Minimizing Maximum Fiber Requirement in Optical Networks

Matthew Andrews [*]   Lisa Zhang [†]

July 15, 2005

[*]Bell Laboratories, Murray Hill, New Jersey. andrews@research.bell-labs.com

[†]Bell Laboratories, Murray Hill, New Jersey. ylz@research.bell-labs.com. Corresponding Author: 600 Mountain Ave 2C-519, Murray Hill, NJ 07974. Tel: (908) 582-5281 Fax: (908) 582-5857.

Abstract

We study wavelength assignment in an optical network where each fiber has a fixed capacity of μ wavelengths. Given demand routes, we aim to minimize the maximum ratio between the number of fibers deployed on a link e and the number of fibers required on the same link e when wavelength assignment is allowed to be fractional. Our main results are negative ones. We show that there is no constant-factor approximation unless NP⊆ZPP. In addition, unless NP $\subseteq$ ZPTIME($n^{\text{polylog } n}$) we show that there is no $\log^{\gamma} \mu$ approximation for any $\gamma \in (0, 1)$ and no $\log^{\gamma} m$ approximation for any $\gamma \in (0, 0.5)$ where m is the number of links in the network. Our analysis is based on the hardness of approximating the chromatic numbers. On the positive side, we present algorithms with approximation ratios $O(\log m + \log \mu)$, $O(\log D_{\max} + \log \mu)$ and $O(D_{\max})$ respectively, where $D_{\max}$ is the length of the longest path.

Keywords: Optical networking, wavelength assignment, fixed capacity fiber, inapproximability.

1   Introduction

We consider the problem of achieving *transparency* in optical networks. A path is said to be routed transparently if it is assigned the same wavelength from its source to its destination. Transparency is desirable since wavelength conversion is expensive and defeats the advantage of all-optical transmission.

More formally, we consider an optical network consisting of vertices and optical links and a set of demands each of which needs to be routed from a source vertex to a destination vertex on a single wavelength. Each optical link has one or multiple parallel fibers deployed. The fundamental constraint is that for each wavelength λ, each fiber can carry at most one demand that is assigned wavelength λ. A common problem is to minimize the number of wavelengths required so that all demands can be routed assuming one fiber per link. However, in reality a more pertinent problem is that the number of wavelengths that each fiber can carry is fixed to some value μ, i.e. the total number of wavelengths is fixed. (For example, [11] lists the fiber capacities from different vendors.) The problem now is to minimize the number of fibers required.

For most service providers, the cost of a fiber on a link can be divided into two components. First, there is the cost of renting the fiber from a "dark-fiber" provider. Second, there is the cost of purchasing optical equipment to "light" the fiber. When networks are being designed, the exact form of these costs are often not well known. For example, the dark-fiber providers may regularly update their rental rates and the cost of optical equipment may be subject to negotiation. Moreover, the service providers may have to rent from different dark-fiber providers in different parts of the country and each may have different pricing strategies. Therefore, over time the fiber cost may vary nonuniformly from link to link.

Despite this, we do know that the number of fibers we use on a link must be at least the total number of demands routed through the link divided by the number of wavelengths per fiber. One robust way to ensure our network cost is low regardless of the exact cost structure is to minimize the ratio between the number of fibers actually used on the link and this lower bound.

In this paper we assume that the path followed by each demand is already fixed. Wavelength assignment is therefore the only problem. In an alternative formulation, routing and wavelength assignment could be performed simultaneously. However, in many practical situations arising in optical network design, routing is determined by some higher-level specifications (e.g. carriers may require min-hop routing, see [10, 8]). Hence, it is important to consider the wavelength assignment

problem in isolation. We also remark that once a demand is assigned a wavelength, which fiber on each link actually carries the demand is not an issue. This is because modern optical devices such as mesh optical add-drop multiplexers allow distinct wavelengths from different fibers to be multiplexed into a new fiber.

Fiber minimization with fixed routing is NP-hard on networks with general topology (by a simple reduction from graph coloring). In this paper we focus on upper and lower bounds for approximating the problem.

1.1 Problem definition and preliminaries

We now describe the basic version of our problem. We consider a network and a set of demands D where each demand i is routed on a given path P_i. We require that each demand is assigned a wavelength λ from the set $\{0, 1, \ldots, \mu - 1\}$. For each link e, if at most r_e demands passing through link e are assigned wavelength λ for each λ, then the number of fibers required on link e is r_e. If ℓ_e is the number of paths that pass through e, then $f_e = \ell_e/\mu$ is clearly a lower bound on r_e.

There are a number of distinct ways to define the objective function. For the reasons mentioned earlier we focus on a variant in which our goal is to minimize the maximum ratio between the number of fibers deployed on a link e and the corresponding lower bound f_e. (We mention some other variants in Section 5.) The problem may be formulated as an integer program. Let variable $C_{i,\lambda}$ indicate whether or not demand i uses wavelength λ. Our problem, which we call MIN-FIBER, can be written as follows for binary $C_{i,\lambda}$.

$$\min z$$

$$s.t.$$

$$\sum_{i:e\in P_i} C_{i,\lambda} \leq z \cdot f_e \qquad \forall e, \lambda \qquad (1)$$

$$\sum_{\lambda} C_{i,\lambda} = 1 \qquad \forall i \qquad (2)$$

We note that the linear relaxation of the above IP always has an optimal solution $z = 1$ and $C_{i,\lambda} = 1/\mu$ for all demands i and wavelengths λ.

1.2 Our Results

- We begin in Section 2 by presenting a negative result. We show that unless NP$\subseteq$ZPP there is no polynomial-time constant-factor approximation algorithm for the MIN-FIBER problem.

