Lighting Fibers in a Dark Network

F. B. Shepherd and A. Vetta

***Abstract*— We consider the problem of network design in transparent, or clear channel, optical networks associated with wavelength division multiplexing (WDM). We focus on the class of traffic engineering models known as *routing, wavelength and capacity assignment* (RWAC) problems. Here, in contrast to traditional networks, traffic flow paths must also be assigned an end-to-end wavelength. This additional requirement means that there can be an increased cost associated with optimal capacity allocations for such *WDM-flows*. In general this can be arbitrarily worse than traditional network designs. We argue that in order to evaluate the benefit of different switch technologies, a good benchmark is to measure the increase in costs purely in terms of link capacity; we call this the *cost of transparency*. Experimental research shows that this cost is small in multi-fiber networks with modest switching functionality at the nodes. We present theoretical justification for why this occurs, and prove that in multi-wavelength multi-fiber transparent networks the cost of transparency all but disappears if there is moderate traffic load. Our arguments are based on efficient heuristics that may also be useful for more complex network optimizations.**

***Index Terms*— WDM, multi-fiber network, transparent optical network, linear programming, randomised rounding.**

I. Introduction

It is widely believed that the continuing growth in data traffic will impose a demand for an efficient optical core network to flexibly handle large numbers of high bandwidth connection requests across it. These connections are embodied as *wavelength paths* and each such all-optical path consists of a sequence of fibers that together provide very high bandwidth between the path's endpoints. A leading technology for offering even more bandwidth is that of *wavelength division multiplexing* (WDM) which enables multiple wavelengths to be supported on each fiber. In a clear-channel or *transparent* optical network, wavelength conversion is not possible along any wavelength path, and so there is a requirement that each communication, or *flow* path be assigned a single wavelength for transmitting data along the path. This *wavelength assignment* is performed in addition to the traditional constraints of network routing, where flow paths are computed to carry pre-specified traffic demands subject to the constraint that no link in the network is loaded beyond its capacity. Since each fiber supports at most one connection at a given wavelength, two flow paths assigned a common wavelength must necessarily be fiber-disjoint. This extra requirement makes the optimal cost of reserved link capacity in a WDM network potentially much greater than in a standard circuit-switched network where wavelength assignment plays no role. We refer to the resulting increase in cost of link capacity as the *cost of transparency*. We also refer to the ratio of the optimal WDM link cost divided by the standard cost for link capacity, as the *transparency gap*.

B. Shepherd is with Bell Laboratories: bshep@research.bell-labs.com
A. Vetta is with McGill University: vetta@math.mcgill.ca

Here we deliberately focus on the minimum link-cost network design version of the problem for a number of reasons. First, there is presently an over-abundance of bandwidth in the core optical networks. In such "dark" networks, links carry a vast number of fibers, most of which have not been lit for use. Thus, a key design issue in such dark or multi-fiber networks is the problem of lighting up fibers from an essentially infinite supply as traffic demand grows. This seems likely to be an operational problem of relevance for some time ahead. The second reason is that the outcome of network designs which incorporate switch costs depends critically on the relative costs of different switch functionalities. These costs are themselves a moving target and are sometimes only speculative. The tradeoffs and thresholds between switch technologies may shift dramatically as, say, the cost of wavelength converters go down. We believe it is therefore useful to use the cost of capacity allocation in traditional networks as a benchmark. We can then get a feel for whether a network is expensive (or not) by comparing its cost to that for the traditional network, using the cost of transparency.

There have been a number of efforts to address and reduce the cost of transparency. Conceptually the most simple approach is to introduce all-optical wavelength converters at the nodes of the network. Such switches enable each link on a flow path to be assigned a distinct wavelength. Hence the cost of transparency disappears altogether. The main drawback to this solution is that the cost of wavelength converting technology is currently extremely high relative to alternatives with less switching functionality. Consequently, one line of attack is to install a limited number of wavelength converters at strategic nodes in the network. These switches can be used to 'break' flow paths into several parts each of which may then be assigned a different wavelength. For instance, in a ring [25], it is known that only one converter is needed to eliminate the cost of transparency.

