

Approximation Algorithms for Wavelength Assignment

Vijay Kumar¹ and Atri Rudra²

¹ Strategic Planning and Optimization Team,
Amazon.com,
Seattle, WA, USA.
email: vijayk@amazon.com

² Department of Computer Science and Engineering,
University of Washington,
Seattle, WA, USA.
email: atri@cs.washington.edu

Abstract. Winkler and Zhang introduced the FIBER MINIMIZATION problem in [10]. They showed that the problem is NP-complete but left the question of approximation algorithms open. We give a simple 2-approximation algorithm for this problem. We also show how ideas from the Dynamic Storage Allocation algorithm of Buchsbaum et al. [4] can be used to give an approximation ratio arbitrarily close to 1 provided the problem instance satisfies certain criteria. We also show that these criteria are necessary to obtain an approximation scheme. Our 2-approximation algorithm achieves its guarantee unconditionally. We also consider the extension of the problem to a ring network and give a $2+o(1)$ -approximation algorithm for this topology. Our techniques also yield a factor-2 approximation for the related problem of PACKING INTERVALS IN INTERVALS, also introduced by Winkler and Zhang in [10].

1 Introduction

The FIBER MINIMIZATION problem on an *optical linesystem* ([10]) relates to the efficient construction of an optical fiber network to meet a specified collection of demands. Consider n links (or edges) connected in a line. Each demand needs to use some consecutive links, and thus the demands can be represented as a set of line intervals. A segment of optical fiber, which is of wavelength μ and spans some consecutive links, can carry a collection of demands such that links required by the demands are contained within the links spanned by the fiber, and no two demands are assigned the same wavelength if there is a link which they both use. The goal is to get a set of fiber segments which can carry all demands such that the total length of fiber used is minimized.

This problem was introduced by Winkler and Zhang in [10]. The problem is motivated by wavelength division multiplexing (WDM), which is used to partition the bandwidth available on an optical fiber into multiple channels each of

which can handle different wavelengths. These technologies have made very fast all-optical transmission physically possible. See [3, 10] for more details on the practical motivations.

1.1 Problem definition

More formally, consider a linesystem of n links e_1, e_2, \dots, e_n . We represent each demand by $d_j = [l_j, r_j]$ for $l_j \leq r_j$ if it requires the links $e_{l_j}, e_{l_j+1}, \dots, e_{r_j}$; the set of demands is denoted by D . We say demands d_j and $d_{j'}$ intersect (or overlap) if either $l_j \leq l_{j'} \leq r_j$ or $l_{j'} \leq l_j \leq r_{j'}$. A fiber interval f is represented by $f = [l_f, r_f]$ for $l_f \leq r_f$ if it spans the edges $e_{l_f}, e_{l_f+1}, \dots, e_{r_f}$. The goal is to construct a set F of fiber intervals (each capable of carrying μ different wavelengths) of minimum total length ($\sum_{F \ni f=[l_f, r_f]} (r_f - l_f + 1)$) such that D can be packed in F . A packing of D in F is an assignment of each demand $d_j = [l_j, r_j]$ to a fiber $f = [l_f, r_f] \in F$, and a wavelength $\omega \in \{1, \dots, \mu\}$ within f , such that $[l_j, r_j] \subseteq [l_f, r_f]$ and no two intersecting demands are assigned the same wavelength in the same fiber.

1.2 Previous work

There is a substantial body of work relating to resource optimization in WDM networks. Most of this work addresses the problem of minimizing the number of wavelengths necessary to satisfy the presented demand for a given network topology. This body of work is too extensive to summarize here — please see [8] for a survey of the field.

In some recent work [2, 3, 5, 9, 10] it is sought to more faithfully represent real-world constraints by assuming that fiber has a fixed capacity. Starting with [10] this vein of research aims to optimize the use of fiber. There are different flavors of this problem depending on what the objective function is —

- Minimizing the ratio of the number of fibers deployed to the minimum number of fibers required for any edge [2].
- Minimizing the maximum over all links the the number of fibers used in a link [2, 5, 6].
- Minimize the total amount of fiber used [3, 6].

