the problem of multiple source and destination nodes in a star topology. We examine the class of wavelength-efficient banding algorithms that use no more than the minimum number of wavelengths required, and provide an optimal banding algorithm that minimizes the total number of wavebands required subject to this minimum-wavelength constraint. We also derive the number of wavebands required as a function of the number of nodes and the number of calls, and analyze the increase in the requirements as a function of the two parameters.

A. System Model

We consider a network consisting of N nodes where each node is connected via a bidirectional fiber link to a single hub. We call this a *star* topology. Note that here the term "star" refers only to the physical link connectivity of the network, and does not imply a broadcast medium or all-to-all traffic. We use the rearrangeably non-blocking P-port model from [5], where each of the N nodes has P ports and is therefore allowed to send and receive at most P calls. The hub node does not send or receive any calls. Any traffic set which satisfies these conditions is termed *admissible*; at any instant in time, the calls in the network could consist of any one of the admissible traffic sets. Call arrivals and departures are modelled as transitions between different admissible sets, capturing the dynamic changes in traffic without making any statistical assumptions about the calls. Sufficient resources must be provisioned in the network so that all calls in any admissible set can be supported. The network thus dynamically supports all admissible traffic sets. The challenge is to design good algorithms that minimize the total cost of the resources that need to be provisioned.

An admissible set is *maximal* if each node sends and receives the maximum number of calls permitted. For simplicity, we allow self-traffic in our analysis; the case where self-traffic is disallowed leads to similar results. Without loss of generality, we

can consider only maximal traffic, since for any non-maximal set, we can insert fictitious calls to obtain a maximal set. These sets are the cases of interest since they send the maximum possible traffic into the network.

[8] showed that for such a P-port N-node topology with no wavebanding (or, alternatively, where all wavebands are of size 1), P wavelengths are both necessary and sufficient to support any maximal traffic set. Therefore a minimum of P wavelengths are also required by any waveband partitioning scheme. We say that a waveband can be *fully utilized* if it is possible to assign enough calls to that waveband such that every wavelength in the waveband is utilized on every link in the network. Note that for a maximal traffic set, if the minimum number of wavelengths are to be used, then every waveband must be fully utilized; otherwise, there will exist a link that does not have sufficiently many wavelengths. This follows from the fact that maximal traffic sets have P calls on each link, and only a total of P wavelengths are available.

In Section II, we ask the question of how large the largest waveband can be before we cannot guarantee that it can be fully utilized by any admissible traffic set. This will be useful in Section III, where we show that once we have an expression for the largest waveband size that can be fully utilized, a greedy algorithm can be used in conjunction with this expression to recursively partition the wavelengths into wavebands. The greedy approach will be shown to be optimal in minimizing the resulting number of wavebands subject to using the minimum number P of total wavelengths. Section IV describes how to assign calls from a traffic set to wavebands given a waveband partition, and provides an upper bound on the number of wavebands required by the greedy algorithm. The bound will show that this number grows only logarithmically with the total number of wavelengths P.

II. MAXIMAL WAVEBAND SIZING

Recall that for maximal traffic sets, every waveband must be fully utilized if we are to use only the minimum number of wavelengths. In this section, we consider the problem of determining $b_{max}(N, P)$, the largest waveband that can be fully utilized by all maximal traffic sets, for general P-port traffic in an N-node star. The approach we will use is to show an equivalence between a maximally utilized waveband and a maximal bipartite graph matching, and then rely on graph theoretic results to provide necessary and sufficient conditions for obtaining maximal matchings.

We can represent any traffic set on a star topology by a bipartite graph. The graph has two sets of N nodes each, denoted by V_1, V_2. Nodes in V_1 represent sources of calls, while nodes in V_2 represent destinations. An edge exists between a node $s_i \in V_1$ and a node $d_j \in V_2$ if there is a call from node s_i to d_j in the star. The edge is given a weight based on the number of calls between that source-destination pair. Denote the set of all edges by $\mathcal{E}$.

