Physical Topology Design for All-Optical Networks

Huan Liu, Fouad A. Tobagi
Department of Electrical Engineering
Stanford University
{huanliu, tobagi}@stanford.edu

***Abstract*— There are many advantages of deploying an all-optical network. Unfortunately, there are still few guidelines on how to properly design the physical topology for such a network. We propose an efficient physical topology design algorithm and we use the asymptotic growth rate of the provisioned capacity as a metric to compare various design alternatives. A higher asymptotic growth rate translates directly into higher deployment cost for large networks. Our study shows that taking fiber length into consideration can lead to lower capacity requirement. We also find that a sufficiently large fiber-to-node ratio is necessary in order to minimize the asymptotic growth in the provisioned capacity, increase capacity utilization and minimize the need for wavelength conversion. We study a real network and find that its fiber-to-node ratio is too low. As a result, large provisioned capacity is required and less than 55% of the capacity is usable. By increasing the ratio, we can reduce the provisioned capacity and achieve close to 80% utilization.**

I. Introduction

Compared to the broadcast-based optical network architectures [1][2][3][4][5], Wavelength Routed All-optical Network (WRAN) utilizing the WDM technology promises to greatly increase the transport capacity at much reduced cost. A connection, also known as a lightpath [6], only occupies one wavelength on each fiber link along the physical route used to connect the two end nodes. Thus, the same wavelength on other fiber links could be reused for other lightpaths to increase the utilization of the provisioned wavelengths.

Besides the increased utilization, WRAN has many other advantages. Since a lightpath is routed transparently through the WRAN, that is, bypassing intermediate nodes without packet processing or costly opto-electronic conversion, much of the queuing delay and electronic equipment cost can be eliminated. The cost savings in electronic equipment will be significant [7], especially when the line speed is very high. Furthermore, WRAN can be easily and cheaply upgraded when the interface speed is increased. This is because optical switching is agnostic to the underlying data rate of an optical channel (up to a certain limit because the channel bandwidth limits the maximum data rate), and thus, the intermediate optical switches do not have to be upgraded when the line speed increases.

Given the high cost of deploying a WRAN, it is important to design the physical (fiber) topology to minimize the total capital investment. The total cost of deploying a WRAN is the sum of two cost components—the link cost and the node cost. The link cost, i.e., the cost of laying down fibers to interconnect nodes, is a function of the total fiber length L. The node cost, i.e., the cost of the all-optical wavelength switch is a function of the number of wavelengths W that are provisioned on each fiber link. In general, there is a tradeoff between L and W—more L will translate into less W and vice versa.

One can design the physical topology to use the minimum amount of fiber by connecting the nodes using a minimum spanning tree. Even though the link (fiber) cost is at the minimum, the node (wavelength) cost will be very high. Alternatively, one can connect all node pairs using direct fibers. The node (wavelength) cost is at its minimum since $W = 1$. However, the link (fiber) cost will be very high. The optimum design with the minimum total (link and node) cost will be between these two extreme solutions.

To pick the best topology with the minimum total cost, one has to solve the problem of designing a physical topology to minimize the number of wavelengths W required given a budget on L. We will present a comprehensive treatment of the design problem, including both a mathematical problem formulation and an efficient heuristic algorithm. Using the design algorithm we proposed, one can design the topology with the minimum total (link and node) cost by repeatedly run the algorithm for different L, comparing the resulting solutions based on the actual cost functions, and then picking the one with the lowest total cost. In order to be independent of the actual cost functions, we study the tradeoff between L and W. In addition, we also derive design principles and guidelines, which are unfortunately nonexistent as of now.

To evaluate the various design alternatives—some use more fiber but fewer wavelengths and some use less fiber but more wavelengths—we propose to use the provisioned capacity $C = LW$ (capacity for short in the following) as a metric. C essentially is the bandwidth-distance product. It is used to measure the amount of network resources that have to be provided for a given set of lightpath demands. Since there is a cost associated with providing the network resources, naturally, it is desirable to *minimize* the provisioned capacity. Note that our use of the term "capacity" may be different from the literature in other contexts, where some fixed network resources are assumed given and the goal is to *maximize* the capacity, i.e., the throughput.

C is a more fair metric compared to other metrics such as MW, where M is the number of fiber links. Consider a sample network as shown in Fig. 1, where 6 nodes lie on a straight line with an internodal distance of 1. Let us assume that the traffic demands are uniform all-to-all, i.e., we need to establish a lightpath between every pair of nodes. Two possible physical topologies are depicted in the figure. Topology (a) is a linear topology, and it requires 9 wavelengths. To see this, we just need to consider link 3–4. Since there are three nodes on either side of this link, there are a total of 9 lightpaths crossing this link; hence, 9 wavelengths are needed. Topology (b) requires 8 wavelengths. To see this, we can consider either link 3–4 or 4–5. Since there are 4 and 2 nodes on either side of these two links respectively, there are a total of 8 lightpaths crossing these two links; hence, 8 wavelengths are needed. Since both topologies have 5 edges, it may be tempting to choose topology (b) because it requires fewer wavelengths. However, the total fiber length in topology (b) is 7 units long, and therefore, 56 units of capacity is required. In contrast, topology (a) only uses 45 units of capacity. The reason for the high capacity in topology (b) is because the demands between nodes 1,2 and 1,3 are not routed along the direct line between the two end nodes. In particular, both demands take a detour—going to node 4 first and then to the destination nodes. Demand 1, 2 wastes 4 units of capacity and demand 1, 3 wastes 2 units of capacity.

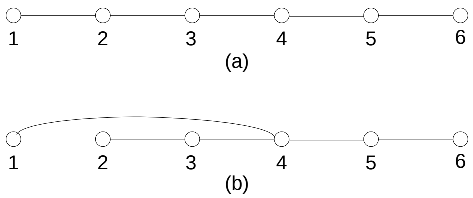

Fig. 1. 6 nodes lie in a line with an inter-nodal distance of 1. Physical topology (a) requires 9 wavelengths and physical topology (b) requires 8 wavelengths.

The capacity metric measures how efficiently the provisioned network resources are utilized—higher capacity for the same set of lightpaths means that the provisioned capacity is less efficiently utilized. Unfortunately, the capacity metric does not directly reflect the cost of deploying a WRAN. A design with a higher capacity may have lower cost than a design with a lower capacity simply because of the differences in the cost functions for L and W.