Much recent work addresses a different flavor of the FIBER MINIMIZATION problem. That version assumes the availability of a device called a *mesh optical add/drop multiplexer* (or MOADM), which allows the signals to be moved from one fiber to another as they pass through the nodes in the optical network under consideration. In the presence of MOADMs, the FIBER MINIMIZATION problem becomes “easier” as we are allowed to “break” a demand into segments which span an edge while routing them through the network. For the case of a linesystem, [10] give a polynomial time algorithm to solve this version of the problem. For ring and tree networks, constant-factor approximation algorithms are known [5, 9]. Recently, Andrews and Zhang have showed that for general

graph topologies, this problem (which they call SUMFIBER-CHOOSEROUTE) is hard to approximate to within a poly-logarithmic factor of the optimal for a general network topology under some complexity assumptions, whether the routing of the demands is performed by the algorithm ([3]), or specified as input ([2]). They also prove ([2]) a similar hardness result under a similar complexity assumption for the problem where the objective function is to minimize the ratio, over all edges, of the number of fibers deployed for that edge to the minimum number of fibers required. These negative results are accompanied in all cases by approximation algorithms with logarithmic approximation ratios.

A related set of problems is tackled in [1], where linesystem design is modeled as a generalized graph coloring problem. A collection of optimal and approximation algorithms are presented for different network scenarios, where the objective is to minimizing costs related to the set of colors employed. It has been pointed out to us that some of their techniques can be employed to substitute Phase 1 of our algorithm in Section 2.

In the setting of [10] where there are no MOADMs, — that is, a demand cannot change fibers at any node — FIBER MINIMIZATION was shown to be NP-complete in [10]. To the best of our knowledge, ours are the first approximation algorithms in this model.

1.3 Our results

We present polynomial time approximation algorithms for the FIBER MINIMIZATION problem on some simple network topologies.

- In Section 2, we give a 2-approximation algorithm for the FIBER MINIMIZATION problem. This problem is similar to the Dynamic Storage Allocation (DSA) problem in some aspects, and we employ techniques from the DSA literature, such as those of Gergov [7] and Buchsbaum et al. [4], to tackle it (Incidentally, the proof of Theorem 1 has essentially the same structure as the proof — to the best of our knowledge, unpublished — of Gergov’s [7] algorithm being a 3-approximation for DSA, and may be interesting on that account).
- We use a result of Buchsbaum et al. [4] to derive an approximation scheme for FIBER MINIMIZATION in Section 3.
- We extend these results to obtain other approximations for related problems. In Section 4 give a factor $2 + o(1)$ approximation algorithm for FIBER MINIMIZATION on a ring.
- We investigate the related problem of PACKING INTERVALS IN INTERVALS — also posed in [10] — in Section 5, and obtain a 2-approximation algorithm.

2 A 2-approximation algorithm for FIBER MINIMIZATION

As specified in Section 1.1, the input D to the FIBER MINIMIZATION problem consists of demands d_j , each of which is an interval of the form $[l_j, r_j]$, where

$l_j \leq r_j \in \{1, \dots, n\}$. For each link e_i , we define $L_D(i) = |\{d_j \in D : i \in [l_j, r_j]\}|$ (WLOG we assume that $L_D(i)$ is a multiple of μ) and $L_D^{max} = \max_i L_D(i)$. The algorithm uses an $L_D^{max} \times n$ matrix H to keep track of wavelength assignments. For each $k \in \{1, \dots, L_D^{max}\}$, the row H_k is referred to as the k^{th} row and finally would correspond to a wavelength in some fiber interval. We say that link e_i is colored c_i in row k if $H_{k,i} = c_i$. For any row k and color C , we call an interval $[l, r]$ in the k th row a C segment if $H_{k,l-1} \neq C$, $H_{k,r+1} \neq C$ and $\forall i \in [l, r]$, $H_{k,i} = C$.

The algorithm, which we refer to as the **Simple** algorithm, is specified in Figures 1 and 2. In **Phase 1**, the algorithm constructs the matrix H and then derives a “packing” P_D from H . We say that demand $d_j \in D$ is packed in P_D if $\langle j, k \rangle \in P_D$ for some $k \in \{1, \dots, L_D^{max}\}$. P_D is not a feasible solution since the assignments of some demands overlap, although only to a limited extent. In **Phase 2**, the overlaps are taken care of and a valid packing derived.

2.1 Algorithm outline

Let us briefly summarize the algorithm before we formally describe it. The basic intuition for the algorithm comes from techniques often used for the Dynamic Storage Allocation (DSA) problem, in particular those used in [7]. FIBER MINIMIZATION has some similarities with DSA, with the important distinction that while the demands in DSA are axis-parallel rectangles, demands in our problem are line intervals.

