Photonic Network Communications, 7:3, 267–278, 2004

A Study of Path Protection in Large-Scale Optical Networks

Yufeng Xin*
Advanced Networking Research, MCNC-RDI, RTP NC 27709
E-mail: yxin@anr.mcnc.org

George N. Rouskas
Department of Computer Science, North Carolina State University, Raleigh, NC 27695-7534
E-mail: rouskas@csc.ncsu.edu

Received July 3, 2003; Revised September 24, 2003 and Accepted January 9, 2004

Abstract. We consider the problem of designing a network of optical cross-connects (OXCs) which provides end-to-end lightpath services to large numbers of client nodes, under the requirement that the network will survive any single-link failure. Our main objective is to quantify the additional resource requirements of implementing path protection schemes over a network with no survivability properties. To this end, we present heuristic routing and wavelength assignment algorithms for dedicated path protection and two variants of shared path protection, and integrate them into the physical and logical topology design framework we developed in an earlier study. We apply our heuristics to networks with up to 1000 client nodes, with a number of lightpaths that is an order of magnitude greater than the number of clients, and for a wide range of values for system parameters such as the number of wavelengths per fiber, the number of optical transceivers per client node, and the number of ports per OXC. Our results provide insight into the relative resource requirements of dedicated and shared path protection schemes. We also find that, using shared path protection schemes, it is possible to build cost-effective survivable networks that provide rich connectivity among client nodes with only a modest additional amount of resources over a network with no survivability properties.

Keywords: survivability, large-scale optical networks, topology design

1 Introduction

One critical challenge in the design and management of large scale optical networks is the survivability of the network. The survivability of a network is defined as its ability to recover service connections that are disrupted by network faults or failures. The need for highly survivable optical networks is obvious considering (1) the vast transport and switching capacity of networks using wavelength division multiplexing (WDM), and the potential data loss as a result of a failure, and (2) the fact that the probability of network failures is not low. For instance, Hermes, a Pan-European carrier, estimates an average of one cable cut every four days on their network, while 136 fiber cuts were reported by various United States carriers to the Federal Communications Commission in 1997 alone [1].

In a survivable network, the original service path is referred to as the primary path. Whenever a failure along a primary path disrupts service connections, the typical approach is to restore the affected connections by rerouting them over an available backup path. The techniques used for survivability in optical networks can be broadly classified into two categories: preplanned protection and dynamic restoration [2]. Preplanned protection relies on dedicated resources (e.g., wavelengths) which are reserved in advance (i.e., during connection setup), and which remain idle while there is no failure. Dynamic restoration, on the other hand, refers to online techniques, invoked immediately after the occurrence of a failure, whose objective is to discover spare capacity to restore the affected services. Restoration is more efficient than protection in terms of resource utilization, but restoration time is typically longer than the time needed to perform protection switching.

The protection and restoration techniques can be

*Corresponding author.

further classified as path-based or link-based [3]. In path-based schemes, upon the occurrence of a failure, a fault notification is sent to the source (and destination) of each affected connection; each such source-destination pair then takes the necessary actions to switch its connection from the primary to the backup path. In link-based schemes, on the other hand, the nodes at the two ends of a failed link are responsible for rerouting all affected traffic around the failure; the source and destination of an affected connection need not become aware of the failure at all. Usually, path-based schemes require less spare capacity than link-based ones. However, path-based recovery may take longer to restore service since it requires the source and destination of a connection to take action, whereas with link-based schemes, the nodes closest to the failure are the ones involved in the recovery process.

Protection and restoration of optical WDM networks is a research area that has attracted considerable attention in recent years. Some of the issues that have been addressed in the literature include: techniques for protection and restoration in networks of general topology [4–6]; implementation issues [7]; the interaction between WDM layer protection and upper layer protection and restoration [8–10]; the tradeoffs between protection and restoration [11]; static versus dynamic establishment of protection paths [12]; the effects of wavelength converters [13]; survivable network design [14–16]; and diverse routing, i.e., the routing of primary/backup paths [17–20], among others.

While survivability is not only desirable, but required in order to support mission-critical applications and services, there are certain costs associated with designing a survivable network. These costs reflect the additional network resources which have to be provisioned to restore any connections disrupted by a failure, including hardware resources (fibers, wavelengths, switches, transceivers, etc.) as well as software resources (e.g., diverse routing algorithms, restoration and protection switching protocols, etc.). Therefore, quantifying the amount of these additional resources is an important issue for network providers deploying survivable networks. Previous studies have taken some initial steps in this direction, by devising techniques to design the physical and logical network topology that satisfies certain survivability properties [11,14,16]. Unfortunately, the topology design problem in these studies is typically formulated as an integer linear programming (ILP) model, which can only be applied to networks of small size.

