

Part III Essay

Routing and wavelength assignment in optical networks

D.J. Wischik

University of Cambridge, Statistical Laboratory,
16 Mill Lane, Cambridge, UK.
djw1005@cam.ac.uk

20th May, 1996

Abstract

In optical networks, the vast bandwidth available in an optical fibre may be utilized by splitting it into several channels, each of a different wavelength. This allows signals to be routed entirely in the optical domain. This essay studies such routings, and examines the number of different wavelengths needed for irregular networks.

We find the number of wavelengths needed in large random networks, in terms of the proportion of edges present. A simple greedy routing algorithm which uses no wavelength conversion is sufficient.

We give an alternative lower bound on the number of wavelengths needed, which performs poorly for large random networks, but which is more appropriate for certain existing networks. It can sometimes be used to show that heuristically-found routings are optimal.

1 Introduction

Contemporary fibre-optic networks transmit data at speeds orders of magnitude faster than standard electronic networks. But they have as yet only realised a small fraction of their potential. If they are to do more than replace copper wires, they must go beyond the limitations imposed by electronic technology.

Currently, optical signals are converted into electronic signals whenever they need to be switched, and then back into optical signals for transmission. Electronic processing becomes a bottleneck in fast networks. Routing a signal using exclusively optical techniques may eventually be both cheaper and faster.

The simplest technique for optical routing, and also one of the most promising, is wavelength division multiplexing. This exploits the ability of an optical fibre to carry a number of independent signals of different wavelengths. These different wavelengths can be separated out and switched independently.

The huge capacities available with optical fibres and the simpler routing technology lead to design considerations different to those for conventional electronic networks [2]. In an electronic circuit-switched network calls are typically routed as they come on-line, in order to fully exploit the limited capacities. But in an optically routed network, there would more commonly be fixed routes for calls to take; and the routing pattern would only change over longer time-scales, to cope with changes in average long-term traffic. This system is called quasi-static routing. It is well-suited for national backbone networks, such as that proposed in [11].




Figure 1: A wavelength-routed optical network.

A wavelength division multiplexed optical network consists of end nodes and routing nodes connected by fibre-optic links, as in Figure 1. Each link can carry a certain number of wavelengths. A routing node can route each wavelength from an input port to an output port, independently of other wavelengths; it may also change the wavelength of a signal with a wavelength converter. An end node consists of one or more tunable optical transmitters and receivers, and is connected to a routing node, usually by a single link, but sometimes by more. We are interested in setting up connections between pairs of end nodes, assigning to each pair of interest a route and a wavelength in such a way that, on any link, each wavelength is used by at most one connection. We call the collection of routes and wavelengths assigned a *routing*. The practicability of a routing will depend on the number of different wavelengths it uses: current limits are of the order of tens of wavelengths.

The *efficiency* of a network is that proportion of the total throughput able to be carried simultaneously. In this context, it is the largest proportion of node pairs that may be connected at any one time. For example, for efficiency 1, each end node must have as many transmitters and receivers as there are other end nodes in the network. If, in addition, there is a single optical link from an end node to its routing node, there must be at least this many wavelengths available. With less than full efficiency, there arises the problem of coordination: the transmitting end node and the receiving end node must both be tuned to the correct wavelength.

1.1 Previous Work

A permutation routing is a routing in which each end node has one transmitter and one receiver, and each transmitter is connected to a single receiver, and each receiver to a single transmitter. Aggarwal *et al* [1] consider the problem of constructing networks which can route all permutations by different tunings of the receivers and transmitters, in the case where each end-node is connected to a routing node by a single link. They find bounds on the number of wavelengths necessary, and also relate the number of wavelengths to the number of routing nodes.

Ramaswami and Sivarajan [9] show how to construct a linear programming problem for an arbitrary network in order to maximize the number of simultaneous connections; in their approach, the number of variables grows exponentially with network size, and the constraint of the number of wavelengths is ignored.

Ramaswami and Svarajan [10] and Zhang and Acampora [12] consider packet-

switched networks, and separate the routing problem into two parts. First they connect certain pairs of nodes to form a *logical topology*, by solving a linear programming problem or by a variety of heuristic methods. Then, they overlay a packet-switched network on top of this logical topology, an optical route in the physical network corresponding to a single logical link. When some node-pairs are not connected by an optical route, a packet may have to pass through several intermediate optical routes to reach its destination; this is called multi-hop routing. Zhang and Acampora simulate various dynamic call admission schemes and find that, although the logical topology is fixed in advance and may not change to accommodate new requests, there is little loss of throughput.

Baroni [3] gives some exact results for standard network architectures, such as the de Bruijn graph and ShuffleNet. But such networks are not easily scalable: their ‘numerology’ permits only certain numbers of nodes and links in specific configurations. If these new technologies are to be applied to practical large-scale networks, it is important to understand in general how the number of wavelengths depends on the network topology. Baroni, Bayvel and Midwinter [4] examine empirically the number of wavelengths needed in random networks. This essay will analytically justify their results.

1.2 The Model

For simplicity, all links will be taken to consist of a single fibre which can transmit signals in both directions. Each routing node will be identified with its end-node. This might correspond to a full efficiency network, in which each end node has many transmitters and receivers and is connected by a large number of fibres to its routing node; or to a network of efficiency $1/(n - 1)$ where n is the number of nodes, in which each end node has a single transmitter and receiver, and a single fibre connecting it to its routing node.

The mathematical formulation is as follows. Let G be a graph, assumed to be connected, with vertex set $V(G)$ and edge set $E(G)$. We define a *lightpath* between vertices v and w in V to be a path in G between v and w , together with an associated colour. A *lightedge* is an edge with an associated colour. We wish to assign lightpaths to all pairs of distinct vertices, so that all lightpaths through a given edge are of different colours.