There are a number of different all-optical switching technologies; we refer the reader to [20] for a comprehensive discussion, and [24] for recent realizations of a number of cross-connects. We now briefly discuss some of the most well-known switch designs. As is to be expected, the cost of a switch tends to increase as it provides more functionality. The standard alternative and one of the least expensive options are switches that are only spatially configurable. These *space switches* allow any terminating fiber at a node to be switched with another terminating fiber (i.e., all wavelengths on a fiber are switched together). A step up in functionality are so-called *wavelength selective* (WSC) switches. These switches allow any incoming wavelength path to be switched onto any outgoing fiber but using the same wavelength. This still results in a transparent network since a path may not change its wavelength. Space switches and WSC switches are both examples of *waveband routing* where each fiber carries k wavelengths that are partitioned into b bands of sizes $k_1, k_2, \ldots, k_b$ respectively, where $k = \sum_i^b k_i$. Each band consists of a subset of the wavelengths. For instance, viewing each wavelength as an integer, the bands could be: $\{1, 2, \ldots k_1\}, \{k_1 + 1, k_1 + 2, \ldots, k_1 + k_2\}, \{k_1 + k_2 + 1, \ldots, k_1 + k_2 + k_3\} \ldots, \{\sum_{i=1}^{b-1} k_i + 1, \ldots, k\}$. Any band of wavelengths on one fiber can be cross-connected to the same band on another fiber. Space switches and WSC switches form the special cases where $b = 1$ and k respectively.

In multi-fiber networks, it is natural to expect the cost of transparency to be reduced if space switches are replaced by WSC switches. The degree of the reduction in empirical observations (for instance, [13] and a related study [18] for dynamic loads) is more than might be anticipated, especially given that WSC technology still only provides quite limited switching between the total set of flow paths. Particularly illuminating is that the cost of such networks is often virtually the same as with wavelength conversion and so the transparency gap is effectively eliminated. The main result of this paper is a theoretical explanation for these observations, given in Section III.

Our theoretical justification of these empirical results is based on randomised techniques for obtaining routing, wavelength and capacity assignment solutions not far from those for the problem without wavelength constraints. The results are especially good for instances where the traffic load (measured in the number of wavelength paths) on network links is not too small; "not too small" means on the order of the logarithm of the number of multiplexed wavelengths (denoted k throughout) per fiber. This situation arises in "mature" networks if capacity allocations can be done from scratch for the total demand on the network. For instance, this is certainly the case for medium-term planning in *dedicated networks* where one is allowed to reroute existing wavelength paths as part of the total optimization. It is also the case for shorter-term planning in networks with dynamically reconfigurable switches.

In combination, these observations impact on how one might evaluate wavelength-conversion assets for dedicated optical networks, i.e., networks where routing and capacity allocations can be reconfigured periodically over time. Even though the cost of wavelength converters is certain to decrease in the future, they will remain substantially more expensive than other technologies. Moreover, our results suggest that the benefits of conversion are greatest in a "young" network, when relatively few fibers are lit up. Unfortunately, in such a network the savings arising in the cost of lighting fibers is quite possibly dwarfed by the initial capital expense of installing wavelength conversion. In contrast, we see that as the network matures, the transparency gap gradually disappears. This simplified argument admittedly depends on our restriction to measuring the transparency gap parameter. We feel, however, that these observations reveal a principle that is worth noting. Of course, in real planning exercises, we must plug in specific numbers for the actual switching costs involved! We discuss this aspect briefly in the conclusions.

In related work, there are a number of alternative approaches to gauging the tradeoffs we have discussed; we briefly mention two. One example is to consider capacitated networks with dynamic loads, and to then estimate the blocking probability. Results in a similar vein appear in [9], where empirical evidence is given that bandwidth utilization is increased by fifty percent with the selective use of wavelength converters. We also mention that it is shown in [23] that the problem of wavelength routing in space-switched multi-fiber networks is mathematically equivalent to that of routing in multi-wavelength TDM networks. Such networks are proposed by [19] and studied in [26], [21].

Finally, in Section IV we explore another area where wavelength routing exhibits unexpectedly desirable behaviour, this time from an algorithmic viewpoint. We consider the core problem of pure wavelength assignment. This is known to be theoretically as hard as the notorious graph colouring problem (e.g., cf. [13]). This latter problem is known for its intractability in theory and often in practice. However, practical and random instances of wavelength assignment are not as difficult to solve as, say, random instances of graph colouring. We discuss this phenomenon and speculate as to why it occurs.

II. A Mathematical Model

In this section we model, via an integer programming (IP) formulation, the *routing, wavelength and capacity assignment problem* (RWAC) associated with a multi-fiber network. We use the resultant formulation to develop an efficient routing algorithm with provable performance guarantees in Section III.

The underlying network may be directed or undirected; for simplicity we represent the network as an undirected graph $G = (V, E)$. Here each edge $e \in E$ represents a topological link in the network. Each link contains many fibers, and each fiber supports a fixed number k of wavelengths[1]. We label these wavelengths $1, 2, \ldots, k$ and we also refer to these as *colours*.