In this paper, we consider the problem of designing large-scale optical WDM networks of optical cross-connects (OXCs) which can survive all single-link failures. The scale of the network is characterized by the number of client devices (e.g., edge routers) using its services, and the number of lightpaths between pairs of clients that it can support. We are interested in typical national or international networks, in which the number of clients can be in the order of hundreds or thousands, and the number of lightpaths that need to be established can be an order of magnitude greater than the number of clients. We address both the physical topology design problem (i.e., determine the number of OXCs required and the fiber links to interconnect them) and the routing and wavelength assignment problem for the survivable WDM network. We assume no wavelength conversion in the network. Since the problem is NP-hard, we present a set of heuristic algorithms to obtain a good solution in terms of the number of OXCs and/or number of wavelengths required. For the physical topology design, we extend the genetic algorithm we developed in our previous study [21] to search the space of physical topologies; that study did not consider any protection requirements. We also present routing and wavelength assignment heuristics for three protection schemes: dedicated path protection, and two variants of shared path protection.

While some of the problems we consider have been addressed in the literature, our work differs from previous studies in several important ways. To the best of our knowledge, this is the first time that heuristics for the survivable physical and logical topology design problems have been applied to networks of realistic size. We provide insight into the design of survivable 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 per fiber and up to 24 optical interfaces per client node. We determine the resource requirements of dedicated and shared path protection, and compare them to those of a design that offers no protection. Our major finding is that, using shared path protection, it is possible to build cost-effective survivable networks which offer rich connectivity among the client nodes with only a moderate additional cost over a network that does not provision spare resources.

The paper is organized as follows. In Section 2, we describe the problem we study, the assumptions we make, and the overall solution approach. In Section 3, we present heuristic algorithms for survivable logical topology design based on dedicated path protection, as well as on two variants of shared path protection. We present numerical results in Section 4 and we conclude the paper in Section 5.

2 Problem Definition

We consider a number N of routers (or other electronic devices, e.g., ATM switches) that are to be interconnected over an optical backbone network which consists of OXC nodes. The service provided by the optical network of OXCs is the establishment of lightpaths among pairs of routers. We assume that each router has Δ optical transceivers, 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 a router is due both to optical hardware and cost limitations (reflected in the number of optical transceivers) and the traffic processing capacity of the router. We also assume that all fiber links, including links between OXCs as well as links between a router and an OXC, can support the same set of W wavelengths, and that each OXC has P input/output wavelength ports, including add/drop ports. Finally, we assume that the OXCs have no wavelength conversion capabilities.

In previous work [21], we considered the problem of designing an optical network of OXCs to interconnect the N routers when each router maintains the maximum possible number of connections, i.e., Δ incoming and Δ outgoing lightpaths. Specifically, we presented a set of heuristic algorithms to address both the physical topology design problem (i.e., determine the number M of OXCs needed to support all the lightpaths and the fiber links among the OXCs) and the logical topology design problem (i.e., determine the routing and wavelength assignment for the lightpaths among the routers). Our heuristics were based on a genetic algorithm, which allowed us to solve the physical and logical topology design problem for very large networks, i.e., for up to $N = 1000$ routers and up to 10 000 lightpaths. Briefly, the solution approach we developed consists of three main steps:

1. A genetic algorithm to search the space of 2-connected physical topologies with exactly M OXCs.
2. A routing heuristic and wavelength assignment heuristic to solve the logical topology problem (i.e., find a path and wavelength for each lightpath).
3. A binary search to determine the least number of OXCs needed to satisfy all constraints (i.e., obtain a feasible solution).

More specifically, the algorithm starts with the assumption that the number M of OXCs in the physical topology is given. We use the genetic algorithm of Step 1 above to search the space of 2-connected physical topologies with M nodes for a topology that minimizes the number of wavelengths required to support the given lightpath set. For each candidate topology, we use the routing and wavelength assignment heuristics of Step 2 to determine the number of wavelengths required. Let $W_M^\star$ denote the smallest number of wavelengths required for any candidate topology with exactly M nodes, as determined by the genetic algorithm. Starting with a value of M equal to the number of routers N, we perform a binary search (as in Step 3) to determine the smallest value of M such that $W_M^\star$ is close to, but less than, the actual number of wavelengths W per fiber.

For the details of each step of this solution approach and the corresponding algorithms, the reader is referred to Xin et al. [21]. Our major finding was 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. Our work, however, did not consider any survivability issues other than requiring that the resulting topology of OXCs be 2-connected. In particular, each lightpath was assigned a single network path and wavelength, and no backup paths were provisioned for protection.