In a nutshell, this is how the algorithm works. During Phase 1 (Figure 1), each matrix entry $H_{k,i}$ either has the value *red* (which means that no demand can be “placed on e_i ” in the k th row), or *green* (which means that at most one demand can be “placed on e_i ” in the k th row), or *blue* (which means overlapping demands can be “placed on e_i ” in the k th row). Each edge e_i is “allotted” $L_D(i)$ rows, which is why the first $L_D^{max} - L_D(i)$ entries in the i th column of H are colored *red* in Step 2. All other entries are colored *green*. Initially all green segments are “available”.

In Step 3, the algorithm iterates over all possible rows and in each iteration looks for an “available” green segment over which some “unpacked” demand can be placed (maybe partially) in the following way: this demand can possibly intersect blue segments in that row but it must not intersect with any other green segment. The placement of a demand can fragment an “available” green segment into smaller “available” green segments (the edges common to the placed demand and the green segment are no longer available). The iteration is complete when no demand can be placed on any available green segment. Edges which were colored blue in the current row or did not have any demand placed on them in the current row are colored blue in the next row. Note that this implies that if an edge becomes blue in one row then it remains blue for all the subsequent iterations (rows). Phase 1 of the algorithm is complete when D_{max} iterations are complete. We will show in Lemma 4 that after the first phase, all demands are “packed”.

Further, in Lemma 5 we show that in a blue segment a demand may intersect, if at all, with no more than one other demand. This suggests Phase 2 (Figure 2) of the algorithm where demands are finally packed into fiber intervals. Consider μ consecutive rows and consider maximal intervals of consecutive edges which are not colored red in any of the μ rows. In each such segment, every one of the μ rows has, by Lemma 5, at most two demands conflicting over any particular link. Thus, creating two fiber intervals corresponding to such a segment is sufficient to accommodate all the demands.

The algorithm is more formally described in Figures 1 and 2.

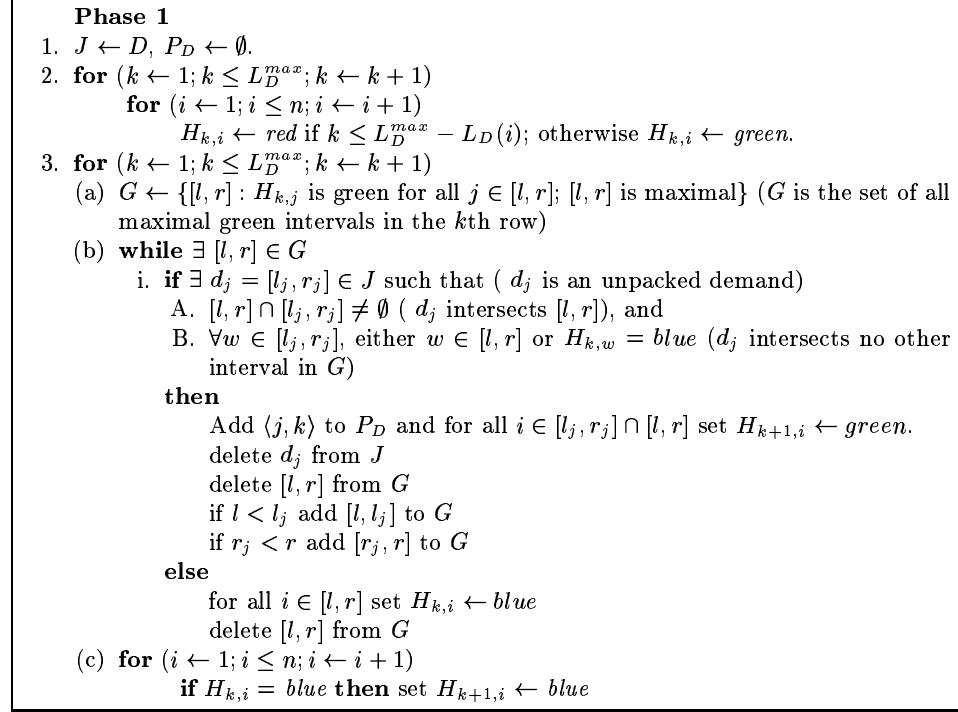


Fig. 1. Phase 1 of the **Simple** algorithm

2.2 Correctness and performance

The following three lemmas follow directly from the way the matrix H is manipulated in the algorithm of Figure 1.

