On the Physical and Logical Topology Design of Large-Scale Optical Networks

Yufeng Xin, George N. Rouskas, *Senior Member, IEEE*, and Harry G. Perros, *Senior Member, IEEE*

***Abstract*—We consider the problem of designing a network of optical cross-connects (OXCs) to provide end-to-end lightpath services to large numbers of label switched routers (LSRs). We present a set of heuristic algorithms to address the combined problem of physical topology design (i.e., determine the number of OXCs required and the fiber links among them) and logical topology design (i.e., determine the routing and wavelength assignment for the lightpaths among the LSRs). Unlike previous studies which were limited to small topologies with a handful of nodes and a few tens of lightpaths, we have applied our algorithms to networks with hundreds or thousands of LSRs and with a number of lightpaths that is an order of magnitude larger than the number of LSRs. In order to characterize the performance of our algorithms, we have developed lower bounds which can be computed efficiently. We present numerical results for up to 1000 LSRs and for a wide range of system parameters such as the number of wavelengths per fiber, the number of transceivers per LSR, and the number of ports per OXC. The results indicate that it is possible to build large-scale optical networks with rich connectivity in a cost-effective manner, using relatively few but properly dimensioned OXCs.**

***Index Terms*—Genetic algorithms, large-scale optical networks, topology design, wavelength-division multiplexing (WDM).**

I. Introduction

THE wide deployment of point-to-point wavelength-division-multiplexing (WDM) transmission systems in the Internet infrastructure has enhanced the need for faster switching at the core of the network. The corresponding massive increase in network bandwidth due to WDM has occurred in conjunction with a growing effort to modify the Internet Protocol to support different levels of quality-of-service (QoS). Label switching routers (LSRs) supporting multiprotocol label switching (MPLS) [1] are being deployed to address these two issues. On one hand, LSRs simplify the forwarding function, thereby making it possible to operate at higher data rates. On the other hand, MPLS enables the Internet architecture, built upon the connectionless Internet Protocol, to behave in a connection-oriented fashion that is more conducive to supporting QoS.

The rapid advancement and evolution of optical technologies makes it possible to move beyond point-to-point WDM transmission systems to an all-optical backbone network that can take full advantage of the available bandwidth by eliminating the need for per-hop packet forwarding. Such a network consists of a number of optical cross connects (OXCs), arranged in some arbitrary topology, and provides interconnection to a number of LSRs. Each OXC can switch the optical signal coming in on a wavelength of an input fiber link to the same wavelength in an output fiber link. The OXC may also be equipped with converters that permit it to switch the optical signal on an incoming wavelength of an input fiber to any wavelength on an output fiber link. The main mechanism of transport in such a network is the lightpath, which is a communication channel established between two OXCs or two LSRs and which may span a number of fiber links (physical hops). If no wavelength converters are used, a lightpath is associated with the same wavelength on each hop. This is the well-known wavelength continuity constraint. Using converters, a different wavelength on each hop may be used to create a lightpath. Thus, a lightpath is an end-to-end optical connection established between two LSRs.

Currently, there is tremendous interest within both the industry and the research community in optical networks of OXCs. The Internet Engineering Task Force (IETF) is investigating the use of generalized MPLS (GMPLS) [2] to set up and teardown lightpaths. GMPLS is an extension of MPLS that supports multiple types of switching, including switching based on wavelengths usually referred to as multiprotocol lambda switching ($\mathrm{MP}\lambda\mathrm{S}$); therefore, we will also use the term "$\mathrm{MP}\lambda\mathrm{S}$ network" to refer to an optical network of OXCs. With GMPLS, the OXC backbone and the network of LSRs will share common functionality in the control plane, making it possible to seamlessly integrate OXC backbones within the overall Internet infrastructure. Also, the Optical Domain Service Interconnection (ODSI) initiative and the Optical Internetworking Forum (OIF) are concerned with the interface between an LSR and the OXC to which it is attached as well as the interface between OXCs, and have several activities to address MPLS over WDM issues [3].

In addition to control plane protocol issues, the problem of designing wavelength-routed networks of OXCs has received considerable attention in the last decade. In order to reduce the complexity of this problem, typically, it is broken down to two sub-problems: *network design* and *routing and wavelength assignment* (RWA). Network design involves *physical topology design* and *configuration design*. The topology design involves the determination of the number of OXCs and their interconnectivity. The network configuration is concerned with the determination of the size of OXCs, the number of fibers and the set of

Manuscript received May 16, 2002; revised September 27, 2002. This work was supported by the NSF under grant ANI-9805016.

Y. Xin is with MCNC, Research Triangle Park, NC 27709-2889 USA.

G. N. Rouskas and H. G. Perros are with the Department of Computer Science, North Carolina State University, Raleigh, NC 27695 USA (e-mail: rouskas@eos.ncsu.edu; hp@eos.ncsu.edu).

Digital Object Identifier 10.1109/JLT.2003.810075

lightpaths. Routing and wavelength assignment involves mapping lightpaths onto the physical topology and assigning wavelengths to these lightpaths. The reader is referred to [4] for a general analysis of various formulations and solution approaches to the above problems. Most of the work in the open literature focuses on the configuration design and the RWA problems under the assumption of a fixed fiber physical topology. That is, given a network traffic demand and a physical network topology, an optimal network configuration, a best virtual topology, and an optimal routing and wavelength assignment are obtained. In the case in which the set of lightpaths is also given, the problem is reduced to a pure RWA problem, which can be further decomposed into a routing sub-problem and a wavelength assignment sub-problem.

The design of virtual topologies has been studied extensively. Constrained by the limited number of available wavelengths and the available number of transceivers, it may not be possible to establish a lightpath between every pair of nodes. Consequently, only a selected set of nodes can be connected by lightpaths, leading to a virtual topology over the given physical network. In a virtual topology, the nodes correspond to actual physical network nodes, while the links correspond to lightpaths. A review of algorithms for virtual topology design can be found in [5].

Typically, the network design problem under a given physical topology can be formulated as an integer programming (IP) problem with the objective of optimizing a performance metric of interest. This IP problem has been shown to be NP-hard, and several heuristic algorithms have been proposed in the literature. These algorithms differ in their assumptions regarding the traffic demands, as well as in the performance metric used. Specifically, for the case of static traffic, i.e., when all the connections are known and they are static, the objective is typically to minimize network resource usage. In the case where the number of fibers is limited, the objective is usually to minimize the total number of wavelengths needed in the network. A linear programming relaxation technique to obtain the optimal solution for the RWA of a WDM ring network was proposed in [6], while a longest-lightpath-first heuristic was developed in [7] as an approximate solution to the RWA problem for a mesh network. In the case of a limited number of wavelengths per fiber, [8] provided an iterative scheme to minimize the costs associated with the working fibers for provisioning a static set of lightpaths on a given WDM network topology. A lightpath accommodation heuristic for a given set of lightpaths was developed to minimize the total number of OXC ports (another important network resource) in [9]. A general IP model for the virtual topology design was presented in [10], and a branch-and-bound algorithm to minimize the average lightpath length was described.

The design of a virtual topology was first formulated as an integer linear programming (ILP) problem in [11], where a heuristic algorithm combining simulation annealing and flow deviation was presented to minimize the delay and the maximum flow in a link. A genetic algorithm to calculate the optimal virtual topology and RWA to minimize the average signal delay was developed in [12]. An upper bound on the maximum carried traffic for any routing and wavelength assignment algorithm was derived in [13]. Most of the routing algorithms are adapted from Dijkstra's shortest path algorithm [4], [12], [14]. The wavelength assignment problem is equivalent to a graph coloring problem, and [15] adopted a known heuristic algorithm to solve it. The RWA problem for a multi-fiber network under dynamic traffic was solved in [14] using the layered-graph heuristic model.

The physical topology design problem has also received some attention. The relationship between the number of wavelengths needed and some topology parameters, such as connectivity, nodal degrees, and average hop distance, was studied in [16], [17] through simulation. A bound on the number of wavelengths given the connectivity requirements of the users and the number of switching states was derived in [18]. Simple approximate equations for the scaling properties of the number of wavelengths, the nodal degree, the total fiber lengths, and the maximum number of transit nodes of a lightpath were given in [19]. An analytical solution of the RWA problem for some regular topologies, such as shuffle and tori was carried out in [20], [21].