To derive design guidelines independent of the actual cost functions, we focus on the asymptotic growth rate of the capacity, which can be used to determine the lower-cost design alternative when N is large, where N is the number of nodes in the network. Consider two design alternatives. Design A uses 1 unit of fiber, and the resulting capacity can be expressed as $a_1 N^{e_1}$. Design B uses 2 units of fiber, and the resulting capacity can be expressed as $a_2 N^{e_2}$. Let us assume $e_1 > e_2$. From the definition of capacity, we can see that Design A uses a factor of $c = 2\frac{a_1 N^{e_1}}{a_2 N^{e_2}}$ more wavelengths than Design B. When N is small, c is small (slightly more than $2\frac{a_1}{a_2}$). If the cost of the fiber is a lot more than the cost of the wavelengths, Design A will be the lower-cost solution. However, when N is large, c will be large because of the exponential term. At some point, the cost of the wavelengths will dominate, and Design B becomes the lower-cost solution. As long as e_1 is greater than e_2, even if e_1 is only slightly larger than e_2, Design B will be a lower-cost solution when the network is large enough. Hence, the exponent is an indicator of the actual cost for large-size networks.

Throughout this paper, we will assume that all fiber links will support the same number of wavelengths. This is a realistic assumption because it is necessary to ensure the maximum interoperability between neighboring nodes. Having different W on each link may make sense currently because WDM systems are used primarily as transmission systems and there are many opto-electronic conversions at a node. However, since there is no opto-electronic conversion in the WRAN, having different W not only will require unsymmetrical optical cross connects, but also will limit the flexibility in routing lightpaths.

A. *Traffic model*

We first consider WRANs that support full-mesh connectivity, i.e., there are T lightpaths to establish between every pair of nodes, where T is a constant that is the same for all node pairs. Our design algorithm could be easily generalized to other lightpath connection patterns. We will present our results on non-uniform lightpath connections at the end.

B. *Prior work*

The physical topology design problem has been studied before in the literature. However, most of these works consider only the case where the cost is proportional to the number of fiber links regardless of their lengths[8] [9]. The work in [2] considered only broadcast-based optical networks, and the topology is restricted to a tree. The work in [10] takes fiber length into consideration. They considered a different problem where W is given and the goal is to minimize the total cost of fiber. The algorithm proposed runs much slower. A problem instance with 100 nodes requires 11 hours. This is not suitable for a tradeoff study like ours, where hundreds of thousands of problem instances have to be solved.

There have been several related studies on the wavelength requirements in an optical network. The first work on wavelength requirements to support full-mesh connectivity was reported in [11]. The authors found that the ensemble average number of wavelengths (W) required is only dependent on α, where $\alpha = \frac{2M}{N(N-1)}$ is the ratio between the number of edges in the network and the number of edges required to fully connect all node pairs. The result only applies when the fiber links are picked randomly, i.e., when short fibers and long fibers have an equal chance of being picked. If many short fibers are picked, then W could be much higher. Also, the work in [11] did not answer the question of what α one should use in designing a network.

In [12], the authors gave an approximate equation for the required number of wavelengths W. Unfortunately, the derivation of the equation was not shown. In section II, we give an equivalent derivation for a more accurate result that is a constant factor away from the one given in [12].

C. *Organization of paper*

The rest of the paper is organized as follows. In section II, we first give a lower bound and an approximate equation for

C. In section III, we present our physical topology design algorithm. In section IV, we evaluate our algorithm and show the tradeoff between L and W. In section V, we consider one real-life network and show how our design guideline could be applied. Lastly, we conclude in section VI.

II. Lower bound and approximate equation

A. Lower bound

The most capacity-efficient way to establish a lightpath between two nodes i and j is to lay down a direct fiber between the two nodes and using one wavelength on that fiber for the connection. The capacity used for this connection is $d_{ij} \cdot 1 = d_{ij}$, where d_{ij} is the physical distance between the two nodes. Another way to establish the lightpath is by hopping through one or more other nodes and use one wavelength on each fiber link along the way. Hopping through other nodes will take strictly more capacity unless the intermediate nodes lie on the direct (straight) line between node i and j. In the following, we say a lightpath uses direct-line routing if the lightpath only hops through zero or more nodes along the direct line. We say the lightpath uses non-direct-line routing if otherwise.

Summing up the minimum capacity required for each lightpath, we can derive the lower bound on the provisioned capacity for full-mesh connectivity as follows:

$$C_{LB} = T \sum_{i=1}^{N} \sum_{j=1}^{N} d_{ij} \tag{1}$$

Note that this lower bound holds regardless of the relative positions of the nodes (e.g., not necessarily uniformly placed). It also holds even if d_{ij} is not the Cartesian distance between the nodes. For example, there may be a physical constraint that forces a fiber link not to be laid along the direct line (e.g., mountains, rivers). In such cases, d_{ij} denotes the actual length of the fiber that has to be laid down.

This lower bound cannot be achieved unless many fiber links are laid down. If only a few fiber links are available, some lightpaths will not be able to use direct-line routing, either because there is no fiber link on the direct line or because there is no wavelength left on the fiber links on the direct line.

B. Approximate equation

In [12], an approximate equation for W was given. Unfortunately, the derivation of the equation was not shown and the steps cannot be easily reproduced. In this section, we give an equivalent derivation for a result that is only a constant factor away from that in [12].

To be consistent with the rest of the paper, we assume that the N nodes of a network are uniformly distributed in a square area, as opposed to a circular disk area used in [12], with unit size as shown in Fig. 2. We note that the choice of a square area is arbitrary, and any bounding area could be chosen. We also assume that the fiber links are uniformly distributed. The assumptions on uniform node placement and link placement are made in order to derive an equation that approximates the average of that of a large ensemble of random networks.

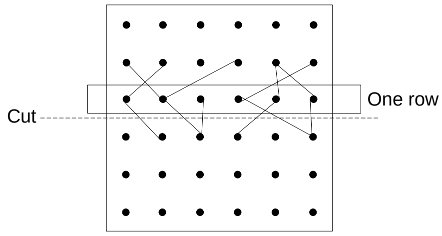

Fig. 2. N nodes uniformly distributed in a square.

Our goal is to compute the number of lightpaths that will cross the cut in the middle of the square and also the number of fiber links that will cross the cut. The minimum number of wavelengths required will simply be a ratio of these two numbers.

