Wavelength Assignment in WDM Rings to Minimize System Cost instead of Number of Wavelengths

Ori Gerstel
Tellabs Operations Inc., Optical Network Group, Hawthorne, NY, ori@tellabs.com

Philip Lin
Tellabs Operations Inc, Tellabs Research Center, Cambridge, MA, pjlin@trc.tellabs.com

Galen Sasaki
Electrical Engineering Department, University of Hawaii, Honolulu, HI, sasaki@spectra.eng.hawaii.edu

Abstract

This work deals with wavelength assignment for lightpaths. We restrict our attention to the near term, in which WDM networks will be in the form of rings and higher level networks will be SONET/SDH self-healing rings. This view changes the goal of wavelength assignment (WLA) vs. previous work on the subject in a number of aspects. First, a pair of SONET add/drop multiplexers (ADMs) terminates each lightpath. These ADMs also terminate adjacent lightpaths to form rings, implying that the WLA has to support this type of sharing. Second, we argue that unless the wavelength limit is exceeded, the first-order optimization goal should be to minimize the overall network cost which is dominated by the number of required ADMs and not the number of wavelengths. We show that these two minimization problems are intrinsically different, and there exist cases where the two minima cannot be simultaneously achieved. We derive a simple lower bound to the number of ADMs. Depending on the given lightpaths, we show that this lower bound is not always achievable. We also show that adding wavelength converters to the system does not improve the cost but that splitting a lightpath and handling each part separately may reduce the total number of ADMs. We develop two heuristics to minimize the number of ADMs: Cut-First, and Assign-First. Both heuristics, using different approaches, attempt to use the smallest number of ADMs possible. We also show that Cut-First always uses the minimum number of wavelengths, but because of the cut, may use more ADMs than necessary. However, the number of extra ADMs is proven to be bounded by the number of supported wavelengths and typically much less. We show instances where Cut-First performs better than Assign-First and vice versa. Finally, we present a set of transformations that take any WLA and improve its cost.

Keywords: Wavelength assignment, wavelength division multiplexing (WDM), optical networks, SONET/SDH, self-healing rings.

1. Introduction

1.1. Background

WDM ring networks are being developed as part of test-beds and commercial products, and will soon be deployed by a growing number of telecom carriers. As these products become real, it is time to revisit one of the most fundamental network design problems for WDM networks that attracted numerous studies – the assignment of wavelengths to a given set of lightpaths [T75, BH96, RS94, KA96, CGK93, RU94, GSR96].

While most of the previous work attempts to minimize the number of wavelengths required or the amount of blocking for the given set of lightpaths, we claim that this is not always the desirable goal due to the following reasons. First, in the near future, many of the WDM systems will be under-deployed since the amount of needed bandwidth will not easily scale up by a factor of 32 (which is the scaling factor current WDM systems provide). Second, a more important goal for many of the customers will be to minimize the entire cost of their system – not just the WDM component but also the TDM line-terminal equipment attached to the optical nodes. In fact, the dominant cost of a system running OC-48 traffic is the TDM component and not the optical one. Third, a factor that was neglected in most of the previous work is the fact that the higher level network cannot necessarily support arbitrary topologies. For example, SONET/SDH self-healing rings, which are the predominant layer above the optical layer, have to be configured in rings for protection purposes.

Minimizing the cost of the entire system is more often than not a hard task, which does not lend itself to analysis. It is therefore reasonable to split it into two parts:

1) Given smaller traffic demands (e.g., T3 or OC-3) concentrate them to form a minimal number of higher speed (e.g., OC-48) lightpaths, and
2) Given the set of lightpaths, route and assign wavelengths to them so that the extra cost is minimized.

The first part tries to minimize the cost of the necessary higher-level line-terminals and is studied in [GRS98]. The second part attempts to minimize the additional overhead resulting from the need of these lightpaths to be configured as rings and is the focus of the current paper.

1.2. Problem statement

Given a set of full-duplex lightpaths, $T=\{(s_i, t_i)\}$, and a wavelength division multiplexed (WDM) ring network comprising optical nodes, the wavelength assignment problem assigns a wavelength to each lightpath. There are N optical nodes on the ring, labeled from 0 to N-1. The forward part of the duplex lightpath (s_i, t_i) traverses downstream from s_i to t_i, and the reverse part traverses upstream from t_i to s_i. For simplicity, call s_i the origin node and t_i the termination node.

The underlying assumptions for this paper are:

- Routing is a hard problem by itself, even for rings [GJMP80, FNSST92]. Therefore we assume here (as most of the earlier work [T75, BH96, KA96, CGK93, RU94, GSR96]) that the lightpaths are already routed, say clockwise on the ring, between s_i and t_i.
- We focus on the static WLA problem here, which the set of lightpaths is given upfront. We expect this to be the main mode of operation for WDM, as is the case for very high-speed pipes in the telecom environment.
- Lightpaths may be split into several segments and the data transferred electrically between the segments. This is a valid option in the SONET case, although it involves higher port and cross-connect costs that are ignored here for sake of simplicity.

