

Island Hopping and Path Colouring with applications to WDM Network Design*

Andrew McGregor
andrewm@cis.upenn.edu

Bruce Shepherd
bshep@research.bell-labs.com

Abstract

Two benefits of optical communication are (a) the possibility of using a single fiber optic cable to carry multiple signals simultaneously if each signal uses light of a distinct wavelength, and (b) the decreased latency for signals that avoid expensive optical-electric-optical (OEO) conversions. This paper studies two natural graph problems that arise when trying to capitalize on these benefits.

Given a supply network we must route multiple demands. We install fibers in the network, each of which may carry λ signals of different wavelengths. We decouple the general problem into the following two sub-problems.

Fiber minimization: We wish to light the cheapest set of fibers, l_e fibers for edge e , such that each demand can be routed monochromatically with at most l_e routes of the same wavelength using an edge e . We first generalize the work of Winkler and Zhang which gave an optimal solution for line networks. Secondly, we present a 3.55-approximation for the single-source problem.

Hop minimization: In general it is not possible to inter-connect all fibers incident on a node. If a signal travels between two fibers that are not connected then an OEO conversion, or *hop*, is necessitated. We wish to minimize the number of hops required when routing the demands subject to the constraint that at most c fibers are connected to each other at each node. For $c = 2$, we show an $O(\log^{1-\epsilon} n)$ -approximation hardness result that complements the $O(\log n)$ -approximation algorithm due to Anshelevich and Zhang. For directed supply graphs we show an $O(n^{1-\epsilon})$ -approximation hardness result. However, for acyclic supply graphs we present an $O(\sqrt{n})$ -approximation algorithm. For strongly connected digraphs, our $O(n^{1-\epsilon})$ -hardness continues to hold for $c = 2$; however, for $c = 3$, we give a $O(\log n)$ -approximation algorithm.

1 Introduction

In this paper we focus on issues that arise in the design of networks in which light travelling along optical fibers is used to transmit information. While such networks have their origins back in the sixties, it has only been in the last twenty years that the technology is becoming fully realized. One development that capitalizes on the inherent advantages of

optical communication is that of wavelength division multiplexing (WDM.) This allows multiple signals using different wavelengths to be transmitted simultaneously on a single optical fiber. Systems that used two different wavelengths first appeared around 1985. Today, *dense* wavelength division multiplexing can use up to 160 different wavelengths thereby providing a substantial increase in the possible transmission rate.

The study of WDM networks includes numerous natural and compelling questions, some of which have already been addressed in the literature. These include minimizing the number of colours needed such that every demand can be routed monochromatically (e.g., [1, 10, 25, 27, 31, 32, 39]). More recently, attention has turned to minimizing the cost of fibers, each being able to carry a fixed number of wavelengths, that can be used to achieve monochromatic routing (e.g., [13, 4, 26, 35, 40]). In another set of problems the possibility of using wavelength converters has been explored (e.g., [39, 23, 21, 12]). These devices allow the wavelength of a signal to be changed mid-route. Unfortunately they are still prohibitively expensive. Other difficult optimization problems arise when building, configuring and operating WDM networks [11]. Very recently, the limits of optically switching has given rise to graph decomposition problems [6].

In the next section we formally describe the parameters of designing WDM networks. Subsequently, we specify the precise problems that we address, the existing results and our improvements. Before we do that however, we give an informal description of two special cases of the problems that we address. We do this in such a way as to highlight that, while motivated by problems of practical interest, there are inherently fundamental problems which may prove pleasant and interesting in and of themselves.

The first problem is to decompose a graph into simple paths T_1, T_2, \dots such that between any two nodes there exists a path P such that P intersects at most t of the paths T_1, T_2, \dots . One can imagine

*This work was partly funded by ONR basic research grant N00014-05-1-0256 and was carried out while the first author was visiting Bell Labs.

numerous pseudo-motivations for such a problem. For example, we may wish to design a subway system (given a set of existing tunnels) such that the number of times that the average commuter needs to change lines is minimized. The second problem starts with a graph $G = (V, E)$ and a set of $s_i t_i$ pairs. For each pair $s_i t_i$ we pick a path in G between the nodes s_i and t_i and assign one of λ colours to this path. We seek to minimize the maximum (over all colours) multiplicity of an edge among the set of all paths with a given colour.

1.1 Design of WDM Networks In general, our optical network design problems take the following inputs. We are given an undirected *supply network* $G = (V, E)$ where E is the set of *edges* uv whose endpoints are the nodes u and v . We are also given a *demand graph* $H = (V, F)$. For our purposes we assume each demand f requests one unit of flow between its endpoints. (We also consider the directed supply graphs which consist of nodes and arcs, denoted (u, v) having tail u and head v .) Edge e of the supply network may have an associated cost c_e . Each edge represents point-to-point links, and within each link there may be many fibers. Typically, there are already many fibers in a link, but one pays a price c_e each time a fiber is lit (which amounts to installing the appropriate line terminating equipment at the link endpoints.) Let l_e be the number of fibers lit on link e . A fiber has a capacity λ which is the number of wavelengths that can simultaneously be used on the fiber.

We now outline the main constraints imposed on routing in an WDM optical network; for a more thorough background the reader is referred to [36].

1. *Routing constraints:* We route each demand in H along a simple path in G .
2. *Capacity constraints:* The number of paths through an edge e does not exceed $l_e \lambda$.
3. *Optical constraints:* Each path is assigned a wavelength $w(P) \in [\lambda]$. Two paths with the same wavelength may not use the same fiber although they may use two different fibers on the same link. At each node there exist switching devices, called *reconfigurable optical add-drop multiplexers* (ROADMs), that optically connect signals on incident fibers. A c -arm ROADM partitions the incident fibers into sets of size at most c . A signal travelling between two fibers may be switched optically if the two fibers are in the same partition. If a signal is not opti-

cally switched then an *optical-electrical-optical* (OEO) conversion is required. Colloquially we say that such a signal must “hop” to the next fiber. Ideally c would be large enough such that signals could be optically switched between all incident fibers. This is called *Full Wavelength Selective Switching* (WSC). In practice, $c = 2$ or 3.

To take advantage of the benefits of optical communication we try to route demands *transparently* as much as possible. Transparency refers to signals that stay in the optical domain without any conversion to electronics. Requiring that each path is monochromatic is a necessary part of ensuring transparency since converting between wavelength requires electronics. Note that without the optical constraints our design problem becomes the same routing and capacity allocation problem as for traditional networks. It is quite easy to construct examples in which the optical constraints necessitate an increase in the cost of fiber required to route the demands. In [35] this factor increase is termed the *transparency gap*. The problem has been studied from an empirical perspective in [29]. They find that there is often very little transparency gap as long as demand is sufficiently large. Theoretical evidence for this is further provided in [35].

1.2 Our Results and Previous Work

Fiber Minimization (Full WSC-Switching) Our first results are applicable when we assume the availability of full WSC-Switching. In this case our sole goal is to minimize the total cost of fiber purchased. This model was introduced by Lowe [28] and has been widely studied in various communities, see for instance [35, 40, 5, 13]. See [22] for a survey.