Let $\Lambda(G)$ be the least number of colours for which such an assignment is possible. We will examine the relationship between the proportion of edges present in the graph and $\Lambda(G)$. Two different sorts of bounds on $\Lambda(G)$ will be presented. Each bound will be defined, motivated by comparison to computer-simulated results, and justified analytically.

1.3 Random graph preliminaries

In what follows, let $n = |V|$ be the order of the graph, $M = |E|$ the size of the graph, $N = \binom{n}{2}$ the maximum number of edges in a graph of order n , and $\alpha = M/N$ be the proportion of edges present, also called the *connectivity* of G .

We will be concerned with two types of random graphs. $\mathcal{G}(n; \alpha)$ is the space of random graphs of order n in which each edge is independently present with probability α or absent with probability $1 - \alpha$. $\mathcal{G}(n; M)$ is the space of random graphs of order n and size M , in which each graph is equally likely. We will write $\mathbb{P}_{n, \alpha}$ and $\mathbb{P}_{n, M}$ for these two cases, and suppress n and α or M when it is clear from the context which is intended.

In either of the two cases, we say that almost all graphs have property Q if $\lim_{n \rightarrow \infty} \mathbb{P}_n(Q) = 1$.

The following result provides a useful link between the two spaces. We say that a property Q of graphs is increasing if $G \subseteq H$ and $Q(G)$ imply $Q(H)$.

PROPOSITION 1 *Let Q be an increasing property. If almost all graphs in $\mathcal{G}(\alpha)$ have Q , then almost all graphs in $\mathcal{G}(\lfloor \alpha N \rfloor)$ have Q .*

Proof. First, note that if $M_1 \leq M_2$ then $\mathbb{P}_{M_1}(Q) \leq \mathbb{P}_{M_2}(Q)$. This can be seen by coupling the two spaces. Consider adding M_2 edges one by one. If Q holds after M_1 edges have been added, then it holds when M_2 have been added.

Suppose it is not the case that almost all graphs in $\mathcal{G}(\lfloor \alpha N \rfloor)$ have Q . Then for some $\epsilon > 0$ there is a sequence of orders n_k such that $\mathbb{P}_{n_k, \lfloor N\alpha \rfloor}(Q) < 1 - \epsilon$ for all k . So

$$\begin{aligned} \mathbb{P}_{n_k, \alpha}(Q) &= \sum_{m=0}^N \mathbb{P}_{n_k, \alpha}(M = m) \mathbb{P}_{n_k, \alpha}(Q | M = m) \\ &= \sum_{m=0}^{\lfloor \alpha N \rfloor} \mathbb{P}_{n_k, \alpha}(M = m) \mathbb{P}_{n_k, m}(Q) + \sum_{m=\lfloor \alpha N \rfloor + 1}^N \mathbb{P}_{n_k, \alpha}(M = m) \mathbb{P}_{n_k, m}(Q) \\ &\leq (1 - \epsilon) \mathbb{P}_{n_k, \alpha}(M \leq \lfloor \alpha N \rfloor) + \mathbb{P}_{n_k, \alpha}(M > \lfloor \alpha N \rfloor) \end{aligned}$$

which contradicts $\mathbb{P}_{n_k, \alpha}(Q) \rightarrow 1$ as $k \rightarrow \infty$. \square

2 The Distance Bound

Given a graph G , all node pairs must be connected by a lightpath. If the shortest path between two vertices v and w is of length $d = d(v, w)$, where the length of a path is the number of constituent edges, then the lightpath between them must consist of at least d lightedges. The sets of lightedges used by different paths are disjoint. If there is an assignment of lightpaths using Λ colours, then there are $\Lambda \cdot M$ lightedges available; and so

$$\Lambda \cdot M \geq \sum_{v \neq w \in V} d(v, w)$$

and in particular

$$\Lambda(G) \geq m_{\text{dist}}(G)$$

where we define the *distance bound* $m_{\text{dist}}(G)$ to be $\frac{1}{M} \sum_{v \neq w} d(v, w)$.

The distance bound has obvious modifications when the network is subject to different constraints. If only a subset of node pairs require to be connected, only the distances between those pairs need be summed. If a link consists of more than one optical fibre, the number of edges can be increased by one. If there are nodes solely for routing, they need not be taken into account in the sum.

2.1 Motivation

This section derives an estimate for m_{dist} and compares it to the computer-simulated results in [4].

Their method of simulation is as follows. They generate random graphs of given size and connectivity, subject to two constraints: firstly that if any one edge is removed the graph remains connected, which is important for the reliability of a

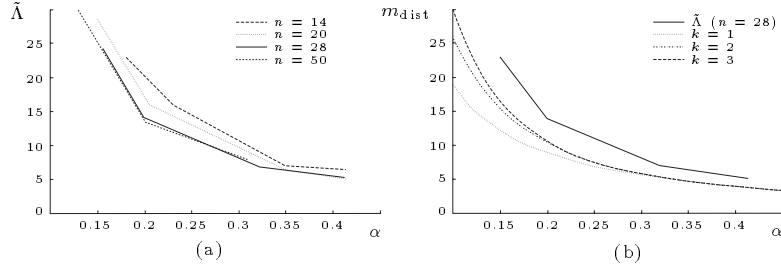


Figure 2: (a) Average number of wavelengths needed by the Minimum Number of Hops algorithm for graphs of various orders n , and varying connectivity α . (b) Approximate expected value of m_{dist} for a random graph in $\mathcal{G}(28; \alpha)$, taking $k = 1, 2, 3$ terms in the expression (1).

network; and secondly that each node has degree less than some specified maximum degree. The graphs are generated by putting down edges at random until the required number of edges is reached, at each step only putting down edges which would not violate the second condition; and finally checking the first condition.