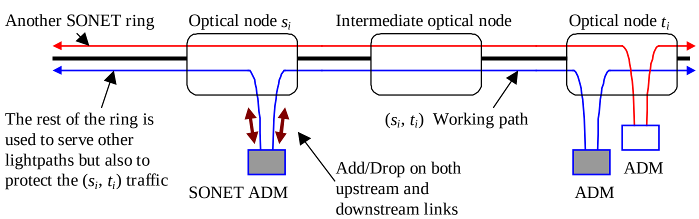

Figure 1. ADMs create protected lightpaths between nodes via wavelength channels

Electrical TDM line-terminals terminate the lightpaths. These nodes are assumed in this paper to be SONET/SDH add/drop multiplexers (ADMs). Figure 1 shows part of a ring configuration comprising 3 optical nodes and 3 ADMs configured in two SONET rings. For clarity, the wavelength channels are separated in the figure although they use a common fiber, and each channel represents wavelength paths traversing in both the upstream and the downstream directions.

Each lightpath uses two ADMs, one at the origin, the other at the termination. Although the origin node needs only the downstream ADM function and the termination node needs only the upstream ADM function, full ADMs will be installed on both nodes to complete the protection path around the ring. The full configuration would result in a number of SONET rings all circumventing a single optical ring. It is easier to think of each wavelength around the ring as providing the connectivity for a single SONET ring and we shall prove it suffices to consider such cases.

It follows that if two adjacent lightpaths are assigned the same wavelength, then they are used by the same SONET ring and the ADM at the common node can be shared by them. Figure 2 depicts two different wavelength assignments for the same set of lightpaths. For clarity, only the working paths are drawn and the paths that serve for protection purposes only are omitted. Each segment represents the working path of a duplex lightpath. Assignment A requires 6 ADMs. In assignment B, lightpaths (0,2) and (2,4) are on the same wavelength and share the same ADM. Hence, Assignment B requires only 5 ADMs total. Assignment B is preferred because ADMs are the costly part of the system.

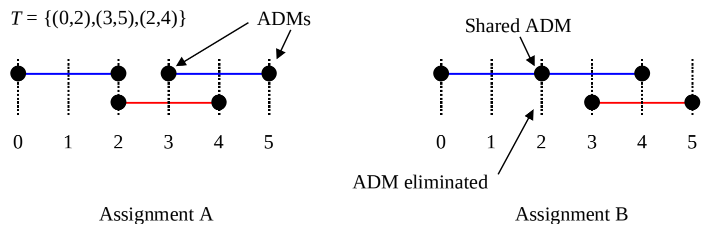

Figure 2. Two wavelength assignments.

The focus of this paper is to assign wavelengths channels to a given set of lightpaths such that the number of ADMs used, α, is minimized. It is important to note that this fact does not necessarily imply a wasteful use of wavelengths. In fact the Cut-First algorithm described below achieves a low number of ADMs while minimizing the number of used wavelengths.

1.3. Summary of results

This paper contains the following results:

1) The wavelength assignment problem with a new goal of minimizing the cost (which is mostly affected by the number of ADMs).
2) Example proving that ADM minimization is different than wavelength minimization and that there exists a tradeoff between the two.
3) A lower bound on the number of ADMs.
4) A worst case example showing that the lower bound may not always be achievable.
5) A lightpath splitting method as a way to reduce the number of ADMs.
6) Two wavelength assignment heuristics and bounds on their performance.
7) Three transformations to further reduce the number of ADMs given any wavelength assignment scheme.

2. Minimum number of wavelengths vs. minimum number of ADMs

The conventional objective in wavelength assignment is to minimize the total number of wavelengths, W. A natural conjecture is that wavelength minimization algorithms may also simultaneously minimize the number of ADMs. The conjecture is sensible because packing the lightpaths into a smaller number of wavelengths results in potentially more ADM sharing. The following example, however, shows that the two problems are intrinsically different from each other.

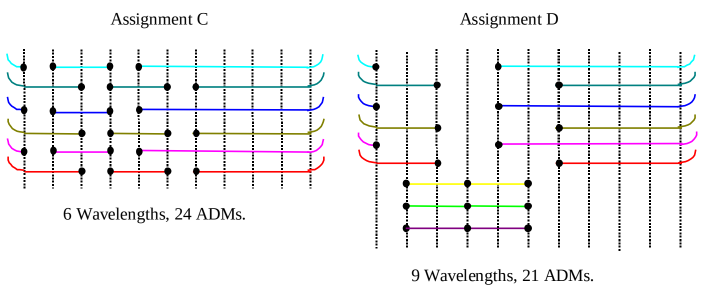

Figure 3. Example to show that wavelength minimization and ADM minimization cannot be simultaneously achieved. The curving lines at the edges of each sketch signify that the lightpath wraps around the ring. The dashed lines symbolize the optical nodes and the dots -- the ADMs.