An instance of our routing problem also comes with a symmetric traffic matrix D. The ij^{th} entry of the matrix, d_{ij}, is the number of connections demanded between nodes i and j. Each connection requires a wavelength-path between nodes i and j. We assume, for each node pair i and j, that $\mathcal{P}_{ij}$ is the set of *feasible* paths joining i and j. We require that the wavelength-paths used to satisfy the demand between i and j are chosen from amongst paths in $\mathcal{P}_{ij}$. We denote by $\mathcal{P}$ the set of all feasible paths in the network, that is $\mathcal{P} = \cup_{ij} \mathcal{P}_{ij}$. We can now present the mathematical model as follows.

$\underline{\texttt{IP}-\texttt{RWAC}}$

$$\min \sum_{e \in E} c_e z_e$$

$$\sum_{P \in \mathcal{P}_{ij}} f(P) = d_{ij} \quad \forall \text{ node pairs } i, j \tag{1}$$

$$\sum_{\lambda=1}^{k} f_\lambda(P) = f(P) \quad \forall \text{ paths } P \in \mathcal{P} \tag{2}$$

[1] In general this number may be link-dependent but for our purposes this simpler version suffices.

$$\sum_{P \in \mathcal{P}: e \in P} f_\lambda(P) \quad \leq \quad z_e \quad \forall \text{ links } e \in E, \ \forall \text{ colours } \lambda \ (3)$$

$$z_e, \, f(P), \, f_\lambda(P) \quad \in \quad Z \qquad (4)$$

We next discuss the different components of this integer programming formulation. For each topological path P in G, there is a variable $f(P)$. This represents the number of wavelength paths routed along P, but does not specify the exact fiber and wavelengths on which these connections are carried. The variables $f_\lambda(P)$ determine how many of the connections are assigned wavelength λ. Constraint (2) ensures that the $f(P)$ wavelength paths on P are all assigned a colour. The constraint (1) ensures that, for each pair i and j, there are enough wavelength paths to satisfy the demand between i and j.

We also have a variable z_e for each topological link $e \in E$. This represents the number of fibers we need to light in the link in order to support the routing and assignments determined by the flow path variables. The constraints (3) ensure that, for each λ, z_e is at least as large as the number of paths through e that are assigned wavelength λ. Alternatively, one could view this as ensuring that the assignment solution may use a wavelength only once in in any fiber. Finally, since c_e represents the cost of lighting and using a fiber in link e, our objective is to minimize the quantity $\sum_e c_e z_e$. Observe that since G is a *multi-fiber network*, we have tacitly assumed that the variables z_e are not bounded. Let WDM-OPT denote the optimal value of the integer program IP-RWAC.

In order to evaluate the transparency gap, we now need to consider routing in non-WDM networks (or, equivalently, in networks with converters at each node). In this case, the same IP works except that we have no need for paths to be assigned a single contiguous wavelength, as nodes may switch any incoming port to any outgoing port. An IP formulation for this problem would drop the variables $f_\lambda(P)$, and the edge constraints (3) are replaced by the simpler $\sum_{P: e \in P} f(P) \leq k z_e$. Thus we obtain the following relaxed problem called the *routing and capacity allocation problem*.

<u>IP − RC</u>

$$\min \sum_{e \in E} c_e z_e$$

$$\sum_{P \in \mathcal{P}_{ij}} f(P) \quad = \quad d_{ij} \qquad \forall \text{ node pairs } i, j \qquad (5)$$

$$\sum_{P \in \mathcal{P}: e \in P} f(P) \quad \leq \quad k z_e \qquad \forall \text{ links } e \in E \qquad (6)$$

$$z_e, \, f(P), \quad \in \quad Z \qquad (7)$$

Let OPT be the optimum value for the integer program IP-RC. Evidently we have OPT $\leq$ WDM-OPT and the quantity $\frac{\text{WDM-OPT}}{\text{OPT}}$ is just the so-called transparency gap.

Due to the computational intractability of integer programming, one common approach to solving a problem such as IP-RWAC is to consider the linear relaxation. This is simply the linear programming problem, LP-RWAC, obtained by dropping the integrality constraints (4) and only requiring that $x_e, f(P), f_\lambda(P) \geq 0$. Such a problem is theoretically tractable and can generally also be solved efficiently in practice. Let LP-OPT denote the optimum for LP-RWAC. We clearly also have that LP-OPT $\leq$ WDM-OPT.