Since we are considering the cut right in the middle of the square, there are exactly $N/2$ nodes on either side of the cut. Therefore, it is easy to see that there are $(N/2)^2 T$ lightpaths crossing the cut.

To compute the number of fibers crossing the cut, we first compute the number of nodes n in the row right above the cut. Since the area of the square is 1, each side of the square is exactly 1 unit long. Each node will take up a square area of roughly $1/N$ with each side of it being $1/\sqrt{N}$ long. Dividing the length of the row, the number of nodes in the row is then $n = 1/(1/\sqrt{N}) = \sqrt{N}$.

If the average node degree is d, there are nd edges originating from these n nodes in the row. Among them, nh edges go to nodes in the same row, and the rest $n(d-h)$ go to nodes in the neighboring rows. Because of the uniform link placement assumption, the number of edges staying in the same row should be proportional to the number of nodes in the row, i.e., $h/d = n/N$. Since the number of edges going to the row below (across the cut) is half the number of edges leaving the row, it can be expressed as:

$$\frac{1}{2}n(d-h) = \frac{1}{2}nd(N-n)/N = \frac{1}{2}d(\sqrt{N}-1)$$

If the average fiber length L_f is more than the average node distance L_n, some fiber links originating from rows above could also cross the cut. Adding these fiber links, the number of edges crossing the cut will increase by a factor of L_f/L_n.

Dividing the number of lightpaths crossing the cut by the number of fiber links crossing the cut, we will get the minimum required number of wavelengths W as follows:

$$\begin{aligned} W &= \frac{N^2 T/4}{\frac{1}{2}d(\sqrt{N}-1)L_f/L_n} \\ &= \frac{1}{2-2/\sqrt{N}} \frac{N^{3/2} T L_n}{dL_f} \\ &= B\frac{N^{3/2} T L_n}{dL_f} \end{aligned} \tag{2}$$

where $B = \frac{1}{2-2/\sqrt{N}}$ is almost a constant, especially when N is large.

Since the total fiber length is $L = NdL_f/2$, we can rearrange the terms in equation (2) to get the expression for the provisioned capacity.

$$C = WNdL_f/2 = \frac{B}{2}N^{5/2}TL_n \tag{3}$$

If a circular disk area is assumed, the same derivation will yield a result that differs from the one in [12] by a constant of $\frac{\sqrt{\pi}}{2}$. We believe our result is more accurate because our result matches the lower bound for the same network (nodes uniformly distributed in a square) almost perfectly. The matching is not surprising because, in the derivation of the approximate equation, we have implicitly assumed that each lightpath will be routed along the direct line, just like we did in deriving the lower bound.

Equation (3) shows that the capacity will scale as $N^{2.5}$ even though the number of lightpaths will only scale as N^2. From the derivation, we can see that the extra 0.5 factor in the exponent comes from the fact that the number of fiber links crossing the cut is on the order of $\sqrt{N}$, but the number of lightpaths crossing the same cut is on the order of N^2. Alternatively, this extra 0.5 factor can be also viewed as coming from the fact that the diameter of the network is on the order of $\sqrt{N}$. Since some lightpaths will have to cross the network to reach their destinations, they will require an order of $\sqrt{N}$ more capacity. This is true whether the lightpaths cross the network using few long fiber links or many short fiber links because only the fiber length comes in the definition of the capacity, not the number of fiber links.

In the derivation, we have assumed that all N nodes are uniformly distributed in the area in order to approximate the average case. However, nodes in real networks are almost never distributed evenly inside the bounding area. But a change in this assumption will only affect the constant term. As long as the physical topology is two-dimensional, the diameter of the network will be on the order of $\sqrt{N}$, and some lightpaths will have to use $\sqrt{N}$ more capacity. We will look at a real network in section V, where nodes are not uniformly distributed. We will see that our observation still applies.

In addition to the non-uniform distribution of nodes, the bounding area of a real network is also seldomly square. If we assume a different bounding area (such as a circle), the same derivation will give a result that differs again only in the constant term. Even for irregular-shaped bounding area, the exponential term remains the same as long as it is a two-dimensional area (as opposed to a line). Since we are more interested in the asymptotic growth rate of the capacity, we will assume a square area throughout this paper without loss of generality.

The average node distance L_n in equation (3) can be thought of as a scaling factor. Consider a topology, if we double the length of each fiber link (and thus push the nodes further apart), the topology still remains the same. The only thing changed is that L_n is doubled. For an arbitrary network, L_n can be determined as follows. We take the area A of the bounding box of the topology and divide it by the number of nodes N, we will get the average area of each node to be A/N. To derive the average node distance, we assume these nodes are uniformly located and each node will take up a square area with each side being $\sqrt{A/N}$ long, and then the distance between two neighboring nodes will be $L_n = \sqrt{A/N}$. If the fiber distance d_{ij} is not the Cartesian distance (e.g., physical constraint forces some fiber to be laid along non-direct lines), L_n can be adjusted proportionally according to the actual fiber distance. Again, using a different value of L_n will only change the constant term, not the exponential term. For simplicity of discussion, we let $L_n = 1$ without loss of generality for the rest of this paper.

Equation (3) predicts two things. First of all, it predicts that the capacity will scale linearly with T. We observe in simulations that this is indeed true. Therefore, this paper will not focus on the scaling as a function of T. In the following, we assume $T = 1$.

The second predication this equation presents is that the asymptotic growth rate of the capacity will remain the same no matter how W and L scale. In other words, W can be traded off with L equally.

Even though the scaling as a function of T (the first prediction) matches very well with our observations in simulation, given the simplicity of the analysis, there are few reasons to believe that the second prediction is true. There are two factors that are not modeled by either the lower bound or the approximate equation. For the first factor (factor F1), we have assumed that each lightpath will be routed along the direct line between the two end nodes. This is rarely the case even if each lightpath goes through the minimum number of fiber links. The reason is because there simply may not be fiber links on the direct line. The second factor (factor F2) not modeled is the utilization of the provisioned capacity. In general, it is not possible to even out the load on each fiber link such that the same number of wavelengths is used on every link, i.e., some wavelengths cannot be utilized because of topological constraints.

III. Physical topology design

The problem of designing a physical topology to minimize W is clearly an NP-complete problem because even if the topology is known, the problem of determining how many wavelengths are needed (the Routing and Wavelength Assignment Problem) is known to be NP-complete.