The number of wavelengths needed for a graph is estimated using a heuristic algorithm, called the Minimum Number of Hops algorithm. First, the set of shortest paths between each node pair is found. Each node pair is arbitrarily assigned one of these paths. Each node-pair is in turn considered, and if the maximum congestion along an alternative path is lower than the maximum congestion along the current path, that alternative path is chosen instead; this is repeated until no more substitutions are possible. The congestion of an edge is the number of paths using that edge. Finally, wavelengths are assigned by ordering the paths in decreasing order of length, and to each in turn assigning the first available wavelength. This provides an estimate $\tilde{\Lambda}$ for Λ . Some results are shown in Figure 2 (a).

The results may be compared to the approximate expected value of m_{dist} for random graphs. They may only be compared approximately, because it is very difficult to deal analytically with a probability space as complicated as the one used above. Instead, we will deal only with $\mathcal{G}(n; M)$ and $\mathcal{G}(n; \alpha)$ and ignore the problem of disconnected graphs – which would make $\mathbb{E}(m_{\text{dist}}) = \infty$.

For a random graph in $\mathcal{G}(n; M)$ of connectivity α ,

$$\mathbb{E}(m_{\text{dist}}(G)) = \mathbb{E}\left(\frac{1}{M} \sum_{v \neq w} d(v, w)\right) = \frac{1}{M} \mathbb{E}\left(\sum_{v \neq w} d(v, w)\right) = \frac{1}{\alpha} \mathbb{E}D$$

where D is the length of the shortest path between an arbitrary pair of nodes in a graph. We may approximate

$$\mathbb{E}D \geq \sum_{d=1}^k d \mathbb{P}(D = d) + (k+1) \mathbb{P}(D > k). \quad (1)$$

When $k = 2$ this gives

$$\mathbb{E}D \geq \alpha + 2p(1 - \alpha) + 3(1 - \alpha)(1 - p),$$

where $p = \mathbb{P}(D = 2 | D > 1)$. We can find p by a simple counting argument, giving

$$\begin{aligned} p &= \sum_{r=0}^{n-2} \mathbb{P}(D = 2 | r \text{ tandem edges present}) \mathbb{P}(r \text{ tandem edges present}) \\ &= \binom{\binom{n}{2} - 1}{M}^{-1} \sum_{r=0}^{n-2} 2^r \binom{n-2}{r} \binom{\binom{n}{2} - 1 - 2\binom{n-2}{r}}{M-r}. \end{aligned}$$

This is still rather unwieldy. It may however be very well approximated in $\mathcal{G}(n; \alpha)$. In this case the number of edges M is a random variable and dependent on the sum of internodal distances, but we approximate:

$$\mathbb{E}\left(\frac{\sum_{v \neq w} d(v, w)}{M}\right) \approx \frac{\mathbb{E}(\sum_{v \neq w} d(v, w))}{\mathbb{E}M} = \frac{1}{\alpha} \mathbb{E}D.$$

In $\mathcal{G}(n; \alpha)$, $\mathbb{P}(D = 1) = \alpha$, and $\mathbb{P}(D > 2 | D > 1) = (1 - \alpha^2)^{n-2}$. These two estimates of m_{dist} are virtually identical for typical α and n .

In fact, we may go a little further. In $\mathcal{G}(n; \alpha)$ it is possible to slightly improve the approximation by finding $\mathbb{P}(D = 3)$. Let B_d be the number of vertices at distance d from some fixed vertex, so that $\mathbb{P}(D = d) = \frac{1}{n} \mathbb{E}(B_d)$. Then $B_0 = 1$; and by considering ‘growing’ a cluster from that vertex, it can be seen that the conditional distribution of B_{d+1} given $B_0 \dots B_d$ is

$$B_{d+1} | B_0 \dots B_d \sim \text{Bin}\left(n - (B_0 + \dots + B_d), 1 - (1 - \alpha)^{B_d}\right).$$

This expression becomes unwieldy for $d > 2$, but it is still simple to calculate $\mathbb{E}(B_3)$ numerically, by summing the following over possible values for B_1 :

$$\mathbb{E}(B_3 | B_1) = (1 - \alpha)^{B_1} (n - 1 - B_1) \left[1 - \left((1 - \alpha)(1 - (1 - \alpha)^{B_1}) + (1 - \alpha)^{B_1} \right)^{n-2-B_1} \right]$$

These estimates are plotted against simulated data in Figure 2 (b). They indicate that m_{dist} might be a reasonable estimate for Λ .

2.2 Bounds

PROPOSITION 2 *If G has connectivity α then $m_{\text{dist}}(G) \geq \frac{2}{\alpha} - 1$.*

Proof. The least m_{dist} is attained from a graph with minimal sum of internodal distances. When the number of edges is $M = n - 1$ so that G is either a tree or disconnected, this minimum is clearly attained for a tree. The optimal tree is a star, by induction on n : when a new node and edge are added, the new node must be distance 1 from one node, and distance at least 2 from all other nodes; only in the case of a star is the distance to all other nodes 2.