In the example, assignment C produces the minimum number of wavelengths, but use more ADMs than necessary. On the other hand, assignment D has the least number of ADMs, 13, but uses more wavelengths than C. Note that with this given set of lightpaths, both minima cannot be achieved simultaneously. Also, intermediate values for the number of wavelengths and ADMs are possible. The two examples below contain the same set of lightpaths as the above, but achieves different numbers of ADMs using different number of wavelengths.

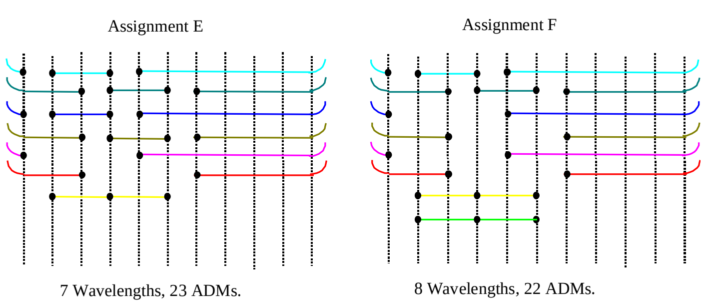

Figure 4. Various assignments that tradeoff between number of ADMs and wavelengths.

The assignments C--F are based on 3 "stripes" of the first upper 4 lightpaths in Assignment C. The lightpaths in each stripe can be assigned 2 wavelengths (configuration I), in which case they require 8 ADMs or 3 wavelengths (configuration II), in which case 7 ADMs suffice for them. By choosing different configurations for different stripes, a tradeoff can be achieved.

The following figure graphically depicts this tradeoff between ADMs and wavelengths. The same technique can be generalized to achieve k tradeoff points instead of just 4. The tradeoff curve is confined to the plane above the lower bounds to the number of wavelength and that of the number of ADMs. These lower bounds may or may not be achievable as will be discussed in a later section.

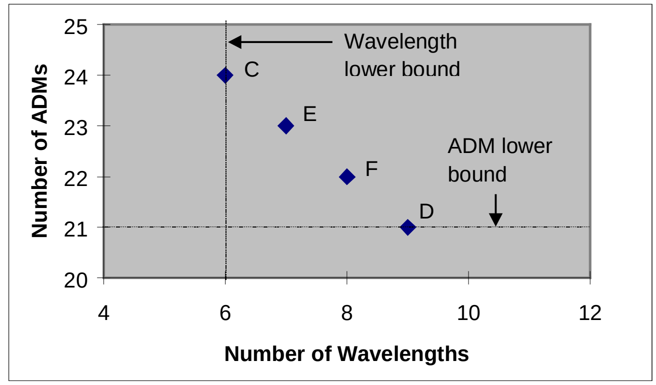

Figure 5. Tradeoff between number of wavelengths and number of ADMs.

The optimal assignment depends on the relative cost of a wavelength and of an ADM. In this study, we assume that the wavelength costs are negligible comparing to the ADM costs.

2.1. Wavelength changing does not help

We argue that, without loss of generality, each self-healing ring can be restricted to operate on a single wavelength. Contrast this property with the case of wavelength minimization, in which wavelength-changing clearly makes a difference. Although wavelength-changers (or ADMs operating on multiple wavelengths) can be used to change the wavelength on portions of a ring, these wavelength-changers do not help in reducing the number of ADMs. To show this, consider an algorithm that utilizes wavelength changers and the assignment results in some multiple wavelength rings. Assume also that this assignment achieves the minimum number of ADMs. Due to protection requirements, one must be able to decompose the network into separate healing rings. One can then reassign the wavelengths such that each separate ring operates on a single wavelength. Therefore, wavelength-changers are unnecessary and can be eliminated, and the single-wavelength ring assumption is not restrictive.

3. Lower bound for number of ADMs

By construction, we have the following claim.

Claim 1:

Two lightpaths, i and j, can share an ADM if and only if the following two conditions are satisfied:

1. i and j are assigned the same wavelength,
2. i and j are adjacent lightpaths, i.e., $s_i = t_j$ or $s_j = t_i$.