Problem: MINFIB

Input: Supply network G , demand graph H , costs $(c_e)_{e \in E}$, fiber capacity λ

Solution: A multiple l_e of fibers bought at link e , simple routing paths $(P_h)_{h \in E(H)}$ and an assignment of one of λ colours to each P_h such that the number of paths of the same colour using any edge e is at most l_e

Goal: Minimize $\sum_{e \in E} l_e c_e$

Winkler and Zhang [40] showed that MINFIB can be solved optimally if G is a path. Our first re-

sult is to extend this to a more general class of directed tree networks. For general graphs, Andrews and Zhang [5] gave a polynomial-time $O(\log n)$ approximation. Moreover, they show that this problem has no $O(\log^{1/4-\epsilon} n)$ -approximation for any $\epsilon > 0$, unless $NP \subset ZPTIME(n^{O(\log^{1/\epsilon} n)})$. In contrast, we show that the single-source multicommodity WDM flow problem has a constant approximation in general graphs. Both of our positive results rely on the *integer decomposition property* of certain polyhedra.

Hop Minimization (Partial WSC-switching)

We now consider the more realistic possibility that full WSC switching is not possible and instead we assume we only have access to c -arm ROADMs. A model that captures this was introduced by Anshelevich and Zhang [6] for $c = 2$. This model assumes that each link consists of a single lit fiber of infinite capacity. The focus thus turns away from the cost of lighting fiber, to minimizing hop costs. To formally define the problem we must first introduce the notion of a *transparent island*.

DEFINITION 1. (TRANSPARENT ISLAND) A proper c -arm transparent island is a set of edges of a subtree of the supply graph that has maximum degree c . A improper c -arm transparent island is defined by a partitioning of the edges incident to each node into sets of size at most c . Let P_v be the partitioning of $\delta(v)$, the edges incident to v . Let $s_v(e) \in P_v$ denote the subset of edges in the same partition as the edge e . We say that $e \sim e'$ if $s_v(e) = s_v(e')$ for some $v \in V$. An improper transparent island is an equivalence class of transitive closure of \sim . A 2-arm transparent island is called a line system.

This definition generalizes that given in [6] for the case of 2-armed ROADMs: a proper (improper) transparent island is a simple path (trail).

DEFINITION 2. (HOPS) Given a set of transparent islands, T_1, \dots, T_l and a path P denoted as a sequence of edges (e_1, e_2, \dots, e_p) , let $t(e)$ denote the transparent island including edge e . We define the number of hops taken by P as $h_{T_1, \dots, T_l}(P) = 1 + |\{i \in [p-1] : T(e_i) \neq T(e_{i+1})\}|$.

Problem: MINHOP $_c$
Input: Supply network G , demand graph H
Solution: c -arm transparent islands T_1, \dots, T_l and simple routing paths $(P_h)_{h \in E(H)}$
Goal: Minimize average number of hops, i.e. $\frac{1}{|E(H)|} \sum_{h \in E(H)} h_{T_1, \dots, T_l}(P_h)$

The main result of [6] relevant here is a pleasingly simple polytime algorithm for MINHOP $_2$ in undirected graphs that achieves a $O(\log n)$ approximation. They also show that MINHOP $_2$ is hard to 2-approximate. Our first result for MINHOP is to improve this hardness result to $\Omega(\log^{1-\epsilon} n)$ and thereby establish the surprising (to us) result that the decomposition algorithm of [6] is close to optimal. We then extend the treatment to directed graphs. Here we show another negative result – it is hard to $O(n^{1-\epsilon})$ -approximate MINHOP $_2$. This is the case even if the supply graph is strongly connected. Note that an $O(n)$ approximation is trivially achieved by routing demands along the shortest paths and designating each edge as a transparent island. On the positive side, we give a polytime $O(\sqrt{n})$ -approximation algorithm when the supply graph is a directed acyclic graph.

In directed graphs, 3-armed ROADMs yield substantial improvements. For example, when the supply graph is strongly connected, we present an algorithm that achieves an $O(\log n)$ approximation for MINHOP $_3$, an exponential improvement over the best possible for 2-arm ROADMs. We also investigate the benefits of 3-arm ROADMs in terms of the cost of optimal solutions. For undirected graphs, we show that if the supply graph G is 3-node connected planar then $\text{MINHOP}_3(G) = 1$ whereas if G is only 2-node connected planar $\text{MINHOP}_3(G)$ can be $\Omega(\log n)$. The algorithmic results are summarized in Table 1.

	2-arm ROADMs		3-arm ROADMs
	Algorithm	Hardness	Algorithm
a)	$O(\log n)$ [6]	$\Omega(\log^{1-\epsilon} n)$	$O(\log n)$ [6]
b)	$O(n)$	$\Omega(n^{1-\epsilon})$	$O(\log n)$
c)	$O(n^{1/2})$	$\Omega(\log n)$	$O(n^{1/2})$

Table 1: Approximation Ratios - Hardness and Guarantees. Rows a), b) and c) correspond to undirected graphs, strongly-connected graphs, and DAGs.

Towards Joint Optimization of Fiber Layout, Routing and ROADM Configuration

The MINHOP $_c$ problem is of interest when we have already decided on where to light fibers. We are often in the case where each link has a large supply of fibers and we can always choose to light extra fibers in order to reduce OEO costs. In Section 4 we make a step towards jointly optimizing fiber design, routing, and ROADM configurations. In particular, we consider the model where fibers again have infinite

capacity, however, we may light as many fibers as needed in a link. We also consider a related fundamental problem of laying down a small number of line systems to reduce the diameter of an acyclic graph to polylogarithmic size. We describe connections to the shortcutting problem and give initial results and describe some possible directions.

2 Fiber Minimization

In this section, we consider the network flow and design problems in WDM networks with full wavelength-selective switching, MINFIB. We show “fractional implies integral” results where given a standard network flow, we are able to transform it into a WDM flow. Underlying both results is an associated class of polyhedra have the *integer decomposition property* (IDP). A polyhedron P has the IDP if for any $x \in P$, and any integer k such that kx is integral, we have $kx = \sum_i^k x^i$ where each x^i is an integral vector in P . Trivially, any polyhedra with the IDP is integral, i.e., each of its vertices is integral. We make use of the following result of Baum and Trotter [8].

THEOREM 2.1. (BAUM, TROTTER [8]) *A matrix A is totally unimodular if and only if $\{x : Ax \leq b, x \geq 0\}$ has the integer decomposition property for every integer vector b .*

Multicommodity WDM Flow in Trees In this section we assume that the supply network is a directed tree $T = (V, A)$ and that the demand graph $H = (V, F)$ has the property that for each $h = (u, v) \in F$, the unique path in T between u and v is directed from u to v . This generalizes fiber minimization on a line as considered by Winkler and Zhang [40].