Empirical studies such as [13] suggest that the cost WDM-OPT is in fact not much more than OPT. In the next section we present theoretical justification for this. Whilst we argue that in most networks the two values are close, this is not helpful from a computational perspective, as the value OPT is itself hard to find. We thus also describe an algorithm for converting a solution to the linear relaxation LP-RWAC into a feasible for IP-RWAC whose cost is not too far away, especially in the case of more "mature" optical networks.

III. Theoretical Foundations behind the High Performance

In this section, we show that very often, the transparency gap is small. Moreover, we describe an efficient algorithm that produces a feasible solution to the RWAC problem in multi-fiber networks with a very good performance (approximation) guarantee.

Suppose that we are given an optimal integral solution f, z for the routing and capacity allocation problem without wavelength assignment (IP-RC). We describe a simple RANDOMISED ASSIGNMENT ALGORITHM for obtaining a solution for (RWAC) as follows. For each P with $f(P) > 0$ randomly assign P a wavelength from $1, 2, 3 \ldots, k$. Clearly this produces a feasible routing and wavelength assignment. We obtain a network design by setting z_e to be the maximum, over all wavelengths λ, of the number of paths through e that were assigned λ by this procedure. Observe that we can view this assignment method in terms of randomly placing "balls into bins", where the paths are balls and the wavelengths are bins. This allows us (see the appendix) to bound the transparency gap via the following result.

Theorem 1 (Transparency Gap Proposition): For any $0 < \delta \leq \frac{1}{2}$, the expected cost of a randomised assignment is at most

$$(1 + \delta)\text{OPT} + \frac{4}{\delta} \log k \sum_{e \in E} c_e.$$

This theorem tells us that if the traffic load per link is not too light (say, at least logarithmic in the number of wavelengths) then the transparency gap is negligible. In particular, for such mature networks the theorem has the very simple form below.

Corollary 1: If OPT $\geq \frac{4}{\delta^2} \log k \sum_{e \in E} c_e$ for some $0 < \delta \leq \frac{1}{2}$, then the expected cost of a solution from RANDOMISED ASSIGNMENT is at most

$$(1 + 2\delta)\text{OPT}.$$

This shows that the cost of transparency will be very small (for example, choosing $\delta = 0.005$ gives a transparency gap of 1%) in optical networks with a sufficiently large number of wavelength requests, and where existing wavelength paths can be rerouted as we augment capacity (light fibers) to handle new traffic growth. This is likely, for instance, in many medium and long-term planning exercises of dedicated optical networks. Similarly, it applies to networks with rapidly reconfigurable switches.

Note also that the gap in Theorem 1 does not depend upon the number m of links in the network, but rather only on k. Thus, in mature networks, the transparency gap all but disappears irrespective of the network size!

Theorem 1, however, is not directly practical in that it may prove difficult to solve the integer program for (IP-RC). One alternative is to simply take the best available solution to the routing capacity problem. We may then run RANDOMISED ASSIGNMENT to efficiently obtain a solution for RWAC. This will be within the prescribed bounds from the cost of the starting solution. If we wish, however, to give guaranteed bounds from optimality, a better heuristic is to apply a different algorithm to the linear relaxation of IP-RWAC. We describe this approach now.

Let f, f_λ, z be a solution to the linear relaxation LP-RWAC of value LP-OPT. We then obtain a feasible solution for IP-RWAC by applying the method of RANDOMISED ROUNDING introduced in [15]. Without loss of generality, we may assume that all of the LP variables have values at most one[2]. Now recall that for each node pair ij we wish to choose exactly d_{ij} paths from $\mathcal{P}_{ij}$. This we achieve by independently selecting d_{ij} paths, where a path $P \in \mathcal{P}_{ij}$ is selected with probability proportional to its weight in the linear programming solution, i.e., with probability $\frac{f(P)}{d_{ij}}$. This produces our flow paths, and now we assign them their colours. Each selected path is assigned a wavelength according to the probabilities again induced by the linear program i.e., $\frac{f_\lambda(P)}{f(P)}$.

Observe that for each link e and each colour λ the expected number of paths of colour λ using link e is at most the load z_e on the link given by the linear program. This allows us to prove (in the appendix) the following result.