Lemma 1. *At the end of Phase 1, if $H_{k,i} = green$ for any k and i , then for all $1 \leq k' < k$, $H_{k',i} = green$ or $H_{k',i} = red$.*

```

 $F_D \leftarrow \emptyset.$ 
for ( $k = 1; k \leq L_D^{max}; k = k + \mu$ )
1.  $\forall i \in \{1, \dots, n\}, c[i] \leftarrow black;$ 
2.  $\forall i \in \{1, \dots, n\},$  if  $\exists j \in [k, k + \mu]$  such that  $H_{j,i} = red$  then  $c[i] \leftarrow red.$ 
3. for each maximal interval  $I = [l, r] \subseteq [1, n]$  such that  $i \in [l, r] \Rightarrow c[i] = black,$ 
   (a) Create two fiber segments  $f_1(I)$  and  $f_2(I)$  and add them to  $F_D.$ 
   (b) for all  $j \in [k, k + \mu),$ 
       i. Let  $S_j \leftarrow \{d_i : d_i \subseteq I, \langle i, j \rangle \in P_D\}.$  Let  $s_{j,1}, s_{j,2}, \dots, s_{j,N}$  be an ordering
          of  $S_j$  such that for any two demands  $s_{j,a} = [l_a, r_a]$  and  $s_{j,b} = [l_b, r_b],$ 
           $a < b \Rightarrow r_a \leq r_b.$ 
       ii. for  $i = 1, 2, \dots, N,$ 
          - Assign  $s_{j,i}$  the wavelength  $j - k + 1.$ 
          - Assign  $s_{j,i}$  the fiber segment  $f_2(I)$  if there exists some demand  $s_{j,u}$ 
            assigned wavelength  $j - k + 1$  and fiber segment  $f_1(I),$  such that
             $s_{j,i} \cap s_{j,u} \neq \emptyset;$  otherwise, assign  $s_{j,i}$  the fiber segment  $f_1(I).$ 

```

Fig. 2. Phase 2 of the **Simple** algorithm

Proof. If $H_{k',i}$ is colored blue, then $H_{j,i}$ gets colored blue for all $j > k'$ due to repeated execution of Step 3(c); in particular $H_{k,i}$ gets colored blue, which contradicts the assumption on $H_{k,i}$ in the lemma. ■

Lemma 2. *At the end of Phase 1, if $H_{k,i} = red$ for any k and i , then for all $1 \leq k' < k, H_{k',i} = red.$*

Proof. This is ensured by Step 2, where $H_{k,i}$ is colored red for all k less than a certain value. The color red is not employed at any other step in the algorithm; nor is it ever replaced by any other color. ■

Lemma 3. *At the end of Phase 1, $L_D(i) = |\{k : H_{k,i} \neq red\}|.$*

Proof. This follows from how the coloring decision is made at Step 2, and the fact that the set $\{(k, i) : H_{k,i} \text{ is red}\}$ is invariant over the later steps of the algorithm. ■

We now show that all demands in D are “packed” in P_D after Phase 1.

Lemma 4. *At the end of Phase 1, J is empty.*

Proof. Assume that this is not the case, and there is a demand $d_t = [l_t, r_t] \in J$ at the end of Phase 1. One of the following three cases must arise:

- **Case 1:** There is one $i \in [l_t, r_t]$ such that $H_{L_D^{max}, i} = red.$ By Lemma 2, for all $k \in [1, L_D^{max}], H_{k,i} = red.$ Lemma 3 implies that $D(i) = 0,$ which contradicts the fact that d_t uses link $i.$
- **Case 2:** There exists $i \in [l_t, r_t]$ such that $H_{L_D^{max}, i} = green.$ It follows from Lemmas 1 and 2 and the coloring criterion of Step 2 that for all $k \in \{1, \dots, L_D^{max} - L_D(i)\}, H_{k,i} = red;$ and for all $k' \in \{L_D^{max} - L_D(i) + 1, \dots, L_D^{max}\}, H_{k',i} = green.$

For each such k' , it must be the case that in the k' th iteration of the **for** loop in Step 3, some demand $d_j \ni e_i$ was placed in row k' , since otherwise $H_{k',i}$ would have been colored blue at Step 3(b). In all there would be $L_D(i)$ such demands, one for each value of k' . Including them, and including d_t , there are at least $L_D(i) + 1$ demands that use link e_i , which contradicts the definition of $L_D(i)$.