In this section, we first formulate the physical topology design problem as an Integer Linear Programming (ILP) problem and then propose a practical heuristic algorithm. The optimization problem takes the budget on fiber (in the form of a fixed fiber-to-node ratio $f = L/N$) as a constraint and designs a topology to minimize the number of wavelengths W required.

We consider the topology design problem both with and without the Wavelength Continuity Constraint (WCC), which requires each lightpath to be assigned a unique wavelength on each fiber link it traverses. The WCC constraint could limit the amount of usable capacity. If a wavelength is used on a link, any other connections that must pass through this link cannot use the same wavelength again.

A. ILP problem formulation

For brevity, we only show the formulation which relaxes the wavelength continuity constraint (WCC). The formulation can be easily modified if the WCC constraint is enforced. We will use the following notations in the formulation:

- z_{ml}: This is the fiber link variable. $z_{ml} = 1$ if node m is connected to node l via a fiber link. $z_{ml} = 0$ if there is no fiber link between node m and l.
- z^{ij}_{ml}: This is the lightpath routing variable. $z^{ij}_{ml} = 1$ if the lightpath between node i and j goes through the fiber link between node m and l.
- d_{ml}: This is a constant that specifies the fiber distance between node m and l, which could be longer than the Cartesian distance (e.g., when physical constraints force a fiber link not to be laid along the direct line).

First, we need to make sure that all demands (lightpaths) are routed. This is the same as a flow conservation constraint.

$$\sum_l z^{ij}_{ml} - \sum_l z^{ij}_{lm} = \begin{cases} 1 & \text{if } m = i \\ -1 & \text{if } m = j \\ 0 & \text{otherwise} \end{cases} \quad \forall m, ij \tag{4}$$

Second, we make sure we only use W wavelengths on each link.

$$\sum_{ij} (z^{ij}_{ml} + z^{ij}_{lm}) \le W z_{ml} \quad \forall ml \tag{5}$$

Note that we assume all lightpaths are bidirectional. Since link ml and lm are considered as separate links mathematically, we need to add the lightpaths crossing both links on the left-hand side. If we make W a variable, then this constraint is no longer linear. We could do a linear search for the right W, making W a constant in each iteration. Alternatively, to make the constraint linear, we can use the following two constraints instead.

$$\sum_{ij} (z^{ij}_{ml} + z^{ij}_{lm}) \le \quad W \qquad \forall ml \tag{6}$$

$$\sum_{ij} (z^{ij}_{ml} + z^{ij}_{lm}) \le \quad Z z_{ml} \qquad \forall ml \tag{7}$$

Z is a large constant. The first constraint makes sure that the number of wavelengths on each link is fewer than W and the second constraint makes sure that no lightpath passes through a fiber link if that fiber link is not installed.

Since the fiber-to-node ratio f is given, we can only use $L = fNL_n = fN$ (since we assume $L_n = 1$) amount of fiber. When $f = 1$, roughly only N fiber links of length L_n can be added to connect all nodes. Therefore, $f = 1$ is the minimum required to guarantee a connected topology. If only short fiber links are used, f roughly corresponds to k, the average edge-to-node ratio. The following constraint limits the total fiber length that can be used.

$$\sum_{ml} z_{ml} d_{ml} < L = fN \tag{8}$$

The objective is to minimize the number of wavelengths.

$$\text{Objective: } \min W \tag{9}$$

B. Heuristic algorithm

The ILP formulations given above can only be used to solve small problems exactly, e.g., for networks with less than 10 nodes. For larger problems, we propose an efficient heuristic algorithm. We compared our algorithm with the ILP for several problem instances, and found it produces near-optimal results. We also compared our algorithm with that in [10] and found that our algorithm not only produces comparable results, but also runs several orders of magnitude faster. For example, for a network with 100 nodes, our algorithm takes less than 0.05 seconds on a Sun Blade 1000 workstation, compared to 11 hours for the algorithm reported in [10]. We call the algorithm TOPO_F to denote that this is a topology design algorithm with the total fiber length as a constraint.

To understand what determines the number of wavelengths needed, consider a set of nodes S and the remaining nodes $\overline{S}$. Let $C(S, \overline{S})$ denote the number of fiber links that cross between the two sets of nodes. Since we want to establish a full-mesh connectivity, there are $|S| \times |\overline{S}|$ lightpaths that will cross between the two sets of nodes. Therefore, we can derive a lower bound on W as follows:

$$W_{LB} = \max_S \frac{|S| \times |\overline{S}|}{C(S, \overline{S})}$$

Such a lower bound was observed in [12][13], and it is also called the flux of a graph in [13]. Note that W_{LB} is the maximum value among all possible sets of nodes. So if only a few fiber links cross between two sets of nodes, W_{LB} will be very high as the denominator is small. To design a topology that needs less capacity, we need to make sure that the number of fiber links crossing any two sets of nodes is sufficiently large. This observation motivates us to propose the following physical topology design algorithm.

The TOPO_F algorithm proceeds in four steps.

- In the first step, a minimum spanning tree is established. This ensures that the topology is connected.
- The second step tries to find the cut that achieves the lower bound W_{LB}, i.e., find the cut such that $|S| \times |\overline{S}|/C(S, \overline{S})$ is the largest. To do so, we pick a starting node s and initialize the set S to contain only the starting node ($S = \{s\}$). Then we add one neighbor node into set S at a time until all nodes are in the set. When adding a node, we pick the neighbor node such that the cut $C(S, \overline{S})$ is the smallest. The reason we do so is because the smallest cut $C(S, \overline{S})$ will give the biggest W_{LB}, which is what we are looking for. When we have a new S (after

adding a node), we compute $|S| \times |\overline{S}|/C(S, \overline{S})$. If it is larger than what we have seen before, we record it. To do a more exhaustive search, we try each node as the starting node in turn and repeat the above process. If there is a tie in the cut, we pick the cut for which we can add a fiber link to bridge the cut and this fiber link is the shortest among all cuts in the tie.