Theorem 2: For any $\delta \leq \frac{1}{2}$, RANDOMISED ROUNDING produces an integral solution to the multi-fiber network problem whose expected cost is at most

$$(1+\delta)\text{LP-OPT} \; + \; \frac{4}{\delta}\log k \sum_{e \in E} c_e.$$

Let us briefly discuss the implications of this result. Again, in mature networks the cost of the linear program solution is large relative to the second term (that is, the term $\frac{4}{\delta}\log k \sum_{e\in E} c_e$). Consequently, using RANDOMISED ROUNDING we will be able to find an assignment that is nearly optimal (say, within a factor $1+2\delta$ of the optimal solution).

We remark that we could also apply a combination of randomised rounding and randomised assignment to the linear relaxation of IP-RC (which is not as tight as that for IP-RWAC). This would lead to a faster algorithm due to the simpler nature of the linear program.

IV. Pure Wavelength Assignment Problems

One natural attack on routing and wavelength assignment problems is as follows. First solve the routing problem so as to minimize the load on the network (e.g., by solving an integer multicommodity flow problem). Second, decompose the solution into flow paths. Finally, colour the resulting collection $\mathcal{C}$ of flow paths so that we obey any wavelength-contention constraints. In this section we focus on the simplest form of this last phase. Namely, where the network only has one fiber in each link (i.e., $z_e = 1$). In this case, the fundamental question is whether we can colour the paths (with k colours, say) of $\mathcal{C}$ so that no two paths with the same colour, cross a common edge. We call this the *path colouring problem.*

Path colouring can be solved as a graph colouring problem. Given a collection of paths in a graph G, the *path intersection graph* $G(\mathcal{C})$ has a node v_P for each path $P \in \mathcal{C}$ and there is an edge $v_{P_1}v_{P_2}$ whenever the paths P_1 and P_2 share an edge in G. Hence there is a feasible wavelength assignment if and only if the path intersection graph $G(\mathcal{C})$ is k-colourable.

The converse reduction also holds. Namely, it is easy to show that for any graph H there is a network G and collection of paths $\mathcal{C}$ such that $H = G(\mathcal{C})$[3]. In this sense, at the heart of RWAC is path colouring, and this is as difficult as graph colouring to solve. Moreover, graph colouring is known to have intractable behaviour in both theory and practice. Despite this, path colouring does not suffer from the same pitfalls as graph colouring, as we review in a moment. We speculate as to why this is so at the end of the section.

The graph colouring problem is a notoriously intractable problem in combinatorial optimization. This is evidenced by empirical experience and by theory. For example, theoretically it is known that, for any choice of $\epsilon > 0$, there is no theoretically efficient (i.e., polynomial time) algorithm that is guaranteed to produce a solution within a factor of $n^{1-\epsilon}$ [7] from the optimum (unless $\mathbf{ZPP} = \mathbf{NP}$). Here, n is the number of nodes in the graph. While certain classes of graphs admit good algorithms for exact or approximate colouring, it is surprising that the DIMACS computational challenge of 1992, gave empirical evidence that optimal colouring of random instances is difficult. Namely, the competition did not produce any exact algorithms which could optimally colour random graphs with 100 nodes where edges were added with probability 0.7.

In contrast, path intersection graph colouring (for example, [12], the full version of [13], and [22]) appears to be much more amenable to computation. It is often possible to exactly colour path intersection graphs with up to 500 nodes. For instance using readily available software and a standard integer programming formulation where for each colouring class we introduce $0-1$ variables that guarantee the variables set to 1 form a *stable set* i.e., a set of nodes, no pair of which is adjacent. We believe one reason for this is that path intersection graphs have compact representations for the collection of stable sets. In particular if we fix a network, then the number of constraints in the integer program is fixed, even while the number of paths grows large.

We discuss this claim in some more detail to explore its relevance. Consider the use of integer programming to solve the stable set problem in a graph H. One standard IP formulation is to just list the *edge inequalities*: $x_u + x_v \leq 1$ for each edge $e = uv$ of H. Clearly this is valid, since any $0-1$

[2] This can be done by "stripping out" an integral sub-solution.

[3] In fact, in the full version of [13], it is argued that you can do this even by choosing a collection of paths all of length at most 6.