4

ZPP is the class of languages that can be recognized using a randomized algorithm that always gives the correct answer and whose expected running time is polynomial in the size of the input. Our result is based on the hardness result for graph coloring of Feige and Kilian [9].

- In Section 3 we further improve the lower bound. Unless NP $\subseteq$ ZPTIME($n^{\text{polylog } n}$), we show that there is no $\log^\gamma \mu$-approximation for any $\gamma \in (0, 1)$ and no $\log^\gamma m$-approximation for any $\gamma \in (0, 0.5)$ where m is the number of links in the network. ZPTIME($n^{\text{polylog } n}$) is the set of languages that have randomized algorithms that always give the correct answer and have expected running time $n^{\text{polylog } n}$.

- In Section 4 we turn our attention to positive results. In Section 4.1 we show that using randomized rounding we can obtain a solution in which the number of fibers required on each link e is at most $2f_e + 6(\log m + \log \mu)$. (All logarithms are to the base e.) This gives us an $O(\log m + \log \mu)$ approximation algorithm. We note that this algorithm can be derandomized using the standard method of conditional expectations.

  In Section 4.2 we apply the path-length rounding scheme of [12] to create a solution in which the number of fibers required on each link e is at most $f_e + D_{\max}$, where $D_{\max}$ is the length of the longest path in the network. This gives us an $O(D_{\max})$ approximation algorithm which is an improvement over the randomized rounding method when the paths are short.

  In Section 4.3 we use the Lovász Local Lemma to show that there always exists a solution in which the number of fibers required on each link e is $2f_e + 6(\log D_{\max} + \log \mu)$.

- In Section 5 we conclude by presenting two variants of the Min-Fiber problem and indicating which of our results still apply.

1.3 Previous Work

For the case in which the number of available wavelengths is not fixed, the problem of minimizing the number of wavelengths used has been much studied, e.g. [1, 3, 4, 20]. Some papers focus on common special topologies such as rings [14, 24] and trees [16, 15, 6]. The work listed here is by no means complete. A good survey on the subject can be found in [13].

Our problem of fiber minimization with a fixed fiber capacity has been introduced more recently. In [25, 17] the authors prove that coloring demands on a *line* only requires the minimum number of fibers per link, i.e. $\lceil f_e \rceil$ fibers on link e. This generalizes the well-known algorithm for coloring

interval graphs. In addition, [17] shows that the problem becomes NP-hard once the network topology is more complicated. The authors provide 2-approximation algorithms for rings and stars. Recent work on trees include [5, 7] and the results in [5] imply a 4-approximation.

In [2] a different objective is studied. The authors aim to minimize the total amount of fiber deployed when the demand routes are unknown *a priori*. They show that for a general network topology there is no $\log^{1/4-\gamma} m$ approximation for any $\gamma > 0$ unless NP $\subseteq$ ZPTIME($n^{\text{polylog } n}$). For total fiber minimization, both [2] and [22] offer approximation algorithms.

2 Basic Lower Bound

In this section we show that there is no constant factor approximation to the MIN-FIBER problem unless NP$\subseteq$ZPP. Our construction is based on hardness of approximation results for graph coloring. For any graph G we use $\chi(G)$ to denote the chromatic number of G and $\alpha(G)$ to denote the size of the maximum independent set of G. Throughout this section we shall use the terms "color" and "wavelength" interchangeably.

Feige and Kilian [9] construct a randomized reduction from 3SAT to graph coloring with the following properties. Given a 3CNF formula φ and a constant ε, they randomly construct an n-node graph G (where n is polynomial in the size of φ) such that,

- If φ is satisfiable then with probability 1, G can be colored with n^ε colors, i.e. $\chi(G) \leq n^\varepsilon$.

- If φ is not satisfiable then with high probability the maximum independent set in G has at most n^ε nodes, i.e. $\alpha(G) \leq n^\varepsilon$ with high probability. Note that since $\alpha(G) \cdot \chi(G) \geq n$ this immediately implies that $\chi(G) \geq n^{1-\varepsilon}$.

Feige and Kilian use this reduction to show that there is no $n^{1-\varepsilon}$ approximation for graph coloring unless NP$\subseteq$ZPP. We shall use it to show that for any constant c there is no c-approximation for MIN-FIBER unless NP$\subseteq$ZPP.

2.1 Constructing an instance of MIN-FIBER

We now demonstrate how to take a graph G and create an instance of MIN-FIBER on a network N. For each node v in G we have a demand d_v. The links in N consist of two sets E_1 and E_2. All links in E_1 are non-adjacent, i.e. no 2 links in E_1 have a vertex in common. The links in E_2 are used to connect up the links in E_1.

More precisely, for each clique Q in G with $c+1$ nodes we create a link e_Q in N and these links form the link set E_1. The demand d_v passes through e_Q for all $v \in Q$. If demand d_v has to pass through links $e_{Q_0}, \ldots, e_{Q_{z-1}}$ then there also exists a link $f_{v,j}$ in E_2 that connects the head of e_{Q_j} with the tail of $e_{Q_{j+1}}$. The full path of d_v is $e_{Q_0}, f_{v,0}, e_{Q_1}, \ldots, e_{Q_{z-2}}, f_{v,z-2}, e_{Q_{z-1}}$. We illustrate the construction of the network N from a graph G in Figure 1. The number of colors in our instance of MIN-FIBER is $\mu = n^\varepsilon$.
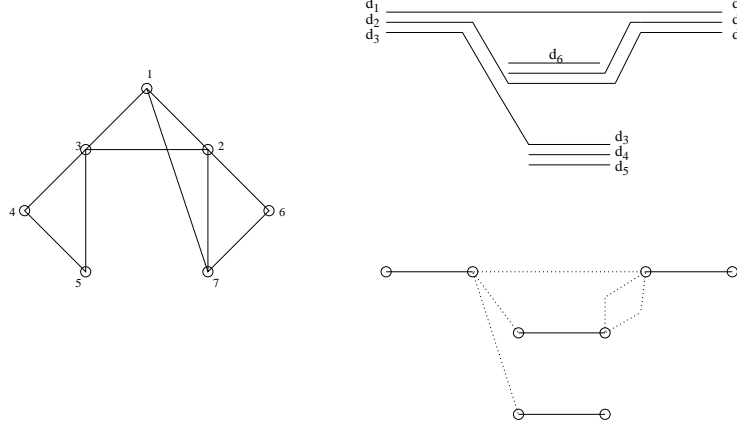