In this paper, we extend the above study [21] to address the physical and logical topology design problem under the additional requirement that the optical network of OXCs be able to survive any single-link failure. Specifically, our objective is to provision both a primary and a backup path (and wavelength) for each lightpath, so that any and all lightpaths affected by the failure of a single link will be restorable. In other words, we will design the optical network to offer a lightpath service that is 100% reliable under single-link failures.

The fundamental question we address in this paper is:

> What is the additional cost of building a network that survives all single-link failures over one that offers no protection?

In our study, we use the number of additional OXCs or wavelengths needed to implement the backup paths as the cost metric. We believe that the answer to this question is of importance to service providers who need to deploy survivable optical networks in a cost-effective manner.

In answering the above question, we use the same three-step solution approach we developed in our previous study [21]. More specifically, we use the same genetic algorithm in Step 1 to search for an appropriate 2-connected physical topology, and the same binary search in Step 3 to determine the least number of OXCs for a feasible solution. The main difference is in the algorithms that we use in Step 2 for obtaining the logical topology. Since our goal is to design a survivable logical topology that can fully recover from any single-link failure, we use new heuristics to route and assign wavelengths to the primary and backup lightpaths. We describe these algorithms in the next section.

3 Survivable Logical Topology Design

We assume that we are given a 2-connected physical topology of M OXCs, as well as a set $\mathscr{L}$ of lightpaths between pairs of OXCs that need to be established. The physical topology is provided by the genetic algorithm (Step 1 of the solution approach we described in the previous section), while the set $\mathscr{L}$ of lightpaths between pairs of OXCs can be derived in a straightforward manner from the set of lightpaths between pairs of routers and information about the OXC to which each router is attached. We represent the physical topology of OXCs as an undirected graph $G = (V, E)$, where V is the set of nodes, $|V| = M$, and E is the set of fiber links. Two paths in this topology are said to be link-disjoint if they have no links in common. As in our previous work [21], we assume that, given the number M of nodes, the number of edges in the physical topology is uniquely determined as $|E| = \alpha M$, where $\alpha, 0 < \alpha \leq 1$, is a parameter that characterizes the degree of connectivity of the network. We have used $\alpha = 0.3$ in this work, since this is the average degree of connectivity for existing backbone networks.

We consider two protection schemes in our study:

1. Dedicated path protection (DPP). This is a $1 + 1$ path protection scheme in which there is no wavelength sharing among backup paths.
2. Shared path protection (SPP). This protection scheme allows backup paths to share a wavelength whenever possible (more details on how this is accomplished will be provided shortly).

Our objective is to route and assign wavelengths to the primary and backup paths under the DPP and SPP schemes in a way that minimizes the number of wavelengths used, while ensuring that the resulting logical topology will survive any single-link failure. It is well known that this routing and wavelength assignment problem is NP-Complete. Therefore, 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. The next subsection discusses the routing algorithm we use for finding link-disjoint paths, and the following three subsections discuss the wavelength assignment heuristics we use for the dedicated and shared path protection schemes.

3.1 Routing of Link-Disjoint Paths

Regardless of the protection scheme used (dedicated or shared path), in order for the network to be able to recover from any single-link failure, each lightpath must be assigned both a primary and a backup path, and these paths must not share a common link. Therefore, the lightpath routing problem is equivalent to finding two link-disjoint paths in the physical topology of OXCs for each lightpath in the set $\mathscr{L}$. We also note that the fact that the physical topology is 2-connected guarantees the existence of a pair of link-disjoint paths between any pair of OXCs.

A straightforward method to find a pair of link-disjoint paths from an OXC s to an OXC d is to use the following two-step approach. First, use some routing algorithm to find the shortest path p from OXC s to OXC d in the physical topology represented by the set of edges E. Next, remove all links of p from E to obtain a new set E', and use the same routing algorithm to find the shortest path b from s to d in the physical topology represented by E'. Obviously, paths p and b are link-disjoint, and can be used as the primary and backup paths, respectively. The advantage of this two-step approach is its simplicity, and we adopt it to determine the primary and backup paths of a lightpath. At each step, we use the routing algorithm we developed in Xin et al. [21] to find the shortest path. This routing algorithm defines the cost metrics in a way that accounts for the number of lightpaths using each link, in order to minimize the number of wavelengths used; for the details of the algorithm, the reader is referred to Xin et al. [21].