Let σ_k be the total number of lightpaths with node k as the origin and τ_k the total number of lightpaths with node k as termination. Then the number of ADMs used at each node is given by,

$$\alpha = \sum_k \left(\sigma_k + \tau_k - \#\text{ of ADMs shared at } k\right). \qquad \textbf{(1)}$$

This is because each origin and termination requires one ADM except in the case of sharing. The number of sharing that can occur is bounded by min{σ_k,σ_k}. Hence, we have the following lower bound:

$$\alpha \geq \alpha_{lb} = \sum_k \left(\sigma_k + \tau_k - \min\{\sigma_k, \tau_k\}\right). \qquad \textbf{(2)}$$

Let M be the number of lightpaths in the configuration. With some algebraic manipulation, the lower bound can be re-written as:

$$\alpha_{lb} = M + \frac{1}{2}\sum_k \left|\sigma_k - \tau_k\right|. \qquad \textbf{(3)}$$

This second form of the lower bound has a nice interpretation: the number of ADMs required equals the number of lightpaths plus the number of “gaps” in the network. A gap is a contiguous span of a wavelength ring that does not contain working traffic. The total number of gaps equals half the number of unmatched ends and hence, our alternative lower bound.

In the lower bound calculation, we ignored condition 1) of claim 1. Therefore, there are cases where the lower bound cannot be achieved. Next, we show one such example.

4. Worst case example

In this section, we will show that even when $\sigma_k = \tau_k$ at each node k, potentially allowing for complete ADM sharing, no such sharing may be possible due to lightpaths that overlap each other. Hence, the total number of ADMs equals twice the total number of lightpaths.

Consider a ring network with N = 8 nodes. (The reader can easily generalize this example to the general case of 4n nodes where n is any positive integer). Number the nodes in the ring from 0 to 7 in the clockwise direction. Suppose there is a collection of $M = N = 8$ lightpaths defined as follows: T={(0,5), (5,2), (2,7), (7,4), (4,1), (1,6), (6,3), (3,0)}. Note that the load L = (N/2) + 1 = 5 on every link.

Each lightpath traverses 5 links, which is more than half way around the ring. Each lightpath will share one or more links with any other lightpath. Thus, all these lightpaths must be assigned different wavelengths. This means no ADMs can be shared. Each lightpath requires 2 ADM, and the total number of ADMs used is $\alpha = 2M = 16$.

Note that in this example, for every termination, s_i, there exist a termination, t_j, such that $s_i = t_j$. According to the lower bound in equation (2), $\alpha_{lb} = M$. Therefore, this example exceeds the lower bound by α-α_{lb} = M = 2(L-1). But this is the best we can do because the overlaps prevent us from any ADM sharing. This example shows that the lower bound cannot always be achieved.

There is a way to reduce the amount of ADMs in excess of the lower bound. This method is called *lightpath splitting* and we describe it next.

5. Lightpath splitting

The previous section demonstrated a worst case situation where the number of ADM exceeds the lower bound by about twice the load. This excess is unavoidable under any wavelength assignment. We can reduce this excess if some of the given lightpaths are allowed to split.

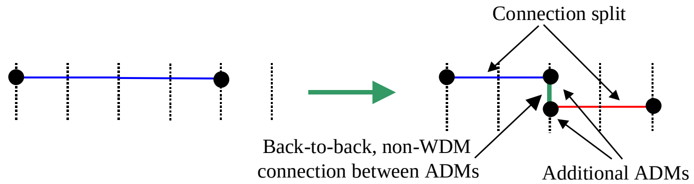

Figure 6. Splitting a lightpath.

When a lightpath is split, two new lightpath endpoints are created, and hence, two additional ADMs may be needed. This operation seems counter-productive because our goal is to minimize ADMs. However, when a lightpath is split into smaller parts, the ADMs of the two parts may be more readily shared with other existing lightpaths.

We demonstrate how splitting can produce more ADM sharing by using the same example as in last section. Take all lightpaths crossing node 0 (there are 4 of them), and split each lightpath into two parts with node 0 as the break point. Then, the resulting lightpaths can fit perfectly into W=L=5 wavelengths:

Given T = {(0,5), (5,2), (2,7), (7,4), (4,1), (1,6), (6,3), (3,0)}

Lightpaths after split (split lightpaths are bold faced)
T' = {(0,5), **(5,0), (0,2)**, (2,7), **(7,0), (0,4)**, **(4,0), (0,1)**, (1,6), **(6,0), (0,3)**, (3,0)}

Assignment:
Wavelength 1 = {(0,5), (5,0)}
Wavelength 2 = {(0,2), (2,7), (7,0)}
Wavelength 3 = {(0,4), (4,0)}
Wavelength 4 = {(0,1), (1,6), (6,0)}
Wavelength 5 = {(0,3), (3,0)}

All the lightpath terminations and origins are paired, and there were $M + L - 1$ lightpaths after splitting. So the number of ADMs is $N = M + L - 1$, an improvement of (L - 1) over the non-splitting method.
Therefore, splitting allows for smaller number of ADMs by increasing ADM sharing. Furthermore, splitting may be necessary in some cases because it achieves a lower number of ADMs than any algorithm.