In this paper we consider the problem of designing large-scale optical WDM networks of OXCs that provide end-to-end lightpath services. The scale of the optical backbone network is characterized by the number of LSRs using its services and the number of lightpaths that it can support. We are interested in typical national or international networks, in which the number of LSRs can be in the hundreds or thousands, and the number of lightpaths that need to be established can be an order of magnitude greater than the number of LSRs. Given the number of LSRs, the number of wavelengths per fiber link, and a set of physical constraints (such as the number of transceivers at each LSR and the number of input–output ports at each OXC), we address both the physical topology design problem (i.e., the number of OXCs required and their interconnectivity) and the routing and wavelength assignment. We assume no wavelength conversion in the network. Since the problem is NP-hard, we present a set of heuristic algorithms to obtain a near-optimal solution in terms of the number of required OXCs, including a genetic algorithm to search the space of physical topologies.

While some of the problems we consider have been studied earlier, our work differs from previous studies in several important ways. To the best of our knowledge, this is the first time that the problem of designing the physical *and* logical topology of a wavelength-routed network is fully formulated and solved. Also, whereas previously published algorithms have been applied to relatively small networks (e.g., 10–20-node topologies such as the NSFNet and the Arpanet) with few (less than 100) lightpaths, we consider large-scale networks of realistic size. We provide new insight into the design of WDM backbone networks by investigating the effect of various system parameters; again, we use realistic ranges for the values of these parameters, e.g., up to 128 wavelengths and up to 24 optical interfaces per LSR. Finally, we obtain lower bounds for the optimization problem and present results which indicate that our heuristics are close to optimal.

We now summarize some of the results of our study. The most important finding is that it is possible to build cost-effective optical backbone networks that provide rich connectivity among large numbers of LSRs with relatively few, but properly dimensioned, OXCs. In particular, the number of OXCs increases linearly with the number of attached LSRs, but the rate of the increase is rather slow. We also find that, in order to take advantage of an increasing number of wavelengths, the number of OXC ports must increase correspondingly. Otherwise, using additional wavelengths has little effect on the number and topology of OXCs required for a given number of LSRs. This result, coupled with the fact that the number of wavelengths in a fiber is expected to continue to increase in the foreseeable future, has implications on the OXC technology and design. For instance, our results indicate that 3-D MEMS switches which can scale to large port sizes would be more appropriate than 2-D MEMS switches which are limited to small port sizes [22]. Finally, we also find that it is possible to enhance the degree of connectivity among the LSRs by a large factor (through a corresponding increase of the number of optical interfaces at each LSR) with a relatively small incremental cost in terms of additional OXCs and fiber links among them.

In Section II we describe the problem we study as well as the assumptions we make, and in Section III we present an overview of our solution approach. In Section IV, we present the heuristic algorithms for generating a 2-connected graph, routing lightpaths, and assigning wavelengths. We also describe the genetic algorithm used for searching the space of physical topologies. We develop lower bounds for the optimization problem in Section V. We present numerical results in Section VI, and we conclude the paper in Section VII.

II. Problem Definition

We consider a number N of LSRs that are to be interconnected over an optical backbone network [23] which consists of OXC nodes supporting GMPLS. The service provided by the $\mathrm{MP\lambda S}$ network of OXCs is the establishment of lightpaths among pairs of LSRs. We assume that each LSR has Δ optical transceivers, and therefore, it may establish at most Δ incoming and at most Δ outgoing lightpaths at any given time. This constraint on the number of simultaneous lightpaths to/from an LSR is due both to optical hardware and cost limitations (reflected in the number of optical transceivers) and the traffic processing capacity of the LSR. We also assume that all fiber links, including links between OXCs as well as links between an LSR and an OXC, can support the same set of W wavelengths, and that each OXC has exactly P input/output ports. We let α, $0 \leq \alpha \leq 1$, be the desired degree of connectivity of the physical topology of OXCs, defined as

$$\alpha = \frac{E}{\frac{M(M-1)}{2}} \tag{1}$$

where E represents the number of fiber links interconnecting the OXCs. Parameter α represents how dense the graph is. For an arbitrary graph, α ranges from 0 to 1, with 0 representing a graph that is totally disconnected and 1 representing a completely connected graph. We note that most of the existing backbone networks have an α value around 0.3.

The fundamental question we address in this paper is: *What is the minimum number M of OXCs required to support the N LSRs, and what is the physical topology of the corresponding* $\mathrm{MP\lambda S}$ *network?*

We believe that the answer to this question is of importance to service providers who need to deploy optical backbone networks in a cost-effective manner. We note that the cost of building an $\mathrm{MP\lambda S}$ network will be mainly determined by: **(i)** the cost of the OXCs (including switch hardware and switch controller software), and **(ii)** the cost of (deploying or leasing) the fiber links between OXCs (including the cost of related equipment, such as optical amplifiers). While in our study we directly model only the OXC cost, we note that the fiber cost is indirectly taken into account through the parameter α: because of (1), for a given value of α, minimizing the number M of OXCs will also minimize the number of fiber links. Since α is an input parameter in our formulation, we believe that α in combination with the number of OXCs is representative of the overall cost of the $\mathrm{MP\lambda S}$ network.

Clearly, in order to determine the number of OXCs in the backbone we need to take into account not only the number N of LSRs but also the traffic requirements (i.e., the number of lightpaths between pairs of LSRs), the survivability properties of the network, etc. In fact, different service providers may well have different and even conflicting requirements for their networks. Rather than trying to account for all possible design requirements, we are instead interested in providing a general framework that can help us answer the above question in a way that can provide practical guidelines for building $\mathrm{MP\lambda S}$ networks. We therefore set the following requirements **R1–R5** that the $\mathrm{MP\lambda S}$ network we design must satisfy. We believe that this list captures the salient features of the network and is sufficiently general to accommodate the requirements of a wide range of network providers.

- **R1.** Each LSR accesses the backbone using two bidirectional fiber links, one to each of two different OXCs; both links are used for carrying traffic to and from each LSR.
- **R2.** The physical topology of the OXCs is 2-connected.
- **R3.** Each LSR maintains 2Δ simultaneous lightpaths to/from other LSRs.
- **R4.** Two neighbor OXCs in the physical topology are interconnected by one bidirectional fiber link.
- **R5.** The OXCs do not have any wavelength conversion capability.

The first two requirements (**R1** and **R2**) ensure that there are at least two edge-disjoint paths between any two LSRs, a necessary condition for a survivable network. **R3–R5** can be viewed as worst case requirements. **R3** ensures that the physical topology of the $\mathrm{MP\lambda S}$ network can support the maximum number of simultaneous lightpaths (recall that no LSR can have more than Δ outgoing and Δ incoming lightpaths). In particular, we use $\mathcal{L}$ to denote the set of ΔN lightpaths that the network must support. Because of **R4**, the resulting network will use single-fiber links between pairs of adjacent OXCs. Finally, **R5** requires that a lightpath be assigned a single wavelength along all the physical links it traverses. **R4** and **R5** can be easily relaxed, but are included here because we believe that there will exist $\mathrm{MP\lambda S}$ networks which will satisfy one or both of these requirements, at least during early deployment. Furthermore, we expect that relaxing either **R4** or **R5** will lead to a topology with a smaller number of OXCs compared to when both are in place, therefore, our results can be used as a worst-case scenario.

III. Solution Approach

Our objective is to determine the optimal physical topology of OXCs for establishing the given set $\mathcal{L}$ of lightpaths among the N LSRs, under the constraint on the number of wavelength W that can be supported in each fiber. This involves determining the minimum number of OXCs required as well as the links interconnecting the OXCs. Because of the difficulty of this problem, we choose an indirect approach to obtaining a near-optimal physical topology. Specifically, we first assume that the number M of OXCs in the physical topology is given and there is no constraint on the number of wavelengths. Consequently, we consider the problem of determining the links of the physical topology such that the number of wavelengths required to establish the set $\mathcal{L}$ of lightpaths is minimized. Later in this section, we show how the solution to this problem can be used to obtain a physical topology with a near-optimal number of OXCs.

Let us start with the assumption that the number M of OXCs is given. Our objective then is to obtain

1) the set of fiber links interconnecting the M OXC nodes (i.e., the physical topology);
2) the routing and wavelength assignment for the lightpaths in the set $\mathcal{L}$ among the N LSRs

such that the required number of wavelengths per fiber link in the physical topology is minimized. The inputs to the problem are: the number of OXCs, M; the number of ports in each OXC, P; and the static traffic matrix $[v_{s,d}]$, where $v_{s,d}$ represents the number of lightpaths that have to be established between OXCs s and d, $s, d = 1, 2, \ldots, M$. This traffic matrix is derived directly from the lightpath set $\mathcal{L}$ as follows: if a lightpath needs to be set up from LSR A connected to OXC s and LSR B connected to OXC d, then we increment $v_{s,d}$.