For any capacity vector $u : A(T) \rightarrow \mathbf{Z}$, we define a polyhedron $P(T, H, u)$ in \mathbf{R}^F of “feasible” flows for our demands: $\{x : B \cdot x \leq u, 0 \leq x \leq 1\}$. Here B denotes the $0, 1$ matrix with a row for each arc of T , and a column for each $h = (u, v) \in F$. There is a 1 in B_{ah} if and only if a lies on the path P_h . The matrix B is a network matrix cf. [34] and hence is totally unimodular. Therefore, the block matrix $[B^T I]^T$ is totally unimodular too. It follows that for any integral u , $P(T, H, u)$ has the integer decomposition property [8].

THEOREM 2.2. *There is a polytime algorithm that solves MINFIB on directed tree instances.*

Proof. It is sufficient to show that in any directed tree instance of MINFIB, if the load on any edge e is at

most $l_e \lambda$, then we may also assign valid wavelengths to the demands F . That is, there is a partition $F = F_1 \cup \dots \cup F_\lambda$ such that the load of any arc of T under F_i is at most l_e . Consider the polyhedron $P(T, H, l)$. By assumption, the vector $\frac{1}{\lambda} \chi^F$ lies in this polyhedron, where χ^F denotes the $0, 1$ incidence vector of the set of edges F . Since P has the IDP, we may decompose the vector $\chi^F = \sum_i^\lambda \chi_i^F$, where $\chi_i^F \in P(T, H, l)$. Since the load of any edge under F_i , is $\leq l_e$, we may route these demands monochromatically.

Single-Source Multicommodity WDM Flows

Now we assume that we have an arbitrary directed supply graph $D = (V, A)$ but our demand graph $H = (V, F)$ is a single-source instance, i.e., the demands are of the form (s, t_i) for some s and t_1, \dots, t_k . We start with another “fractional implies integral” result for WDM flows.

THEOREM 2.3. *If there is a single source network integral flow such that the load on any arc e is at most $l_e \lambda$, then there is a wavelength-labeled such flow that uses at most l_e fibers on each edge e .*

Proof. Consider the node-arc incidence matrix A for D : for each node v and arc a , $D_{v,a}$ is 1 if a has tail v , -1 if it has head v , and otherwise it is 0. Since A is totally unimodular, we have that $P = \{f : Af \leq l, f \geq 0\}$ has the IDP. Thus if f' is a flow vector whose total load on any edge e is at most λl_e , then $\frac{f'}{\lambda} \in P$, and so once again we may decompose f' as $\sum_{i=1}^\lambda f^i$ where each f^i is an integral vector P . As the flow paths determined by any f^i place a load of at most l_e on any link, there are enough fibers to route these demands on the same wavelength.

The above theorem can also be proved using an efficient combinatorial algorithm by employing standard flow computations. To see this, we call a flow k -ready (for l) if the total load on any link e is at most kl_e . We call an arc e *critical* for a k -ready flow f if its load is $(k-1)l_e + r$ for some $r > 0$. We also call r its *requirement*. If e is not critical, then its requirement is 0. We create a typical auxiliary graph for flows as follows. For each arc with a flow of x on it, we include the reverse arc with capacity x . For each arc with a load of $(k-1)l_e - r$ for some $r > 0$, we include the forward arc with capacity r . In addition, from each terminal t_i , receiving d_i units of flow from s , we include an arc from t_i back to s with capacity 1 (if there are d_i demands terminating we give it capacity d_i .) Finally, for each critical arc we include this arc with a lower bound equal to its requirement.

One easily checks (we omit the details) that if f is a k -ready flow, then Hoffman's Circulation Conditions [20] hold in this auxiliary graph. Moreover, if f' is such a circulation, the flow f'' , produced in the standard symmetric difference manner from f, f' , is $(k - 1)$ -ready.

Note that the circulation f' determines the set of flow paths to be routed on wavelength k . We remove these demands, and then repeat with the flow f'' . The circulation can be obtained by repeated solving of a shortest path problem. This method also gives a simple algorithm for the computation of so-called p -multi-route flow vectors [2] of size say kp by taking $l_e = 1$ and then finding a k -ready flow of size kp .

We turn to the problem of computing the optimal choice for lighting fibers. Theorem 2.3 implies that it is sufficient to determine values l_e such that the minimum cut condition holds. That is, for each set S containing s , $\lambda \cdot l(\delta^+(S)) = \lambda \cdot \sum_{e \in \delta^+(S)} l_e \geq f(S)$ where $f(S) = \sum_{t_i \in V-S} 1$. This single-sink edge installation problem is hard, even when each fibre has infinite capacity. For directed graphs, there is an approximation preserving reduction from set cover. However, for undirected graphs there is a factor 3.55 approximation algorithm [18]. Combining this with Theorem 2.3 yields the following result.

THEOREM 2.4. *In polytime we can 3.55-approximate MINFIB on single source instances.*

Theorem 2.4 may be of practical use for empirical studies, by giving formulations for multicommodity WDM flows based on setting up n single source multicommodity flow vectors, as opposed to $\Omega(n^2)$ source-destination flow vectors. Capacity sharing between distinct source flows remains a complication.

3 Hop Minimization

2-arm ROADMS (Undirected Supply Graphs) In [6] it was shown that routing along an arbitrary spanning tree of G and then setting ROADMS optimally yielded an $O(\log n)$ approximation to MINHOP₂. They also prove that it is NP-hard to approximate MINHOP₂ by better than a factor 2. Given that the routing is reasonably arbitrary it is perhaps surprising that their algorithm is almost optimal. Indeed, we show that, unless $NP \subset DTIME(2^{O(\log^{1/\epsilon} n)})$, it is hard to approximate MINHOP₂ to within a factor of $\log^{1-\epsilon} n$. Our reduction is from LONGPATH, the problem of finding long paths in 3-regular Hamiltonian graphs. It is known that, unless $P = NP$, LONGPATH cannot be approximated to within a constant factor [9].

THEOREM 3.1. *For $\epsilon > 0$, unless $NP \subset DTIME(2^{O(\log^{1/\epsilon} n)})$, there exists no polytime algorithm that $O(\log^{1-\epsilon} n)$ -approximates MINHOP₂ for undirected graphs.*

Proof. Take an instance L of LONGPATH on $t + 1$ nodes. We construct an instance of MINHOP₂ on $n = 2^{O(\log^{1/\epsilon} t)}$ nodes such that the optimal cost is 1 but finding a solution with cost less than $\Omega(\frac{\log n}{\log t})$ yields a path of length greater than $t/50$ in L . The construction is as follows.