- At the end of the second step, we have identified a S_{max} such that $|S_{max}| \times |\overline{S_{max}}|/C(S_{max}, \overline{S_{max}})$ is the largest. Then, in the third step, we pick the shortest fiber link that can bridge the cut (one end of the fiber link in S_{max} and the other end in $\overline{S_{max}}$). If adding the fiber link will exceed the total fiber length budget, we proceed to step four. Otherwise, we add the fiber link, then return back to step two to find a new limiting cut. Note that we could have added a parallel fiber in this step, even if a fiber link has already been added between the two end nodes.
- In the last step, we check each remaining node pair in turn in increasing order of distance. If adding a fiber link between the node pair will not violate the total fiber length budget, the fiber link will be added to the topology.

The TOPO_F algorithm tends to use short fiber links so that more fiber links could be added. Because the topologies generated by the TOPO_F algorithm have many short fiber links, we call these topologies the short-fiber topologies.

IV. Numerical results

Using the physical topology design algorithm, we can now study the tradeoff between L and W and see how the provisioned capacity scales as a function of N.

We consider networks of practical sizes with up to a few hundred nodes. The backbone topologies of most Internet Service Providers (ISP) in the US have fewer than 100 nodes. Networks with more than a few hundred nodes are less likely to be practical because of the difficulty involved in design and management. A hierarchical network architecture might then be more appropriate. Note that, unlike some recent work on understanding the Internet topologies [14] [15], we are focusing on the backbone fiber (physical) topology of a single ISP. The backbone network is much smaller compared to the Internet, and it does not necessarily follow the characteristics of the Internet, such as a power-law distribution.

We consider a practical range of fiber-to-node ratio f, from 1 up to 4, to include networks that are currently deployed or to be deployed in the near future. The fiber-to-node ratio of the current fiber networks of most ISPs is below $f = 2$. Note that the fiber-to-node ratio is roughly the same as the edge-to-node ratio k if short fibers are used, and the edge-to-node ratio is half of the average node degree, i.e., $k = d/2$.

For a network with N nodes, we assume these N nodes are uniform-randomly distributed within a square of area N ($\sqrt{N}$ on each side) to ensure $L_n = 1$. To generate a network, we randomly place each one of the N nodes at a point in the square with equal probability. Once the network is generated, we then apply the topology design algorithm.

To determine W after the physical topology is designed, we use a Routing and Wavelength Assignment (RWA) algorithm to fit all lightpaths into as few wavelengths as possible. Several heuristic RWA algorithms have been proposed in the literature [16] [17]. The algorithms we use are reported in a technical report [18], one with the Wavelength Continuity Constraint (WCC) and one without. Since this paper focuses on physical topology design and capacity scaling, we will not discuss the details of the RWA algorithms here.

In order for the experiments to be statistically significant, we repeat the above procedure 100 times and then take the average before reporting the data. In other words, for each N, we randomly generate 100 separate networks and then take the average result from the 100 separate networks.

A. Heuristic algorithm and design approach evaluation

The results from the TOPO_F algorithm (along with the RWA algorithms) are very close to that from solving the ILP formulation directly for small-size networks (<5 nodes). Unfortunately, because of the high computation time, we are not able to evaluate our heuristic algorithm against the optimal for larger-sized networks. Instead, we will evaluate its performance against the lower bound.

In Fig. 3, we plot the average capacity of short-fiber topologies (from the TOPO_F algorithm) assuming $f = 4$, and the average lower bound (equation 1) from the same set of networks. For comparison purpose, we also plot the average and minimum capacity from 1000 random topologies with an edge-to-node ratio $k = 4$. These random topologies are generated by randomly picking a pair of nodes and then placing an edge across them until the desired number of edges are generated. Since edges are picked without regard to their lengths, many long fiber links are used.

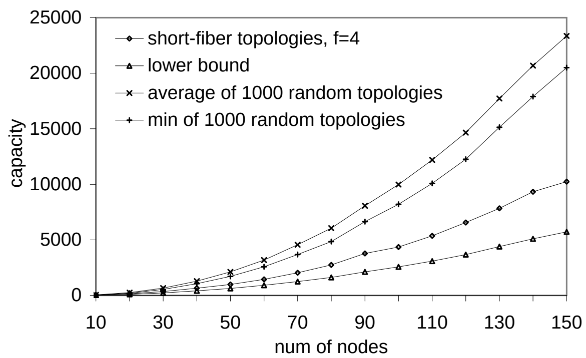

Fig. 3. Comparison of the results between the heuristic algorithm, the lower bound and the average and minimum of many random topologies. With WCC.

As shown in the figure, our topology design algorithm can achieve capacity that is less than twice of the lower bound. Considering that the lower bound is not achievable unless all lightpaths use direct-line routing, we suspect that our algorithm is not too far away from the optimal. Our physical topology design algorithm can greatly reduce the required capacity, not only compared to the average but also compared to the minimum of the 1000 random topologies.

Compared to random topologies, short-fiber topologies with $f = k$ use much less fiber at the cost of more wavelengths. In Fig. 4, we show the wavelength requirement in random topologies with $k = 4$. For short-fiber topologies, we show two results. One is the wavelength requirement if we set $f = k = 4$; the other is the wavelength requirement if we set f to a value such that it will use the same amount of fiber as in the random topologies. We can see that the short-fiber topologies require much fewer wavelengths if the same amount of fiber is used. This is because the TOPO_F algorithm is conscious about the fiber usage and, therefore, it is able to establish more fiber links with the limited budget. This result suggests that taking fiber length into consideration can lead to better designs.

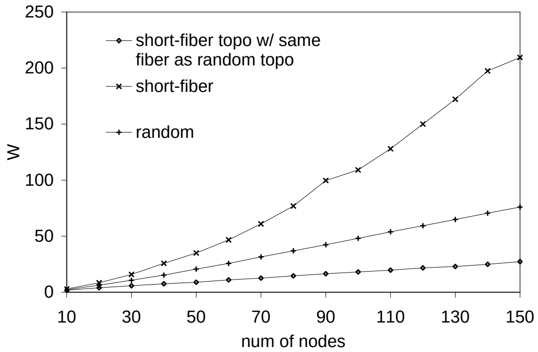

Fig. 4. Wavelength requirement comparisons between random topologies and short-fiber topologies. With WCC.

B. Capacity scaling

In Fig. 5, we show the tradeoff between L and W. Clearly, W decreases as we increase L. Although hard to see in this figure, W increases by a larger proportion as N gets larger when f is small compared to the case when f is large. This trend can be captured by the growth rate in W (or C). It is higher when f is small compared to the case when f is large.