Figure 1: An example of the construction for $c = 2$. (Left) Graph G with 4 cliques of size 3. (Upper right) Demands and routes created from G. (Lower right) Network N, solid lines represent links in E_1 and dotted lines represent those in E_2.

2.2 Reduction from 3SAT to MIN-FIBER

Given a 3CNF formula φ we first choose a constant ε such that $\varepsilon < \frac{1}{c+1}$. We then construct a random n-node graph G according to the method of Feige and Kilian [9] for this parameter ε. Finally, we convert the graph G into an instance of MIN-FIBER on a network N according to the method of the previous section. Note that since c is a constant, the number of demands and links in N are both polynomial in n which is in turn polynomial in the size of φ.

Lemma 1 *If φ is satisfiable then with probability 1 the demands in N can be colored such that at most one fiber is required on each link. If φ is not satisfiable then with high probability, for any coloring of the demands in N, some link requires $c+1$ fibers.*

Proof: Suppose that φ is satisfiable. Then with probability 1 the graph G is colorable with $\mu = n^\varepsilon$ colors. For any such coloring, we color the demands in N such that demand d_v receives the same

7

color as node v. Clearly, for any clique Q in G and any color λ, there is at most one node in Q that receives color λ. Hence for any link e_Q in E_1, there is at most one demand passing through link e_Q that receives color λ. Therefore each link in E_1 requires only one fiber in order to carry all its demands. The links in E_2 have only one demand and so they trivially require one fiber only. Hence at most one fiber is required on any link in N.

To prove the other direction, suppose that φ is unsatisfiable. Then with high probability $\alpha(G) \leq n^\varepsilon$. Suppose for the purpose of contradiction that we can color the demands in N with $\mu = n^\varepsilon$ colors such that each link requires at most c fibers. This implies that for any link e_Q in E_1, not all the demands passing through e_Q receive the same color. Consider now the corresponding coloring of the nodes in G.[1] By the construction of our network N, for any clique Q with $c + 1$ nodes, not every node in Q receives the same color.

Let X be the induced subgraph of G on the set of nodes that constitutes the largest color class. We have just shown that X does not contain a clique of size $c+1$. Moreover, since X is contained in G, $\alpha(X) \leq \alpha(G) \leq n^\varepsilon$. Ramsey's theorem (see e.g. [18]) immediately implies that,

$$|X| \leq \binom{(\alpha(G)+1)+(c+1)-2}{(c+1)-1} \leq \alpha(G)^c. \tag{3}$$

Since X constitutes the largest color class and there are n^ε colors, $|X|n^\varepsilon \geq n$. Hence,

$$|X| \geq n^{1-\varepsilon}$$
$$\Rightarrow \quad \alpha(G)^c \geq n^{1-\varepsilon}$$
$$\Rightarrow \quad \alpha(G) \geq n^{\frac{1-\varepsilon}{c}} > n^\varepsilon,$$

since $\varepsilon < \frac{1}{c+1}$. This contradicts the fact that $\alpha(G) \leq n^\varepsilon$. $\qquad\square$

Theorem 2 *There is no c-approximation to* MIN-FIBER *for any constant c unless NP$\subseteq$ZPP.*

Proof: Suppose for the purpose of contradiction that C is a polynomial time c-approximation algorithm. We use this to construct a randomized algorithm B for 3SAT. For each instance φ, algorithm B creates a random graph G and then converts it to an instance of MIN-FIBER on a network N as described above. It then runs algorithm C on the instance of MIN-FIBER. If the solution returned by algorithm C is at most c then algorithm B returns "satisfiable", otherwise algorithm B returns "unsatisfiable". Lemma 1 implies that,

---

[1]Note that this is not necessarily a proper coloring. Some edges in G may have both endpoints assigned the same color.

- If φ is satisfiable then the optimal solution to the instance of MIN-FIBER is 1. Since algorithm C is a c-approximation algorithm, it returns a value of at most c. Therefore algorithm B outputs "satisfiable".

- If φ is unsatisfiable then with high probability the optimal solution to the instance of MIN-FIBER is $c+1$. Therefore algorithm C returns a solution of at least $c+1$. Therefore algorithm B outputs "unsatisfiable".

Note that algorithm B has one-sided error. Hence 3SAT$\in$coRP and so NP $\subseteq$ coRP. This implies RP $\subseteq$ NP $\subseteq$ coRP $\subseteq$ coNP which in turn implies NP = coNP = RP = coRP = RP$\cap$coRP = ZPP.