Let the nodes of L be $\{u_1, u_2, \dots, u_{t+1}\}$. Replace each node u of L by nodes $V(u) = (u^a, u^b, u^c)$ where each node becomes an endpoint for one of the edges originally incident on u . We call these nodes the “split-triple” for u . We also add two nodes $I(u) = (v_u^a, v_u^b)$ which we call the “interfacing-pair” for u . Next add a complete set of edges between the nodes in $V(u)$ and $I(u)$. Call the resulting graph L' . Let $G' = (V', E')$ be a balanced t -ary tree on n' nodes. Call the root node r . Double each edge e to produce an “in-and-out” edge pair (e^a, e^b) . Replace each node v of degree $2(t + 1)$ by a copy of L' such that the two nodes of each interfacing-pair in L become the endpoints of an in-and-out pair of edges. We call the resulting graph $G = (V, E)$. Let the demand graph H consist of all edges from r to degree two nodes of G (the nodes from G' that were originally leaves.) See Figure 1. By analogy to the original tree structure of G' we consider copies of L' to be super-nodes in a tree-like structure and define “parent”, “height”, “sub-tree”, “leaves” in the natural way. For example, the “depth” of G is $d = \log_t(n'(t - 1) + 1) - 1$.

Solve MINHOP₂ on G, H . Consider a copy X of L' and the $t + 1$ in-and-out pairs of edges, $E_X = (e_1^a, e_1^b, e_2^a, e_2^b, \dots, e_{t+1}^a, e_{t+1}^b)$ where (e_1^a, e_1^b) is the in-and-out pair of edges that join this copy of L' to its parent copy. Our solution has partitioned these edges into line systems. Let P_a be the line system including e_1^a and let $E_a = E_X \cap P_a$. Similarly define P_b and E_b .

We claim that $|E_a| > t/10$ gives rise to a simple path of length at least $t/50$ in L . To see this, consider the sequence $S = (V(u_{i_1}), V(u_{i_2}), \dots, V(u_{i_l}))$ of split-triples visited by P_a . The sequence must be of length at least $|E_a|/2 - 1$. Furthermore all terms of the sequence occur only once with the possible exceptions of $V(u_{i_1})$ and $V(u_{i_l})$ which may appear twice. Hence there is a (consecutive) sub-sequence S' of S of length $(|E_a|/2 - 2)/2$ that gives rise to simple path in L' of length $t/50$.

Hence assume that $|E_a|$ (and $|E_b|$ by an identical argument) is $\leq t/10$. Since every demand in the subtree rooted at v must use either e_1^a or e_1^b and each




Figure 1: Reduction from LONGESTPATH to MINHOP₂ (undirected).

is on the same line system as at most $t/10 - 2$ of the edges in $E_X \setminus \{e_1^a, e_2^b\}$, at least $(1 - 2(t/10 - 2)/t) > 1/2$ of the leaves require to hop between line systems at X . Hence the average number of hops required by a demand is $\Omega(d)$. But since L is Hamiltonian it is straightforward to see that each demand can be routed with only the initial hop.

2-arm ROADMS (Directed Supply Graphs) We now turn our attention to directed graphs. We first give the proverbial bad news: it is hard to approximate MINHOP₂ up to a factor $O(n^{1-\epsilon})$ even in a strongly connected graph. Our reduction is from the problem 2DIRPATH. This decision problem is known to be NP-complete [14].

Problem: 2DIRPATH
Input: A directed graph $H = (V, A)$, distinct vertices $x_1, x_2, y_1, y_2 \in V$.
Question: Are there edge-disjoint directed paths, from x_1 to y_1 from y_1 to y_2 , in H ?

Note that an $O(n)$ approximation for MINHOP₂ is trivially achieved by routing each demand along the shortest path and letting each edge be a separate transparent island.

THEOREM 3.2. *For any constant $\epsilon > 0$, unless $P = NP$, there exists no polynomial algorithm that $O(n^{1-\epsilon})$ -approximates MINHOP₂ for strongly connected graphs.*

Proof. Consider a graph on nodes $v_1, v_2, \dots, v_{n'}$. There are arcs (v_i, v_{i+1}) for $i = 1, \dots, n' - 1$ and arcs (v_i, v_{i-3}) for $i = n', n - 2, n - 4, n - 6, \dots, 4$. We wish to route 1 unit of flow from v_1 to $v_{n'}$ and one unit from $v_{n'}$ to v_1 . This costs $\Omega(n')$ since at each v_i ($1 < i < n'$) there is at least one hop necessitates are every (there is a path along all edges of the graph because removal of any edge disconnects v_1 from $v_{n'}$ or vice versa. Hence we pay 1 at each degree 3 node.) We now prove

a hardness result by embedding numerous copies of an instance $(H, in_1, in_2, out_1, out_2)$ of 2DIRPATH as follows. For each $i = n' - 3, n' - 5, \dots$, we remove nodes v_i and v_{i+1} and incident arcs. We add H and arcs $(v_{i+3}, in_1), (v_{i-1}, in_2), (out_1, v_{i-2})$ and (out_2, v_{i+2}) . Let the resulting graph be G . Now, H needs to be of size at least n^ϵ for some constant ϵ if we are to argue that there are no polynomial (in n) algorithm for solving 2DIRPATH. Hence $n' = n^{1-\epsilon}$. Clearly if $(H, in_1, in_2, out_1, out_2)$ is a YES instance of 2DIRPATH then $MINHOP_2(G) = 1$. However if it is a NO instance, then $MINHOP_2(G) = \Omega(n^{1-\epsilon})$. There is a slight detail in that the hard instances of 2DIRPATH as constructed in [14] are not strongly connected. However, this problem is easily overcome by adding some extra nodes in the graph that will not aid the routings in the case of a NO instance (details omitted.)

We now give some good news: if the supply network is acyclic then we can do better. We start with a preliminary lemma.

LEMMA 3.1. *We call a length l sequence of natural numbers a_1, a_2, \dots, a_l boosted if $a_i \neq a_{i+1}$ for all $i < l$, and for any pair $i < j$ with $a_i = a_j$, we have $a_p \leq a_i$ for all $p \in (i, j)$. Then the length of a boosted sequence $a = (a_1, a_2, \dots, a_l)$ with a support of size k is at most $2k$.*

Proof. Without loss of generality the support, $\{a_i : i = 1, 2, \dots, l\}$, consists of distinct elements $1, 2, \dots, k$. The proof is by induction on k , the case $k = 1$ being trivial. If there is no repeated element in the sequence, then the length is at most k , so assume that q is the minimum repeated number, and let i_1, i_2, \dots, i_n be the only positions such that $a_{i_1} = a_{i_2} = \dots = a_{i_n} = q$. For each $j = 1, 2, \dots, n - 1$, let $I_j = \{a_p : i_j < p < i_{j+1}\}$. By definition and minimality of q , the elements of each I_j are distinct. Moreover, these sets are themselves disjoint, for otherwise some a_{i_j} is surrounded by a pair of equal,

but smaller elements in the sequence, contradicting the fact that the sequence was boosted. Define $r = \sum_{j=1}^{n-1} |I_j|$. It follows that the total length of the sequence is at most $2r$ plus the length of the sequence $a' = (a_1, a_2, \dots, a_{i_n}, a_{i_n+1}, \dots, a_l)$. The sequence a' has a support of size $k - r$ and so by induction (it is also boosted), it has length at most $2(k - r)$, and the result follows.