One drawback of the above two-step method is that sometimes it may fail to find a pair of link-disjoint paths even when such a pair actually exists [17]. Therefore, if this method fails when we apply it to the 2-connected topology, then we run the algorithm developed in Bhandari [17] which is guaranteed to find a pair of link-disjoint paths whenever one exists.

3.2 Dedicated Path Protection (DPP)

As we mentioned earlier, the dedicated path protection scheme is a $1+1$ protection scheme, in which we assume that the service demand is transported along the primary and backup paths simultaneously, and the receiver selects the best signal. As a result, it is not possible for backup paths to share wavelengths. We adopt the following method to perform routing and wavelength assignment in the DPP scheme. First, a pair of link-disjoint primary and backup paths is obtained for each lightpath in set $\mathscr{L}$, as we explained in the previous subsection. The result of the routing step is a set of $2|\mathscr{L}|$ paths, and these paths need to be assigned wavelengths such that if two paths 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 on an induced simple graph [22]. The induced graph has $2|\mathscr{L}|$ vertices, each corresponding to a path on the physical topology. Two vertices of the induced graph are joined by an edge only if the two corresponding paths share the same fiber link. A heuristic algorithm was developed in Zhang and Acampora [22] to solve this vertex coloring problem. The algorithm uses a greedy heuristic to assign wavelengths (colors) to the paths (vertices). As in Xin et al. [21], we adopt this algorithm to perform wavelength assignment, since it has been shown to have good accuracy. As we explain in Xin et al. [21], the complexity of the routing and wavelength assignment in this manner takes time $O(\Delta NM^2 + \Delta^2 N^2)$, where Δ is the number of optical transceivers per router, N is the number of routers, M is the number of OXCs, and ΔN is the size of the given lightpath set $\mathscr{L}$.

We note that this approach does not differentiate between primary and backup paths for the purposes of wavelength assignment. As a result, there is no wavelength sharing among backup paths, and we expect the dedicated path protection scheme to require more wavelengths than the shared path protection schemes we discuss next.

3.3 Shared Path Protection (SPP)

In SPP, the spare (backup) capacity is shared among several lightpaths. In this case, the backup paths are pre-decided and the spare capacity (wavelength) is reserved but not used. When a particular primary path fails, the service is switched to the reserved backup path. To survive from a single-link failure, all the primary paths that share the same spare capacity must be mutually link-disjoint. Obviously, SPP requires less spare capacity than the dedicated protection scheme we discussed in the previous subsection. On the other hand, shared protection requires a sophisticated protection switching scheme to configure the backup path for a particular service path upon a failure, whereas dedicated protection is simpler and faster since the only action required when a failure occurs is for the receiver to switch to the backup path.

We consider two approaches to solve the routing and wavelength assignment problem for primary and backup paths in shared path protection. The main difference in the two approaches is in whether the backup paths are considered jointly with, or separately from, the primary paths when assigning wavelengths.

3.3.1 Shared Path Protection with Joint Wavelength Assignment (SPP-JWA)

In the first approach, we first determine the primary and backup paths for each lightpath in set $\mathscr{L}$, by finding a pair of link-disjoint paths as we explained in Section 3.1. We then convert these paths into an induced graph Q, and we use the vertex coloring algorithm in Zhang and Acampora [22] to color the graph Q and hence assign a wavelength to the paths. Each of the $2|\mathscr{L}|$ vertices of graph Q represents one of the primary or backup paths determined by the routing algorithm, as in the dedicated path protection scheme we discussed in Section 3.2. However, since backup

Input: The physical topology $G(V, E)$ of the network of OXCs, and the set of lightpaths, $\mathscr{L}$
Output: Routing (i.e., a pair of primary/backup paths for each lightpath in $\mathscr{L}$) and wavelength assignment such that backup paths share wavelengths whenever possible, and the logical topology can survive any single-link failure

1. **begin**
2. **while** ($\mathscr{L}$ is not empty) **do**
3. Remove a lightpath l from $\mathscr{L}$
4. Find a pair of link-disjoint paths for l using the algorithm in Section 3.1
5. Let the shorter one be the primary path p, and the longer one be the backup path b
6. **end** // of the while loop
7. Convert all the primary and backup paths into an induced graph Q using the rules listed in Section 3.3.1
8. Color the vertices of graph Q using the algorithm in [22]
9. **end** // of the algorithm

Fig. 1. Algorithm for shared path protection with joint wavelength assignment (SPP-JWA).

paths may now share wavelengths, we use a different set of rules than in the DPP scheme to determine when two vertices of the induced graph are joined by an edge. Let l_1 and l_2 be two lightpaths, $l_1, l_2 \in \mathscr{L}$, and let p_i and b_i denote the primary and backup path, respectively, for lightpath $l_i, i = 1, 2$. Note that p_i and $b_i, i = 1, 2$, are vertices of the induced graph Q. We use the following rules to determine the edges of Q:

- Since by definition p_1 and b_1 are link-disjoint, they may use the same wavelength, hence there is no edge between them in graph Q; similarly for p_2 and b_2.
- If p_1 and p_2 are link-disjoint, then the two paths may use the same wavelength, and the corresponding vertices are not connected in Q. Since, under the single-link failure scenario, paths p_1 and p_2 will not fail simultaneously, their corresponding backup paths b_1 and b_2 may also use the same wavelength, even if they are not link-disjoint. Therefore, there is no edge between vertices b_1 and b_2 in Q.
- If p_1 and p_2 overlap (i.e., they are not link-disjoint), then the two paths must use different wavelengths, and vertices p_1 and p_2 are connected in Q. We now consider two cases depending on the relationship between b_1 and b_2:
  - If b_1 and b_2 also overlap, then these two paths must also use different wavelengths, since a single-link failure may affect both primary paths, in which case b_1 and b_2 will become active simultaneously; hence vertices b_1 and b_2 are connected in Q.
  - If b_1 and b_2 are link-disjoint, then there is no edge between vertices b_1 and b_2 in graph Q.
- If p_1 and b_2 overlap, then the two paths must use different wavelengths, and vertices p_1 and b_2 are connected in Q; otherwise, there is no edge between p_1 and b_2 in Q. Similar observations apply to p_2 and b_1.

The complete routing and wavelength assignment algorithm is shown in Fig. 1. We will refer to this algorithm as shared path protection with joint wavelength assignment (SPP-JWA).

For a particular number N of routers and number M of OXCs, the **while** loop in Steps 2–6 takes $O(\Delta N M^2)$ time since we use Dijkstra's algorithm for the lightpath routing. Step 7 takes $O(\Delta^2 N^2)$ time for the construction of the induced graph Q, and Step 8 takes $O(\Delta^2 N^2)$ for the graph coloring since we have a total of $2\Delta N$ (primary or backup) lightpaths. Therefore, the time taken by the SPP-JWA algorithm is $O(\Delta N M^2 + \Delta^2 N^2)$.

3.3.2 Shared Path Protection with Separate Wavelength Assignment (SPP-SWA)

In this approach, we again use the routing algorithm of Section 3.1 to obtain a pair of link-disjoint primary and backup paths for each lightpath in set $\mathscr{L}$. However, we use a different method for assigning wavelengths to these paths. Specifically, let $\mathscr{L}_p$ and $\mathscr{L}_b$ denote the set of primary and backup paths, respectively, determined by the routing algorithm. We assign wavelengths in two steps, as follows.

In the first step, we assign wavelengths to the primary paths in $\mathscr{L}_p$ by coloring an induced graph Q_p. In Q_p, each vertex represents a primary path, and two vertices in Q_p are connected by an edge if the corresponding paths have at least a physical link in common in the physical topology of OXCs. We use the algorithm in Zhang and Acampora [22] to color the graph Q_p.

In the second step, we assign wavelengths to the backup paths in set $\mathscr{L}_b$ by coloring another induced graph Q_b. Each vertex of Q_b represents a backup path in $\mathscr{L}_b$, but we use the following rules to determine the set of edges in Q_b; as before, we use p_i and b_i represent the primary and backup paths, respectively, of lightpaths $l_i \in \mathscr{L}, i = 1, 2$.

- If p_1 and p_2 are link-disjoint, then b_1 and b_2 may use the same wavelength, hence no edge exists between vertices b_1 and b_2 in Q_b, regardless of whether or not paths b_1 and b_2 have a common link.
- If p_1 and p_2 overlap, we consider two cases:
  - If b_1 and b_2 also overlap, then b_1 and b_2 must use different wavelengths; therefore, vertices b_1 and b_2 are connected in Q_b.
  - If b_1 and b_2 are link-disjoint, then b_1 and b_2 may use the same wavelength; hence, no edge exists between b_1 and b_2 in Q_b.

Note that we cannot use the vertex coloring algorithm in Zhang and Acampora [22] directly to color graph Q_b. This is because, when assigning colors to the backup paths, we must take into account a set of constraints arising from the fact that primary paths have already been assigned wavelengths. In particular, the coloring algorithm must ensure that wavelengths assigned to the backup paths do not conflict with the wavelengths assigned to the primary paths.