$\square$

3 Improved Lower Bound

In this section we derive more general hardness results by examining the construction of Feige and Kilian in more detail. In particular, given a 3CNF formula φ and a constant ε, they construct a random graph G on n nodes with parameters a, ρ, A and k. (As an aside, the parameters a, ρ and A are associated with a randomized Probabilistically Checkable Proof for NP and k is associated with a random graph product on a graph generated from the PCP. However, these interpretations are not important for our purposes.) The parameters are chosen so that the following relationships hold. More specifically, the parameters a and ρ are fixed to some constants such that Eq. (5) holds. The parameter A is polynomial in the size of φ and k is polylogarithmic in the size of φ. In particular, k is chosen sufficiently large such that Lemma 3 holds.

$$n \;=\; a^k \tag{4}$$

$$1 \;\geq\; \frac{\log a\rho}{\log a} \geq 1 - \varepsilon \tag{5}$$

$$A \;=\; \mathrm{poly}(|\varphi|) \tag{6}$$

$$k \;=\; \Theta(\log^{1/\delta}|\varphi|) \qquad \text{for any } \delta \in (0,1) \text{ of our choice} \tag{7}$$

Feige and Kilian show that graph G has the following properties.

1. If φ is satisfiable then with probability 1, G can be colored with $(1 + \log n)/\rho^k$ colors, i.e. $\chi(G) \leq (1 + \log n)/\rho^k$.

2. If φ is not satisfiable then $\alpha(G) \leq kA$ with high probability, which implies $\chi(G) \geq n/(kA)$.

From the graph G we construct an instance of MIN-FIBER in the same manner as in the previous section. We set,

$$\mu \;=\; (1 + \log n)/\rho^k \tag{8}$$

$$c \;=\; \log^{1-\delta} n \qquad \text{for any } \delta \in (0,1) \text{ of our choice} \tag{9}$$

Lemma 3 *We can choose $k = \Theta(\log^{1/\delta}|\varphi|)$ such that $\frac{(1+\log n)(kA)^c}{\rho^k} < n$.*

Proof: Immediate from the parameter definitions. $\qquad\qquad\square$

The following is analogous to Lemma 1.

Lemma 4 *If φ is satisfiable then with probability 1 the demands in N can be colored with μ colors such that at most 1 fiber is required on each link. If φ is not satisfiable then with high probability, for any coloring of the demands in N, some link requires $c+1$ fibers.*

Proof: For the case in which φ is satisfiable, the proof is identical to Lemma 1.

For the other direction, suppose that φ is unsatisfiable but we color the demands in N with μ colors such that each link requires at most c fibers. Consider the corresponding coloring of G and let X be the induced subgraph of G on the set of nodes that constitutes the largest color class. As in Eq. (3) in the proof of Lemma 1, $|X| \leq \alpha(G)^c$. By the construction of G, with high probability $\alpha(G) \leq kA$, which implies $|X| \leq (kA)^c$. Since X constitutes the largest color class and there are $\mu = (1 + \log n)/\rho^k$ colors, $|X| \geq n/\mu = n\rho^k/(1 + \log n)$. These inequalities imply that $(kA)^c \geq n\rho^k/(1 + \log n)$ which contradicts Lemma 3. $\qquad\qquad\square$

Note that since we have a link in the network N for each subset of $c+1$ nodes in G, the size of the instance of MIN-FIBER is polynomial in n^c. The following is analogous to Theorem 2.

Theorem 5 *Unless 3SAT has a randomized algorithm with expected running time $O(|\varphi|^{\mathrm{polylog}(|\varphi|)})$, there is no $\log^\gamma \mu$-approximation to MIN-FIBER for any $\gamma \in (0,1)$, and there is no $\Theta(\log^\gamma m)$-approximation for any $\gamma \in (0, 0.5)$. Here, μ is the number of colors per fiber and m is the number of links in MIN-FIBER.*

Proof: As in Theorem 2 we assume for the purpose of contradiction that C is a polynomial time c-approximation algorithm where c is defined in Eq. (9). From C we can construct a randomized algorithm B for 3SAT such that if φ is satisfiable then B outputs "satisfiable"; if φ is unsatisfiable then with high probability B outputs "unsatisfiable".

The correctness of B is identical to Theorem 2. The running time of B is $O(|\varphi|^{\text{polylog}(|\varphi|)})$ since both k and c are polylogarithmic in $|\varphi|$. Since $\mu \leq n$ and $m = O(n^c)$ we can show that $c > (\log \mu)^{1-\delta}$ and $c = \Omega((\log m)^{1-1/(2-\delta)})$. We note that B can give an incorrect answer with low probability. However, in the same way that NP$\subseteq$coRP implies NP$\subseteq$ZPP we can convert B into a randomized algorithm that always gives the correct answer and whose expected running time is $O(|\varphi|^{\text{polylog}(|\varphi|)})$.

$\square$

4 Upper Bounds

4.1 Randomized Rounding

Recall that the linear relaxation of our MIN-FIBER problem always has an optimal solution $z = 1$ and $C_{i,\lambda} = 1/\mu$ for all demands i and wavelengths λ. We adopt the technique of randomized rounding introduced in [19]. For each demand i we choose a number x_i uniformly at random in the range $[0,1]$. If $x_i \in [k/\mu, (k+1)/\mu)$ then we round $C_{i,\lambda}$ to 1 for $\lambda = k$ and round $C_{i,\lambda}$ to 0 for $\lambda \neq k$. After rounding the constraint in Eq. (2) still holds. We use the Chernoff Bound from Theorem 3.35 in [21] to see how much the constraint in Eq. (1) is violated. Let $\hat{C}_{i,\lambda}$ denote the rounded solution.