As more edges are added, optimal graphs are obtained from a star by adding edges between any unconnected spokes; call some increasing sequence of these graphs H_M . In this way, $\sum_{v \neq w} d(v, w)$ decreases by 1 as each edge is added, from $n - 1$ for when H is a star down to $\frac{1}{2}n(n - 1)$ when H is complete. They are optimal, for suppose they were not: that is, that there is a G with M edges for which $m_{\text{dist}}(G) < m_{\text{dist}}(H_M)$. Then m_{dist} must decrease by at least one as each edge is added to G , until it is less than $m_{\text{dist}}(H_N)$ when G has been made complete, which is impossible. \square

THEOREM 3 *If $\Lambda > \frac{2}{\alpha} - 1$ then almost all graphs in $\mathcal{G}(\alpha)$ have a routing using Λ or fewer colours.*

The proof is rather long, and is left for the appendix. A simple greedy algorithm which uses only lightpaths of lengths one and two is sufficient to construct the routing, suggesting that for practical networks equally simple routing strategies might behave well. Note that in the construction of the routing there was no wavelength conversion. This shows that for large random networks, there is nothing to be gained by using wavelength converters.

This problem is closely related to Kelly's Triangle Problem. The problem as originally posed is as follows. Consider a complete graph with n nodes. Let $0 < p < 1/3$ and independently colour each edge of the graph red with probability p and white with probability $1 - p$. A triangle is a set of three distinct edges, each pair of which has a distinct common node; and a triangle is called good if it contains exactly one red edge. How many disjoint good triangles is it possible to construct?

In the notation of this essay, the red edges correspond to links which are not present. All pairs of nodes joined by links (white edges) are connected by a lightpath along the edge between them, all using the same colour. We seek to connect all node pairs not directly connected (connected by a red edge) by a lightpath through a tandem pair of existing edges (form a good triangle) using a second colour. We require all these lightpaths to be disjoint (all the good triangles to be disjoint). The Triangle Problem asks: if $\alpha > 2/3$, is it possible to connect all node pairs using 2 colours?

Hajek [7] proves that the answer is almost always yes. The result is established using techniques from stochastic approximation theory. It provides the basis for the proof of theorem 3.

This result has close connections with Dynamic Alternative Routing [6]. Calls, arriving dynamically between nodes of a circuit-switched network, may be routed over indirect routes when direct routes have insufficient free capacity. DAR is a simple and efficient method of assigning alternative routes using only direct paths and two-hop indirect paths, in effect solving the dynamic version of the Triangle Problem.

We summarise the results of this section in the following corollary:

COROLLARY 4 *For almost all graphs in $\mathcal{G}(\lfloor \alpha N \rfloor)$,*

$$\lceil \frac{2}{\alpha} - 1 \rceil \leq \Lambda(G) \leq \lfloor \frac{2}{\alpha} \rfloor$$

Proof. The first inequality is from proposition 2, the second from theorem 3 and proposition 1. \square

3 The Partition Bound

The last section showed that random graphs have very good properties for optical routing. In this section, we look at another bound suggested by F.P. Kelly [8], which turns out to be weaker than the distance bound for random graphs, but stronger when applied to some existing networks.

Given a graph G , let $A \subset V(G)$. Let $c(A)$ be the set of edges linking vertices in A to vertices in $G \setminus A$. The cut $c(A)$ must carry connections between $|A| \cdot |G \setminus A|$

pairs of vertices. If G can be coloured with Λ colours, then the cut has capacity $\Lambda|c(A)|$ lightedges. So,

$$\Lambda \cdot |c(A)| \geq |A| \cdot |G \setminus A|$$

and in particular

$$\Lambda(G) \geq \lceil m_{\text{cut}}(G) \rceil$$

where we define the partition bound $m_{\text{cut}}(G)$ to be $\sup_{A \subset G} \frac{|A| \cdot |G \setminus A|}{|c(A)|}$.

3.1 Motivation

Enumerating over all subsets $A \subset V(G)$ is $O(2^n - 1)$ which becomes impractical for networks of even moderate size. A heuristic algorithm has been developed to find an estimate \tilde{m}_{cut} of m_{cut} . The algorithm works as follows.

1. Let the cost of traversing an edge be 1.
2. Determine a cheapest path between each node pair.
3. Calculate the number of paths through each edge, i.e. the *congestion* of that edge.
4. Increase the cost of the edges which are maximally congested by ε .
5. Repeat steps 2-4 until a cut $c(A)$ is observed on which costs are being repeatedly increased.
6. Use this cut to calculate \tilde{m}_{cut} (It is clear that a cut is obtained. If the cut partitions the network into more than two pieces, select the piece which gives rise to the best \tilde{m}_{cut}).





	n	$ E $	α	$\lceil \frac{2}{\alpha} - 1 \rceil$	m_{dist}	$\lceil \frac{1}{\alpha} \rceil$	\tilde{m}_{cut}	$\tilde{\Lambda}$
Arpanet 	20	31	0.16	11	17	7	33	33
EON 	20	37	0.19	10	13	6	18	18
National 	24	48	0.17	11	16	6	40	40
NSF 	14	21	0.23	8	10	5	13	13

Table 1: Results for some planned or implemented networks. \tilde{m}_{cut} gives the value of m_{cut} found by the heuristic algorithm, and $\tilde{\Lambda}$ gives the number of colours used in the heuristic wavelength assignment algorithm. In all cases, the lower bounds and the feasible solutions agree, showing that $\tilde{\Lambda}$ is optimal.

Table 1 shows the results of this algorithm and the earlier algorithm for assigning lightpaths when applied to various existing or planned networks. It can be seen that

in all cases the partition bound matches precisely the number of colours needed, and performs significantly better than the distance bound. This indicates that these networks were not designed for the sort of optical routing presented in this essay: random networks tend to need fewer colours.