The problem of assigning wavelengths to the set of backup paths $\mathscr{L}_b$ can be expressed as a constrained vertex coloring problem. In this problem, each vertex (backup path) b_i of the induced graph Q_b is associated with a color constraint set C_i, which is determined by the wavelength assignment for the set of the primary paths. More specifically, the color constraint set C_i associated with vertex b_i is the set of wavelengths (colors) which are used by some primary path that has one or more links in common with the corresponding backup path b_i. Therefore, a coloring of graph Q_b, in addition to satisfying the usual coloring constraints, must also be such that no vertex b_i is ever assigned a color in its constraint set C_i.

This constrained vertex coloring problem can be formally defined as follows:

Definition 3.1: *Let $Q_b = (V, E)$ be a simple graph. Each vertex $v \in V$ is associated with a set of colors C_v that may not be used in coloring the vertex. What is the minimum number of colors to cover all the vertices in V such that (1) any two vertices u and v are not assigned the same color if there is an edge $(u, v) \in E$, and (2) the color assigned to vertex v, $c_v \notin C_v, \forall v$?*

It is easy to see that the constrained vertex coloring problem is NP-Complete, since by letting the constraint sets $C_v = \phi$ for all vertices v, it reduces to the well-known NP-Complete vertex coloring problem. In order to solve this constrained vertex coloring problem, we propose an efficient greedy algorithm that combines the construction of the induced graph Q_b with wavelength assignment. The main steps of the algorithm are as follows. We begin with an empty induced graph Q_b, the set of backup paths $\mathscr{L}_b$ that need to be assigned wavelengths, and the set of primary paths $\mathscr{L}_p$ which have been colored. At each iteration, we first remove a backup lightpath b_i from the set $\mathscr{L}_b$ and determine the associated color constraint set C_i. C_i is determined by checking each physical link along the route of b_i and adding all the wavelengths used by primary paths on this link into C_i. We then add vertex b_i to the induced graph Q_b, and we also add an edge from vertex b_i to some vertex (if any) b_j already in the Q_b if backup paths b_i and b_j share a physical link and p_i and p_j also have a common link. Let c be the color with the smallest index that (1) does not belong to set C_i, and (2) has not been assigned to any neighbor of b_i in graph Q_b. We assign color c to vertex (backup path) b_i; we note that this selection corresponds to a first-fit assignment. This process continues until all backup paths in $\mathscr{L}_b$ have been considered; when the process terminates, the constrained induced graph Q_b is complete and all vertices (backup paths) have been assigned colors.

The complete routing and wavelength assignment algorithm is shown in Fig. 2. We will refer this algorithm as shared path protection with separate wavelength assignment (SPP-SWA).

It is not difficult to see that, for a given number N of routers and a number M of OXCs, the routing and wavelength assignment of primary lightpaths take $O(\Delta N M^2 + \Delta^2 N^2)$ time, since we use Dijkstra's algorithm for routing and there are ΔN primary

Input: The physical topology $G(V, E)$ of the network of OXCs, and the set of lightpaths, $\mathscr{L}$
Output: Routing (i.e., a pair of primary/backup paths for each lightpath in $\mathscr{L}$) and wavelength assignment such that backup paths share wavelengths whenever possible, and the logical topology can survive any single-link failure

begin
  while ($\mathscr{L}$ is not empty) **do**
    Remove a lightpath l from $\mathscr{L}$
    Find a pair of link-disjoint paths for l using the algorithm in Section 3.1
    Let the shorter one be the primary path p, and the longer one be the backup path b
  end (of the while loop)
  $\mathscr{L}_p \leftarrow$ set of all primary paths
  $\mathscr{L}_b \leftarrow$ set of all backup paths
  Convert all the primary paths in $\mathscr{L}_p$ into the induced graph Q_p
  Color the vertices of graph Q_p using the algorithm in [22]
  Initialize the induced graph Q_b of backup paths to an empty graph
  for all backup paths $b_i \in \mathscr{L}_b$ **do**
    $C_i \leftarrow \phi$ // C_i is the color constraint set,
      // i.e., the set of wavelengths that cannot be used by path b_i
    for all primary paths $p_j \in \mathscr{L}_p$ **do**
      $c \leftarrow$ the color of path p_j
      if (paths b_i and p_j have a common link) **then** $C_i \leftarrow C_i \cup \{c\}$
    end // of the inner for loop
    Add vertex b_i and the appropriate edges to the graph Q_b, as described in Section 3.3.2
    $c \leftarrow$ color of smallest index not in C_i *and* not assigned to a neighbor of b_i in Q_b
    Assign color c to the vertex (backup path) b_i
  end // of the outer for loop
end // of the algorithm

Fig. 2. Algorithm for shared path protection with separate wavelength assignment (SPP-SWA).

lightpaths. The routing and wavelength assignment of backup lightpaths has a similar complexity, hence the overall complexity of SPP-SWA is also $O(\Delta NM^2 + \Delta^2 N^2)$. We note that, while all three algorithms (DPP, SPP-JWA, and SPP-SWA) have the same worst-case complexity, the detailed analysis presented in Xin [23] indicates that the constant coefficient of DPP is the largest of the three, the coefficient of SPP-SWA is the smallest, and the coefficient of SPP-JWA lies between those of the other two algorithms.