In Fig. 6, we show W as a function of N (with the WCC constraint) in the short-fiber topologies. Note that the graph is on a log-log scale. Indeed, when f is small, W grows very quickly, as evident from the steeper slope, and close to 20000 wavelengths are needed when $N = 500$ and $f = 1$. This growth drops quickly as f is increased, resulting in a large reduction in W. Since the fiber-to-node ratio is fixed, the fiber length is a linear function of N, so the capacity will follow the same trend as the wavelengths.

The data in Fig. 6 appears to fall on a straight line, suggesting that W is a power function of N. If we assume that C, L and W are all power functions of N in the form of aN^e, then we can use the least square method to estimate the parameters a and e on a log-log plot. The results for short-fiber topologies are shown in Table I. All curve fittings have a correlation coefficient of more than 0.99, confirming that a power function is a good fit.

$e(L)$, the exponent in L, is simply 1 because we have fixed the fiber-to-node ratio. $e(W)$, the exponent in W, and $e(C)$, the exponent in C, again drop quickly as f increases. When $f = 1$, $e(C)$ is more than 2.75. As f increases beyond 2.5, $e(C)$ quickly drops and it is between 2.6 and 2.64. When f is small, the topology is barely connected. Many lightpaths are forced to be routed through long detours. When f is large, there is a better chance for a lightpath to be routed close to the direct line. $e(C)$ is smaller than that of the random topologies that we studied, which is around 2.8, suggesting that our physical topology design algorithm is effective at reducing the capacity requirement.

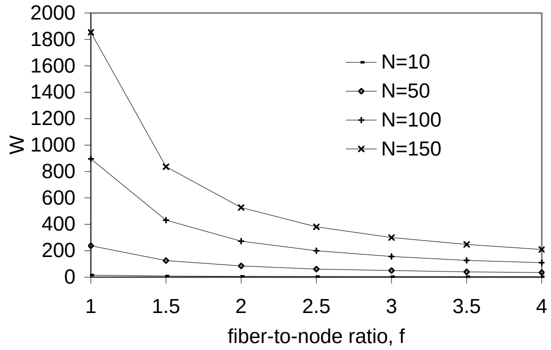

Fig. 5. The tradeoff between L and W. With WCC.

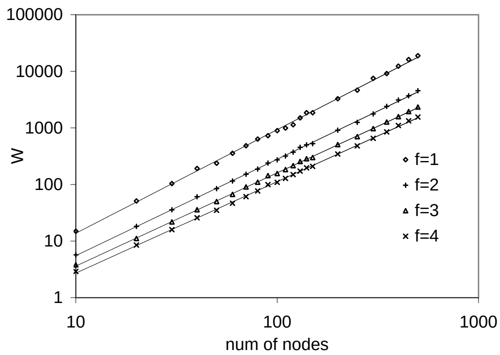

Fig. 6. W as a function of N. With WCC.

The constant $a(L)$ of L is simply f, the fiber-to-node ratio. $a(C)$, the constant in C, is almost the same regardless of the parameter f. In other words, an increase in f (and therefore $a(L)$) would result in a corresponding decrease in $a(W)$, the constant in W. This result is consistent with our theoretical analysis (equation 3).

The results in Table I are generated from networks with up to 500 nodes, large enough to cover networks to be deployed in the near future.

C. Effects of the WCC constraint

In Fig. 7, we plot the wavelength requirement as a function of N for short-fiber topologies, for both the case of with WCC constraint and without. When f is small ($f = 1$) and N is large ($N \geq 100$), the differences are noticeable. The differences are up to 9% for $f = 1$, and they are up to 7% for $f = 2$. In all other cases ($f > 2$), the differences are very small, less than 5%. This suggests that the WCC constraint makes little

TABLE I
CONSTANT PARAMETERS FROM LEAST SQUARE CURVE FITTING FOR SHORT-FIBER TOPOLOGIES. WITH WCC.

f	a			e		
	W	L	C	W	L	C
1	0.27	1	0.27	1.76	1	2.76
1.5	0.15	1.5	0.23	1.73	1	2.73
2	0.12	2	0.23	1.69	1	2.69
2.5	0.10	2.5	0.24	1.66	1	2.66
3	0.09	3	0.26	1.64	1	2.64
3.5	0.08	3.5	0.28	1.61	1	2.61
4	0.07	4	0.28	1.60	1	2.60

difference especially when the fiber-to-node ratio f is high enough. Therefore, the costly wavelength converters could be avoided by simply increasing f. This observation is consistent with that in [11], where they found that WCC has little effect on two-connected topologies.

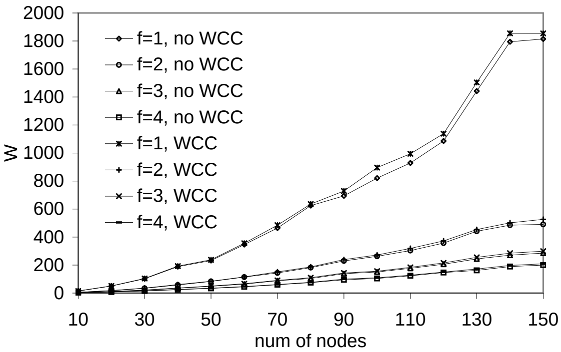

Fig. 7. Number of wavelengths with or without the WCC constraint.

D. Capacity utilization

In this section, we look at the effects of factor F2: the utilization of the provisioned capacity. Let f_i denote the number of wavelengths that are used on fiber link i, then the capacity utilization ratio can be defined as $\frac{\sum_i f_i}{WM}$, i.e., the ratio between the sum of the number of used wavelengths on each fiber link and the product of the number of wavelengths (W) and the number of fiber links (M). The capacity utilization ratio for short-fiber topologies (with the WCC constraint) is shown as a function of N in Fig. 8. When f is small, the utilization ratio is very low. For example, when $f = 1$, at most 70% utilization can be achieved and it is as low as 40% when N is large. The situation quickly improves as f gets bigger. Roughly 70% or 80% utilization is possible when $f > 2.5$. The utilization ratio decreases for large N. This is because the average hop count of lightpaths increases as N increases, making it less likely to fully utilize the provisioned capacity. The decrease in the utilization ratio partly contributes to the higher asymptotic growth rate (the other contributor is factor F1, non-direct-line routing of lightpaths). We see similar trend on the utilization ratio when the WCC constraint is relaxed.