THEOREM 3.3. *There is a polytime $O(\sqrt{n})$ -approximation algorithm for MINHOP_2 for DAGs.*

Proof. The algorithm routes demands in two phases. In the first phase we find paths which we subsequently use to define our line systems. We find these paths in iterations. In iteration i , we have already found some paths P_1, P_2, \dots, P_{i-1} for $i - 1$ of the demands. We define G^i to be the graph obtained by adding a clique of edges, each with a label j , between the nodes $V(P_j)$ for each $1 \leq j \leq i - 1$. We call these *short arcs* and assign them length $1/n^2$. The original edges of G are given length 1. If each demand is at distance at most $3\sqrt{n}$ apart in G^i , then we stop this phase of the algorithm. Otherwise, we find a demand with a shortest path R_i in G^i whose length is at least $3\sqrt{n}$. By minimality and the fact that G is acyclic, note that any such path contains at most 2 nodes from any path P_j with $1 \leq j \leq i - 1$, otherwise we can “shortcut” (and any such pair is adjacent in G^i .) Therefore R_i contains at least $3\sqrt{n} - 2(i - 1)$ nodes not contained in any P_j , $1 \leq j \leq i - 1$. Let P_i be the directed (simple by acyclicity) path obtained by replacing any short arcs in R_i , say (x, y) of label $j < i$, by the subpath of P_j from x to y . Note that we route $O(\sqrt{n})$ demands since if we kept on finding demands at a distance at least $3\sqrt{n}$ then after at most \sqrt{n} stages, every node would appear in some path. Let P_{i^*} be the last route determined in this phase.

In the second phase, we create our line systems using the P_i 's as follows. Starting with $H^0 = G$, for each $i = 1, 2, \dots, i^*$, we set each maximal subpath of P_i in H^{i-1} to be a line system (with label i). We then delete all edges of P_i to obtain H^i . Finally, edges not contained in $\bigcup_{1 \leq i \leq i^*} P_i$, form single edge line systems. The result follows from the next claim.

Claim. For any pair of nodes u, v , there is a $u - v$ path that hops onto line systems created up to the H^i stage at most $O(i)$ times.

Proof. [Proof of Claim] Let R be the shortest $u - v$ path in G^i and let P be the path in G formed by replacing any short arcs by the corresponding subpath of one of the routes determined in the first phase. The

path P can be decomposed into fresh segments and non-fresh segments as $R = F_0, N_0, F_1, N_1, \dots, F_p, N_p$ (where F_0 and N_p are potentially empty.) Each line system used in P has an associated label $j \leq i$. We consider the sequence $\bar{a} = (a_1, a_2, \dots, a_l)$ of these labels as we traverse P from left to right - we only list a label once for a contiguous set of edges in P that all belong to the same line system. We ignore edges in free segments as they have not been assigned a label. We claim that \bar{a} is boosted and hence, by Lemma 3.1, the path that hops onto line systems created up to the H^i stage at most $O(i)$ times.

Firstly the only way in which $a_j = a_{j+1} = a$ for some $1 \leq j \leq l - 1$ is if the last arc (v_1, v_2) of N_q and the first arc (v_3, v_4) of N_{q+1} (for some q) are from line systems with label a . But in this case the path R in G^i was not the shortest - it incurred a cost at least 1 (the edges in the segment N_{q+1} each cost 1) to join v_2 to v_3 whereas there was an edge (v_2, v_3) of length $1/n^2$ in G^a since (v_1, v_2) and (v_3, v_4) are both in P^a .

Secondly suppose that $a_j = a_{j'} = a$ with $j < j'$, yet there is some $j'' \in (j, j')$ such that $a_{j''} < a_j$. By minimality of R in G^i , we know that the whole sub-path of P “from” a_j to $a_{j'}$ must be contained in P_{a_j} . For this to be not so, R would have had to share more than two nodes with P_{a_j} . But then after stage a_j , the segment associated with $a_{j''}$ should have become a line system of label at most a_j . This is a contradiction.

This claim implies two things. For demands that do not get routed in the first phase, it immediately implies that there is a path that uses at most $O(\sqrt{n})$ existing systems, but in addition its total length in G^{i^*} is at most $3\sqrt{n}$, and hence it uses at most $O(\sqrt{n})$ edges not in some P_i . Therefore, this is routed with $O(\sqrt{n})$ hops. Second, any P_i determined in the first stage, alternates between segments of previously created line systems, and “fresh” subpaths. Moreover, the number of fresh subpaths is clearly bounded by the number of line systems, and by the claim the latter is $O(i)$. Since each fresh subpath becomes a new line system, P_i traverses at most $O(\sqrt{n})$ line systems from end to end.

A special case of Lemma 3.3 will show that the above upperbound on $\text{MINHOP}_2(G)$ is essentially tight.

LEMMA 3.2. *There exists a directed acyclic graph G such that $\text{MINHOP}_2(G) = \Omega(\sqrt{n})$.*

While it would appear that the construction used to prove the above lemma could be used to construct




Figure 2: Reduction from set cover.

a hardness of approximation result for acyclic graphs using a reduction from 2DIRPATH, this is not the case. This is because, for acyclic graphs, 2DIRPATH is not NP-hard. The following hardness result uses a reduction from SETCOVER [33].

THEOREM 3.4. *Unless $P = NP$, there is no polytime algorithm that $O(\log n)$ -approximates MINHOP_2 on directed acyclic graphs.*

Proof. Our hardness result is based on a reduction from set cover. Consider a universe $U = \{u_1, \dots, u_n\}$ and a set of sets $\mathcal{S} = \{S_1, \dots, S_m\}$. Wlog. we assume that $m \leq n$. The set cover problem to find a subset $\mathcal{S}' \subseteq \mathcal{S}$ of minimum cardinality such that for each $u_j \in U$ there is a $S_j \in \mathcal{S}'$ such that $u_j \in S_i$. It is known that this is $\Omega(\log n)$ hard to approximate.

We construct a directed network as follows. Consider a graph as pictured in Figure 2. Note that all vertical edges are directed upwards and all horizontal edges are directed to the right. We associate elements of U with the nodes at the bottom level. Nodes in the level above correspond to the sets in \mathcal{S} . For each u_j in S_i we direct an edge from u_j to S_i . For each S_i we have a path P_i of length $2B + 1$ up to a node d . For each P_i there are B separate paths that share an edge with P_i (as shown.) There are $n + mB$ demands. For each u_j we need to route a demand to d - we call these *set cover demands*. For each P_i there are B paths intersecting with P_i 's upon which we need to route a demand - we call these *nuisance demands*. (In the diagram $B = 1$.)

Consider a solution of the routing that contains paths that use s nodes $\mathcal{S}' \subseteq \mathcal{S}$. Now if $S_i \in \mathcal{S}'$, the remaining routes incur a cost of exactly $2B$ along path P_i . This is because for each intersection of one of the B paths, the cost is at least two and if we place the line system on P_i then we incur 2 at each intersection. For $S_i \notin \mathcal{S}'$ no cost is incurred along P_i .