Define a maximal matching to be a set of edges such that exactly one edge is incident on every node in the bipartite graph. All calls in a maximal matching can use the same waveband, since every call leaving a given source node goes to the same destination, and vice versa. Calls from the maximal matching
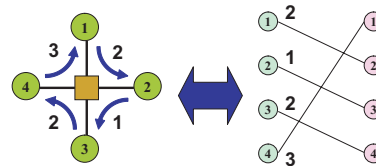

Fig. 2. A 4-node star with a total of 8 calls, as denoted by the weighted arrows. The calls correspond to a maximal matching on the bipartite graph, and could all fit on a single waveband. The waveband could consist of at most a single wavelength, since more calls from node 1 to 3 cannot be found to fully utilize a larger waveband.
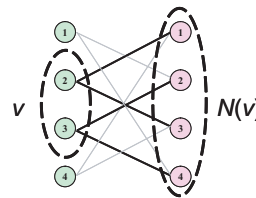

Fig. 3. To apply the test given by Hall's Theorem, a subset v of the source nodes V_1 is first chosen. In this case, v consists of 2 nodes, and the neighborhood $N(v)$ contains 4 nodes. Therefore the test is passed for this choice of v. This test must be repeated for all possible choices of v.

can be assigned to fully utilize a waveband of size at most equal to the smallest edge weight in the matching. (If all edges are not of the same weight, some calls in the matching will not be assigned, and must be allocated to subsequent wavebands.) This is illustrated in Figure 2. Therefore, the problem of finding $b_{max}(N, P)$ can be reduced to finding the largest waveband size for which we can still be guaranteed to find a maximal matching where the smallest edge is at least equal to that size.

Conveniently, a theorem exists which provides necessary and sufficient conditions for the existence of maximal matchings.

Hall's Theorem [12]*:* In a bipartite graph $(V_1, V_2, \mathcal{E})$, define the neighborhood of a subset $v \in V_1$ to be those nodes in V_2 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 $v \in V_1$, its neighborhood $N(v)$ has size $|N(v)| \geq |v|$.

Hall's Theorem therefore provides the basis for determining the existence of maximal matchings. The following test is applied: a subset v of the source nodes V_1 is chosen. If the neighborhood of the subset $N(v)$ is of size greater than or equal to the size of the subset itself, the test is passed. This is shown in Figure 3. The test is then repeated for all possible subsets v of V_1. If the test is passed for all subsets, then a maximal matching exists. If at least one test is failed, then no maximal matching exists.

We can determine if a waveband of a given size b can be fully utilized by a given traffic set as follows. Determine the bipartite graph corresponding to the traffic set, and delete any edges with weight less than b. (This guarantees that any maximal matching found will have minimum edge weight at least b.) The tests given by Hall's Theorem can then be applied to this graph to determine if a maximal matching exists that is sufficiently large to fully utilize the waveband. If the test fails, then a waveband of size b is too large to be sufficiently utilized. This test should be applied to all maximal traffic sets to guarantee that *any* maximal set can fully utilize the waveband.
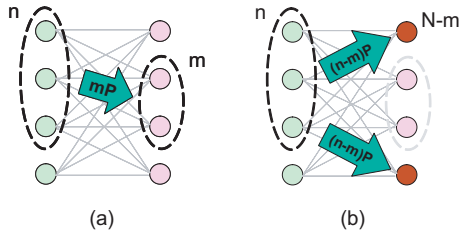
Fig. 4. At most mP calls can be sent to nodes in $N(v)$, and hence $(n-m)P$ calls must go to non-neighborhood nodes.

In principle, the preceding approach could be used to determine $b_{max}(N, P)$ numerically by brute force. However, we will see that a closed-form solution can be obtained. The method for obtaining the closed-form expression for $b_{max}(N, P)$ relies on attempting to construct a bipartite graph which causes the test given by Hall's Theorem to be failed. (In a slight abuse of notation, we call such a bipartite matching a "counterexample".) $b_{max}(N, P)$ is then the largest waveband size for which no counterexample exists.