[**Chernoff Bound**] *If $X_1, \ldots, X_n$ are independent binary random variables where the expectation $x = E\left[\sum_i X_i\right]$, then it holds for all $\delta > 0$ that,*

$$\Pr\left[\ \sum_i X_i \geq (1+\delta)x\right] \leq e^{-\min(\delta^2, \delta) \cdot x/3}.$$

Lemma 6 *For a particular link e and wavelength λ,*

$$\Pr\left[\ \sum_{i: e \in P_i} \hat{C}_{i,\lambda} \geq 2f_e\ \right] \leq m^{-2}\mu^{-2} \qquad \text{if } f_e \geq 6(\log m + \log \mu),$$

$$\Pr\left[\ \sum_{i: e \in P_i} \hat{C}_{i,\lambda} \geq f_e + 6(\log m + \log \mu)\ \right] \leq m^{-2}\mu^{-2} \qquad \text{if } f_e < 6(\log m + \log \mu).$$

Proof: By definition, the expected value of $E[\hat{C}_{i,\lambda}]$ is $1/\mu$. Hence, $E\left[\sum_{i: e \in P_i} \hat{C}_{i,\lambda}\right] = f_e$. Note that for a fixed link e and wavelength λ, the rounding of variables $C_{i,\lambda}$ for demands i that go through e are independent events. We can therefore apply the Chernoff Bound.

If $f_e \geq 6(\log m + \log \mu)$, then

$$\Pr\left[\textstyle\sum_{i:e\in P_i} \hat{C}_{i,\lambda} \geq (1+1)f_e \right] \leq e^{-f_e/3} \leq e^{-2\log m - 2\log \mu} = m^{-2}\mu^{-2}.$$

If $f_e < 6(\log m + \log \mu)$, then

$$\Pr\left[\textstyle\sum_{i:e\in P_i} \hat{C}_{i,\lambda} \geq \left(1 + \tfrac{1}{f_e} \cdot 6(\log m + \log \mu)\right) f_e \right] \leq e^{-2\log m - 2\log \mu} = m^{-2}\mu^{-2}.$$

$\square$

By applying the union bound over all links and wavelengths, we obtain the following.

Theorem 7 *We can round the fractional optimal solution such that with high probability the number of fibers deployed on each link e is at most $2f_e + O(\log m + \log \mu)$. This implies an $O(\log m + \log \mu)$ approximation algorithm.*

We note that for large values of f_e the approximation ratio approaches 2. We also note that by using the slightly tighter Chernoff bound $Pr[\sum_i X_i \geq (1+\delta)x] \leq (e^\delta/(1+\delta)^{1+\delta})^x$, the approximation ratio can be marginally improved to $O(\frac{\log m}{\log\log m} + \frac{\log \mu}{\log\log \mu})$. However, for ease of exposition we ignore "log log" factors in this paper.

4.2 Path Length Rounding

The following lemma is a variation of the rounding theorem in [12], due to Karp, Leighton, Rivest, Thompson, Vazirani and Vazirani.

Lemma 8 *Let A be a $0/1$ matrix whose column sum is at most Δ; and let x be a vector of fractional indicator variables where $x^{k\ell} \in [0,1]$ for each component indexed by k and ℓ, and $\sum_\ell x^{k\ell} = 1$ for each k; and let vector $b = Ax$. We can compute a vector of integral indicator variables $\hat{x}$ in polynomial time such that*

1. *$\hat{x}^{k\ell} = \{0,1\}$ for each component indexed by k and ℓ, and $\sum_\ell \hat{x}^{k\ell} = 1$ for each k;*

2. *$\hat{b}_i - b_i \leq \Delta$ for each i where vector $\hat{b} = A\hat{x}$.*

Proof: The rounding algorithm proceeds in iterations. At the beginning of each iteration, let r be the number of remaining constraints in the system $Ax = b$, let p be the number of remaining constraints in the system $Cx = \vec{1}$, which represents $\sum_\ell x^{k\ell} = 1$ for each k; and let s be the number of remaining indicator variables in vector x. (We reindex the remaining constraints and variables

so that they are in the ranges of $1 \ldots p$, $1 \ldots r$ and $1 \ldots s$ respectively.) In brief, when $r + p < s$ we round one or more variables; when $r + p \geq s$ we eliminate a constraint in the system $Ax = b$. We repeat the iterations described below until either all variables are rounded or all constraints in $Ax = b$ are eliminated.

1. $r + p \geq s$: *Reduction in the number of constraints in the system $Ax = b$.* Let S be the set of $0/1$ vectors obtained by rounding each remaining $x^{k\ell}$ either up to 1 or down to 0. We show below that for some i, $1 \leq i \leq r$, $(Ay)_i \leq (Ax)_i + \Delta$ for all $y \in S$. This means we can drop the ith constraint in $Ax = b$ since $\hat{b}_i - b_i \leq \Delta$ no matter how the rounding is done. The existence of a redundant constraint i is shown below.

$$
\begin{aligned}
&\sum_{1 \leq i \leq r} \max_{y \in S} \left((Ay)_i - (Ax)_i \right) \\
=\ & \sum_{1 \leq i \leq r} \left(\sum_{j : a_{ij} = 1} a_{ij}(1 - x_j) \right) \\
=\ & \sum_{1 \leq j \leq s} \left(\sum_{i : a_{ij} = 1} a_{ij}(1 - x_j) \right) \\
\leq\ & \sum_{1 \leq j \leq s} \left(\Delta(1 - x_j) \right) \\
=\ & (s - p)\Delta \\
\leq\ & r\Delta.
\end{aligned}
$$

The first equality holds since A is a $0/1$ matrix and so the maximum value of $(Ay)_i - (Ax)_i$ is achieved when $y = \vec{1}$. The last equality follows from $\sum_j x_j = \sum_{1 \leq k \leq p} \sum_\ell x^{k\ell} = p$. (Here we index the vector x in two ways, x_j and $x^{k\ell}$.) We can now conclude that there exists an i, $1 \leq i \leq r$, such that $(Ay)_i \leq (Ax)_i + \Delta$ for all $y \in S$. We can identify and discard such a constraint since it satisfies the above property if and only if $(A\vec{1})_i - (Ax)_i \leq \Delta$.