Note that the set of nodes used in the second level corresponds to a valid set cover. Hence, once we have decided to use \mathcal{S}' the optimal cost is,

$$(n - s) + 2sB + mB + n$$

where $n + mB$ correspond to costs to jump onto the first line segment - call these *start-up costs*.

Consider t copies of this graph. We link together the paths that intersect P_i in each copy. Note that the route of the nuisance demands are fixed. Furthermore the path routing any set cover demands remains in a single copy. Consider choosing the same set \mathcal{S}' in each - there is no reason to do otherwise. Hence the total cost is,

$$(n - s)t + 2sBt + mB + nt$$

Crucial is the fact that there are still only mB nuisance demands. Let $t = n^2$ and $B = n^3$ and this becomes $s(2n^5 - n^2) + 2n^3 + n^4$. For large values of n this is proportional to s . Hence if we find an α approximation to our problem then we get an $\alpha + o(1)$ approximation to the set cover instance. We deduce that our problem is $\Omega(\log n)$ hard.

Multi-arm ROADMs: Benefits of an Extra Arm

We start with a result that shows the considerable utility of 3-arm ROADMs. This result contrasts the $\Omega(n^{1-\epsilon})$ -hardness result of Theorem 3.2.

THEOREM 3.5. *There is a polytime $O(\log n)$ -approximation algorithm for MINHOP_3 (strongly connected).*

Proof. Pick an arbitrary node r . Consider an in-arborescence and an out-arborescence. As in [6], we may define transparent islands by doing a caterpillar decomposition of the in-arborescence. Subsequently, each demand can route to r and incur at most $O(\log n)$ hops. Similarly, for the out-arborescence. Now consider any node v and suppose the in-arborescence joined its unique outgoing link a with an incoming link b . Similarly, suppose that the unique incoming link a' of the out-arborescence was paired with some outgoing link b' . If a, b, a', b' are disjoint or consist of precisely the same pair of links then there is no problem. Otherwise, without loss of generality $a = b'$ and $b \neq a'$, and use a 3-armed ROADM to pair up a, a' and b at v . Note that we may easily ensure routing paths are simple by removing any cycles in the routing paths - the number of hops required can only go down.

For DAGs the situation is more complex.

LEMMA 3.3. For $c \geq 2$, there exists a DAG G with $\text{MINHOP}_c(G) = \Omega(n^{1/m})$ but $\text{MINHOP}_{c+1}(G) = 1$.

It is known that $\text{MINHOP}_2(G) = O(\log n)$ for all undirected graphs G [6]. They also note that for a balanced binary tree T (with demands between the root and all leaves) $\text{MINHOP}_2(T) = \Omega(\log n)$. Clearly $\text{MINHOP}_3(T) = 1$. However there also exist graphs, for example the balanced tertiary tree, such that $\text{MINHOP}_3(\cdot) = \Omega(\log n)$.¹ There is one large class of graphs where 3-arm ROADMs give a strong improvement. This follows from a classical result due to Barnette [7].²

LEMMA 3.4. (BARNETTE [7]) *Every 3 node connected planar graph has a spanning tree of max degree at most 3.*

The consequence is that, for every 3 node connected planar graph G , $\text{MINHOP}_3(G) = 1$. Furthermore, Barnette’s proof gives a polytime algorithm that will find the necessary spanning tree. For “all-pairs” instances, $\text{MINHOP}_2(G) > 1$ if G is non-Hamiltonian (and sufficiently large). Also of note is the result by Fürer and Raghavachari [16] that gives a polytime algorithm that finds a degree 3 spanning tree in a Hamiltonian graph G . The bounded degree spanning tree approach seems promising in general [24]. In contrast to Barnette’s algorithm, the next result shows that if the supply graph is only 2-node connected planar then $\text{MINHOP}_3(\cdot)$ can still be $\Omega(\log n)$.

THEOREM 3.6. *There exists a planar 2-node connected graph G such that $\text{MINHOP}_3(G) = \Omega(\log n)$.*

We next highlight the significance of restricting our attention to simple routing paths. Our proof uses a special case of a result due to Mader [30]. The *splitting off* of edges uv and uv' at u consists of deleting these edges and adding the edge vv' (possibly resulting in a double edge). The *local connectivity* between a pair of nodes x, y is the minimum cut separating x, y and is denoted $\lambda(x, y)$.

THEOREM 3.7. *For all 2 edge connected graphs G , $\text{MINHOP}_3(G) = 1$ if non-simple routings are allowed.*

Proof. The proof uses Mader’s Theorem [30] that says that for any node s in an undirected graph, that

¹The analysis in [6] can be extended to show $\text{MINHOP}_c(\cdot) = O(\log_c n)$ but for small values of c the change is negligible.

²For comparison, all 4-node-connected planar graphs are Hamiltonian by a celebrated theorem of Tutte [38].

is not incident to a bridge or of degree 3, there is some pair of edges incident to s that can be split off without changing any local connectivity $\lambda(x, y)$ ($x, y \neq s$).

If all nodes are of degree at most 3, we are done. So consider any node v of degree 4 or more. One may argue (we omit the details) that it is possible to “split off” a pair of edges $uv, u'v$ from v to obtain a new 2-edge-connected graph. By induction, the smaller graph has a single transparent island. If it does not use the new edge, uu' , then it is also a transparent island for G , and if it does, replacing uu' by the $uv, u'v$ and a new ROADM at v , gives the desired island.

4 Future Directions

The above two sections have dealt with fiber minimization and hop minimization separately. The simplifying assumptions in each problem, *full WSC* and *existing unbounded capacity single fibers* respectively, are unfortunately somewhat orthogonal. In this final section, we propose a hybrid model that more completely captures the full network design problem.

WDM network design is essentially bi-criteria: we wish to pay little for fibers and minimize the total OEO conversion. In the new model, rather than assuming that there already exists lit fibers of unbounded capacity as in Section 3, we insist that each fiber used is purchased at the appropriate cost. We still assume that each fiber has unbounded capacity but allow more than one fiber to be purchased in a link. The utility of having multiple unbounded capacity fibers in a link can be seen by considering the problem of using 2-arm ROADMs for routing between the leaves of a star graph with 3 leaves.

There are numerous ways in which the bi-criteria problem can be approached. For example, we can define an (α, β) -*approximation* as a solution which pays at most α times the cost of the minimum Steiner forest (the cheapest set of fibers that allow each demand to be routed) and allows the average (over demands) number of hops to be β . For undirected graphs, using existing results on approximating the minimum Steiner forest [3, 17] and ideas from the previous sections, it is straightforward to find algorithms that achieve a $(4, 1)$ -approximation (if paths are allowed to be non-simple) or a $(2, 2 \log n)$ -approximation (if paths must be simple.)