vector x that satisfies these constraints determines a stable set $\{v : x_v = 1\}$. While this is valid, one can obtain formulations with fewer inequalities by using stronger constraints. The most standard approach is to use clique inequalities (a clique is a set of mutually adjacent nodes). Clearly, we could include a collection of constraints of the following form: $\sum_{v \in Q} x_v \leq 1$ for each $Q \in \mathcal{M}$, where $\mathcal{M}$ denotes the collection of maximal cliques in the graph H. This *clique formulation* is also valid for the stable set problem. A major drawback is that graphs generally may have exponentially many maximal cliques, and so even finding the cliques is an intractable problem. But we do not need the whole collection $\mathcal{M}$ to have a valid formulation. In fact, a collection $\mathcal{M}'$ yields a formulation for the stable set problem if and only if each edge of H has both endpoints inside some clique of $\mathcal{M}'$. Let $\nu(H)$ denote the minimum number of cliques needed to cover the edges of a graph (this is called the *clique-edge cover number*; its asymptotic behaviour is studied in [8]). The above remarks show that the minimum number of constraints in a valid IP formulation for the stable set problem in H is precisely the value $\nu(H)$. It is natural to expect that having a smaller formulation may be computationally advantageous. In particular, for large cliques Q, the setting of one variable $x_v = 1$ already forces many nodes u to satisfy $x_u = 0$ (that is, all the other nodes in Q).

Path intersection graphs arising from a network have precisely this property that their clique-edge cover number is small. To see this, note that if $H = G(\mathcal{C})$ is a path intersection graph arising from a collection paths $\mathcal{C}$, then each edge e of G gives rise to a clique in H. Namely, the set of paths which contain the edge e induce a clique. This gives a clique-edge cover for H whose size depends only on the original network G and not on how many paths we use in G. That is, it does not depend upon the size of H!

These arguments are made in terms of the pure colouring problem, but may be a partial explanation for their relative computational ease. For general RWAC, one can also visualize the impact of having WSC switches, in terms of these IP formulations for stable sets. Namely, they are equivalent to solving stable set problems where for each clique Q, we have a capacity z_Q. We then require each colour class to have the property that $\sum_{v \in Q} x_v \leq z_Q$. Such relaxations are studied in [4] where they show that taking $z_Q = O(1)$ everywhere, does not greatly improve the ability to find theoretically efficient algorithms (again, in contrast to what is found for path intersection graphs).

V. Conclusion

We have discussed that one concern in deploying WDM technology in a transparent optical core network is the potential increase of cost imposed by the extra wavelength contention constraints. It has been observed, however, that cost-effective designs are possible with wavelength selective switches as long as there are many parallel fibers (multi-fiber networks). We gave a theoretical justification for why this is the case, via a result that shows that the cost of transparency is muted in multi-fiber networks with WSC switches. Our arguments were based on simple randomised heuristics that can be used themselves to generate solutions for the routing, wavelength and capacity assignment problem. In fact, in mature networks where the total demand is relatively large, we showed that these heuristics produce solutions very close to the optimal cost of link capacity in a network with converters. Thus in the case of designing dedicated transparent networks, it appears that the use of wavelength converters might be avoided altogether.

Our study focused on link capacity costs under the presence of wavelength selective switches (WSC). An interesting future direction is to consider designs that incorporate a variety of waveband switches. Intuitively, the cost of such a switch decreases as the size of wavebands increases. This *cost reduction function* has its maximum occurring for WSC switches, since they have waveband size of one. It would be interesting to consider the *waveband optimization problem*. That is to consider the cost of the network optimization with switch costs included, under different choices for the cost reduction function described above.

In the final section, we discussed wavelength assignment in single fiber networks. This gives rise to colouring of path intersection graphs, a problem inherent to general routing and wavelength assignment. This problem is theoretically equivalent to the fundamental optimization problem of node colouring in general graphs. While exact node colouring is known to be intractable in theory and practice, much larger instances of path colouring can be solved in general. We speculated as to why path colouring exhibits such desirable behaviour in practice.

Finally, we mention that recently, Matthew Andrews and Lisa Zhang have derived hardness of approximation results for the problems considered here.

Appendix I
Analysis of Randomised Rounding of LP Solution

Here we present the technical details involved in proving the performance guarantee of the algorithm. We begin with some background. The following theorem can be obtained via techniques due to Chernoff ([5]).

Theorem 3: Let $X_1, \ldots, X_r$ be independent Poisson trials ($0-1$ random variables) such that $E(X_i) = p_i$. If $X = \sum_{i=1}^r X_i$, then for $0 < \beta < 2e - 1$, we have