2. $r + p < s$: *Reduction in the number of variables.* Let matrix M consist of 2 parts, the top p rows representing matrix A and the bottom r rows representing matrix C. Since matrix M is singular, there exists a non-zero vector z in the null space of M. Let $\lambda^* = \min\{\lambda > 0 : x + \lambda z \text{ has an integer component}\}$. We update x to $x + \lambda^* z$. This new vector satisfies $Ax = b$ and $Cx = \vec{1}$, each of its component remains between 0 and 1 and some component is integral. We now delete any integer component of x^{kl} from x, its corresponding column A_m from A, its corresponding column C_m from C, and update b to $b - x^{kl} A_m$. For a given k, if $x^{k\ell}$ is

13

integral for all ℓ we also remove the constraint $\sum_\ell x^{k\ell} = 1$ from $Cx = \vec{1}$.

It is easy to see that we maintain the following invariants at the beginning of each iteration.

- The remaining constraints and variables satisfy the systems $Ax = b$ and $Cx = \vec{1}$.

- When constraint i in $Ax = b$ is eliminated, then $\hat{b}_i - b_i \leq \Delta$ holds no matter how the remaining variables are rounded.

- For a given k, when $x^{k\ell}$ is rounded for each ℓ then $\sum_\ell \hat{x}^{k\ell} = 1$.

At the end of all iterations, if some variables are not rounded we consider all the unrounded $x^{k\ell}$. For each k, we round $x^{k\ell^*}$ to 1 for one arbitrary ℓ^* and the rest to 0. Since we maintain the invariants, we have therefore rounded all variables and our lemma holds. If every variable is rounded, then each remaining constraint i in $Ax = b$ is a null constraint due to the first invariant and $\hat{b}_i - b_i \leq \Delta$ holds automatically.

$\square$

It is easy to see that matrix A in the LP formulation of MIN-FIBER has 0/1 entries and its column sum is upper bounded by the longest path length plus 1. Recall we denote the longest path length by $D_{\max}$. By applying Lemma 8, we obtain,

Theorem 9 *We can round the fractional optimal solution such that the number of fibers deployed on each link e is at most $f_e + D_{\max}$.*

4.3 Rounding with the Lovász Local Lemma

In this section we use the Lovász Local Lemma to show that there exists a solution in which the number of fibers deployed on each link e is at most $2f_e + 6(\log D_{\max} + \log \mu)$. This bound is always an improvement over the bound of Theorem 7 and is an improvement over the bound of Theorem 9 in some situations.

Our argument is based on the rounding algorithm of Section 4.1. In Theorem 7 the additive term needs to be $O(\log m + \log \mu)$ due to a union bound that is applied to *all links* and all wavelengths. However, if two links carry no common demand then rounding the demands going through one link is independent of rounding those going through the other. In the following we break down the dependencies among the links and apply the Lovász Local Lemma [23, pages 57-58] to show that no link requires too many fibers.

[**Lovász Local Lemma**]   *Let $E_1, \ldots, E_n$ be a set of "bad events" each occurring with probability p and with dependence at most d (i.e. every bad event is mutually independent of some set of $n - d$ other bad events). If $4pd < 1$, then with probability greater than zero no bad event occurs.*

We say that a *bad event* $E_{\{e,\lambda\}}$ occurs on a link e and wavelength λ whenever too many demands in the rounded solution are routed through e on wavelength λ. (We specify "too many" later.) Two bad events $E_{\{e,\lambda\}}$ and $E_{\{e',\lambda'\}}$ are dependent only if there is a demand that goes through both e and e'. If we let $f_{\max}$ denote $\max_e f_e$ then $f_{\max} \cdot \mu$ is the maximum number of demands that go through any link. Since $D_{\max}$ is the maximum path length, each bad event is dependent on at most $D_{\max} \cdot f_{\max} \cdot \mu$ other bad events.

If $E_{\{e,\lambda\}}$ happens when more than $2f_e + 6(\log D_{\max} + \log \mu + \log f_{\max})$ demands are routed through e on wavelength λ, then $E_{\{e,\lambda\}}$ happens with probability $1/\text{poly}(D_{\max}, \mu, f_{\max})$ by a Chernoff bound argument as in Lemma 6. Since the dependency is at most $D_{\max} \cdot f_{\max} \cdot \mu$, no bad event happens with a positive probability.

In the following we provide a construction to ensure that $f_{\max}$ is small thereby reducing the dependency on a bad event and hence removing $\log f_{\max}$ from the rounding error. For each link e we create $\lceil x_e \rceil$ parallel links $c_e^1, c_e^2, \ldots, c_e^{\lceil x_e \rceil}$, where $x_e = f_e/6(\log D_{\max} + \log \mu)$. In addition to choosing a wavelength, each demand also chooses which parallel link c_e^k to go through for each link e along its given path. It is easy to see that the following fractional solution is feasible. As before, each demand is carried over all μ wavelengths, a $1/\mu$ fraction over each wavelength. In addition, each demand is only carried on *one* out of $\lceil x_e \rceil$ parallel links on e along its given path, i.e. a demand does not split among the parallel links. In the fractional solution the total fractional demand going through any parallel link over any wavelength is at most $6(\log D_{\max} + \log \mu)$. (One parallel link carries less demand if x_e is not integral.)