- Case 3: For all $i \in [l_t, r_t]$, $H_{L_D^{max},i}$ = blue. We complete the proof by showing that d_t would have been placed by the algorithm in some row. Consider $k^* = \min\{k : \forall j \in [l_t, r_t], H_{k,j} = \text{blue}\}$. By the choice of k^* , there exists an interval $[l, r] \subseteq [l_t, r_t]$ such that in some iteration of the **while** loop of Step 3(b),

- $[l, r] \in G$, and
- for all $i \in [l_t, r_t] - [l, r]$, $H_{k^*,i}$ = blue, and
- for all $j \in [l, r]$, $H_{k^*,j}$ is colored blue by the **else** clause of Step 3(b)(i).

This can only happen when there is no demand d_j which is suitable for placing over $[l, r]$ in row k^* : as indicated by conditions A and B in the **if** statement. However, d_t is precisely such a demand; it is unpacked, it intersects $[l, r]$, and (by the choice of k^*) it does not anymore intersect any other interval in G . Thus, instead of coloring $[l, r]$ blue, d_t should have been placed over it; and this completes our proof by contradiction. ■

We now show that the packing P_D has a nice property— no link is used by more than two demands that are placed in the same row. In other words, no three demands conflict simultaneously.

Lemma 5. $\forall(i, k), |\{j : \langle j, k \rangle \ni e_i\}| \leq 2$.

Proof. It is easy to see from the way intervals are added to and deleted from G that only one demand is placed on a green segment, that is, demands do not overlap over green segments. Thus, it follows that overlaps can only take place over blue segments (as no demands are placed over red segments).

Consider such a segment $[l, r]$ in row k . Among the demands that are placed over this segment, there can be at most one demand that contains $H[k, l-1]$, and at most one demand that contains $H[k, r+1]$ — that is because both $H[k, l-1]$ and $H[k, r+1]$ are non-blue and thus can not have overlapping demands placed on them. By placement rules [3(b)i.A.] and [3(b)i.B.], no demand is contained within a blue segment. Thus, no more than two demands can be placed over $[l, r]$ in row k . ■

We next show that Phase 2 outputs a valid solution.

Lemma 6. *Phase 2 (Figure 2) produces a valid packing of D in the set of F_D .*

Proof. Consider a demand $d_i = [l_i, r_i]$. Let d_i be placed in row j where $j \in [(h-1)\mu, h\mu)$ for some j and h — that is, $\langle i, j \rangle \in P_D$.

First of all, let us verify that d_i is assigned a fiber interval in Phase 2. Consider the h th iteration of the **for** loop (Figure 2) in Phase 2. At Step 3(a), an interval $I \ni [l_i, r_i]$ would indeed be created, except if $c[u] = \text{red}$ for some $u \in [l_i, r_i]$. Is

that possible? For that to be the case, there must exist some $j' \in [(h-1)\mu, h\mu)$ such that $H_{j',u} = \text{red}$.

Recall how certain elements of H are colored *red* at Step 2 in Phase 1 (Figure 1). If $H_{j',u}$ is *red*, then $H_{v,u}$ is *red* for all $v \leq L_D^{max} - L_D(u)$, and it is not *red* for any other v . Since $L_D^{max} - L_D(u)$ is a multiple of μ (we have assumed the load at any link to be a multiple of μ), it follows that $H_{v,u}$ is *red* either for all $v \in [(h-1)\mu, h\mu)$, or for none of those values of v . In the latter case, the desired fiber interval I is indeed created; while the former case is easy to rule out, since it implies that $H_{j,u}$ is *red* as well, which is not possible given that $\langle i, j \rangle \in P_D$ and $u \in [l_i, r_i]$ (no d_i can be placed over a *red* interval in Phase 1).