4 Numerical Results

In this section, we present results to compare the various protection schemes in terms of the resources they require to ensure that the logical topology of the network of OXCs can survive all single-link failures. The three protection schemes we consider are dedicated path protection (DPP), SPP-JWA and SPP-SWA. For comparison purposes, we also show the results from our previous study [21], in which the logical topology consisted of a single path for each lightpath connection, with no backup path for protection. We expect that the logical topology with no protection will require the least amount of network resources, the DPP scheme will require the most resources, while the resources needed by the SPP schemes will be somewhere in between these two extremes. However, to the best of our knowledge, no study has quantified the additional amount of resources needed for protection in large-scale optical networks that can support hundreds of routers and thousands of lightpaths. The results we present in this section provide significant insight into this important issue.

In our study, we vary the network parameters as follows: the number N of routers interconnected by the optical network of OXCs varies between 100 and 1000; the number W of wavelengths per link takes values from 8 to 128; the number P of ports per OXC varies from 16 to 64; and the number Δ of optical transceivers per router takes values from 4 to 24. The set $\mathscr{L}$ of lightpaths to be established is selected so that each router has exactly Δ incoming and Δ outgoing lightpaths to a random set of other routers. For the graphs shown here, we have assigned the N routers to the OXCs in a round-robin manner. That is, the first router is attached to the first and second OXC, the second router to the third and fourth OXC, and so on (as in Xin et al. [21], we require that each router attach to two OXCS). This assignment is made for convenience only, and is not inherent to our approach; in fact, our algorithms can accommodate any arbitrary assignment of routers to OXCs. We use the algorithm we developed in Xin et al. [21] to obtain the physical and logical topology of OXCs, but we use the routing and wavelength assignment algorithms we described in Section 3 for the various protection schemes instead of the corresponding algorithms of Xin et al. [21]; refer also to Section 2. We should point out that the above ranges of system parameter values are based on realistic assumptions regarding the state of the technology and the size of emerging optical networks; our results, therefore, illustrate that our algorithm can be applied to networks of size between one and two orders of magnitude greater than that of the small networks considered in previous studies on protection. We also emphasize that the algorithm in Xin et al. [21] computes the complete physical network topology (i.e., the number of OXCs needed and the fiber links that interconnect them), as well as the routing and wavelength assignment for the primary/backup path pair for each lightpath in set $\mathscr{L}$. 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. Finally, the experimental setup was identical to that in our previous study, in order to provide a meaningful comparison of our new results to our earlier ‘‘no protection’’ results in Xin et al. [21].

In Fig. 3, we plot the number W of wavelengths required by each protection scheme (i.e., DPP, SPP-JWA, SPP-SWA, and no protection) against the number N of routers that use the services of the optical network. For these results, we have let $\Delta = 4$, $P = 64$, and $M = N/20$, i.e., the number of OXCs in the optical network is (1/20)-th of the number of routers. From the figure, we can clearly see that shared path protection is significantly more efficient in terms of wavelength usage compared to the DPP scheme. We also observe that the gap between the shared and dedicated schemes becomes larger as the number N of routers increases. When N equals 1000, the number of wavelengths required by the DPP scheme is almost twice that required by the shared protection schemes. On the other hand, the shared path protection schemes require only a moderately larger number of wavelengths than the no-protection scheme which does not provision any backup paths. Between the two shared protection schemes we study, the one based on joint wavelength assignment of primary/backup paths (SPP-JWA) outperforms the one based on separate wavelength assignment (SPP-SWA). As we shall show later, however, the SPP-SWA algorithm is faster than SPP-JWA.