For the rounded solution, we define a bad event $E_{c,\lambda}$ on each parallel link c and wavelength λ. Suppose c carries $y \leq 6(\log D_{\max} + \log \mu)$ fractional demands over wavelength λ before rounding. Then the bad event $E_{c,\lambda}$ happens if c carries more than $y + 6(\log D_{\max} + \log \mu)$ integral demands over wavelength λ after rounding.

Lemma 10 *With positive probability no bad event happens.*

Proof: We first show that a particular bad event $E_{\{c,\lambda\}}$ happens with probability $1/\text{poly}(D_{\max}, \mu)$. By the construction of the fractional solution the expected number of demands that go through a

parallel link c over wavelength λ is at most $6(\log D_{\max} + \log \mu)$. By applying the Chernoff bound as in Lemma 6 we have

$$\Pr\left[\, E_{\{c,\lambda\}} \text{ happens} \,\right] \le e^{-2\log D_{\max} - 2\log \mu} = D_{\max}^{-2}\mu^{-2}.$$

The bad event $E_{\{c,\lambda\}}$ is dependent on another bad event $E_{\{c',\lambda'\}}$ only if there exists a demand that goes through both parallel links c and c'. By construction we know a demand does not split among parallel links. Moreover, each parallel link carries at most $6(\log D_{\max} + \log \mu)$ demands over each wavelength. Hence, each bad event depends on at most $D_{\max} \cdot 6(\log D_{\max} + \log \mu) \cdot \mu$ other bad events. The Lovász Local Lemma therefore implies our lemma. $\qquad\square$

Theorem 11 *There exists a rounding of the fractional solution such that the number of fibers deployed on each link e is at most $2f_e + 6(\log D_{\max} + \log \mu)$ where $D_{\max}$ is the maximum number of links along any demand path.*

Proof: If f_e is more than $6(\log D_{\max} + \log \mu)$ then $\lceil x_e \rceil$ parallel links are created for link e. Of these parallel links $\lceil x_e \rceil - 1$ are "full", i.e. each carries $6(\log D_{\max} + \log \mu)$ fractional demand over each wavelength, and one possibly carries less demand and is "not full". By Lemma 10, if no bad event happens then after the rounding the number of demands that go through each full parallel link over any wavelength is at most doubled and the number of demands that go through an unfull parallel link increases by at most $6(\log D_{\max} + \log \mu)$.

If f_e is at most $6(\log D_{\max} + \log \mu)$ then one parallel link is created for link e. By Lemma 10, if no bad event happens, then at most $6(\log D_{\max} + \log \mu)$ more demands go through e over any wavelength λ after the rounding.

Hence, the total number of demands that go through link e over any wavelength λ is at most $2f_e + 6(\log D_{\max} + \log \mu)$ after the rounding.

$\qquad\square$

5   Conclusions

In this paper we have presented positive and negative results for approximating the MIN-FIBER problem. We conclude by briefly discussing two variants of MIN-FIBER with different objective functions and seeing how our results apply. In the basic MIN-FIBER problem the objective is to

minimize the ratio between the number of fibers deployed on link e and the lower bound f_e. For the sake of comparison, we restate the integer program.

Basic Version

$$\min z$$

$$s.t. \quad \sum_{i:e\in P_i} C_{i,\lambda} \leq z \cdot f_e \qquad \forall e, \lambda$$

$$\sum_\lambda C_{i,\lambda} = 1 \qquad \forall i$$

In the first variant the new objective is to minimize the maximum, over all links e, of the number of fibers used on link e. As an integer program, this variant may be written as,

Variant 1

$$\min z$$

$$s.t. \quad \sum_{i:e\in P_i} C_{i,\lambda} \leq z \qquad \forall e, \lambda$$

$$\sum_\lambda C_{i,\lambda} = 1 \qquad \forall i$$

We note that the hardness results of Sections 2 and 3 follow through, e.g. there is no constant-factor approximation for Variant 1 of MIN-FIBER unless NP$\subseteq$ZPP. This is because for the case in which the 3CNF formula ϕ is satisfiable, all links in the network require exactly 1 fiber. We also note that a lower bound on the optimal value of this problem is $\max_e f_e$. By concentrating on the link e with the maximum value of f_e, the same approximation ratios proved in Section 4 also hold for this variant.

In the second variant we assume that we somehow know the cost per fiber on link e. We denote this cost by L_e. Our objective is to minimize the total cost of fiber needed to carry all the demands. We can formulate this variant as the following integer program.

Variant 2

$$\min \sum_e z_e L_e$$

$$s.t. \quad \sum_{i:e\in P_i} C_{i,\lambda} \leq z_e \qquad \forall e, \lambda$$

$$\sum_\lambda C_{i,\lambda} = 1 \qquad \forall i$$

Once again, the approximation ratios proved in Section 4 apply to this variant. However, our hardness results of Sections 2 and 3 no longer apply. Indeed, for the instances constructed in

our reductions, randomized rounding gives a constant factor approximation for the problem of minimizing the total fiber length.

We remark that all the results in this paper assume demand routes are given. If demand routes are not fixed, [2] shows that minimizing the objective of $\sum_e z_e$ has a logarithmic hardness bound.

Acknowledgement

The authors wish to thank Chandra Chekuri, Bruce Shepherd and anonymous referees for many helpful comments.

References

[1] A. Aggarwal, A. Bar-Noy, D. Coppersmith, R. Ramaswami, B. Schieber, and M. Sudan. Efficient routing and scheduling algorithms for optical networks. In *Proceedings of the 5th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 412 – 423, January 1994.

[2] M. Andrews and L. Zhang. Bounds on fiber minimization in optical networks with fixed fiber capacity. In *Proceedings of IEEE INFOCOM '05*, Miami, FL, March 2005.