Alternatively we could seek to minimize fiber costs given an upper bound on the hops necessary for any demand. This problem has similarities to the problem of “shortcutting” a directed graph. In shortcutting, we are given a directed graph, and wish

to add a small number of arcs (u, v) in the closure (shortcuts), so that the maximum distance (between related pairs) becomes polylogarithmic. Thorup [37] proved that in planar graphs, this was always possible with $O(m)$ shortcuts and he conjectured this to hold in general acyclic graphs. This was recently disproved [19]. In our version, the shortcuts are more like *turnpikes*, in the sense that nodes in the middle of the turnpike may use it to decrease the shortest path lengths to other nodes. Unfortunately even the weaker analogue of Thorup's conjecture is false. The proof uses the same (lovely) counterexample of [19].

Acknowledgments: The authors thank Chandra Chekuri, Steve Fortune, Luis Goddyn, Peter Winzer, and Francis Zane for stimulating discussions. We are grateful to Lisa Zhang for interesting conversations, and especially for driving the many fundamental research problems in this area, during her work on the OCUBE optical design tool [15].

References

- [1] A. AGGARWAL, A. BAR-NOY, D. COPPERSMITH, R. RAMASWAMI, B. SCHIEBER, AND M. SUDAN, *Efficient routing in optical networks.*, J. ACM, 43 (1996), pp. 973–1001.
- [2] C. AGGARWAL AND J. ORLIN, *On multi-route maximum flows in networks*, Networks, 39 (2002), pp. 43–52.
- [3] A. AGRAWAL, P. N. KLEIN, AND R. RAVI, *When trees collide: An approximation algorithm for the generalized steiner problem on networks.*, SIAM J. Comput., 24 (1995), pp. 440–456.
- [4] M. ANDREWS AND L. ZHANG, *Wavelength assignment in optical networks with fixed fiber capacity.*, in ICALP, 2004, pp. 134–145.
- [5] ———, *Bounds on fiber minimization in optical networks.*, in INFOCOM, 2005.
- [6] E. ANSHELEVICH AND L. ZHANG, *Path decomposition under a new cost measure with applications to optical network design.*, in ESA, 2004, pp. 28–39.
- [7] D. BARNETTE, *Trees in polyhedral graphs.*, Canadian Journal of Mathematics, 18 (1966), pp. 731–736.
- [8] S. BAUM AND L. E. TROTTER JR., *Integer rounding and polyhedral decomposition for totally unimodular systems.*, Optimization and Operations Research, (1978), pp. 15–23.
- [9] C. BAZGAN, M. SANTHA, AND Z. TUZA, *On the approximation of finding a(nother) hamiltonian cycle in cubic hamiltonian graphs.*, J. Algorithms, 31 (1999), pp. 249–268.
- [10] I. CARAGIANNIS, A. FERREIRA, C. KAKLAMANIS, S. PERENNES, AND H. RIVANO, *Fractional path coloring with applications to WDM networks.*, in ICALP, 2001, pp. 732–743.
- [11] C. CHEKURI, P. CLAISSE, R. ESSIAMBRE, S. FORTUNE, D. KILPER, W. LEE, K. NITHI, I. SANIEE, B. SHEPHERD, C. WHITE, AND G. WILFONG., *Design tools for transparent optical networks.*, Bell Labs Technical Journal.
- [12] L.-W. CHEN AND E. MODIANO, *Efficient routing and wavelength assignment for reconfigurable WDM networks with wavelength converters.*, in INFOCOM, 2003.
- [13] T. ERLEBACH, A. PAGOURTZIS, K. POTIKA, AND S. STEFANAKOS, *Resource allocation problems in multifiber WDM tree networks*, in WG 2003, 2003, pp. 218–229.
- [14] S. FORTUNE, J. E. HOPCROFT, AND J. WYLLIE, *The directed subgraph homeomorphism problem.*, Theor. Comput. Sci., 10 (1980), pp. 111–121.
- [15] S. FORTUNE, W. SWELDENS, AND L. ZHANG, *Line system design for DWDM networks.*, in Networks 2004, 2004, pp. 315–320.
- [16] M. FÜRER AND B. RAGHAVACHARI, *Approximating the minimum-degree steiner tree to within one of optimal.*, J. Algorithms, 17 (1994), pp. 409–423.
- [17] M. X. GOEMANS AND D. P. WILLIAMSON, *A general approximation technique for constrained forest problems.*, SIAM J. Comput., 24 (1995), pp. 296–317.
- [18] R. HASSIN, R. RAVI, AND F. S. SALMAN, *Approximation algorithms for a capacitated network design problem.*, Algorithmica, 38 (2003), pp. 417–431.
- [19] W. HESSE, *Directed graphs requiring large numbers of shortcuts.*, in SODA, 2003, pp. 665–669.
- [20] A. HOFFMAN, *Some recent applications of the theory of linear inequalities to extremal combinatorial analysis*, in Combinatorial Analysis (Proceedings of Symposia in Applied Mathematics, Volume X(R. Bellman and M. Hall, Jr, eds.), Providence, R.I., 1960, American Mathematical Society, pp. 113–127.
- [21] X. JIA, D.-Z. DU, X.-D. HU, H. HUANG, AND D. LI, *Placement of wavelength converters for minimal wavelength usage in wdm networks.*, in INFOCOM, 2002.
- [22] R. KLASING, *Methods and problems of wavelength-routing in all-optical networks*, tech. rep., Department of Computer Science, University of Warwick, 1998.
- [23] J. M. KLEINBERG AND A. KUMAR, *Wavelength conversion in optical networks.*, J. Algorithms, 38 (2001), pp. 25–50.
- [24] J. KÖNEMANN AND R. RAVI, *A matter of degree: Improved approximation algorithms for degree-bounded minimum spanning trees.*, SIAM J. Comput., 31 (2002), pp. 1783–1793.
- [25] V. KUMAR, *Approximating circular arc colouring and bandwidth allocation in all-optical ring networks.*, in APPROX, 1998, pp. 147–158.
- [26] V. KUMAR AND A. RUDRA., *Approximation algorithms for wavelength assignment*, in UW CSE Tech Report No. 2004-08-01., August 2004.

- [27] V. KUMAR AND E. J. SCHWABE, *Improved access to optimal bandwidth in trees.*, in SODA, 1997, pp. 437–444.
- [28] E. LOWE, *Performance and dimensioning of optical transport networks*, PhD thesis, University of Essex, 1995.
- [29] E. LOWE, B. SHEPHERD, AND N. WALKER, *Routing and configuration of static path optical transport networks*, IEE Electronic Letters, 32 (1996), pp. 1913–1914.
- [30] W. MADER., *A reduction method for edge-connectivity in graphs.*, Annals of Discrete Mathematics 3, 22 (1978), pp. 145–164.
- [31] M. MIHAIL, C. KAKLAMANIS, AND S. RAO, *Efficient access to optical bandwidth - wavelength routing on directed fiber trees, rings, and trees of rings.*, in FOCS, 1995, pp. 548–557.
- [32] P. RAGHAVAN AND E. UPFAL, *Efficient routing in all-optical networks.*, in STOC, 1994, pp. 134–143.
- [33] R. RAZ AND S. SAFRA, *A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP.*, in STOC, 1997, pp. 475–484.
- [34] A. SCHRIJVER., *Theory of Linear and Integer Programming.*, John Wiley and Sons, 1987.
- [35] B. SHEPHERD AND A. VETTA., *Lighting fibers in a dark network.*, Journal on Selected Areas in Communication, 22 (2004), pp. 1583–1588.
- [36] T. STERN AND K. BALA., *Multiwavelength Optical Networks: A Layered Approach.*, Addison-Wesley, Reading, MA, 1999.
- [37] M. THORUP, *On shortcutting digraphs.*, in WG, 1992, pp. 205–211.
- [38] T. TUTTE, *A theorem on planar graphs.*, Trans. Amer. Math. Soc., 82 (1956), pp. 99–116.
- [39] G. WILFONG AND P. WINKLER, *Ring routing and wavelength translation*, in SODA, 1998, pp. 333–341.
- [40] P. WINKLER AND L. ZHANG, *Wavelength assignment and generalized interval graph coloring.*, in SODA, 2003, pp. 830–831.