This problem can be formulated as an integer linear programming problem (ILP). The formulation, which is reported in [24], is similar to the ILP formulation of the logical topology design problem developed in [25], with additional constraints that correspond to the physical topology design subproblem (while [25] considers the physical topology as part of the input, in our case the physical topology design is part of the problem). The solution to this ILP gives an optimal topology for M OXC nodes, as well as the optimal lightpath routing and wavelength assignment that minimizes the number of wavelengths used in any link.

Let W^* be the number of wavelengths in the optimal solution to the ILP. We note that W^* can be greater than, equal to, or less than the number of wavelengths W actually supported by the fiber links of the MPλS network. If $W^* = W$, then the solution is not only optimal, but it is also feasible given the available number of wavelengths W. However, if $W^* > W$, then this optimal solution is not feasible. In this case, we may have to increase the number M of OXCs that was given as input to the problem. By solving the same problem with a larger value of M, we will obtain a new optimal solution requiring a smaller number of wavelengths. On the other hand, if $W^* < W$, the solution is feasible, but it may also be possible that another solution exists, one in which the physical topology consists of a smaller number of OXCs and which requires no more than W wavelengths. Thus, we can solve the same problem with a smaller value for M as input in the hope of finding such a solution.[1]

The above observations naturally lead to a binary search approach to obtaining a solution that requires no more than W wavelengths *and* minimizes the number of OXCs in the MPλS network. The binary search is illustrated in the overall algorithm shown in Fig. 3.

IV. The Heuristic Algorithm

The physical and logical topology design problem we consider in this work is NP-complete since it includes the wavelength assignment problem which has been shown to be NP-complete. The ILP model we discussed in Section III, and which is described in detail in [24], can only be solved for very small size networks in reasonable time because the number of variables and constraints increases much faster than the size of the network. In this section, we present a set of heuristic algorithms for this problem that can be applied to medium and large size networks. We divide the problem into the following tasks: 1) generation of a feasible physical topology, i.e., one that is 2-connected and in which no OXC has an out-degree or in-degree greater than P, 2) routing of lightpaths, and 3) assignment of wavelengths to lightpaths. Each of these tasks is solved using a heuristic algorithm. We then use a genetic algorithm (GA) to generate additional feasible physical topologies, and we iterate in order to obtain a near-optimal solution with a minimum total number of wavelengths.

A. Generation of a Random Feasible Physical Topology

Recall that a feasible OXC network is at least 2-connected. In order to get a feasible topology, we first generate a random tree, then grow a 2-connected graph from it. If we number the leaves of the tree as $i, \ldots, Y$, we can sequentially connect pairs of leaves with edges $(i, i+1), \ldots, (Y-1, Y)$ to obtain a 2-connected graph. Thus, the only question is how to generate a random tree of M nodes, where each node represents one of the OXCs.

We have adopted the method from [26] to generate a random tree of M nodes. Each tree of M nodes has an one-to-one relationship with a Prufer number that has $(M-2)$ digits. The digits are integers between 1 and M. Consider a tree $\mathcal{T}$ of M nodes numbered 1 to M in some manner. To obtain its Prufer number $P(\mathcal{T})$, we start with a null Prufer number (one with no digits) and we repeat the following steps to build $P(\mathcal{T})$ by appending one digit at a time to the right of the current Prufer number. Let i be the lowest numbered leaf in the tree, and let j be the parent of i. Then, j becomes the rightmost digit of $P(\mathcal{T})$. We remove i and edge (i, j) from $\mathcal{T}$. If i was the only child of j, then j becomes a leaf. If only two nodes remain in the tree, we stop; $P(\mathcal{T})$ has been formed. Otherwise, we repeat the above process with the new lowest numbered leaf.

[1]If W^* is less than, but close to, the number W of available wavelengths (i.e., $W - \epsilon \le W^* < W$) it is reasonable to expect only a small decrease in the number of M of OXCs in any optimal solution that uses a larger number (less than W) of wavelengths. After extensive numerical experiments, we have found that $\epsilon = 5$ provides a good tradeoff between running time and quality of the final solution, and thus we have used this value in Step 17 of Fig. 3.

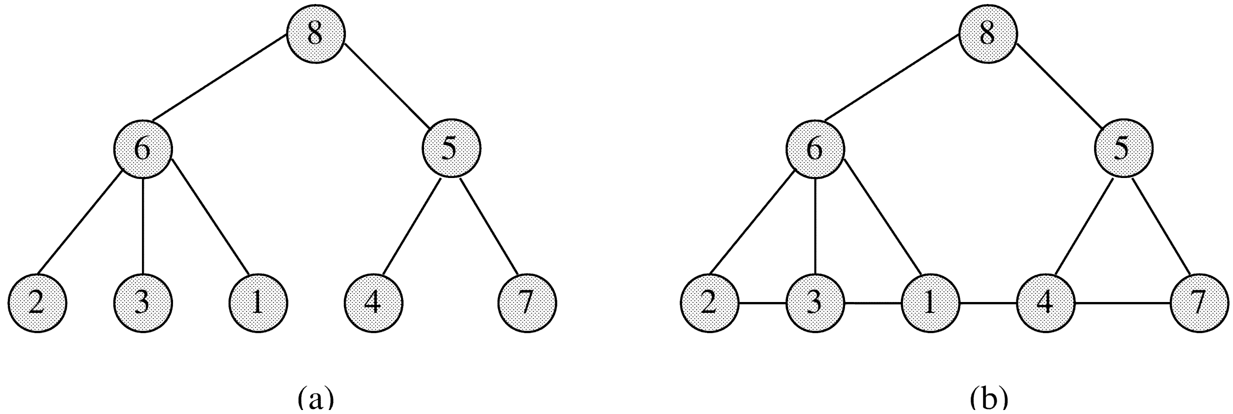

Fig. 1. (a) Tree of 8 nodes corresponding to the Prufer number 666585. (b) Resulting 2-connected graph.

The reverse process of obtaining a tree from a Prufer number $P(\mathcal{T})$ with $(M-2)$ digits in the range 1 to M consists of these steps. First, designate all nodes whose number does not appear in $P(\mathcal{T})$ as *eligible*. Let i be the lowest numbered eligible node, and let j be the leftmost digit of $P(\mathcal{T})$. Add the edge (i,j) to $\mathcal{T}$ and remove the leftmost digit j from $P(\mathcal{T})$. Designate i as no longer eligible. If j does not occur anywhere in what remains of $P(\mathcal{T})$, then designate j as eligible. Repeat the above procedure until no digits remain in $P(\mathcal{T})$, in which case there are exactly two nodes, say, i and j, still eligible. Add (i,j) to $\mathcal{T}$ and stop. Since there are exactly $(M-2)$ digits in the Prufer number and we remove a digit each time we add an edge to $\mathcal{T}$, the final graph $\mathcal{T}$ has exactly $M-1$ edges; it is also shown in [26] that $\mathcal{T}$ has no cycles, thus it is a tree.

As an example, suppose that $M = 8$, and that the Prufer number is 666 585 (consisting of $M-2=6$ digits in the range 1–8). Nodes 1, 2, 3, 4, and 7 do not appear in $P(\mathcal{T})$ and are designated as eligible for consideration. Node 1 is the lowest numbered eligible node, and digit 6 is the leftmost digit of the Prufer number. Consequently, we add the edge (6,1) to the tree, we make node 1 ineligible, and we remove the leftmost digit 6 from the Prufer number. The remaining Prufer number is thus 66585, and since digit 6 appears in it, node 6 is not made eligible. In the second step, node 2 is the lowest numbered eligible node and digit 6 is the leftmost digit in the remaining Prufer number. We add edge (6,2) to the tree $\mathcal{T}$, we make node 2 ineligible, and we remove digit 6 from the Prufer number. The number which remains is 6585, so node 6 remains ineligible. In the third iteration, node 3 is the lowest numbered eligible node and the leftmost digit in the Prufer number is once again digit 6. Thus, we add edge (6,3) to the tree, we make node 3 ineligible, and we remove digit 6 to obtain the new number 585. Since the digit 6 does not appear in the new number, we make node 6 eligible. We continue in this manner, and we add edges (5,4), (8,6), and (5,7) to the tree, at which point the Prufer number becomes null (all its digits have been removed). At this stage there are only two nodes eligible, nodes 5 and 8. We next add edge (8,5) to obtain the tree shown in Fig. 1(a).

Based on the above discussion, the following steps summarize the algorithm for generating a random feasible physical topology.

1) Given the number M of OXCs, randomly generate $(M-2)$ digits in the range of 1 to M to form $P(\mathcal{T})$.
2) Generate a tree $\mathcal{T}$ using the Prufer number, as above.
3) Construct a 2-connected graph by adding edges sequentially to connect the leaves of the tree $\mathcal{T}$.
4) After connecting the leaves of the tree in Fig. 1(a), we get the 2-connected graph shown in Fig. 1(b).