$$Pr(X > (1+\beta)E(X)) \leq \exp\left(-\frac{1}{4}\beta^2 E(X)\right).$$

We may also replace the $\frac{1}{4}$ term by $\frac{1}{3}$ in the case in which $\beta < 1$ [11]. We now prove Theorem 1. Let $X_\lambda(P)$ denote the random variable that is 1 if P is assigned wavelength λ, and 0 otherwise. Thus, via our "balls into bins" approach, $Pr[X_\lambda(P) = 1] = \frac{1}{k}$ for each λ and each P with $f(P) > 0$. For each link e, we may consider the paths $P_1, P_2, \ldots, P_{t_e}$ that contained e and were assigned flow under f. Let X_λ^e denote the random variable $\sum_{j=1}^{t_e} X_\lambda(P_j)$ which counts the number of flow paths assigned a given λ. How many fibers do we need on link e to be able to route our solution obtained by *randomised assignment*? This is simply $Z_e = \max_\lambda X_\lambda^e$. Note here that $E(X_\lambda^e) = \frac{t_e}{k}$. However, the expected cost of our randomised assignment is $E(\sum_{e \in E} c_e Z_e)$. In order to evaluate this, by linearity of expectation, it suffices to bound $E(Z_e)$ for each link e.

Towards this goal, we wish to show that, with large probability, Z_e does not get too much bigger than $\frac{t_e}{k} \leq z_e$. To do this we will apply Chernoff's Theorem; note that we are in the special case in which $p_i = \frac{1}{k}$, for all $1 \leq i \leq k$ and $r = t_e \leq kz_e$. First suppose $z_e = c \log k$ where $c \geq 4$. Then by Theorem 3,

$$Pr(X^e_\lambda > 3z_e) \leq \exp(-z_e) \leq \frac{1}{k^4}$$

Consequently, using the fact that the probability of the union of events is at most the sum of their individual probabilities, we have $Z_e \geq 3z_e$ with probability at most $\frac{1}{k^3}$. Since $Z_e \leq kz_e$, we see that the contribution to its expectation when Z_e is larger than $3z_e$ is $o(1)z_e$. Moreover, we also have

$$Pr(3z_e \geq X^e_\lambda \geq z_e + \sqrt{3(1+\epsilon)c}\log k) \leq \frac{1}{k^{1+\epsilon}}$$

It then follows that

$$E(Z_e) \leq z_e + 2\sqrt{c}\log k = (1 + \frac{2}{\sqrt{c}})z_e$$

On the other hand, if $z_e \leq c \log k$ where $c \geq 4$ then, by a similar argument, we obtain that $E(Z_e) \leq 2c \log k$. In addition, $2c \log k \geq \frac{2}{\sqrt{\gamma}}\gamma \log k$ for all $\gamma \leq c^2$. Putting this together we have that, for any $c \geq 4$,

$$E(Z_e) \leq (1 + \frac{2}{c})z_e + 2c \log k$$

Our main result then follows.

Theorem 4: For any $\delta \leq \frac{1}{2}$, the expected cost of a randomised assignment is at most

$$(1+\delta)\text{OPT} + \frac{4}{\delta}\log k \sum_{e \in E} c_e.$$

Next we prove Theorem 2. We consider some LP solution f, z. We have m edges and k colours. For each edge e, we are interested in the maximum number of fibers required by any wavelength on that link. Recall that, from the linear program, $\sum_{P \in \mathcal{P}: e \in P} f_\lambda(P) = \sum_{ij} \sum_{P \in \mathcal{P}_{ij}: e \in P} f_\lambda(P) \leq z_e$ for any link/wavelength pairing (e, λ). Now take P a wavelength path using link e, and set $p_P = f_\lambda(P)$. Then we can upper bound the probability that the number of fibers required by wavelength λ exceeds a specified value by considering the following random variable. Let $X^e_\lambda = \sum_{P: e \in P} X_P$, where the X_P are independent Poisson trials with $E(X_P) = p_P$. Recall that $0 \leq p_P \leq 1$ for all P. Then we obtain our result using Chernoff bounds.

Theorem 5: For any $\delta \leq \frac{1}{2}$, randomised rounding produces an integral solution to the multi-fiber network problem whose expected cost is at most

$$(1+\delta)\text{LP-OPT} \; + \; \frac{4}{\delta}\log k \sum_{e \in E} c_e.$$

Finally, using sampling techniques we can, with high probability, efficiently obtain a solution whose cost follows the bound guaranteed by the theorem.

ACKNOWLEDGMENTS

We are extremely grateful to Chandra Chekuri for numerous comments and suggestions that greatly improved this paper. We also thank Luc Devroye, Debasis Mitra and Iraj Saniee for helpful suggestions and inspiring discussions. Part of this work was carried out while the second author took part in the Bell Labs Summer Intern Program.

REFERENCES

[1] S. Baroni, S.K. Korotky, and P. Bayvel, "Wavelength interchange in multi-wavelength optical transport networks", in *Tech. Dig. 23rd European Conf. on Opt. Commun.*, Edinburgh, Vol. **3**, pp164-167, 1997.