A Omitted Proofs

Proof. [Proof of Theorem 3.6] We define a class of graphs as follows. Given an edge e , *splitting* e refers to making some k parallel copies of e and then subdividing each copy to become a path of length three. In each such path $P = e_1, e_2, e_3$ we refer to e_2 as the *new* edge, and e_1, e_3 as the *attachment* edges for e_2 .

We create a series parallel graph from an edge $e = uv$ as follows. Apply a k -splitting to e . Recurse on each of the new edges produced. One can thus view this as a k -ary tree in a sense, and all demands will originate at the top-level i.e., at the original edge e , and must be routed to the “leaf” edges. In any splitting operation, the k -branches have an associated 2-edge cut determined by the previous level’s “non-new” edges e_1, e_3 . We consider any such 2-cut associated with a pair of edges e_1, e_3 a path $P = e_1, e_2, e_3$, where the new edge e_2 goes onto be a “complicated” subgraph – called the *bubble graph* for this path – but e_1, e_3 remain untouched. We argue that if $k \geq 13$, then some fraction of e_2 ’s children (i.e., subgraphs obtained from splitting e_2) are not in the same component as either e_1 or e_3 . Hence some fraction ($\frac{1}{13}$ th) must incur an OEO cost when they cross this cut $\{e_1, e_3\}$ to get to their destination.

To see this, we first note the following. Consider a $K_{2,13}$ where the small side consists of two nodes u, v and X is the large side. Each of u, v is then made adjacent to a new pair of degree 1 nodes u', v' . Suppose F is an island in this graph containing uu' and for which there exist a simple path from u' to any node of F . Then $|V(F) \cap X| \leq 6$. First note that there is at most one ROADM, call it R , at u . For otherwise, at least one of these is not incident to the edge uu' . But then any node of X that is adjacent, in F , to this second ROADM would only have a non-simple path to u' in F . Since R is 3-armed, at most two nodes of X are adjacent to R . So let $x \in X$ be adjacent to R in F , that is, $xu \in F$. We claim that at most two other nodes of X are both non-adjacent to R , but have a simple path to x in $F - xu$. Any path between x and such nodes must use an edge xv incident to some ROADM R' at v . If there are 3 or more such nodes, then one of these is not incident to R' . But then since it is not adjacent to R , its path to v in F must traverse v twice, a contradiction. It now follows, that in any such $K_{2,13}$, there is at least one node that is not in the same island as either uu' or vv' .

Now consider any 2-cut at some level. We claim that at least one of the next-level bubblegraphs has

none of its nodes in the same transparent island as either e_1, e_3 (if everyone has a simple path across the 2-cut). To see this, it is enough to note, that any such configuration would yield such a set of islands in the $K_{2,13}$ graph obtained by shrinking the bubble graphs.

THEOREM A.1. *Adding $O(m)$ optical fibers to a graph is not sufficient to ensure that each demand can be routed with at most poly-logarithmic OEO conversions.*

Proof. Hesse [19] showed that there exists a graph with the following properties:

1. Vertices are partitioned into $2n^{1/17}$ layers L_1, L_2, \dots such that each arc (u, v) has $u \in L_i$ and $v \in L_{i+1}$ for some i .
2. There are $m = \Theta(n^{19/17})$ arcs.
3. There are $\Theta(mn^{1/17})$ pairs of vertices $p = \{u, v\}$ with $u \in L_1$ and $v \in L_{2n^{1/17}}$ such that
 - (a) There is a unique path in G between each pair of vertices.
 - (b) The paths between any two pairs share at most one edge.

Consider such a graph. Consider the modest goal of ensuring that the number of OEO conversions between each pair is at most $n^{1/17}$. Note that before any fibers are laid there are effectively $2n^{1/17} - 1$ OEO conversions required. Let there be a set S of k optical fibers. Let h_{sp} be the reduction of OEO conversions required in the path between pair p by fiber s . Let the unique routing between pair p be $r(p)$. To reduce the number of OEO conversion to at most $n^{1/17}$ it is necessary that for each path $p \in P$, $\sum_{s \in S} h_{sp} \geq n^{1/17}$. And hence that $\sum_{s \in S, p \in P} h_{sp} \geq \Theta(mn^{1/17})n^{1/17}$. Unfortunately we have that $\sum_{p \in P} h_{sp} \leq 2n^{1/17}$. This is because for all $s \in S, p, p' \in P$ we have $h_{sp} = |r(p) \cap s| - 1$, $|r(p) \cap r(p')| = 1$ and therefore $\sum_{p \in P} h_{sp} \leq |s| - 1$. We therefore deduce that $\sum_{s,p} h_{sp} \leq k2n^{1/17}$. Hence we need that $k = \Omega(mn^{1/17})$.

Proof. [Proof of Lemma 3.3] Construct the graph with nodes corresponding to vectors (a_1, \dots, a_c) where $a_i \in [n^{1/c}]$. Call these nodes the *awkward* nodes. Add nodes of the form (d_1, \dots, d_c) where $d_j = 0$ or $n^{1/c} + 1$ for some j and $a_i \in [n^{1/c}]$ for $i \neq j$. Call these the *source/sink* nodes. There is an arc in the supply graph of the form $((a_1, \dots, a_c), (b_1, \dots, b_c))$ if, for some j , $b_j = a_j + 1$ and




Figure 3: DAG G with $\text{MINHOP}_2(G) = \Omega(\sqrt{n})$.

$b_i = a_i$ for $i \neq j$. The demand arcs are of the form $((d_1, \dots, 0, \dots, d_c), (d_1, \dots, n^{1/c} + 1, \dots, d_c))$. Now replace each awkward node v of this graph by an “a” node v_a and a “b” node v_b and an arc (v_a, v_b) such that all arcs into v are now into v_a and arcs out from v are now out from v_b . It is straight forward to see that $\text{MINHOP}_c(\cdot)$ for this supply graph and demand graph is $\Omega(n^{1/c})$. See Figure. 3 for a diagram for the case $c = 2$ where $k = n^{1/2}$.