In the figure, there is an anomaly. The utilization ratio is low when f is large and N is small. This is because of the "rounding effect:" when W is small, adding one additional wavelength adds a large percentage of capacity, which in turn greatly reduces the utilization ratio.

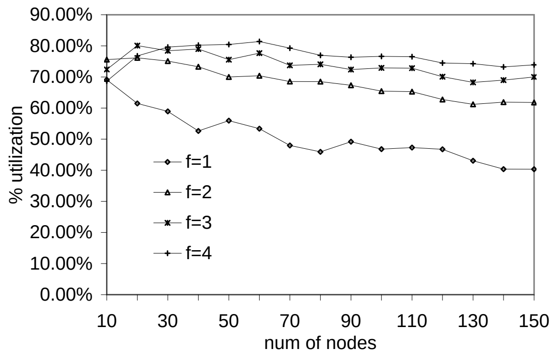

Fig. 8. Utilization ratio as a function of N for short-fiber topologies. With WCC.

As mentioned in the beginning, we notice that C scales almost perfectly linearly as T in our simulation study, as suggested by equation (3). It suggests that the increased number of connection requests (lightpaths) cannot improve the capacity utilization. The low capacity utilization ratio seems to be a fundamental limit of the topology. The only way to improve the utilization ratio seems to be redesigning a better topology using a higher f.

Recall that f_i denotes the number of wavelengths that is used on fiber link i. Let h_{ij} denote the shortest path (in number of physical hops) between node i and j in a given topology. Then we can define *load* as $\sum_i f_i$, i.e., the total number of wavelengths that are utilized. We can also define the *shortest path load* as $\sum_{ij} h_{ij}$, i.e., the load if all lightpaths are routed using the shortest path. The difference between *load* and *shortest path load* represents the extra number of wavelengths needed to route lightpaths along non-shortest paths in order to fully utilize the provisioned capacity. In Fig. 9, we plot the exponent e from least square data fitting to a power function (aN^e) for C, the load and the shortest path load as a function of f.

As shown in the figure, the gap between C and the load is decreasing as f increases. This suggests that the capacity utilization keeps on improving as f increases. In addition, the gap between the load and the shortest path load is increasing. This suggests that as f gets bigger, it becomes increasingly easier to find alternative paths to route a lightpath if the shortest path is congested. Therefore, there is greater chance to even out the routing of lightpaths onto all fiber links and thus reduce the required number of wavelengths W.

E. Non-uniform lightpath demands

We have so far only considered the full-mesh connectivity because it makes the analysis easier. However, the observation

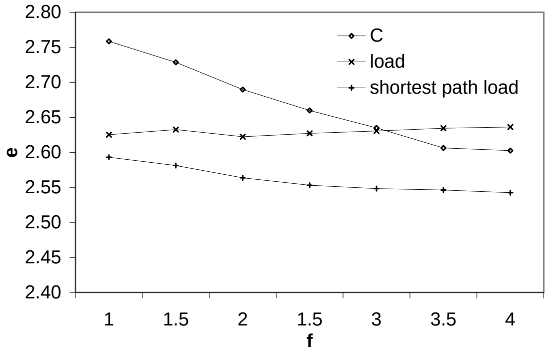

Fig. 9. Growth exponent for C, the load and the shortest path load as a function of f for short-fiber topologies. With WCC.

TABLE II

ASYMPTOTIC GROWTH RATE OF THE CAPACITY FOR ONE PARTICULAR SET OF LIGHTPATH DEMANDS IN SHORT-FIBER TOPOLOGIES

f	1	1.5	2	2.5	3	3.5	4
$e(C)$	1.78	1.74	1.65	1.64	1.62	1.58	1.59

is not strongly dependent on the set of lightpaths to support. Let us consider a different set of lightpaths which is generated as follows. From each node, we randomly pick m other nodes and establish lightpaths to them. The number of lightpaths generated is mN, i.e., the number of lightpaths is on the order of N instead of N^2 as in the full mesh. We pick $m = 4$ in our experiment so that the resulting W is large enough to avoid the rounding effect.

We use the TOPO_F algorithm to design the topology and use curve fitting to determine the asymptotic growth rate of the capacity. The results are shown in Table II.

$e(C)$ should be 1.5 theoretically. But in simulation, it ranges from 1.58 to 1.78. $e(C)$ again decreases quickly as f increases, which is consistent with our observation under the full-mesh connectivity.

V. REAL-LIFE NETWORKS

We have seen that increasing the fiber-to-node ratio f can reduce the capacity growth, increase the capacity utilization and avoid the use of wavelength converters. Therefore, as a design guideline, it seems to be a good idea to have a high enough f. In this section, we examine real-life networks and see how this design guideline can be applied. We studied several real-life networks. Even though the bounding areas are not regular and the node placements are not uniform, the observations we derive are very similar, therefore, we only report our study on one real-life network here.

We consider the backbone network from Level 3 (an Internet Service Provider). According to the Rocketfuel project [19], which maps ISP topologies as seen by the electronic routers (the logical topology), Level 3 attempts to establish a full-mesh connectivity, i.e., a lightpath between every pair of nodes.

Even though the logical topology is a full mesh, the underlying physical fiber topology is far from a full mesh, as shown in Fig. 10. This fiber map could be downloaded from their website directly[20].

Fig. 10. Level3's backbone network

This network has 56 nodes and 63 edges. The edge-to-node ratio is only $k = 1.125$. Only short fiber links that connect neighboring nodes are used, so it is a short-fiber topology. In fact, in most of the topologies we studied, long transcontinental fibers are almost never used.

Using our RWA algorithms, we determine that 374 wavelengths are needed to support the full-mesh connectivity with the WCC constraint. The capacity utilization ratio is only 55%.

We manually measure the distance between every node pair and assume a direct fiber of that length could be laid down to connect the node pair. We then apply the TOPO_F topology design algorithm and RWA algorithms. Using the same total fiber length, the TOPO_F algorithm designed a new topology with 73 edges and the RWA algorithm found that only 363 wavelengths are needed to support the full-mesh connectivity. Our topology design algorithm is able to design a topology that requires less capacity than that of the manually designed topology. This is a confirmation that our algorithm performs well.

We also applied the TOPO_F and RWA algorithms for different values of f. The results are shown in Fig. 11.