3.2 Bounds

PROPOSITION 5 *If G has connectivity α then $m_{cut}(G) \geq 1/\alpha$.*

Proof. Let $v(A) = |A| \cdot |G \setminus A|/|c(A)|$ be the value of the cut $c(A)$. Consider the collection of cut values. By a simple counting argument, their sum is fixed. And by looking at them as a collection of positive real numbers of fixed sum, their maximum must be at least their average. But the average is what a cut value would be if all cut values were equal, which is $\alpha|A| \cdot |G \setminus A|$, giving the result. \square

PROPOSITION 6 *If $C > 1/\alpha$ then for almost all graphs in $\mathcal{G}(\alpha)$, $m_{cut}(G) < C$.*

Proof. Suppose not. Then there exists $\epsilon > 0$ and a sequence n_k such that $\mathbb{P}_{n_k}(m_{cut} \geq C) > \epsilon$. But if $m_{cut} \geq C$, then for some cut $A \subset V(G)$, $v(A) \geq C$.

For a given cut A of order a , let X be the number of edges between A and $V(G) \setminus A$. Then $X \sim \text{Bin}(a(n-a), \alpha)$ and

$$\begin{aligned} \mathbb{P}_{n_k}(v(A) \geq C) &= \mathbb{P}_{n_k}(X \leq a(n_k - a)/C) \\ &\leq \mathbb{E}_{n_k}(x^{-(X - a(n_k - a)/C)}) \\ &= x^{a(n_k - a)/C} \left(\frac{\alpha}{x} + (1 - \alpha)^{a(n_k - a)} \right) \quad \text{for all } x > 1 \end{aligned}$$

Setting $x = \frac{\alpha}{1 - \alpha}(C - 1) > 1$, we see that $\mathbb{P}_{n_k}(v(A) \geq C) \leq p^{a(n_k - a)}$ where

$$p = (1 - \alpha)^{1 - \frac{1}{C}} \alpha^{\frac{1}{C}} (C - 1)^{\frac{1}{C} - 1} C < 1$$

so that by a union bound,

$$\mathbb{P}_{n_k}(m_{cut} \geq C) \leq \sum_{a=1}^{n_k-1} \binom{n_k}{a} p^{a(n_k - a)}$$

which tends to 0 as $k \rightarrow \infty$. \square

COROLLARY 7 *For almost all graphs in $\mathcal{G}(\lfloor \alpha N \rfloor)$ and for all $\epsilon > 0$,*

$$\frac{1}{\alpha} \leq m_{cut}(G) < \frac{1}{\alpha} + \epsilon$$

Proof. The first inequality is from proposition 5, the second from propositions 6 and 1. \square

4 Conclusion

We have presented the distance bound, which gives exactly the number of wavelengths needed for wavelength routing in random networks of order tending to infinity. A simple greedy routing algorithm which uses no wavelength conversion has been shown to be sufficient. It is generalisable to various different routing regimes, such as allowing multiple links between nodes: we suspect that if a regime were to be suitably framed, the distance bound would again prove to be exact. The predicted number of wavelengths is in reasonable accord with the results of a heuristic routing algorithm applied to computer-simulated random networks.

The partition bound is a lower bound on the number of wavelengths needed. It is far worse on large random networks, but works very well on certain practical networks: for these networks, the lower bound agrees with a heuristically-found feasible solution, and is therefore exact. This suggests that there are factors which have not been taken into account which influence the design of networks: they may have been designed for a particular non-uniform traffic matrix, or they may rely upon different links having different capacities. There may be constraints on the topology of the graph: arising from reliability requirements for the network, such as stipulating that the graph remain connected when a single edge is removed; or from engineering limitations, such as limiting the node degrees; or from the physical geography of the network.

An important unanswered question is that of how an existing network might best be modified to improve the number of wavelengths needed: Which links are superfluous? Where is the best place for a new link?

5 Appendix

Proof of Theorem 3.

The proof involves constructing a simple greedy lightpath-assignment algorithm, and finding a deterministic approximation for its behaviour. It is based on the proof for the Triangle Problem given by Hajek [7]. This method is more generally applicable: in the same paper he analyses the Independent Set Problem using a similar technique.

We may assume that $0 < \alpha < 1$.

The Algorithm.

Let $G \in \mathcal{G}(n; \alpha)$. The edges that are present correspond to optical links, each capable of carrying colours $\{1 \dots \Lambda\}$.

First, to all node-pairs which are directly joined by an edge we assign a direct lightpath of colour Λ . We then seek to connect all remaining node-pairs by assigning them an indirect lightpath of length 2. We call the node-pairs not joined by an edge *virtual edges*.

Suppose some such lightpaths have been assigned, and that the lightedges and virtual edges involved have been marked as used. If there are no unused virtual edges remaining, declare the algorithm successful and stop. Otherwise, select an unused virtual edge at random. Examine all lightpaths of length two between its endpoints. If both lightedges of a lightpath are unused, call the lightpath available. If none of the lightpaths are available, declare the algorithm unsuccessful and stop. Otherwise, select one at random. Assign that lightpath to the virtual edge, and mark those two lightedges and the virtual edge as used. Repeat until the algorithm stops.

Let E be the set of edges in G , \bar{E} the set of lightedges, and R the set of virtual edges. Write $e(\lambda)$ for the lightedge through edge e of colour λ . We say a *triangle of colour* λ for a set of two lightedges of colour λ and a virtual edge, which join each pair from a set of three nodes; and call it *good* if the two lightedges and the virtual edge are unused.