In order for the test to be failed, a maximal traffic set must be found for which we can choose a v such that the size of the neighborhood $N(v)$ is smaller than the size of v. We therefore wish to construct a counterexample where, if $|v| = n$, then $|N(v)| = m$, where $m < n$.

Under the P-port model, the nodes in v can send at most mP calls to nodes in $N(v)$. The remaining *residual traffic* is therefore at least $(n - m)P$. These calls are sent to nodes outside the neighborhood, and hence must belong to edges adjacent to a non-neighborhood node. Call these edges *non-neighborhood edges*. Non-neighborhood edges have weight less than $b_{max}(N, P)$ and are removed from the graph before the search for maximal matchings begins, since they do not contain enough calls to fully utilize the waveband. This is illustrated in Figure 4. There are at most $n \cdot (N - m)$ non-neighborhood edges.

We can therefore construct a counterexample if and only if the residual traffic can be divided among the non-neighborhood edges such that no non-neighborhood edge has weight greater than or equal to $b_{max}(N, P)$. Since there are $(n-m)P$ residual calls and $n(N - m)$ non-neighborhood edges, there is at least one non-neighborhood edge with weight at least

$$\left\lceil \frac{(n-m)P}{n(N-m)} \right\rceil$$

If $b_{max}(N, P)$ is chosen to be at most this number, then no counterexample exists for the given values of n and m. We can choose $b_{max}(N, P)$ to guarantee that *no* counterexample exists by minimizing over n and m and choosing $b_{max}(N, P)$ at most this minimum:

$$
b_{max} \leq \min_{n,m} \left\lceil \frac{(n-m)P}{n(N-m)} \right\rceil
= \min_{n,m} \left\lceil \frac{\left(1 - \frac{m}{n}\right)P}{N-m} \right\rceil \quad (1)
$$

We fix for the moment m and consider the minimization over n. Since $n > m$, the minimization is subject to the constraint

$0 < m < n < N$. Equation 1 is minimized by choosing n as small as possible. Since n and m are both integer quantities, we should choose $n = m + 1$; conversely, $m = n - 1$.

Making this substitution, the minimization becomes:

$$
b_{max} \leq \min_n \left\lceil \frac{P}{n[N-(n-1)]} \right\rceil \quad (2)
$$

Since the ceiling function is monotonic,

$$
\min_n \left\lceil \frac{P}{n[N-(n-1)]} \right\rceil = \left\lceil \min_n \frac{P}{n[N-(n-1)]} \right\rceil
$$

Ignoring integrality constraints, the right-hand size is easily shown to be minimized at $n^* = \frac{N+1}{2}$. If N is odd, then $\frac{N+1}{2}$ is an integer and hence is a valid choice for n^*. We subsequently obtain a value for b_{max} of

$$
b_{max} = \left\lceil \frac{P}{\left(\frac{N+1}{2}\right)\left(\frac{N+1}{2}\right)} \right\rceil
$$
$$
= \left\lceil \frac{4P}{(N+1)^2} \right\rceil \quad , \text{ if } N \text{ odd}
$$

If N is even, then since $\frac{P/n}{N-(n-1)}$ is convex, the minimizing value of n must be one of the integers adjacent to $\frac{N+1}{2}$, namely either $\lfloor \frac{N+1}{2} \rfloor = \frac{N}{2}$ or $\lceil \frac{N+1}{2} \rceil = \frac{N}{2} + 1$. It is easy to verify that either case results in the same value of

$$
b_{max} = \left\lceil \frac{P}{\left(\frac{N}{2}\right)\left(\frac{N}{2}+1\right)} \right\rceil
$$
$$
= \left\lceil \frac{4P}{N(N+2)} \right\rceil \quad , \text{ if } N \text{ even}
$$

In summary,

$$
b_{max} = \begin{cases} \left\lceil \frac{4P}{N(N+2)} \right\rceil & , N \text{ even} \\[2mm] \left\lceil \frac{4P}{(N+1)^2} \right\rceil & , N \text{ odd} \end{cases} \quad (3)
$$

III. Optimal Waveband Partitioning

The goal of waveband partitioning is to assign each wavelength to a waveband such that (1) for any admissible traffic set, all wavelengths in a given waveband can be routed together without being demultiplexed, and (2) the total number of wavebands (subject to the minimum-wavelength constraint) is minimized. We will show in this section that the optimal method of accomplishing these goals is to use a greedy algorithm.