B. Routing and Wavelength Assignment

We now assume that we are given a 2-connected physical topology of M OXCs, as well as a set of lightpaths between pairs of OXCs that need to be established. We now present two algorithms, one to route each lightpath over a physical path of fiber links, and another to assign wavelengths to the lightpaths. Note that we treat the routing and wavelength assignment subproblems independently; this approach may require a larger number of wavelengths than a combined solution, but the latter is intractable while our approach can be applied directly to networks of realistic size. Furthermore, as we shall see, the routing algorithm takes into account the number of lightpaths using each link in order to minimize the number of wavelengths needed.

We use Dijkstra's shortest path algorithm to route the set of lightpaths over the given physical topology. In order to minimize the number of wavelengths used on a physical link, we use two heuristic approaches. First, the link weight used in Dijkstra's algorithm is dynamically adjusted to reflect the number of wavelengths already allocated on each link. Consider physical link ℓ and let C_ℓ be the actual link cost (C_ℓ is a constant) and w_ℓ be the number of lightpaths already using this link. Then, each time we run Dijkstra's algorithm to find a path for a certain lightpath, we use the quantity $L_\ell = C_\ell + Hw_\ell$ as the cost of link ℓ, where H is a tunable weight parameter. This cost function forces new lightpaths to be routed over less congested links in the physical topology, reducing the total number of wavelengths used in the network. For the next application of the algorithm, quantities w_ℓ are incremented for all links ℓ along the path of the just routed lightpath.

The second heuristic approach has to do with the order in which we consider the given lightpaths for routing. Specifically, we first sort the OXC nodes in an ascending order according to their degree (ties are broken arbitrarily). Starting with the first OXC node (the one with the smallest degree), we apply Dijkstra's algorithm to route all lightpaths that have this node as source or destination. We proceed in this manner by considering nodes with higher degrees. This method results to a considerably lower wavelength usage than when selecting the nodes randomly. This is because nodes with smaller degrees have fewer alternative links to route their lightpaths. In view of this, routing lightpaths originating or terminating at these nodes first will increase the wavelength use of their links. Because of the cost function described above, later lightpaths will tend to avoid the links around these nodes. On the other hand, if a node with a small degree was considered late in the process, its lightpaths would have to use one of its links regardless of how congested these links were, potentially increasing the overall number of wavelength required. We refer to this scheme as the smallest-degree-first-routing (SDFR) algorithm. The following steps summarize the heuristic algorithm for routing lightpaths.

1) Use the link weight function $L_\ell = C_\ell + Hw_\ell$.
2) Sort the nodes in an ascending order of their degree.
3) Consider each node in this order and use the Dijkstra's algorithm to build the shortest path for its lightpaths.

Once the physical links for each lightpath have been obtained, we need to assign wavelengths such that if two lightpaths share the same link then they are assigned a different wavelength. This wavelength assignment problem can be shown to be equivalent to the vertex coloring problem of an induced simple graph [15]. The induced graph is such that its vertices correspond to lightpaths in the original network, and vertices of the induced graph are linked by an edge only if the two corresponding lightpaths share the same physical link. A heuristic algorithm was developed in [15] to solve the vertex coloring problem. The algorithm uses a greedy approach to assign wavelengths (color) to the lightpath (vertex). We adopt this algorithm to perform wavelength assignment, since it has been shown to have good accuracy and to run in polynomial time. We also note that the upper bound of the number of distinct colors (wavelengths) used is equal to the maximum degree of the induced connection graph plus one.

C. *The Genetic Algorithm (GA)*

In the last decade, genetic algorithms (GAs) [27], [28] have proved to be a practical and robust optimization and search tool. These algorithms are based on the mechanisms of evolution and natural genetics that lead to the survival of the fittest by the process of natural search and selection. A GA generates a sequence of populations using a selection mechanism, and then applies crossover and mutation as search mechanisms. A GA is a global random search technique in the solution space of the problem and it usually avoids entrapping into a local optimization; for further details on the properties of GAs the reader is referred to [27], [28] and the references therein. The steps involved in a GA are as follows.

1) Design an efficient encoding scheme (chromosome) of the solution. Usually, this is a one-to-one bit string mapping to a solution.
2) Generate an initial set of feasible solutions. This set is referred to as a *generation*. The population size of the generation is determined by the variable G_s. This initial generation becomes the current generation.
3) If the stop criterion has been met, return the best solution in the current generation as the near-optimal solution and stop; otherwise, continue to Step 4 for another iteration. Usually the stop criterion is a predetermined number of iterations.
4) Randomly select solutions from the current generation based on their fitness values given by the value of the objective function. The selection probability for a solution is generally proportional to its fitness value. The selected solutions form the basis of the *offspring* generation.
5) Perform crossover on the selected solutions. This step includes picking up pairs of strings at random, randomly choosing a crossover point, and switching the two strings after the point. This crossover is controlled by the *crossover rate* R_c. The algorithm invokes crossover only if a randomly generated number is less than R_c.
6) Perform mutation on the solutions obtained from Step 5, i.e., flip the bits of each string. Mutation is controlled by the *mutation rate* R_m, which is the probability that a bit will be flipped.
7) Calculate the fitness value for the new individuals if the new individual solutions are feasible. If an individual is not feasible, it is either made feasible or it is dropped.
8) Repeat Steps 4–6 until a new generation of the population size is generated. Set this as the current generation. Add the best solution from the last generation to the current generation. Go back to Step 3.

We use a GA algorithm to generate feasible physical topologies starting from the initial physical topology we obtained in Section IV-A. The objective is to search for physical topologies that will improve on the number of wavelengths required to establish the given set of lightpaths. Since we assume that the number M of OXCs is given, we define our solution space as the set of feasible physical topologies on M nodes. We now proceed to describe the encoding of the solution, the calculation of the fitness value, the selection, crossover, and mutation strategies, and the handling of infeasible solutions.

Encoding. Consider a graph with M vertices. We number the vertices from 0 to $M-1$ and the edges of the graph from 0 to $(M(M-1)/2)-1$. We read the string from left to right. Each edge (i,j) is numbered using an index k, which is defined according to the index of the two endpoints of the edge, i and j. Specifically, $k = (M(M-1)/2) - ((M-i)(M-i-1)/2) + j - i - 1$, $0 \le i < j < M$. The solution is encoded into a chromosome (i.e., a bit string of length $M(M-1)/2$) as follows. Each edge is represented by a bit in the bit string. If the edge exists, the bit is set to 1; otherwise it is set to 0. The position of the bit representing an edge is given by the edge index k.

As an example, let us consider the graph shown in Fig. 2(a). The code of this graph is 101001000110011, which is simply the concatenation of the five rows above the diagonal of the adjacency matrix for this graph, shown in the shaded area of Fig. 2(b).

Fitness value. The fitness value of a feasible individual (physical topology) is the negative of the number of wavelengths required to support all the lightpaths in the set $\mathcal{L}$ (i.e., the smaller the number of wavelengths, the higher the fitness value). This number of wavelengths is obtained by running the routing and wavelength assignment heuristic algorithms on this individual.

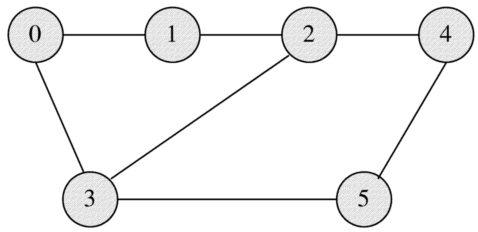

(a)

Adjacency Matrix

$$\begin{array}{c|c} & \begin{matrix} 0 & 1 & 2 & 3 & 4 & 5 \end{matrix} \\ \begin{matrix} 0 \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{matrix} & \begin{bmatrix} 0 & 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 & 0 \end{bmatrix} \end{array}$$

(b)

Fig. 2. (a) Physical topology with six nodes. (b) Its adjacency matrix.

GA parameters. The performance of a GA is determined by the right choice of control parameters: the crossover rate R_c, the mutation rate R_m, and the population size, G_s. In order to shorten the running time of the algorithm, we choose a relatively small population size and therefore a high level of string disruption. We set $R_c = 0.8$, $R_m = 0.1$, and $G_s = 25$. We use the following operations to generate an offspring from the parent generation until we get a new generation of the population size G_s. These operations work on the edge string encoding a graph.