6. Re-routing Lightpaths

Consider the following set of lightpaths: L lightpaths of (1,2), L lightpaths of (3,4), L lightpaths of (5,6) … etc. Assume both L, the load, and N, the number of nodes, are even. Then there are a total of $M = LN/2$ lightpaths. No ADM sharing is possible because no origin matches any termination. Furthermore, no splitting can be done because every lightpath spans only one link. Therefore, $2M = LN$ ADMs are required for this set of lightpaths. Note that L wavelengths are used in this case.

If the lightpaths can be re-routed, then the number of ADMs can be reduced. Re-route half of the lightpaths (1,2) so they become (2,1), and similarly half of (3,4) into (4,3), and so on. Then each pair of (x, y), (y, x) lightpaths can share and use only 2 ADMs. This way, every ADM is shared, and the total number of ADMs required is reduced by half to $M = LN/2$. However, this method is wavelength intensive. Every pair of (x, y) (y, x) lightpaths occupies a full wavelength. Therefore, the number of wavelength used is increased from L to $M/2 = LN/4$.

We mention re-routing as another possible method of reducing ADMs. However, this paper concentrates on wavelength assignment and not on routing. Therefore, we will not consider incorporating re-route in our algorithms.

7. Heuristics

7.1. *Cut-First heuristic*

The benefit of lightpath splitting leads to the following wavelength-assignment heuristics, which we term the Cut-First heuristics. This wavelength assignment algorithm guarantees that the number of ADMs $\alpha \le \alpha_{lb} + L - 1$. This heuristic follows closely to the cut-and-color heuristics described in [T75].

Let σ_i = the total number of lightpath origins at node i.
τ_i = the total number of lightpath terminations at node i.
x_i = the number of lightpath crossovers at node i.
L = the maximum load on a link.

Step 1: Choose a node, i, such that
$$i = \arg\min_i \{2x_i + \min\{\sigma_i, \tau_i\}, L - \max\{\sigma_i, \tau_i\}\}.$$
Label that node as node 0 and label the rest of the nodes clockwise in increasing order. Label also all wavelengths from 1 to W.

Step 2: For all the lightpaths crossing node 0, split the lightpaths into two parts, one on each side of node 0. Treat the two parts as independent lightpaths.

Step 3: Let $k = 0$

Step 4: Wavelength w is preferred if there exists a lightpath i such that $t_i = k$ and lightpath i is assigned wavelength w. For all lightpaths j such that $s_j = k$, assign lightpath j to the lowest numbered preferred wavelength. If no preferred wavelength is left, then assign j to the lowest numbered wavelength available.

Step 5: Increment k by 1.

Step 6: If $k < N$ then go to step 4.

Step 7: All lightpaths are now assigned. Check for possible merge of the lightpath splits done in step 2. Note that the two halves of a split lightpath can be merged only if they were assigned the same wavelength.

Like Cut-and-Color [T75], this algorithm starts at some node and cuts all the lightpaths that pass through it into two segments. Then (in Steps 3—6) it sweeps the ring in a clockwise direction and whenever it encounters a lightpath that originates at some node k it strives to allocate a wavelength that has been used by another lightpath which had terminated before reaching node k. The differences between Cut-and-Color and the current cut-first heuristics are in Step 1, where we are more careful in choosing the initial node, and in step 4, where our algorithm is careful in achieving maximum sharing of ADMs between lightpaths. No such attempt is made in [T75]. The reader can easily verify that, similar to Cut-and-Color, the Cut-First algorithm also achieves the minimum number of wavelengths. The number of wavelength used is L, where L is the maximum load on a link in the ring.

7.1.1. *Cut-First performance*

In step 4, the heuristic attempts to pair up origins with terminations wherever possible. In fact, the pairing is always successful except at node 0. At node 0, the pairing is not guaranteed. Since the algorithm performs assignment on a node by node basis without looking ahead, there may exist situations where a pairing was possible but not found by the heuristic.

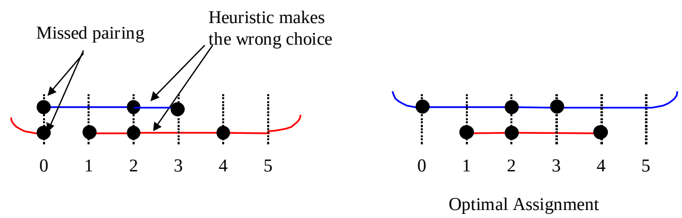

Figure 7. A mismatch could happen at node 0.

Therefore, Cut-First almost achieves the optimal assignment except for two factors: 1) Additional ADMs created due to the cut. 2) Mismatches at node 0. We can now bound the number of ADMs used by the Cut-First heuristics. We have,

$$\alpha \le \alpha_{lb} + \text{number of additional ADMs at node 0} \quad \textbf{(4)}$$

The additional ADMs at node 0 were created due to lightpath splitting and mismatch of origins and terminations at node 0. There are two ways to calculate the number of additional ADMs.