[2] S. Baroni, P. Bayvel, R.J. Gibbens, and S.K. Korotky, "Analysis and design of resilient multi-fiber wavelength-routed optical transport networks", *Journal Lightwave Technology*, **17**, pp743-758, 1999.

[3] P. Bayvel, "Future high-capacity optical telecommunication networks", *Phil. Trans. Royal Society London, A*, **358**, pp303-329, 2000.

[4] C. Chekuri and S. Khanna, "On Multi-dimensional Packing Problems", *10th SODA*, January 1999.

[5] H. Chernoff, "A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations", *Annals of Math. Stat.*, **23**, pp493-509, 1952.

[6] L. Devroye, *Non-Uniform Random Variate Generation*, Springer-Verlag, 1986.

[7] U. Feige and J. Kilian, "Zero knowledge and the chromatic number", *J. Comput. System Sci.*, **57**, pp187-199, 1998.

[8] A. Frieze and B. Reed, "Covering the edges of a random graph by cliques", *Combinatorica*, **15**, pp489-497, 1995.

[9] R. Giles, K. Kumaran, D. Mitra, C. Nuzman and S. Saniee, "Selective transparency in optical networks", *Proc. of ECOC*, Copenhagen, 2002.

[10] V. Guruswami, S. Khanna, B. Shepherd, R. Rajaraman and M. Yannakakis, "Near-optimal hardness results and approximation algorithms for edge-disjoint paths and related problems", Submitted to *JCSS*; preliminary version in *Proc. of 31st Annual ACM Symposium on Theory of Computing (STOC)*, pp19-28, 1999.

[11] T. Hagerup and C. Rüb, "A guided tour of Chernoff bounds", *Information Processing Letters*, **33**, pp305-308, 1990.

[12] K. Kumaran and I. Saniee, "Shared recovery in transparent optical networks", *Proc. of 39th Annual Allerton Conference on Communication, Control, and Computing*, 2001.

[13] E. Lowe, B. Shepherd and N. Walker, "Routing and configuration of static path optical transport networks", *BT Labs technical report: Networks* **54(3)**, July 1995; short version appeared in *IEE Electronic Letters*, **32**, pp1913-1914, 1996.

[14] M. Raab and A. Steger,. "Balls into Bins - A Simple and Tight Analysis", *Proceedings of RANDOM'98, LNCS 1518*, Springer Verlag, pp159-170, 1998.

[15] P. Raghavan and C. Thompson, "Randomized rounding: A technique for provably good algorithms and algorithmic proofs", *Combinatorica*, **7(4)**, pp365-74, 1987.

[16] R. Ramaswami and K. Sivarajan, "Routing and wavelength assignment in all-optical networks", *IEEE/ACM Trans. Netw.*, **3(5)**, pp489-500, 1995.

[17] A. Mehrotra and M. Trick, "A clique generation approach to graph coloring", *INFORMS Journal on Computing*, **8**, pp344-354, 1996.

[18] D. Mitra, C. Nuzman, I. Saniee and P. Whiting, "Optical crossconnect with shared wavelength conversion under dynamic load", *IEEE/OFC Technical Digest*, pp160-162, March 2002.

[19] M. Sabry and J. Midwinter, "Towards an Optical Ether", *IEE Proc. Optoelectron.*, **141**, pp327-335, 1994.

[20] T.E. Stern and K. Bala, *Multiwavelength optical networks: a layered approach*, Addison Wesley, 1999.

[21] S. Subramaniam and R.A. Barry, "Wavelength assignment in fixed routing WDM networks", *Proc. IEEE Int. Conf. on Communications ICC '97*, Montréal, pp406-410, 1997.

[22] M. Trick, *Personal Communication*, 1995.

[23] N. Wauters and P. Demeester, "Wavelength translation in optical multi-wavelength multi-fibre transport networks", *Int. J. Optoelectronics,* **11(1)**, pp53-70, 1997.

[24] G. Wilfong, B. Mikkelsen, C. Doerr and M. Zirngibl, "WDM cross-connect architectures with reduced complexity", *Journal of Lightwave Technology*, **17(10)**, pp1732-1741, 1999.

[25] G. Wilfong and P. Winkler, ‘Ring routing and wavelength translation”, *Proc. of Symposium on Discrete Algorithms*, 1998.
[26] J. Yates, J. Lacey and D. Everitt, “Blocking in multiwavelength TDM networks”, *Telecommunication Systems*, **6**, pp1-22, 1998.