Note that demand d_i is assigned wavelength $j - (h-1)\mu + 1$, and one of the fibers $f_1(I)$ and $f_2(I)$. It remains to be verified that no other demand is assigned the same wavelength and the same fiber. We do this by pointing out that the set of demands that have been assigned wavelength $j - (h-1)\mu + 1$ on $f_1(I)$ or $f_2(I)$ is exactly the set of demands d_k for which $\langle k, j \rangle \in P_D$. This set of demands has been characterized by Lemma 5—any given link is contained in no more than two of these demands. This set of demand can be thought of as a collection of line intervals. Packing them into $f_1(I)$ and $f_2(I)$ without conflict is akin to 2-coloring the corresponding interval graph (whose clique number is 2). It is well-known and easy to see that a legal 2-coloring can be obtained simply by scanning the intervals left-to-right, and greedily assigning them one of two colors. Note that this is precisely what Step 3(b)ii of Phase 2 (Figure 2) attempts to do. ■

Let L be defined as $\frac{\sum_{i=1}^n L_D(i)}{\mu}$. For any set F of fiber intervals, let L_F be the total length of fiber used in F .

The following lemma shows that F_D is a pretty good solution.

Lemma 7. $L_{F_D} = 2L$.

Proof. Let us denote a fiber interval f by $[l_f, r_f]$, and for any link e_i , let $L_{F_D}(i) = |\{f \in F_D : i \in [l_f, r_f]\}|$. We will prove a stronger claim— for all e_i , $L_{F_D}(i) = 2\frac{L_D(i)}{\mu}$. The lemma follows by summing over all links.

Consider Step 3 of Phase 2 (Figure 2), where fiber intervals are created in pairs $(f_1(I)$ and $f_2(I))$. This step is repeated in the L_D^{max}/μ iterations of the **for** loop. A link e_i will not be contained in these fiber intervals for iterations $1, 2, \dots, \frac{L_D^{max} - L_D(i)}{\mu}$, and included in both intervals of a pair $(f_1(I), f_2(I))$ in all subsequent iterations. This is because $H_{v,i}$ is *red* for all $v \leq L_D^{max} - L_D(i)$, and not *red* for any other v , as we saw in the proof of Lemma 6. This implies that during the first $\frac{L_D^{max} - L_D(i)}{\mu}$ iterations, $c[i]$ is *red*, and thus the fiber intervals created do not include link e_i . In other words, link e_i is contained in exactly $2\frac{L_D(i)}{\mu}$ fiber intervals. ■

Clearly, L is a trivial lower bound on the length of the optimal set of fiber intervals for D . Thus, we have the following result.

Theorem 1. *Simple algorithm is a 2-approximation algorithm for the FIBER MINIMIZATION problem.* ■

3 An Approximation Scheme

In this section we employ the *boxing* technique of Buchsbaum et al. [4] to the FIBER MINIMIZATION problem. The application of a result of [4] yields an approximation scheme for our problem. Let us first briefly describe *boxing*, which is applied in [4] to the Dynamic Storage Allocation problem. Let Z be a set of *jobs*, where each job i is a triple of start time l_i , end time r_i and height h_i . To *box* Z means placing the jobs in a box b starting from time $l_b = \min\{l_j : j \in Z\}$, ending at time $r_b = \max\{r_j : j \in Z\}$ and of height $h_b \geq \sum_{j \in Z} h_j$. A *boxing* of Z into a set B of boxes is a partition of Z into $|B|$ subsets, each of which is then boxed into a distinct $b \in B$. At any time t , let $L_Z(t)$ denote $\sum_{j \in Z: l_j \leq t \leq r_j} h_j$, and let $L_B(t) = \sum_{b \in B: l_b \leq t \leq r_b} h_b$.

We will be working with jobs of unit height, for which the algorithm in Section 2.1 of [4] comes with the following performance guarantee:

Theorem 2. [4] *Given a set Z of jobs, each of height 1, an integer box-height parameter H , and a sufficiently small positive ϵ , there exists a set B of boxes, each of height H , and a boxing of Z into B such that for all time t :*

$$L_B(t) \leq (1 + 4\epsilon)L_Z(t) + O\left(\frac{H \log H}{\epsilon^2} \log \frac{1}{\epsilon}\right).$$

The key insight here is that each demand in the FIBER MINIMIZATION problem can be viewed as a job of unit height, and packing of demands into fiber intervals is analogous to the boxing of a collection of jobs. This leads us to the following bound, where $L_F(t)$ denotes as before the number of fibers in F containing the link e_t and $L_D(t)$ denotes the number of demands in D using link e_t .