- *Selection:* Select two candidates from the parent generation using the roulette wheel selection mechanism. Individuals with higher fitness value have higher probability to be selected.
- *Crossover:* We crossover the two candidates obtained by the above selection operation. For example, let us assume that the two selected individuals from the parent generation are 1100001101 and 0110110011, and that the crossing point is the sixth digit. In this case, we obtain the following two new codes: 1100000011 and 0110111101.
- *Mutation:* We mutate bit-by-bit the two individuals obtained from the above crossover operation. For example, after mutating the above two codes we may obtain: 1101000011 and 0111111101.

Infeasible individuals. It is possible that an offspring generated by the above process corresponds to an infeasible graph. That is, the graph may not be 2-connected or it may exceed the constraint on the connectivity α. Therefore, we run the following four tests for each generated offspring.

1) Check the nodal degree of each node to make sure that the constraint on the number of ports per OXC is satisfied.
2) Check if the total number of edges of the generated graph exceeds the given number of edges determined by the connectivity α.
3) Check that the degree of each node is greater than 2.
4) Use Depth-First-Search to ensure that the graph is 2-connected. Steps 3 and 4 combine to check if the network is 2-connected.

If tests 1 or 2 fail, then we drop the individual. If tests 3 or 4 fail, the graph is not 2-connected. We add edges between the disconnected OXCs until the graph becomes 2-connected. If the solution passes the above four tests, we run the RWA heuristics to obtain the minimum number of wavelengths needed in the network.

Complete Algorithm for Designing an MPλS Network
Input: The number N of LSRs, the set $\mathcal{L}$ of lightpaths, the number P of ports per OXC, the number of wavelengths W in each fiber, the OXC connectivity α, and the number of generations G for the genetic algorithm.
Output: A physical topology of M OXCs, and the routing and wavelength assignment for the set $\mathcal{L}$ of lightpaths.

1. begin
2. $M_{max} \leftarrow N/2$, $M_{min} \leftarrow 1$
3. $num\_gen \leftarrow 0$
4. $M \leftarrow (M_{max} + M_{min})/2$ // Binary search
5. $phy\_top \leftarrow$ initial physical topology generated using the algorithm described in Section IV-A
6. For the physical topology phy_top:
7. Solve the RWA problem (see Section IV-B)
8. Calculate the objective function (fitness value) W_f
9. While $num\_gen < G$: // Genetic algorithm
10. $num\_gen \leftarrow num\_gen + 1$
11. Use the genetic algorithm in Section IV-C to create a new generation of physical topologies
12. For each individual phy_top in the new generation calculate the fitness value W_f using Steps 7 and 8
13. $W^\star \leftarrow$ smallest fitness value W_f from last generation
14. $phy\_top^\star \leftarrow$ physical topology corresponding to $W^\star$
15. If $W^\star > W$ then
16. $M_{min} \leftarrow M$; go back to Step 3
17. Else if $W^\star < W - 5$ then
18. $M_{max} \leftarrow M$; go back to Step 3
19. Else return $M, phy\_top^\star$
20. end of the algorithm

Fig. 3. Algorithm for the design of MPλS networks.

The complete algorithm is shown in Fig. 3. A detailed complexity analysis is presented in [24], but Steps 6–18 of the algorithm take time $O(GG_s(\Delta N M^2 + \Delta^2 N^2))$ for an iteration of the binary search in which the number of OXCs is M; recall that G is the number of generations in Step 9 of the algorithm, G_s is the population size, N is the number of LSRs, and Δ is the number of transceivers per LSR. Also, the size of the lightpath set is ΔN, therefore, the above function expresses the running time of the algorithm as a function of the amount of traffic demands.

V. Lower Bounds on the Number of OXCs

We now present three different lower bounds on the number M of OXCs required given: the number N of LSRs, the number W of wavelengths, the degree of connectivity α of the optical network, the number P of ports of each OXC, and the number Δ of optical interfaces at each router. The maximum of the three values is a lower bound on the number of OXCs for the given problem instance.

We obtain the first bound simply by counting the number of OXC ports required in the network. Recall that E denotes the number of links in the optical network. Now, $2N$ OXC ports are needed to connect the N LSRs to the optical network (since each LSR is connected to two OXCs), while $2E$ ports are needed for the fiber links interconnecting the OXCs. Since the total number of OXC ports is MP, we have

$$MP \geq 2N + 2E = 2N + \alpha M(M-1). \tag{2}$$

Solving this equation for M, we obtain

$$M \geq \left\lceil \frac{P+\alpha-\sqrt{(P+\alpha)^2-8\alpha N}}{2\alpha} \right\rceil \tag{3}$$

where input parameters N, P, and α should be such that $(P+\alpha)^2 \geq 8\alpha N$.

The above result only considers the constraint on the number of ports per OXC and the degree of connectivity of the OXC network, and, in general, it is not a tight lower bound. We now present a better lower bound that takes the available number W of wavelengths into consideration. Recall from requirement **R3** that the network must support $N\Delta$ lightpaths, where Δ is the number of optical interfaces at each of the N LSRs. Let D denote the average number of hops (fiber links traversed) over all lightpaths. Then, we have

$$N\Delta D \leq 2WE \tag{4}$$

where $2WE$ is the total number of link-wavelengths in the network, assuming that each fiber link between any pair of OXCs consists of two unidirectional fibers. Combining this result with (2), we obtain

$$M \geq \frac{N\Delta D}{WP} + \frac{2N}{P} = \frac{N}{P}\left(\frac{\Delta D}{W} + 2\right). \tag{5}$$

From (5) we see that the lower bound on M depends on the lower bound on the average hop length D over all lightpaths. An immediate lower bound on D is 1, so we obtain a second lower bound on M as

$$M \geq \frac{N}{P}\left(\frac{\Delta}{W} + 2\right). \tag{6}$$

To obtain a better lower bound on D, we note that, in a network with M nodes, at least $M/2$ nodes are at distance $\lfloor \log_d M/2 \rfloor$ or more from any given node [29], where d is the maximum nodal degree. In our case, every OXC has P ports, among which $2N/M$ are used to interconnect the LSRs assigned to this OXC, so the maximum nodal degree in the core OXC network is $d = P - 2N/M$. Therefore, we have that $D \geq ((\log_d M/2) + 1)/2$. If we use this value of D as the average hop length in (5), we have

$$W(MP-2N) \geq N\Delta \frac{\left(\log_d \frac{M}{2}\right)+1}{2}, \quad d = P - \frac{2N}{M}. \tag{7}$$

Using iteration over M from 0 to N on both sides of the above inequality until it is satisfied, we can obtain the minimum feasible value of M.

The final lower bound on M is the maximum value obtained from expressions (3), (6), and (7).

VI. Numerical Results

In this section we present results that illustrate how the different design parameters affect the number of OXCs required to interconnect a set of LSRs. We vary the design parameters as follows: the number N of LSRs is varied between 100–1000; the number W of wavelengths per link takes the values 32, 64, 128; the number P of ports per OXC varies from 16–64; and the number Δ of transceivers per LSR takes the values 4–24. We also set the upper bound on the degree of connectivity α of the $\mathrm{MP}\lambda\mathrm{S}$ network of OXCs to 0.4. For the graphs shown here, we have assigned the N LSRs to the OXCs in a round-robin manner. That is, the first LSR is attached to the first and second OXC, the second LSR to the third and fourth OXC, and so on (recall that, by requirement **R1**, each LSR must attach to at least two OXCS). This assignment is made for convenience only, and is not inherent to our approach; in fact, the algorithm in Fig. 3 can accommodate any arbitrary assignment of LSRs to OXCs. The set $\mathcal{L}$ of lightpaths is chosen so that each LSR has exactly Δ incoming and Δ outgoing lightpaths (see requirement **R3**) to a random set of other LSRs. Again, however, our algorithm will accommodate any set of lightpaths.

We should point out that the above ranges of the design parameters are based on realistic assumptions regarding the state of the technology and the size of the $\mathrm{MP}\lambda\mathrm{S}$ networks, and go far beyond the small networks to which previous virtual topology algorithms were limited. The results presented here illustrate that our algorithm can be applied to networks of size between one and two orders of magnitude greater than that of the networks studied previously. We also emphasize that the algorithm in Fig. 3 computes not just the number of OXCs, but also the physical topology of the $\mathrm{MP}\lambda\mathrm{S}$ network (i.e., the physical links between the OXCs), as well as the routing and wavelength assignment for the set $\mathcal{L}$ of lightpaths between the LSRs. However, due to the large size of the resulting networks, it is not possible to draw the physical topology of fiber links or the logical topology of lightpaths here.