[3] D. Banerjee and B. Mukherjee. A practical approach for routing and wavelength assignment in large wavelength-routed optical networks. *IEEE Journal on Selected Areas in Communications*, 14(5):903 – 908, 1996.

[4] I. Caragiannis, A. Ferreira, C. Kaklamanis, S. Perennes, and H. Rivano. Fractional path coloring with applications to WDM networks. In *Proceedings of the 28th International Colloqium on Automata, Languages, and Programming (ICALP '01)*, pages 732 – 743, 2001.

[5] C. Chekuri, M. Mydlarz, and F. B. Shepherd. Multicommodity demand flow in a tree. In *Proceedings of the 30th International Colloqium on Automata, Languages (ICALP)*, pages 410–425, 2003.

[6] T. Erlebach, K. Jansen, C. Kaklamanis, M. Mihail, and P. Persiano. Optimal wavelength routing in directed fiber trees. *Theoretical Computer Science*, 221(1–2):119 – 137, 1999.

[7] T. Erlebach, A. Pagourtzis, K. Potika, and S. Stefanakos. Resource allocation problems in multifiber WDM tree networks. In *Proceedings of the 29th International Workshop on Graph Theoretic Concepts in Computer Science*, pages 218 – 229, 2003.

[8] T. Erlebach and S. Stefanakos. On shortest-path all-optical networks without wavelength conversion requirements. In *Proceedings of the 20th International Symposium on theoretical aspects of computer science (STACS)*, pages 133–144, 2003.

[9] Uriel Feige and Joe Kilian. Zero knowledge and the chromatic number. *Journal of Computer System Sciences*, 57(2):187–199, 1998.

[10] S. Fortune, W. Sweldens, and L. Zhang. Line system design for DWDM networks. In *Proceedings of the 11th International Telecommunications Network Strategy and Planning Symposium (Networks)*, Vienna, Austria, 2004.

[11] Alan Gnauck. Digital transmission. Post-OFC2004 Reviews, 2004.

[12] R. M. Karp, F. T. Leighton, R. L. Rivest, C. D. Thompson, U. V. Vazirani, and V. V. Vazirani. Global wire routing in two-dimensional arrays. *Algorithmica*, 2:113 – 129, 1987.

[13] R. Klasing. Methods and problems of wavelength-routing in all-optical networks. Technical Report CS-RR-348, Department of Computer Science, University of Warwick, 1998.

[14] V. Kumar. Approximating circular arc coloring and bandwidth allocation in all-optical ring networks. In *Proceedings of the International Workshop on Approximation Algorithms for Combinatorial Optimization (APPROX '98)*, pages 147–158, 1998.

[15] V. Kumar and E. Schwabe. Improved access to optical bandwidth in trees. In *Proceedings of the 8th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 437 – 444, January 1997.

[16] M. Mihail, C. Kaklamanis, and S. Rao. Efficient access to optical bandwidth in trees. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*, pages 548 – 557, 1995.

[17] C. Nomikos, A. Pagourtzis, and S. Zachos. Routing and path multi-coloring. *Information Processing Letters*, 80(5):249–256, 2001.

[18] G. Polya, R. Tarjan, and D. Woods. *Notes on Introductory Combinatorics, Progress in Computer Science, No. 4*. Birkhauser, Boston, Basel, Stuttgart, 1983.

[19] P. Raghavan and C.D. Thompson. Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica*, 7:365 – 374, 1991.

[20] P. Raghavan and E. Upfal. Efficient routing in all-optical networks. In *Proceedings of the 26th Annual ACM Symposium on Theory of Computing*, pages 134–143, 1994.

[21] C. Scheideler. *Probabilistic Methods for Coordination Problems.* Habilitation thesis, Paderborn University, 2000.

[22] B. Shepherd and A. Vetta. Lighting fibers in a dark network. *Journal on Selected Areas in Communications*, 22(9):1583 – 1588, 2004.

[23] J. Spencer. *Ten Lectures on the Probabilistic Methods.* Capital City Press, Philadelphia, Pennsylvania, 1994.

[24] G. Wilfong and P. Winkler. Ring routing and wavelength translation. In *Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 333 – 341, January 1998.

[25] P. Winkler and L. Zhang. Wavelength assignment and generalized interval graph coloring. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 830–831, January 2003.
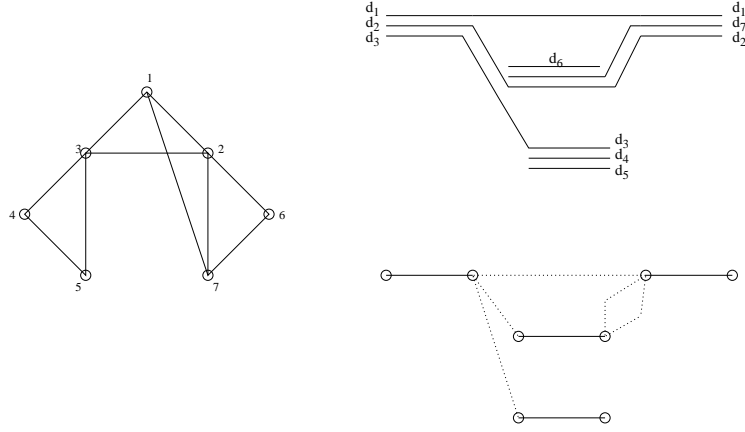
Figure 1: An example of the construction for $c = 2$. (Left) Graph G with 4 cliques of size 3. (Upper right) Demands and routes created from G. (Lower right) Network N, solid lines represent links in E_1 and dotted lines represent those in E_2.