Lemma 8. *Given a set D of demands and a sufficiently small positive ϵ , there exists a set of fiber intervals F and a packing of the demands of D into F such that for all links e_t :*

$$L_F(t) \leq (1 + 4\epsilon) \frac{L_D(t)}{\mu} + O\left(\frac{\log \mu}{\epsilon^2} \log \frac{1}{\epsilon}\right).$$

Proof. A straightforward reduction maps an instance D of the FIBER MINIMIZATION problem to an instance of boxing. Corresponding to every demand $[l, r]$ in D , let there be a job with a start time of $l - 1$ and end time of r . Each link e_i is mapped to the time interval $[i - 1, i]$, and fiber intervals of wavelength $\mu = H$ map to boxes of height $H = \mu$.

Consider the boxing algorithm in Section 2.1 of [4]. As observed above, a demand $[l, r]$ in D corresponds to a job $(l - 1, r)$ in the DSA setting. Further note that a box of height μ in the DSA setting corresponds to a fiber interval. The set of fiber intervals F and set of demands D map directly to the set B of boxes and set Z of jobs, respectively, in the boxing instance. Observe that $L_D(t) = L_Z(t)$ and $L_F(t) = \frac{L_B(t)}{\mu}$.

The lemma follows directly from an application of Theorem 2 to the boxing instance Z with $H = \mu$. ■

Noting that $L_F = \sum_{i=1}^n L_F(i)$ is the total length of fiber used by the algorithm, and $L_D = \frac{\sum_{i=1}^n L_D(i)}{\mu}$, we have

Theorem 3. $L_F \leq (1 + 4\epsilon)L_D + O(n)$ for sufficiently small positive constant ϵ .

Proof. Using Lemma 8 and summing over all links. The last term on the right hand side in Lemma 8 is a constant, which leads to the $O(n)$ upon summation. ■

As in Section 2, we observe that for F^* , the optimal set of fiber intervals, $L_{F^*} \geq L_D$. Thus, algorithm of [4] has a competitive ratio arbitrarily close to 1 provided $L_D = \Omega(n)$.

Next, we look at the $O(n)$ term in Theorem 3 more carefully.

3.1 A Lower Bound

It is easy to see that the $O(n)$ additive term in the statement of Theorem 3 cannot be done away with for any approximation scheme. Consider the set D_{bad} of demands over $2n + 1$ links. D_{bad} contains one copy of demand $[1, 2n + 1]$, and $\mu - 1$ copies each of demands $[1, n + 1]$ and $[n + 1, 2n + 1]$.

Clearly, $L_{D_{bad}} \leq 2n + 2$ while L_{F^*} , the value of the optimal solution, is $3n + 2$. That is, there can be no positive constant $\delta < \frac{1}{2}$ such that $L_{F^*} < (1 + \delta)L_{D_{bad}} + o(n)$ for all n .

4 FIBER MINIMIZATION in a Ring

Next, we look at what is perhaps the most natural generalization of the FIBER MINIMIZATION problem. Ring topologies are very commonly encountered and widely studied in optical routing literature. Consider a ring of n links where each demand d_j in the demand set D is an arc $[l_j, r_j]$ which requires links $e_{l_j}, e_{l_j+1 \bmod(n)}, \dots, e_{r_j}$. Fiber intervals are now arcs each of which can support μ different wavelengths. The goal is to find the set of fiber arcs with the minimum total length.

The straightforward technique of partitioning an arc coloring problem into two interval coloring problems and taking the union of the two solutions would seem to directly obtain an approximation ratio of twice of that of the approximation ratio for the problem on a line system (that is, a ratio of four in this case). However, a small tweak gives a $(2 + \epsilon)$ -approximation ratio algorithm with two invocations of the **Simple** algorithm of Figures 1 and 2. Arbitrarily pick a link e_i and consider the set of demands using link e_i , $D^i = \{d_j \in D : i \in [l_j, l_j + 1 \bmod(n), \dots, r_j]\}$. Now run the **Simple** algorithm on both D^i and $D - D^i$ to get sets of fiber intervals (arcs) F_{D^i} and F_{D-D^i} . Define $F_D = F_{D^i} \cup F_{D-D^i}$. Due to (possible) rounding³ “errors” we now have for each link e_i , $L_{F_D}(i) \leq 2L_D(i) + 1$. Thus we have:

Theorem 4. *The combined algorithm gives $L_{F_D} \leq 2L_{F_D^*} + n$, where F_D^* is the optimal set of fibers arcs for D .*

³ For each arc e_i at most 2 fiber intervals containing e_i from the two solutions can be merged in the solution for the original problem.