We define, for $r \in R$ and $e \in E$ and $x \in R \cup \bar{E}$,

$$S_x = \begin{cases} k & \text{if } x \text{ is used at step } k \\ \infty & \text{if } x \text{ is never used} \end{cases}$$

$$X_r^\lambda(k) = \begin{cases} \text{number of good triangles of colour } \lambda \text{ containing } r, & \text{if } k < S_r \\ X_r^\lambda(S_r - 1) & \text{otherwise} \end{cases}$$

$$Y_e^\lambda(k) = \begin{cases} \text{number of good triangles of colour } \lambda \text{ containing } e(\lambda), & \text{if } k < S_{e(\lambda)} \\ Y_e^\lambda(S_{e(\lambda)}) & \text{otherwise} \end{cases}$$

For any process $(X(k))_{k \in \mathbb{Z}^+}$ write $\Delta X(k) = X(k+1) - X(k)$. Let \mathcal{F}_k be the σ -algebra generated by R and $X_r^\lambda(1) \dots X_r^\lambda(k)$, $Y_e^\lambda(1) \dots Y_e^\lambda(k)$ for $1 \leq \lambda \leq \Lambda - 1$ and $e \in E$, $r \in R$. Note that for a virtual edge r , if $k < S_r$ then

$$\mathbb{E}(\Delta X_r^\lambda(k) | \mathcal{F}_k) = - \sum_f \sum_s \frac{1}{|R| - k} \left(\sum_{\gamma=1}^{\Lambda-1} X_s^\gamma(k) \right)^{-1} \quad (2)$$

where f ranges over the $2X_r^\lambda(k)$ lightedges in available lightpaths of colour λ between the endpoints of r , and s ranges over the $Y_f^\lambda(k) - 1$ virtual edges apart from r which might be routed through s . Similarly, for an edge e , if $k < S_{e(\lambda)}$ then

$$\begin{aligned} \mathbb{E}(\Delta Y_e^\lambda(k) | \mathcal{F}_k) &= - \sum_f \sum_s \frac{1}{|R| - k} \left(\sum_{\gamma=1}^{\Lambda-1} X_s^\gamma(k) \right)^{-1} \\ &\quad + \sum_t \frac{1}{|R| - k} \left[1 - \left(\sum_{\gamma=1}^{\Lambda-1} X_t^\gamma(k) \right)^{-1} \right] \end{aligned} \quad (3)$$

where f ranges over the $Y_e^\lambda(k)$ edges that may form a lightpath with $e(\lambda)$, s ranges over the $Y_f^\lambda(k) - 1$ virtual edges which may be routed via a , except for the one containing f , and t ranges over the virtual edges that may be routed via e .

The Approximation.

Choose α_0 with $\alpha < \alpha_0 < 1$. We will approximate the behaviour of the algorithm over the first $\frac{1}{2}n^2(1 - \alpha_0)$ steps.

The idea behind the approximation is that for large n the edges are nearly independent. If a certain number of lightedges have been used, then the probability that a given lightedge is unused is nearly binomial and independent of other lightedges. There are roughly $\frac{1}{2}n^2\alpha(\Lambda - 1)$ available lightedges altogether, and after k steps we have used $2k$ of them. There are roughly $\frac{1}{2}n^2(1 - \alpha)$ virtual edges altogether, and after k steps we have used k of them. We define

$$\begin{aligned} X(k) &= n\alpha^2(1 - q)^2 \\ Y(k) &= 2n(1 - \alpha)(1 - r)\alpha(1 - q) \end{aligned}$$

where $q = 2k/(\frac{1}{2}n^2\alpha(\Lambda - 1))$ and $r = k/(\frac{1}{2}n^2(1 - \alpha))$.

LEMMA 8 *There exist constants L and δ such that whenever $\epsilon_1 \dots \epsilon_5$ satisfy $\epsilon = \max(|\epsilon_1|, \dots, |\epsilon_5|) \leq n\delta$ then for sufficiently large n and $k \leq \frac{1}{2}n^2(1 - \alpha_0)$,*

$$\left| \Delta X(k) + \frac{2}{\frac{1}{2}n^2(1 - \alpha) - k + \frac{1}{2}n\epsilon_1} \frac{(X(k) + \epsilon_2)(Y(k) + \epsilon_3)}{(\Lambda - 1)(X(k) + \epsilon_4)} \right| \leq \frac{L}{n^2}(1 + \epsilon)$$

$$\left| \Delta Y(k) + \frac{1}{\frac{1}{2}n^2(1 - \alpha) - k + \frac{1}{2}n\epsilon_1} \left((Y(k) + \epsilon_2) + \frac{(Y(k) + \epsilon_3)(Y(k) + \epsilon_4)}{(\Lambda - 1)(X(k) + \epsilon_5)} \right) \right| \leq \frac{L}{n^2}(1 + \epsilon)$$

Proof. Define

$$x(t) = \alpha^2(1 - q)^2$$

$$y(t) = 2(1 - \alpha)(1 - r)\alpha(1 - q)$$

where $q = 2t/(\alpha(\Lambda - 1))$, and $r = t/\alpha$. Then

$$\dot{x}(t) = \frac{-2}{1 - \alpha - t} \left(\frac{y(t)}{\Lambda - 1} \right)$$

$$\ddot{x}(t) = \frac{8}{(\Lambda - 1)^2}$$

Since $X(k) = nx(k/(\frac{1}{2}n^2))$, Taylor's Theorem gives

$$\Delta X(k) - \frac{2}{n} \dot{x}\left(\frac{k}{\frac{1}{2}n^2}\right) = \frac{16}{n^3(\Lambda - 1)^2} \quad (4)$$

Now, $\dot{x}(t)$ is differentiable with constant derivative, and for $0 \leq t \leq 1 - \alpha_0$,

$$0 < \alpha_0 - \alpha < 1 - \alpha - t \leq 1 - \alpha$$

$$0 < x(1 - \alpha_0) \leq x(t) \leq 1$$

$$0 \leq y(t) \leq 1.$$

Therefore there exist constants L and δ for which whenever $\eta, \eta_1 \dots \eta_4$ are such that $\eta = \max(|\eta_i|) < \delta$, then

$$\left| \dot{x}(t) + \frac{2}{1 - \alpha - t - \eta_1} \frac{(x(t) + \eta_2)(y(t) + \eta_3)}{(\Lambda - 1)(x(t) + \eta_4)} \right| \leq \frac{1}{2}L\eta$$

The triangle inequality applied to this and equation 4 yields the desired equation for $\Delta X(k)$, when $n \geq \frac{16}{L(\Lambda - 1)^2}$. The equation for $\Delta Y(k)$ is obtained similarly. \square

Relation between algorithm and approximation.