The greedy algorithm for waveband assignment looks at the P wavelengths required to support P-port traffic and determines the largest waveband size that can always be fully utilized by any admissible traffic set. It then creates a waveband of this largest size, and repeats this process recursively using the remaining wavelengths.

Suppose the size of the largest waveband is b_{max}. Since the largest waveband can always be allocated without waste, it handles b_{max} calls to and from each node. The remaining traffic
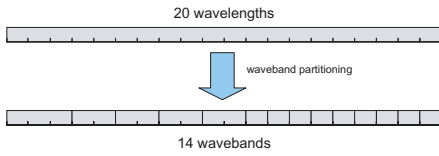
20 wavelengths

waveband partitioning

14 wavebands

Fig. 5.   The waveband partitioning for a 20-port 5-node star.

therefore forms a $(P - b_{max})$-port traffic set. The greedy algorithm determines the size of the largest waveband for this residual traffic set by repeating the above process, proceeding recursively until all wavelengths have been partitioned into wavebands.

A formal statement of the algorithm is as follows:

1) Initialize P to be the number of ports and N to be the number of nodes.
2) Using (3), determine b_{max}, the maximum size of waveband for a P-port N-node star such that no wavelengths in the band are wasted regardless of the traffic set.
3) Let $P = P - b_{max}$.
4) If $P > 0$ go to 1.

Note that the number of wavebands $B_{N,G}(P)$ used by the greedy algorithm is nondecreasing in P; that is, $B_{N,G}(P_1) \leq B_{N,G}(P_2)$ if $P_1 < P_2$. This property will prove useful later.

Example 1: Consider a 5-node star with $P = 20$. Using the greedy algorithm, we would determine that the largest waveband should be $\left\lceil \frac{4P}{(N+1)^2} \right\rceil = \left\lceil \frac{(4)(20)}{(6)^2} \right\rceil = 3$.

After this step, 17 wavelengths remain to be partitioned. Repeating, we determine the next-largest waveband to be $\left\lceil \frac{(4)(17)}{(6)^2} \right\rceil = 2$. This leaves 15 wavelengths to be partitioned.

Iterating through this procedure produces a waveband partitioning of $\{3, 2, 2, 2, 2, 1, 1, 1, 1, 1, 1, 1, 1, 1\}$. Notice that there are $\frac{(N+1)^2}{4} = 9$ wavebands consisting of a single wavelength. There are always $\frac{(N+1)^2}{4}$ single-wavelength bands for large P since $b_{max} = 1$ for $P \leq \frac{(N+1)^2}{4}$. The final waveband partitioning is shown in Figure 5.

We will next show that the approach of the greedy algorithm is optimal, in that it results in the creation of the minimal number of wavebands subject to the minimum-wavelength constraint. We do so via an induction proof. For a P-port N-node star, let the minimum number of wavebands required be denoted by $B_N^*(P)$.

Step 1: Base Case

Consider the case of $P = 1$. Then there is only a single wavelength, and the maximum band size b_{max} is trivially equal to 1. Therefore the greedy algorithm, which would produce a single band of size 1 in this case, is optimal.

Step 2: Induction Step

In the induction step, we consider a P_i-port star and assume that for any value of $P \leq P_i - 1$, we are given that the greedy algorithm is optimal. We then wish to prove that the greedy algorithm is optimal for $P = P_i$.