Theorem 4 implies that if $L_{F_D^*} \geq \frac{1}{\epsilon}n$, then the combined algorithm achieves an approximation ratio of $2 + \epsilon$.

The ideas developed in the paper so far can be applied to the related problem of PACKING INTERVALS IN INTERVALS [10].

5 PACKING INTERVALS IN INTERVALS

Winkler and Zhang also introduced the PACKING INTERVALS IN INTERVALS problem in [10]. Here we are given a set of demands D and a set of fiber intervals F , and the goal is to determine if D can be packed in F . Consider the optimization version of this decision problem. What is the smallest value of μ such that F can accommodate D ? Our techniques imply a 2-approximation algorithm for this problem in the following sense:

Theorem 5. *If D can be packed in F using no more than $\frac{\mu}{2}$ wavelengths in each fiber, then there exists an algorithm which can pack D in F utilizing no more than μ wavelengths in each fiber.*

Proof (Sketch): A small modification to the **Simple** algorithm is required. Let F be the set of fiber intervals and for each edge e_i define $L_D(i) = \frac{\mu}{2} |\{f : F \ni f = [l_f, r_f] \text{ and } i \in [l_f, r_f]\}|$. As in Section 2, L_D^{max} is defined as $\max_i L_D(i)$. Run Step 1 and 2 of Figure 1 with these values of $L_D(i)$ and L_D^{max} . Run Step 3 of Figure 1 with the given set of demands D . Execute Phase 2 (Figure 2) using $\frac{\mu}{2}$ in place of μ , and in the output F_D merge each $(f1(I), f2(I))$ pair.

Using arguments similar to ones used to prove Lemma 4, one can show that if D can be packed in F using no more than $\frac{\mu}{2}$ wavelengths in each fiber then J is empty after the execution of Phase 1. Similarly, analogues of Lemmas 5, 6 and 7 can be proved if the assumption of the theorem statement is valid. The definition of $L_D(i)$ and the fact that the constructed F_D is actually F completes the proof. A detailed proof is omitted due to space restrictions. ■

6 Conclusions

We presented a clean 2-approximation algorithm for the FIBER MINIMIZATION problem on a linesystem. We also apply techniques from [4] to give an approximation scheme for this problem. Based upon our 2-approximation algorithm, we obtain good approximations for the related problems of FIBER MINIMIZATION on a ring and PACKING INTERVALS IN INTERVALS.

Interesting open problems include investigating the FIBER MINIMIZATION problem on other network topologies, particularly those common in optical fiber networks, such as trees and meshes.

References

1. M. Alicherry and R. Bhatia. Line system design and a generalized coloring problem. In *Proc. of ESA 03*, 2003.

2. M. Andrews and L. Zhang. Wavelength assignment in optical networks with fixed fiber capacity. In *Proc. of ICALP*, 2004.
3. M. Andrews and L. Zhang. Bounds on fiber minimization in optical networks. In *Proc. of IEEE INFOCOM 05*, 2005.
4. A. Buchsbaum, H. Karloff, C. Kenyon, N. Reingold, and M. Thorup. Opt versus load in dynamic storage allocation. In *Proc. of STOC 03*, 2003.
5. C. Chekuri, M. Mydlarz, and F. B. Shepard. Multicommodity demand flow in a tree. In *Proc. of ICALP 03*, 2003.
6. T. Erlebach, A. Pagourtzis, K. Potika, and S. Stefanakos. Resource allocation problems in multifiber WDM tree networks. In *Proc. of the 29th Workshop on Graph Theoretic Concepts in Computer Science*, pages 128–229, 2003.
7. J. Gergov. Algorithms for compile-time memory optimization. In *Proc. of 10th SODA*, 1999.
8. R. Klasing. Methods and problems of wavelength-routing in all-optical networks. In *Tech rep. CS-RR-348, Department of Computer Science, University of Warwick*, 1998.
9. C. Nomikos, A. Pagourtzis, and S. Zachos. Routing and path-multicoloring. *Information Processing Letters*, 2001.
10. P. Winkler and L. Zhang. Wavelength assignment and generalized interval graph coloring. In *Proc. of SODA 03*, 2003.