We first prove a small lemma which will be used later.

LEMMA 9 *If X is a random variable such that $|X| \leq c$, $\mathbb{E}X \leq 0$, and $\mathbb{E}(X^2) \leq \frac{c}{n}$, then for ζ sufficiently small, depending only on c ,*

$$\mathbb{E}e^{\zeta X} \leq e^{\zeta^2 c/n}$$

Proof. Since $|X| \leq c$, we may use Dominated Convergence to show

$$\begin{aligned}
\mathbb{E}(e^{\zeta X}) &= 1 + \zeta \mathbb{E}X + \sum_{k=2}^{\infty} \frac{\zeta^k}{k!} \mathbb{E}X^k \\
&\leq 1 + \mathbb{E}X^2 \sum_{k=2}^{\infty} \frac{\zeta^k}{k!} c^{k-2} \\
&= 1 + \mathbb{E}X^2 (e^{\zeta c} - 1 - \zeta c)/c^2 \\
&\leq \zeta^2 \mathbb{E}X^2 \quad \text{for } \zeta \text{ sufficiently small} \\
&\leq 1 + \zeta^2 c/n \\
&\leq e^{\zeta^2 c/n}
\end{aligned}$$

as required. \square

We define

$$\begin{aligned}
Z_e^\lambda(k) &= \begin{cases} X_e^\lambda(k) - X(k) & \text{for } e \in R \\ Y_e^\lambda(k) - Y(k) & \text{for } e \in E \end{cases} \\
\Phi_k &= \max_{\substack{1 \leq \lambda < \Lambda \\ e: k < S_e(\lambda)}} |Z_e^\lambda(k)| \wedge \frac{2}{n} (|R| - \frac{1}{2}n^2(1 - \alpha))
\end{aligned}$$

LEMMA 10 *Fix ξ with $0 < \xi \leq \delta$. Then*

$$\mathbb{P}(\Phi_k \leq \xi n \text{ for all } k \leq \frac{1}{2}n^2(1 - \alpha_0)) \rightarrow 0$$

Proof. Fix $e \in E$ and $\lambda < \Lambda - 1$. Suppose that $k \leq S_{e(\lambda)}$. Equations 2 and 3 and lemma 8 show that if $\Phi_k < n\delta$ then

$$\left| \mathbb{E}(\epsilon_e^\lambda(k) | \mathcal{F}_k) \right| \leq \frac{L}{n^2} (\Phi_k + 1)$$

We will show that the random values $X_e^\lambda(k)$ and $Y_e^\lambda(k)$ are likely to lie within a certain distance r_k of their expected values. The distance r_k will be chosen sufficiently large to be likely, but not so large as to be uninformative. Let

$$\begin{aligned}
r_0 &= n^a, \text{ for some fixed } a \in (\frac{1}{2}, 1) \\
r_k &= r_0 + \sum_{j=0}^{k-1} \frac{L}{n^2} (r_j + 1)
\end{aligned}$$

Note that $r_k = (n^a + 1)(1 + L/n^2)^k - 1$ so that $r_k \leq r_{n^2} \leq (n^a + 1)e^L$. Suppose that n is sufficiently large that $r_k < n\xi$. Let

$$\begin{aligned}
D &= \min\{k \geq 0 : \Phi_k \geq r_k\} \wedge \frac{1}{2}n^2(1 - \alpha_0) \\
\epsilon_e^\lambda(k) &= \begin{cases} Z_e^\lambda(k+1) - Z_e^\lambda(k) & \text{for } k < D \\ \epsilon_e^\lambda(D-1) & \text{for } k \geq D \end{cases}
\end{aligned}$$

Then $|\mathbb{E}(\epsilon_e^\lambda(k) | \mathcal{F}_k)| \leq \frac{L}{n^2} (r_k + 1)$. Also, by writing ϵ_e^λ in terms of X_e^λ and Y_e^λ , we can find a constant c independent of n such that

$$\begin{aligned}
|\epsilon_e^\lambda(k)| &\leq c \\
\mathbb{E}(\epsilon_e^\lambda(k)^2 | \mathcal{F}_k) &\leq c/n
\end{aligned}$$

Now for a given λ and e (which will be suppressed), defining $\eta_i^\pm = \pm c_e^\lambda(i) - \frac{L}{n^2}(r_i + 1)$,

$$\begin{aligned} \{|Z_e^\lambda(k)| \geq r_k\} &\subseteq \{|Z(0)| \geq \frac{1}{2}n^a\} \cup \{|\epsilon_0 + \dots + \epsilon_{k-1}| \geq r_k - \frac{1}{2}n^a\} \\ &= \{|Z(0)| \geq \frac{1}{2}n^a\} \cup \{\eta_0^+ + \dots + \eta_{k-1}^+ \geq \frac{1}{2}n^a\} \cup \{\eta_0^- + \dots + \eta_{k-1}^- \geq \frac{1}{2}n^a\}. \end{aligned}$$

From the rate function for the binomial, the probability of the first term decreases exponentially in n^{2a-1} . For the remaining terms, note that from the corresponding conditions on ϵ_e^λ there is a constant d such that $\mathbb{E}(\eta_k|\mathcal{F}_k) \leq 0$ and

$$\begin{aligned} |\eta_k| &\leq d \\ \mathbb{E}(\eta_k^2|\mathcal{F}_k) &\leq d/n \end{aligned}$$