Let the number of wavelengths assigned to the first waveband by the greedy algorithm be b_{max}, and let the number of wavelengths assigned by any other algorithm be b. Then we can express the total number of wavebands created by the two

algorithms as $B_{N,G}(P_i)$ and $B_{N,O}(P_i)$, respectively. The first quantity can be written as:

$$
\begin{aligned}
B_{N,G}(P_i) &= 1 + B_{N,G}(P_i - b_{max}) \\
&= 1 + B_N^*(P_i - b_{max}) \quad (4)
\end{aligned}
$$

with the second equality resulting because by the induction hypothesis the greedy algorithm is optimal for $P \leq P_i - 1$. The second quantity is:

$$
\begin{aligned}
B_{N,O}(P_i) &= 1 + B_{N,O}(P_i - b) \\
&\geq 1 + B_N^*(P_i - b) \quad (5)
\end{aligned}
$$

where the second inequality results from the fact that B_N^* is optimal.

By the induction hypothesis, since the greedy algorithm is optimal for $P \leq P_i - 1$, and the greedy algorithm is non-decreasing in P, $B_N^*(P)$ is also non-decreasing for $P \leq P_i - 1$. Comparing equations 4 and 5 and noting that $b_{max} \geq b$, we conclude that since $B_N^*(P)$ is non-decreasing, $B_{N,G}(P_i) \leq B_{N,O}(P_i)$ and hence the greedy algorithm is optimal for $P = P_i$ as well, concluding the proof.

IV. RWA USING WAVEBANDS

A. Waveband RWA Algorithm

The partitioning of wavelengths into wavebands is a process that is done only once for any given network with fixed P and N. The wavebands are reused for any admissible traffic set simply by reconfiguring the switch at the hub. What now remains is the problem of wavelength assignment and the subsequent switching of the wavebands for each admissible traffic set. Once this problem is solved, when there is a new call arrival, the wavelength assignment and switch settings are simply reconfigured to support the resulting new traffic set.

We can use the same approach as in the development and proof of the greedy algorithm for wavelength assignment. Given the waveband partition, we choose the largest waveband and find a set of calls that fully utilize that waveband. We are guaranteed to find such a set by our choice of b_{max}. Removing those calls from consideration, we proceed to the next-largest waveband, and repeat recursively until all calls are assigned.

Formally, given the waveband partition, the procedure for assigning calls from any admissible traffic set to wavelengths in those bands is as follows:

1) *Obtain a set of calls which fully utilizes the largest free waveband.* This can be accomplished by drawing the bipartite graph corresponding to all traffic carrying more than b_{max} calls. From Section II, we know b_{max} is sufficiently small that a maximal matching with minimum edge weight at least b_{max} can always be found. Assign calls from that matching to fully utilize the waveband.
2) *Remove assigned calls, and repeat.* Remove all calls which have been assigned wavelengths from the traffic set. If the traffic set is now empty, stop. Otherwise, return to Step 1.

We emphasize that the problem of partitioning the wavebands is disjoint from the wavelength assignment for calls in the traffic set. The waveband partition is **fixed** for a given P and N, and any admissible traffic set uses the same waveband partition; only the wavelength assignment and switching changes.

B. Upper Bound on Number of Wavebands

Having obtained an optimal algorithm for allocating wavebands, it is natural to ask how many wavebands are required for given values of N and P. The question is relevant because if routing is done at the waveband level, then the number of wavebands determines the switching cost. In principle, the number of wavebands could be determined numerically by iterating through the greedy algorithm (a procedure that is not computationally difficult) for each set of N and P. In this section, an alternate approach based on relaxing the integrality constraints associated with b_{max} is used to obtain a closed-form upper bound on the optimal number of wavebands.