First, we calculate ADM creation directly. Each crossover at node 0 is split, and that gives rise to at most 2 ADMs per crossover. Furthermore, each origin (in the original set of lightpaths) may or may not be matched to a termination and vice versa. This gives rise to at most another $\min\{\sigma_0, \tau_0\}$ ADMs.

Therefore, the maximum number of ADM added is $2x_0 + \min\{\sigma_0, \tau_0\}$.

The second method of calculation is as follows: There cannot be more than L ADMs at node 0 because the total number of wavelengths used is L and no more than a single ADM is needed per wavelength. Because there are σ_0 origins and τ_0 terminations at node 0, any algorithm will result in at least $\max\{\sigma_0, \tau_0\}$ ADMs at node 0. At most L is needed out of which at least $\max\{\sigma_0, \tau_0\}$ are accounted for by the lower bound, and therefore, the number of additional ADMs at node 0 is at most $L - \max\{\sigma_0, \tau_0\}$.

Combining the two results, the number of additional ADMs is at most

$$\min\{2x_0 + \min\{\sigma_0, \tau_0\}, L - \max\{\sigma_0, \tau_0\}\} \quad \textbf{(5)}$$

In step 1, we minimize the number of additional ADMs by choosing node 0 such that the above equation is minimized. Therefore, the Cut-First algorithm achieves,

$$\alpha \le \alpha_{lb} + \min_i\{\min\{2x_i + \min\{\sigma_i, \tau_i\}, L - \max\{\sigma_i, \tau_i\}\}\} \quad \textbf{(6)}$$

The right hand side of equation (6) is at most $\alpha_{lb} + (L\text{-}1)$. We have already shown a worst case example in section 4 where it is necessary for $\alpha = \alpha_{lb} + (L\text{-}1)$. Therefore, the Cut-First algorithm is the best possible for that worst case example. However, it is not in general the optimal algorithm. The Cut-First algorithm often produces extra ADMs because of lightpath splitting at node 0. We next introduce the Assign-First heuristic in an attempt to find solutions with $\alpha = \alpha_{lb}$ whenever possible.

7.2. Assign-First heuristic

The Assign-First wavelength assignment heuristics avoids ADM creation by skipping the cut step in the Cut-First algorithm. Initially, assign all lightpaths that originate at or crossover node 0. Then greedily assign the rest of the lightpath node by node, from node 1 to node N - 1, as in the Cut-First heuristic.

The Assign-First heuristic also attempts to pair up origins with terminations whenever possible. The only place where it may fail are the assignments for those lightpaths that crossover node 0 and those that originate or terminate at node 0. Each crossover lightpath may cause 1 additional ADM because its origin may be unable to match to an available termination. At node 0, each non-matching origin/termination adds 1 additional ADM. Therefore, the maximum number of additional ADMs is at most $x_0 + \min\{\sigma_0, \tau_0\}$. Again, we can choose node zero so as to minimize this. Therefore, the performance of Assign-First is,

$$\alpha \le \alpha_{lb} + \min_i\{x_i + \min\{\sigma_i, \tau_i\}\} \quad \textbf{(7)}$$

7.3. Heuristic comparison

Each of the heuristics described has its advantage. To better illustrate the differences between the heuristics and understand the type of traffic that each heuristic is suited for, we have, in this section, two examples, each favoring a different heuristic. The favored heuristic will achieve the minimum number of ADMs while the other heuristic has worse performance.

Before delving into the comparisons, some explanation is needed in the tie-breaking criteria within heuristics. In both heuristics, more than one lightpath may be eligible to connect to the termination of an assigned lightpath (Step 4). In addition, an eligible lightpath may have more than one place to which it may connect. For a particular set of lightpaths, the heuristic performance depends on the tie-breaking criteria used. However, there is not one criterion that is generally good for all possible lightpath sets. In fact, for any given criteria, one can always find a set of lightpaths such that the tiebreaker works poorly. Therefore, we will assume the following tie-breaking criteria without loss of generality: always assign the shortest lightpath first, and assign it to the lowest numbered eligible wavelength. The examples in this section are created with this tie-breaking criterion in mind.