In Fig. 4 we plot the number M of OXCs in the $\mathrm{MP}\lambda\mathrm{S}$ network against the number N of LSRs. Three plots are given for three different values of the number of wavelengths per link, $W = 32$, 64, 128. For these results, we have let $\Delta = 12$ and $P = 64$. We make two important observations. First, the number M of OXCs increases almost linearly with the number N of LSRs, but the slope of the curves is moderate. In particular, an increase by a factor of ten in the number of LSRs results in an increase in the number of OXCs by a factor between four (for $W = 32$) and seven (for $W = 128$). Since we have kept the degree of connectivity at around 0.4 for all physical topologies, the

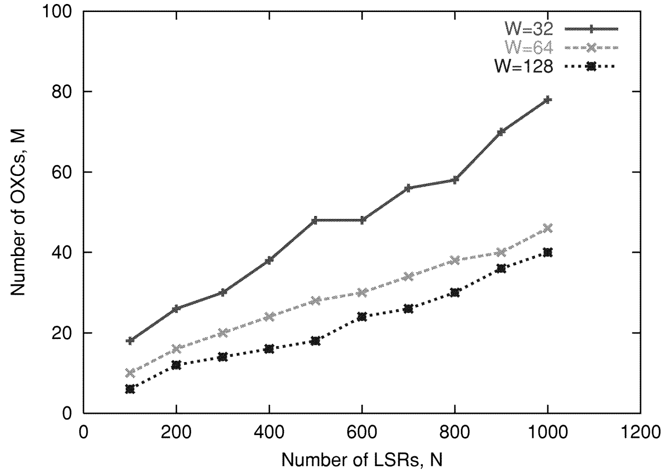

Fig. 4. Number of OXCs in the physical topology ($\Delta = 12, P = 64$).

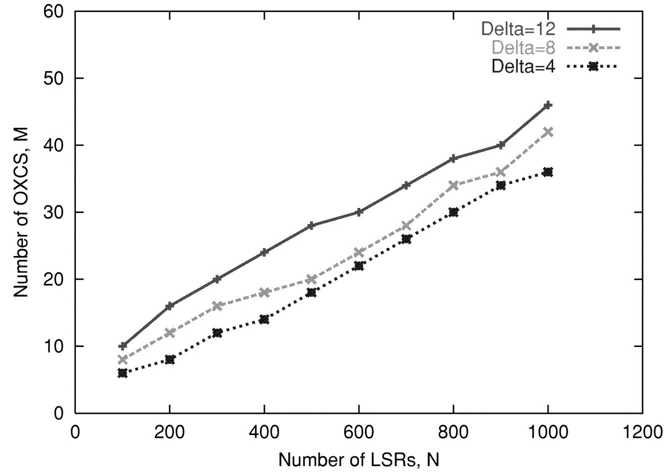

Fig. 5. Number of OXCs in the physical topology ($W = 64, P = 64$).

corresponding increase in the number of links in the topology, as N increases, is similar. This result implies that OXC networks to interconnect very large number of LSRs can be built cost-effectively. The second observation is that the larger the number of wavelengths available at each fiber link, the smaller the number of OXCs required for a given number N of LSRs. This result is expected, however, we note that a two-fold increase in the number of wavelengths (from 32 to 64) reduces the number of OXCs by less than 1/2 (between 44% for $N = 100$ and 41% for $N = 1000$). We also see the effect of diminishing returns, since a second two-fold increase in the number of wavelengths (from 64 to 128) results in a smaller reduction in the number of OXCs (between 40% for $N = 100$ and 13% for $N = 1000$). Note that, for a given value of N, the number of lightpaths in the set $\mathcal{L}$ remains constant at ΔN across the three curves in Fig. 4, so one would expect a larger decrease in the number of OXCs as the number of wavelengths increases. However, recall that $2N$ OXC ports are needed to attach the LSRs to the OXC network. Since the number of ports per OXC remains constant at $P = 64$, as W increases, the number of OXCs needed is constrained by the number of ports required rather than the number of lightpaths that need to be established. In other words, in order to take full advantage of the larger number of wavelengths in the fiber, OXCs with a larger number of ports must be employed.

In Fig. 5 we fix the number of wavelengths to $W = 64$ and the number of ports per OXC to $P = 64$, and we plot the number of OXCs against the number N of LSRs. Three curves are shown, one for a different value of the number of transceivers per LSR, $\Delta = 4$, 8, 12. Note that the curve for $\Delta = 12$ is identical to the middle curve of Fig. 4 for $W = 64$, although the scale in the two figures is different. Again, we see that the number of OXCs increases linearly with the number of LSRs. Recall that the number ΔN of lightpaths that must be established for a given value of N increases linearly with Δ. However, the curves in Fig. 5 show that the number of OXCs for a given N value needed to support the larger number of lightpaths increases much more slowly than Δ. For instance, for $N = 100$, 6 OXCs are needed for $\Delta = 4$, 8 OXCs for $\Delta = 8$, and 10 OXCs for $\Delta = 12$. For $N = 1000$, the corresponding number of OXCs are 36, 42, and 46, respectively. These results indicate that a relatively small

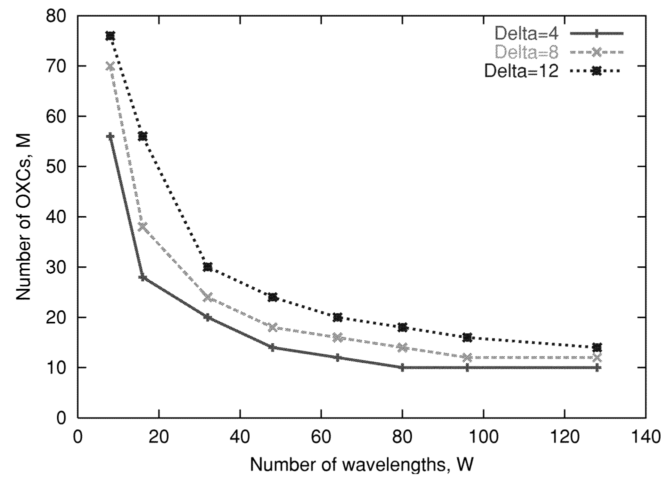

Fig. 6. Number of OXCs in the physical topology ($N = 300, P = 64$).

incremental cost (in terms of additional OXCS and fiber links to interconnect them) can provide a significantly richer connectivity among the LSRs.

In Fig. 6 we plot the number of OXCs against the number of wavelengths when the number of LSRs is constant at $N = 300$ and the number of ports per OXC is $P = 64$. Three curves for different numbers of transceivers per LSR are shown, $\Delta = 4$, 8, 12. The results are as expected. Specifically, the number of OXCs needed decreases as the number W of wavelengths per fiber increases, but the curves flatten out once $W > 80$. As we mentioned above, this reflects the fact that a larger number of ports per OXC is needed to take full advantage of the large number of wavelengths. Also, more OXCs are required as Δ increases, but the results are consistent with the previous figure in that the increase in the number of OXCs is significantly slower than the increase in Δ (and the corresponding increase in the number of lightpaths to be established). Similar observations can be made from Fig. 7, where we let $N = 300$ and $P = 64$, and we plot the number of OXCs against the number Δ of transceivers per LSR. Note that Δ (and, consequently, the number of lightpaths) increases by a factor of six from 4 to 24, the number of OXCs required increases much slower, from 20 to 48 when

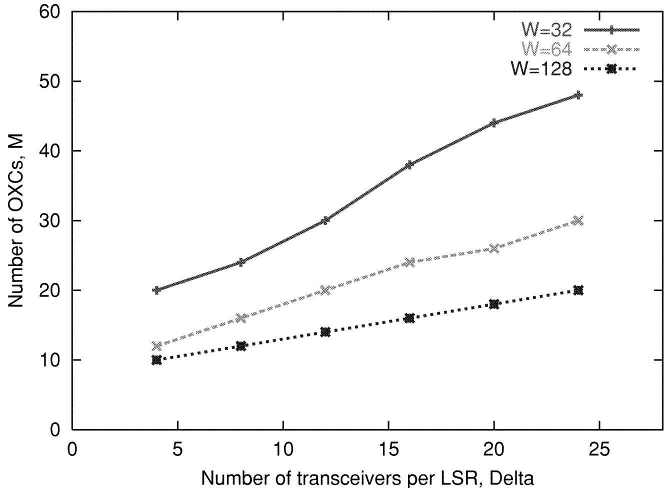

Fig. 7. Number of OXCs in the physical topology ($N = 300, P = 64$).