Applying lemma 9 with $\zeta = \frac{1}{4d}n^{a-1}$, it follows that for n sufficiently large and $k \leq n^2$,

$$\begin{aligned} \mathbb{P}(\eta_0^\pm + \dots + \eta_{k-1}^\pm \geq \frac{1}{2}n^a) &\leq \mathbb{E}\left[\exp\left(\zeta\left(\sum_i \eta_i - \frac{1}{2}n^a\right)\right)\right] \\ &\leq \exp(\zeta^2 dk/n - \zeta n^a) = \exp\left(-\frac{1}{16c}n^{2a-1}\right) \end{aligned}$$

Thus the probability that $|Z_e^\lambda| \geq r_k$ for some e and λ and $k \leq \min(D, S_{e(\lambda)})$ tends to 0 as $n \rightarrow \infty$. Since also $\mathbb{P}(|R| - \frac{1}{2}n^2(1 - \alpha) > \frac{1}{2}n^{a+1})$ tends to zero, we have

$$\mathbb{P}(\Phi_k \geq r_k \text{ for some } k \leq D) \rightarrow 0$$

and so

$$\mathbb{P}(\Phi_k \leq \xi n \text{ for all } k \leq \frac{1}{2}n^2(1 - \alpha_0)) \rightarrow 0$$

as desired. \square

Almost certain success.

Since $\Lambda > \frac{2}{\alpha} - 1$, we can choose ξ sufficiently small that $\xi \leq \delta$ and $X(k) > \xi n$ for $k \leq \frac{1}{2}n^2(1 - \alpha_0)$. Now, if $\Phi_k \leq \xi n$, and k is such that $k \leq \frac{1}{2}n^2(1 - \alpha_0)$ and there are some virtual edges left unused, then for some unused virtual edge r and colour λ , $X_r^\lambda(k) > 0$. Hence the algorithm cannot fail at step k . Since, by lemma 10, $\mathbb{P}(\Phi_k \leq \xi n \text{ for all } k \leq \frac{1}{2}n^2(1 - \alpha_0)) \rightarrow 0$, the probability that the algorithm fails before $\frac{1}{2}n^2(1 - \alpha_0)$ virtual edges have been assigned tends to zero as $n \rightarrow \infty$.

This does not assign lightpaths to all virtual edges. However, we are now in position to prove the main theorem, by considering a modified algorithm. Choose ϵ with $\epsilon < \alpha - \frac{2}{\Lambda+1}$. Independently delete each link with probability ϵ . Then run the original algorithm on the modified graph, though only using undeleted links.

LEMMA 11 *This modified algorithm is almost surely successful.*

Proof. Call the deleted edges false virtual edges. To an observer who could not distinguish between true and false edges, the algorithm would seem statistically identical to the original algorithm, although run for a larger set of edges, up until the time when all the original virtual edges have been assigned lightpaths. The probability that an undeleted link is present is $\alpha(1 - \epsilon)$. Replacing α by $\alpha(1 - \epsilon)$ and setting α_0 equal to $\alpha(1 - \frac{1}{2}\epsilon)$, we see that the modified algorithm almost surely

does not fail within the first $\frac{1}{2}n^2(1 - \alpha_0)$ steps. But the number of true virtual edges is almost surely less than $\frac{1}{2}n^2(1 - \alpha_0)$. Therefore, all the true virtual edges may be assigned lightpaths, almost surely. \square

This lemma completes the proof of theorem 3.

References

- [1] Alok Aggarwal, Amotz Bar-Noy, Don Coppersmith, Rajiv Ramaswami, Baruch Schieber, and Madhu Sudan. Efficient routing in optical networks. Technical report, IBM, 1994.
- [2] Joseph Bannister, Mario Gerla, and Milan Kovacević. Routing in optical networks. In Martha E. Steenstrup, editor, *Routing in Communications Networks*. Prentice-Hall, New Jersey, 1995.
- [3] Stefano Baroni. Analysis and design of wavelength routing and allocation algorithms for application in advanced high capacity optically amplified communication networks. UCL Dept. of Electronic and Electrical Engineering, 1995.
- [4] Stefano Baroni, Polina Bayvel, and John E. Midwinter. Wavelength requirements in dense wavelength-routed optical transport networks with variable physical connectivity. UCL Dept. of Electronic and Electrical Engineering, 1995.
- [5] Vincent W. S. Chan. All-optical networks. *Scientific American*, September 1995.
- [6] Richard J. Gibbens. *Dynamic Routing in Circuit-Switched Networks: The Dynamic Alternative Routing Strategy*. PhD thesis, University of Cambridge, July 1988. Chapter 5.
- [7] Bruce Hajek. Average case analysis of greedy algorithms for Kelly’s Triangle Problem and the Independent Set Problem. In *IEEE Conference on Decision and Control*, 1987.
- [8] Frank P. Kelly. Wavelength routing in optical networks. Private communication, July 1995.
- [9] Rajiv Ramaswami and Kumar N. Sivarajan. Routing and wavelength assignment in all-optical networks. Technical Report RC 19592, IBM, 1994.
- [10] Rajiv Ramaswami and Kumar N. Sivarajan. Design of logical topologies for wavelength-routed all-optical networks. In *Proc. INFOCOM*, pages 1316–1325, 1995.
- [11] M. Sabry and J. E. Midwinter. Towards an optical ether. *IEE Proc.-Optoelectron.*, 141(5), 1994.
- [12] Zhensheng Zhang and Anthony Acampora. A heuristic wavelength assignment algorithm for multihop wdm networks with wavelength routing and wavelength reuse. In *Proc. IEEE INFOCOM*, pages 534–543, 1994.