7.3.1. Cut-First advantage

The motivation behind the Cut-First heuristics is based on the sharing of ADMs promoted by lightpath splitting. This advantage at times overcomes the additional ADM created from the split. In fact, splitting may be necessary to reduce the number of ADMs in certain cases. Here is an example:

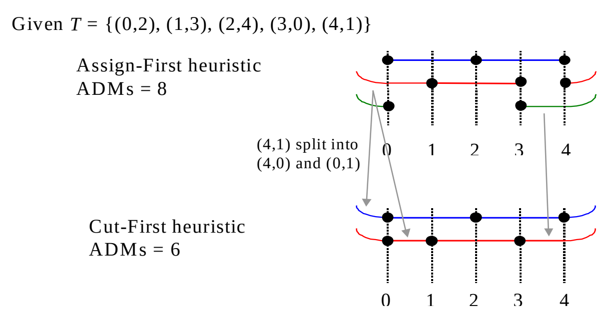

Figure 8. Example showing the advantage of Cut-First

Cut-First splits lightpath (4,1) into (4,0) and (0,1), and hence, allows (3,0) to fit on the second wavelength. Consequently, the total number of ADMs is reduced.

An additional advantage of Cut-First is that it uses the minimum number of wavelengths W=L.

7.3.2. Assign-First advantage

The Assign-First heuristic was developed because at times Cut-First would unnecessarily split a lightpath and hence create more ADMs. Assign-First does not allow lightpath splitting, and hence, it prevents unnecessary ADM creation.

The example below demonstrates the Assign-First advantage. As shown, lightpath (7,2) is unnecessarily split by Cut-First.

Given T = {(0,1), (0,5), (1,3), (1,8), (4,0), (6,8), (7,2), (8,0)}

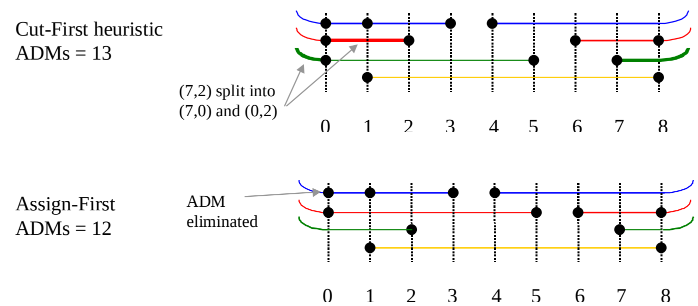

Figure 9. Example showing the advantage of Assign-First

Therefore, Assign-First may wins over Cut-First because it prevents unnecessary splitting.

8. Post-assignment transformations

In this section, we introduce three transformations that can reduce the number of ADMs. These transformations can be applied after the set of lightpaths has been assigned wavelengths under *any* algorithm. The three transformations are called: merging, combining, and further splitting.

8.1. Merging

If the assignment contain lightpaths that have been split, e.g. the output assignment resulted of Cut-First, then it may be possible to merge back the lightpaths and eliminate the ADMs produced from splitting.

Define a *valid segment* to be a contiguous subset of wavelength channels on a single wavelength that contain zero or more integral lightpaths and partial or integral gaps. See the following figure for examples.

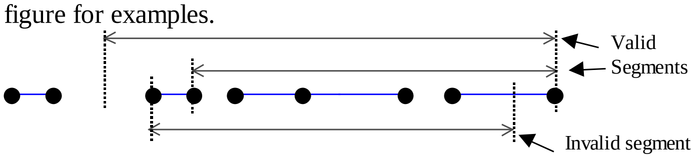

Figure 10. Example of valid and invalid segments.

Two valid segments are *swappable* if they use different wavelengths and start and end at the same nodes. When two segments are swapped, the lightpaths within one segment are re-assigned to the wavelength of the other segment and vice versa.

The goal of the merging transformation is to swap some swappable segments so as to merge back parts of a lightpath that was previously split. The following figure shows an example where merging is helpful.

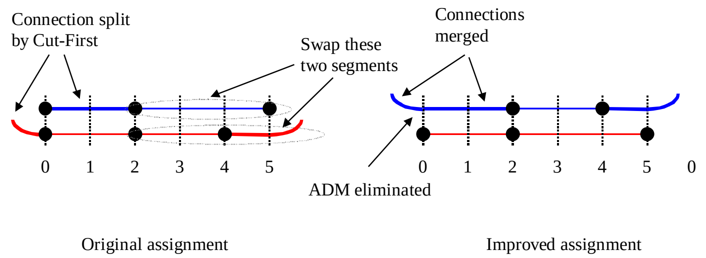

Figure 11. Merging created by segment swapping

The example shows only two wavelengths of the ring. The Cut-First algorithm split the lightpath (4,2) into (4,0) and (0,2), and the wavelength assignment created the situation depicted on the left of the figure. The assignment can be improved by swapping segments as shown. With the segment swapped, the split lightpath can be merged to eliminate the extra ADMs.

In general, merging can reduce up to 2 ADMs because each half of the split lightpath can eliminate one ADM.

8.2. Combining

Segment swapping not only allows split lightpaths to merge, it can also allow two lightpaths to share an ADM. This is best described with an example.