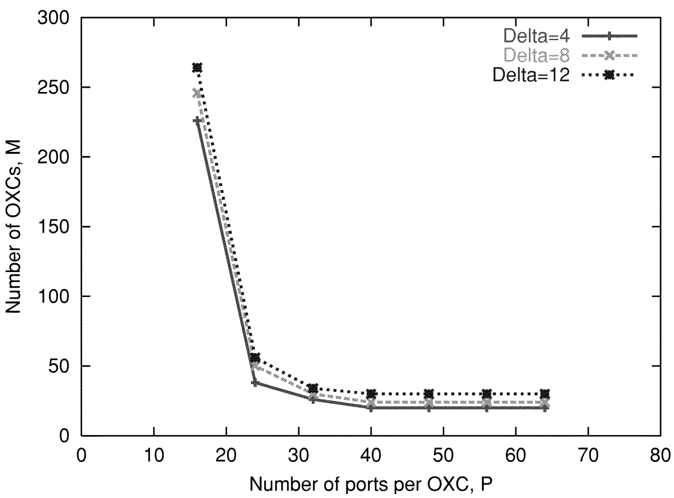

Fig. 8. Number of OXCs in the physical topology ($N = 300, W = 32$).

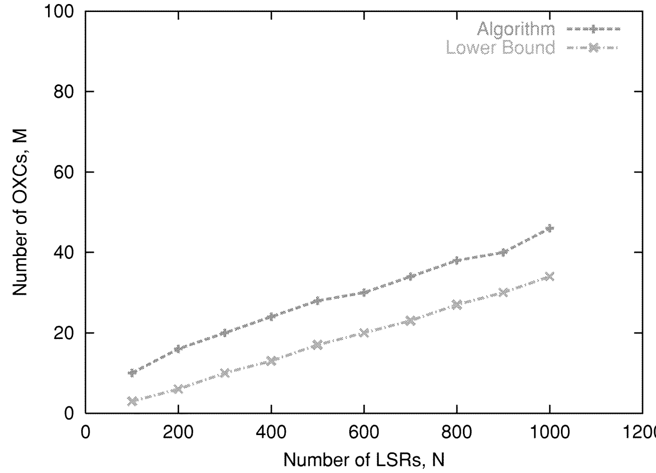

Fig. 9. Heuristic versus lower bound ($\Delta = 12, P = 64, W = 64$).

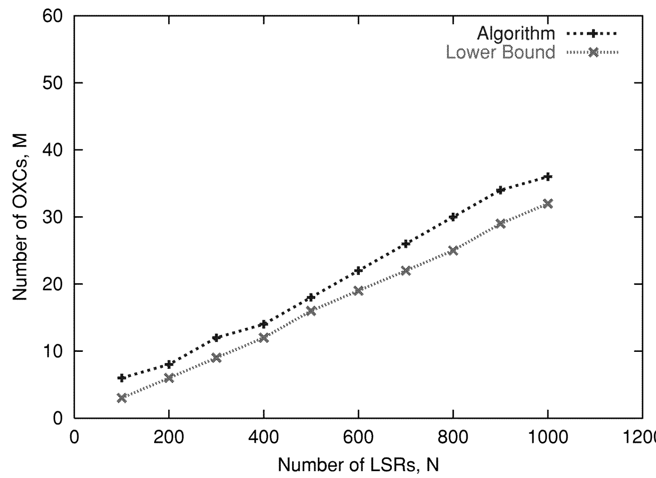

Fig. 10. Heuristic versus lower bound ($W = 64, P = 64, \Delta = 4$).

$W = 32$. In fact, the increase in the number of OXCs is even slower for larger number of wavelengths, from 10 to only 20 for $W = 128$.

In Fig. 8 we let $N = 300$ and $W = 32$ and we plot the number of OXCs against the number P of ports, for three different values of Δ, $\Delta = 4$, 8, 12. There is a sharp drop in the number of OXCs initially, as P increases from 16 to 24, but the curves level off after that. Also, we again see that the value of Δ does not significantly affect the number of OXCs. These results indicate that by employing OXCs of medium size (in terms of P) can have a dramatic effect in the number of OXCs required. While OXCs with many ports are expected to be more expensive than those with few ports, the dramatic drop in the curves of Fig. 8 indicates that it may be cost-effective to employ the former; also, for the same degree of connectivity α, fewer OXCs implies significant savings in fiber links.

Finally, in Figs. 9 and 10 we compare the results of our heuristic algorithm to the lower bound we presented in Section V. Fig. 9 plots the middle curve of Fig. 4 and the corresponding lower bound; thus, these plots correspond to the following values of the input parameters: $\Delta = 12$, $P = 64$, and $W = 64$. As we can see, the two curves have very similar behavior, and the results from our heuristic are close to the lower bound. We emphasize that the lower bound is obtained simply by counting the number of network resources that are absolutely necessary to support the given set of lightpaths. In other words, there is no guarantee that there exists a physical topology with a number of OXCs equal to the lower bound, such that it is possible to route and assign wavelengths to the given set of lightpaths. Therefore, the optimal number of OXCs lies somewhere between the two curves in Fig. 9.

Fig. 10 is similar to Fig. 9, but it plots the bottom curve of Fig. 5 and the corresponding lower bound, for $W = 64, P = 64$, and $\Delta = 4$. Again, we see our heuristic returns a number of OXCs that is close to the lower bound, and this result is consistent across the range of the number N of LSRs shown in the figure. Very similar results regarding the relative performance of our heuristic and the lower bound have been obtained for a wide range of the input parameters.

Finally, we investigate the sensitivity of our algorithm to the set of lightpaths in the set $\mathcal{L}$ that is used as input for designing the physical and logical topology of the optical network (refer to Fig. 3). To this end, we used the same parameters as in Fig. 4,

TABLE I
MINIMUM, MAXIMUM, AND AVERAGE NUMBER OF OXCS IN THE PHYSICAL TOPOLOGY OVER TEN DIFFERENT LIGHTPATH SETS FOR THE STATED VALUES OF THE SYSTEM PARAMETERS ($\Delta = 12, P = 64$)

No. of LSRs	No. OXCs in physical topology								
	$W = 32$			$W = 64$			$W = 128$		
	Min	Max	Avg	Min	Max	Avg	Min	Max	Avg
100	16	18	18	10	12	11	6	8	7
200	24	28	26	16	18	16	10	12	11
300	30	34	32	18	22	20	12	14	14
400	38	38	38	22	26	24	16	16	16
500	44	52	47	26	30	28	18	20	19
600	46	52	49	28	34	30	22	24	23
700	52	62	57	32	34	34	26	28	26
800	56	66	61	36	42	37	30	32	30
900	62	72	68	40	42	40	34	36	35
1000	70	80	75	44	46	45	38	40	40

and we run the algorithm for 10 different lightpath sets. The lightpath sets were randomly selected such that each LSR has exactly Δ incoming and outgoing connections. Our results are summarized in Table I, where we present the minimum, maximum, and average number of OXCs obtained by our algorithm for each set of lightpaths for a given set of values for the system parameters. As we can see, the results are not particularly sensitive to the input set $\mathcal{L}$ of lightpaths: when all other parameters are the same, the difference between the minimum and maximum number of OXCs for the different lightpath sets is within 10–15%. Recall that each lightpath set represents a worst-case scenario, in that the optical interfaces at each LSR are all assumed to be active. However, in a typical network operation, the number of simultaneous connections will in general be less than in a worst-case scenario. Therefore, we believe that the networks designed by our algorithms will be capable of supporting a wide range of typical dynamic traffic demands.

VII. CONCLUDING REMARKS

We have described a set of heuristic algorithms for the physical and logical topology design of large-scale optical networks of OXCs. Our objective has been to minimize the number of OXCs given a constraint on the number of wavelengths per fiber link and certain constraints (i.e., biconnectivity) on the physical topology. We presented routing and wavelength assignment heuristics, as well as a genetic algorithm to iterate over the set of physical topologies. We have applied our algorithms to design networks that can accommodate hundreds of LSRs and several thousands of lightpaths. Our results are close to optimal, and they have shed new light into the design of $\mathrm{MP}\lambda\mathrm{S}$ networks. The most important finding is that it is possible to build cost-effective networks that provide rich connectivity among the LSRs with relatively few, but properly dimensioned, OXCs.

Our work can be extended in several directions. First, our algorithm does not take into account protection, while one of the requirements of current and future networks is to provide backup paths in the case of failures. Therefore, it would be interesting to investigate how the results and conclusions we presented in Section VI are affected by different protection schemes, such as dedicated or shared protecttion; we anticipate that new design algorithms will be needed for such a study. Also, given the feasibility and potential of performing multicast in the optical domain through light splitting, the design of physical and logical topologies for optical networks supporting multicast is emerging as an important field of study. These problems will be the focus of our future research efforts.