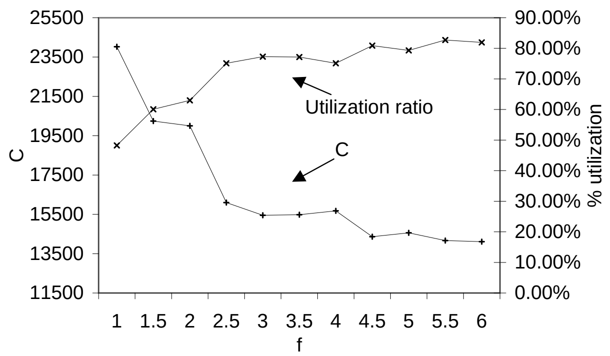

Fig. 11. C and the utilization ratio as a function of f for Level3's network. With WCC.

We can see that C rapidly decreases as f increases. The slope of decrease flattens out when f is at least 2 or 3. At the same time, the utilization ratio quickly improves. When

f is small, the utilization ratio is only 50%, and it quickly improves to nearly 80%. This result suggests that the current fiber-to-node ratio in Level 3's network is too low. It should be increased in order to lower the capacity requirement and increase the utilization ratio.

VI. Conclusion

We proposed an efficient algorithm to design the fiber topology with the goal of minimizing the number of wavelengths under the constraint of a fixed fiber-to-node ratio. Using the algorithm, we studied the tradeoff between fiber (link cost) and wavelengths (node cost). Understanding this tradeoff allows network designers to design cost effective transport networks.

We evaluated and compared designs under different fiber-to-node ratios using the provisioned capacity as a metric, which is an indicator of how efficient the provisioned resources are utilized. The asymptotic growth rate of the capacity not only captures the tradeoff between fiber and wavelengths independent of the network size, but it also indirectly translates into the deployment cost regardless of the actual cost functions for fiber and wavelengths.

We showed that, compared to random topologies and the lower bound, our physical topology design algorithm is very effective not only at reducing the capacity requirement but also at reducing the asymptotic growth rate of the capacity. We found that taking fiber length into consideration can reduce the capacity requirement. We showed that having a large fiber-to-node ratio can greatly reduce the asymptotic growth rate and can lead to lower cost when N is large.

On studying several real-life topologies, we find that most of them have too low a fiber-to-node ratio. By increasing it, we can greatly reduce the capacity requirement and increase the capacity utilization.

References

[1] G. Kramer, B. Mukherjee, and A. Maislos, "Ethernet PON: Design and analysis of an optical access network," *Photonic Network Commun.*, vol. 3, no. 3, pp. 307–319, 2001.

[2] J. A. Bannister, L. Fratta, and M. Gerla, "Topological design of the wavelength-division optical network," in *Proc. IEEE Infocom*, San Francisco, CA, June 1990.

[3] J. Labourdette and A. Acampora, "Logically rearrangeable multihop lightwave networks," *IEEE Trans. Commun.*, vol. 39, no. 8, Aug. 1991.

[4] I. White, M. Rogge, K. Shrikhande, and L. Kazovsky, "A summary of the HORNET project: A next-generation metropolitan area network," *IEEE J. Select. Areas Commun.*, vol. 21, no. 9, pp. 1478–94, 2003.

[5] I. Widjaja, I. Saniee, R. Giles, and D. Mitra, "Light core and intelligent edge for a flexible, thin-layered and cost-effective optical transport network," *IEEE Commun. Mag.*, vol. 41, pp. s30–s36, 2003.

[6] I. Chlamtac, A. Ganz, and G. Karmi, "Lightpath communications: an approach to high bandwidth optical wans," *IEEE Trans. Commun.*, vol. 40, no. 7, pp. 1171–82, July 1992.

[7] P. Roorda, C. Lu, and T. Boutilier, "Benefits of all-optical routing in transport networks," in *Proc. OFC'95*, 1995, pp. 164–165.

[8] C. Guan and V. Chan, "Efficient physical topologies for regular WDM networks," in *Proc. Optical Fiber Commun. Conf.*, vol. 1, Feb. 2004, pp. 23–27.

[9] Y. Xin, G. N. Rouskas, and H. G. Perros, "On the physical and logical topology design of large-scale optical networks," *J. Lightwave Technol.*, vol. 21, no. 4, pp. 904–15, Apr. 2003.

[10] G. Xiao, Y. wing Leung, and K.-W. Hung, "Two-stage cut saturation algorithm for designing all-optical networks," *IEEE Trans. Commun.*, vol. 49, no. 6, pp. 1102–15, June 2001.

[11] S. Baroni and P. Bayvel, "Wavelength requirements in arbitrarily connected wavelength-routed optical networks," *J. Lightwave Technol.*, vol. 15, no. 2, pp. 242–251, 1997.

[12] D. Hjelme, A. Royset, and B. Slagsvold, "How many wavelengths does it take to build a wavelength routed optical network?" in *Proc. 22nd European Conference on Optical Communication*, vol. 1, 1996, pp. 27–30.

[13] F. T. Leighton and S. Rao, "An approximate max-flow min-cut theorem for uniform multicommodity flow problems with applications to approximation algorithms," in *Proc. FOCS'88*, 1988, pp. 422–431.

[14] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the Internet topology," in *Proc. ACM SIGCOMM*, 1999.

[15] A. Medina, I. Matta, and J. Byers, "On the origin of power laws in internet topologies," *ACM Computer Communication Review*, Apr. 2000.

[16] S. Ohta and A. Greca, "Comparison of routing and wavelength assignment algorithms for optical networks," in *Proc. IEEE Workshop on HPRS*, Texas, May 2001.

[17] D. Banerjee and B. Mukherjee, "A practical approach for routing and wavelength assignment in large wavelength routed optical networks," *IEEE J. Select. Areas Commun.*, vol. 14, no. 5, pp. 903–908, June 1996.

[18] H. Liu and F. A. Tobagi, "Designing physical topology for all-optical networks with minimal provisioned capacity," Stanford University, Tech. Rep., 2006.

[19] N. Spring, R. Mahajan, and D. Wetherall, "Measuring ISP topologies with rocketfuel," in *Proc. SIGCOMM*, 2002.

[20] "Level 3's fiber map in the US." [Online]. Available: http://www.level3.com/userimages/dotcom/en_US/images/ir_us.jpg