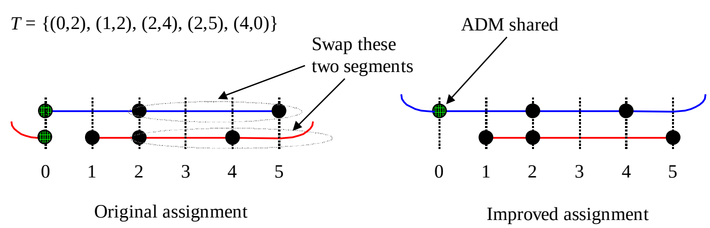

Since two different ADMs are combined, this method reduces the number of ADMs by one.

8.3. Splitting

The splitting transformation tries to split a lightpath so that the two parts of the split lightpath fit perfectly into gaps on other wavelengths. Although this transformation often creates more ADMs, if the two parts can use existing ADMs, then splitting will reduce the total number of ADMs. Below is an example:

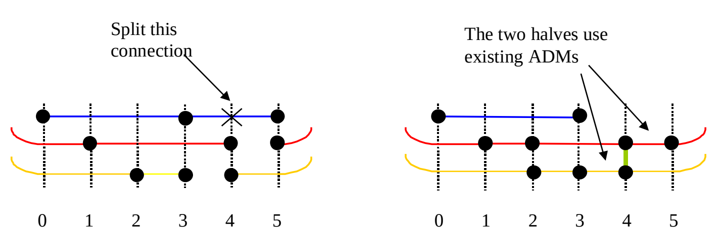

Because the split lightpath is now using existing ADMs, up to two ADMs may be eliminated using this method.

8.4. Transformation summary

The following table summarizes the three transformations:

Table 1. Transformation summary

Transformation	Applicable situation	ADMs reduced
Merging	Swapping segments merges previously split lightpaths	up to 2
Combining	Swapping segments increases the sharing of ADMs	1
Splitting	Split parts of a lightpath use existing ADMs	up to 2

Each of these three methods provides a stepwise reduction in the total number of ADMs. The overall algorithm is to first use one of the two heuristics and then apply these transformations. The order in which these transformations are applied is crucial to the final results. However, no consistent ordering methodology is known at this time.

9. Conclusions

When assigning wavelengths to lightpaths in a WDM ring, it is important not only to minimize the number of wavelengths but also the number of ADMs. We developed a lower bound on the number of ADMs necessary, and two algorithms that attempt to achieve the lower bound. The Cut-First algorithm uses at most L-1 more ADMs than the lower bound, and simultaneously achieves the minimum number of wavelength. The Assign-First algorithm does not allow splitting, therefore, could perform better than Cut-First, but generally uses more wavelengths than necessary. Both the Cut-First and Assign-First algorithms perform greedy assignments node-by-node. We also developed three transformation methods that improve the result of any assignment.

Acknowledgement. We would like to thank Rajiv Ramaswami for helpful discussions and comments.

References

[BH96] R. Barry and P. Humblet, "Models of blocking probability in all-optical networks with and without wavelength changers", *IEEE JSAC / IEEE-OSA JLT: Special Issue on Optical Networks*, vol. 14, no. 5, pp. 858-867, 1996.

[CGK93] I. Chlamtac, A. Ganz, and G. Karmi, "Lightnets: Topologies for high-speed optical networks", *IEEE/OSA Journal on Lightwave Technology*, vol. 11, pp. 951-961, 1993.

[FNSST92] A. Frank, T. Nishizeki, N. Saito, H. Suzuki, and E. Tardos, "Algorithms for routing around a rectangle", Discrete Applied Math, vol. 40, pp. 363-378, 1992.

[GJMP80] M. Garey, D. Johnson, G. Miller, and C. Papadimitriou, "The complexity of coloring circular arc graphs and chords", *SIAM Journal of Discrete Math*, vol. 1, no. 2, pp. 216-227, 1980.

[GSR96] O. Gerstel, G. Sasaki, and R. Ramaswami, "Dynamic wavelength allocation in WDM ring networks with little or no wavelength conversion", *The 34th Allerton conf. On Communications, Control, and Computing*, 1996.

[GSR98] O. Gerstel, G. Sasaki, and R. Ramaswami, "Cost Effective Traffic Grooming in WDM Rings", *Infocom'98*.

[KA96] M. Kovacevic and A. Acampora, "Benefits of wavelength translation in all-optical clear-channel networks", *IEEE JSAC / IEEE-OSA JLT: Special Issue on Optical Networks*, vol. 14, no. 5, pp. 868-880, 1996.

[RU94] P. Raghavan and E. Upfal, "Efficient routing in all-optical networks", *in Proc. 26th ACM Symp. Theory of Computing*, pp. 134-143, 1994.

[T75] A. Tucker, "Coloring a family of circular arcs", *SIAM Journal of Applied Math*, vol. 29, no. 3, pp. 493-502, 1975.