REFERENCES

[1] B. Davie and Y. Rekhter, *MPLS Technology and Applications*. San Diego, CA: Morgan Kaufmann, 2000.
[2] P. Ashwood-Smith *et al.*, "Generalized MPLS—Signaling Functional Description, Work in Progress," IETF Draft draft-ietf-mpls-generalized-signaling-06.txt, 2001.
[3] D. H. Su and D. W. Griffith, "Standards activities for MPLS over WDM networks," *Optical Networks*, vol. 1, no. 3, July 2000.
[4] E. Karasan and E. Ayanoglu, "Effects of wavelength routing and selection algorithms on wavelength conversion gain in WDM optical networks," *IEEE J. Select. Areas Commun.*, vol. 16, pp. 1081–1096, Sept. 1998.
[5] R. Dutta and G. N. Rouskas, "A survey of virtual topology design algorithms for wavelength routed optical networks," *Optical Networks*, vol. 1, no. 1, pp. 73–89, Jan. 2000.
[6] T. Lee, K. Lee, and S. Park, "Optimal routing and wavelength assignment in WDM ring networks," *IEEE J. Select. Areas Commun.*, vol. 18, pp. 2146–2154, Oct. 2000.
[7] I. Chlamtac, A. Ganz, and G. Karmi, "Lightpath communications: An approach to high bandwidth optical WANS," *IEEE Trans. Commun.*, vol. 40, pp. 1171–1182, July 1992.
[8] M. Alanyali and E. Ayanoglu, "Provisioning algorithms for WDM optical networks," *IEEE/ACM Trans. Networking*, vol. 7, pp. 767–778, Oct. 1999.
[9] N. Nagatsu, S. Okamoto, and K. Sato, "Optical path cross-connect system scale evaluation using path accommodation design for restricted wavelength multiplexing," *IEEE J. Select. Areas Commun.*, vol. 14, no. 5, pp. 893–902, June 1996.
[10] D. Banerjee and B. Mukherjee, "Wavelength-routed optical networks: Linear formulation, resource budgeting, and a reconfiguration study," *IEEE/ACM Trans. Networking*, vol. 8, pp. 598–607, Oct. 2000.
[11] B. Mukherjee *et al.*, "Some principles for designing a wide-area WDM optical network," *IEEE/ACM Trans. Networking*, vol. 4, pp. 684–696, Oct. 1996.
[12] D. Saha, M. D. Purkayastha, and B. Mukherjee, "An approach to wide area WDM optical network design using genetic algorithms," *Comput. Commun.*, vol. 22, pp. 156–172, 1999.
[13] R. Ramaswami and K. Sivarajan, "Routing and wavelength assignment in all-optical networks," *IEEE/ACM Trans. Networking*, vol. 3, pp. 489–500, Oct. 1995.
[14] S. Xu, L. Li, and S. Wang, "Dynamic routing and assignment of wavelength algorithms in multifiber wavelength division multiplexing networks," *IEEE J. Select. Areas Commun.*, vol. 18, pp. 2130–2137, Oct. 2000.
[15] Z. Zhang and A. Acampora, "A heuristic wavelength assignment algorithm for multihop WDM networks with wavelength routing and wavelength reuse," *IEEE/ACM Trans. Networking*, vol. 3, pp. 281–288, June 1995.
[16] S. Baroni and P. Bayvel, "Wavelength requirements in arbitrarily connected wavelength-routed optical networks," *J. Lightwave Technol.*, vol. 15, pp. 242–251, Feb. 1997.
[17] ——, "Key topological parameters for the wavelength-routed optical networks," in *Proc. ECOC 1996*, 1996.
[18] R. A. Barry and P. A. Humblet, "On the number of wavelengths and switches in all-optical networks," *IEEE Trans. Commun.*, vol. 42, pp. 583–591, Feb./Mar./Apr. 1996.
[19] D. R. Hjelme, A. Royset, and B. J. Slagsvold, "How many wavelengths does it take to build a wavelength routed optical network," in *Proc. ECOC 1996*, 1996.

[20] B. Beauquier, "All-to-all communication for some wavelength routed all-optical networks," *Networks*, vol. 30, 1999.

[21] M. A. Marsan, A. Bianco, and E. Leonardi, "Topologies for wavelength-routing all-optical networks," *IEEE/ACM Trans. Networking*, vol. 1, pp. 534–546, Oct. 1993.

[22] P. B. Chu, S.-S. Lee, and S. Park, "MEMS: The path to large optical crossconnects," *IEEE Commun. Mag.*, vol. 40, pp. 80–87, Mar. 2002.

[23] S. Baroni, J. O. Eaves, M. Kumar, and M. A. Qureshi, "Analysis and design of backbone architecture alternatives for IP optical networks," *IEEE J. Select. Areas Commun.*, vol. 18, pp. 1980–1994, Oct. 2000.

[24] Y. Xin, "Topology Design of Large-Scale Optical Networks," Ph.D. dissertation, North Carolina State University, 2002.

[25] R. M. Krishnaswami and K. N. Sivarajan, "Design of logical topologies: A linear formulation for wavelength-routed optical networks," *IEEE/ACM Trans. Networking*, vol. 9, pp. 186–198, Apr. 2001.

[26] C. C. Palmer and A. Kershenbaum, "An approach to a problem in network design using genetic algorithms," *Networks*, vol. 26, pp. 151–163, 1995.

[27] M. Gen and R. Cheng, *Genetic Algorithms and Engineering Optimization*. New York: Wiley, 2000.

[28] M. Srinivas and M. Patnaik, "Genetic algorithms: A survey," *IEEE Comput.*, vol. 27, pp. 17–26, June 1994.

[29] R. K. Pankaj and R. G. Gallager, "Wavelength requirements of all-optical networks," *IEEE/ACM Trans. Networking*, vol. 3, pp. 269–280, June 1995.

Yufeng Xin received the M.E. and B.E. degrees from Northwestern Polytechnical University, Xi'an, China, in 1993 and 1990, respectively, and the Master's degree in computer science and the Ph.D. degree in operations research from North Carolina State University, Raleigh, in 2001 and 2002, respectively.

He is currently a Technical Specialist with MCNC, Research Triangle Park, NC.

George N. Rouskas (S'92–M'95–SM'01) received the diploma in electrical engineering from the National Technical University of Athens (NTUA), Athens, Greece, in 1989, and the M.S. and Ph.D. degrees in computer science from the College of Computing, Georgia Institute of Technology (Georgia Tech), Atlanta, in 1991 and 1994, respectively.

He joined the Department of Computer Science, North Carolina State University, Raleigh, in August 1994 as an Assistant Professor, and he has been a Professor since July 2002. During the 2000–2001 academic year, he spent a sabbatical term at Vitesse Semiconductor, Morrisville, NC, and in May and June 2000, he was an Invited Professor at the University of Evry, France. His research interests include network architectures and protocols, optical networks, multicast communication, and performance evaluation.

Dr. Rouskas received a 2003 NCSU Alumni Outstanding Research Award. He is a recipient of the 1997 NSF Faculty Early Career Development (CAREER) Award, and a coauthor of a paper that received the Best Paper Award at the 1998 SPIE conference on All-Optical Networking. He also received the 1995 Outstanding New Teacher Award from the Department of Computer Science, NCSU, and the 1994 Graduate Research Assistant Award from the College of Computing, Georgia Tech. He was a co-Guest Editor for the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS's Special Issue on Protocols and Architectures for Next-Generation Optical WDM Networks, published in October 2000. He is also on the editorial boards of the IEEE/ACM TRANSACTIONS ON NETWORKING, *Computer Networks*, and *Optical Networks*. He is a Member of the ACM and of the Technical Chamber of Greece.

Harry G. Perros (M'87–SM'97) received the B.Sc. degree in mathematics from Athens University, Athens, Greece, in 1970, the M.Sc. degree in operational research with computing from Leeds University, U.K., in 1971, and the Ph.D. degree in operations research from Trinity College, Dublin, Ireland, U.K., in 1975.

He has held visiting faculty positions at INRIA, Rocquencourt, France (1979) and the University of Paris 6, France (1995–1996, 2000, and 2002). He has also spent two sabbatical terms at Nortel, Research Triangle Park, NC (1988–1989 and 1995–1996). He is currently a Professor of Computer Science and a Program Coordinator of the M.S. degree in Computer Networks at North Carolina State University, Raleigh. He has published extensively in the area of performance modeling of computer and communication systems, and he has organized several national and international conferences. He published a monograph entitled *Queueing Networks With Blocking: Exact and Approximate Solutions* (New York: Oxford Press, 1994) and a textbook entitled *An Introduction to ATM Networks* (New York: Wiley, 2001). His current research interests are in the areas of optical networks and satellites.

Dr. Perros is Chairman of the IFIP Working Group 6.3 on the Performance of Communication Systems and the IFIP Working Groups 7.3 and 6.2.