Consider the case of choosing the largest waveband to be $b = \frac{4P}{(N+1)^2} \leq b_{max}$, where the integrality constraints on b have been relaxed. We then use this to obtain an upper bound on the value of the largest waveband in the greedy algorithm, and track the value of P_k through each iteration. Let P_k be the value of P after running the k^{th} iteration of the greedy algorithm. It can be shown that the series progressed according to the relation:

$$P_k = \left[1 - \frac{4}{(N+1)^2}\right]^k \cdot P$$

If $P \leq \frac{(N+1)^2}{4}$, then the number of bands B is simply equal to $\frac{(N+1)^2}{4}$ since each band is composed of only a single wavelength. Therefore consider $P > \frac{(N+1)^2}{4}$ and determine the number of bands k required to reduce the number of unassigned wavelengths to $\frac{(N+1)^2}{4}$. Then the total number of wavebands would be $k + \frac{(N+1)^2}{4}$. We first solve for k as follows:

$$P_k = \frac{(N+1)^2}{4}$$

$$\left[1 - \frac{4}{(N+1)^2}\right]^k \cdot P = \frac{(N+1)^2}{4}$$

$$k = \frac{\log\left[\frac{(N+1)^2}{4P}\right]}{\log\left[1 - \frac{4}{(N+1)^2}\right]}$$

This gives an upper bound on the number of wavebands B, namely:

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

To obtain a sense of how quickly the number of wavebands increases with P, we use the approximation that, for large N,

$$\frac{1}{\log\left[1 - \frac{4}{(N+1)^2}\right]} \approx -\frac{(N+1)^2}{4}$$

which gives us

$$B \approx \frac{(N+1)^2}{4}\left(1 + \log\left[\frac{4P}{(N+1)^2}\right]\right)$$
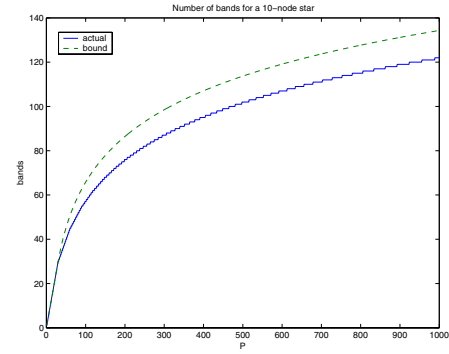

Fig. 6. A comparison of the actual number of wavebands required for a 10-node star to the upper bound.

From this, we can see that the number of wavebands grows with order $O(N^2 \log(P/N^2))$. Figure 6 plots the bound compared to the exact number of wavebands required.

V. CONCLUSIONS

We considered the problem of partitioning wavelengths into wavebands for star networks. We provide a greedy algorithm for waveband partitioning and show that it is optimal in that it requires the minimum number of wavebands subject to using the minimum possible number of wavelengths. We also give an algorithm for allocating calls from any admissible traffic set to the wavebands in a non-blocking manner. Finally, we show that number of wavebands required grows as $O(N^2 \log(P/N^2))$, where P is the number of ports per node and N is the number of nodes.

REFERENCES

[1] R. Ramaswami and K. N. Sivarajan, *Optical Networks: A Practical Perspective*. Morgan Kaufmann, 1998.
[2] ——, "Routing and wavelength assignment in all-optical networks," *IEEE/ACM Trans. Networking*, vol. 3, pp. 489–500, October 1995.
[3] 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.
[4] 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.
[5] 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.
[6] 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.
[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] ——, "Dynamic wavelength assignment for wdm all-optical tree networks," in *Proc. Allerton*, September 2003.
[9] L.-W. Chen and E. Modiano, "Efficient routing and wavelength assignment for reconfigurable WDM networks with wavelength converters," in *Proc. IEEE INFOCOM*, April 2003.
[10] X. Cao, V. Anand, Y. Xiong, and C. Qiao, "Performance evaluation of wavelength band switching in multi-fiber all-optical networks," in *Proc. IEEE INFOCOM*, April 2003.
[11] R. Izmailov, S. Ganguly, V. Kleptsyn, and A. C. Varsou, "Non-uniform waveband hierarchy in hybrid optical networks," in *Proc. IEEE INFOCOM*, April 2003.
[12] M. N. S. Swamy and K. Thulasiraman, *Graphs, Networks, and Algorithms*. New York: Wiley, 1981.