Fig. 4 presents a different perspective on the resource requirements of the different schemes. In this figure, we fix the number of wavelengths to $W = 64$, the number of ports per OXC to $P = 64$, and the number of transceivers per router to $\Delta = 4$, and we plot the number of OXCs needed to support all lightpaths against the number of client routers, N. We make two important observations. First, for all four

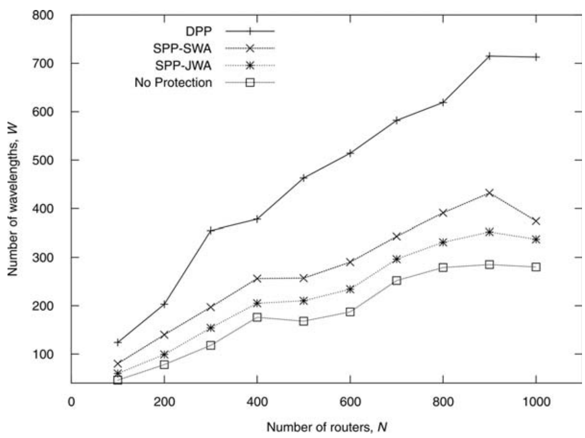

Fig. 3. Wavelength requirements of the various protection schemes ($M = N/20, P = 64, \Delta = 4$).

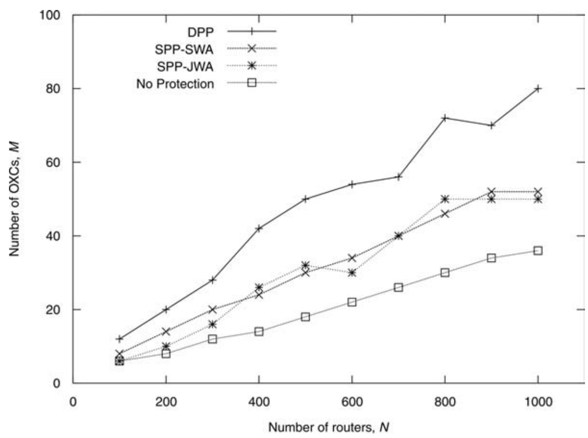

Fig. 4. OXC requirements of the various protection schemes ($\Delta = 4, P = 64, W = 64$).

schemes, the number M of OXCs increases almost linearly with the number N of routers, but the slope is moderate. Second, both shared path protection schemes require a similar number M of OXCs, which is about 1/3 higher than that required by a network offering no protection. Also, DPP is the worst scheme since it requires twice as many OXCs as the no-protection one. The results in Figs 3 and 4 indicate that optical networks to interconnect very large numbers of routers can be built cost-effectively even under the requirement that the network must survive all single-link failures. In particular, the shared path protection schemes are a good choice since their additional resource requirements, in terms of wavelengths or OXCs, are only modestly higher than when no protection is provided.

In Fig. 5, we plot the number of OXCs against the number W of wavelengths when the number of routers is constant at $N = 300$, the number of transceivers per router is $\Delta = 8$, and the number of ports per OXC is $P = 64$. As expected, the number of OXCs needed decreases as the number of wavelengths per fiber increases; also, the DPP scheme requires more OXCs than the SPP-based schemes, which in turn have higher requirements than the no-protection scheme. However, we observe that the difference among the number of OXCs required by the different protection schemes tends to diminish as W increases. This result is important, as it implies that as WDM technology improves and makes more wavelengths available, the incremental cost for adding protection to the network will decrease. We also note that for $W = 8$ wavelengths, the DPP and SPP algorithms have not found a feasible solution with a number of OXCs $M \leq 300$; this implies that the number of wavelengths is too small to support the given lightpath set. While a feasible solution may exist with $M \leq 300$, an optimal algorithm for finding it might take an enormous amount of time to complete. Our heuristic algorithms would be able to find a feasible solution with $M > 300$, but we have explicitly limited the number of OXCs to be at most equal to the number N of routers (in this case, $N = 300$); otherwise, the optical network would require more OXCs than routers, an inefficient and expensive solution.

In Fig. 6, we let $N = 300, P = 64$, and $W = 64$, and we investigate the effect of the number Δ of transceivers on the OXC requirements of the various protection schemes. Note that, as Δ increases, the number of lightpaths in set $\mathscr{L}$ also increases linearly, since the total number of lightpaths to be established is equal to ΔN; for the parameters used in this experiment, the number of lightpaths increases from 1200 when $\Delta = 4$ to 7200 when $\Delta = 24$. As expected, the number of OXCs required increases with Δ. While this increase is significantly slower than the increase in Δ for the no-protection case, it is faster for the protection schemes. However, we note that as Δ (and, consequently, the number of lightpaths) increases by a factor of six from 4 to 24, the number of OXCs required increases slower for both SPP-JWA (from 18 to 90) and for SPP-SWA (from 18 to 86).

In Fig. 7, we let $N = 300, W = 32$, and $\Delta = 8$, and we plot the number M of OXCs against the number P of ports per OXC. When $P = 16$, a feasible solution could not be found even for $M = 300$ for any of the

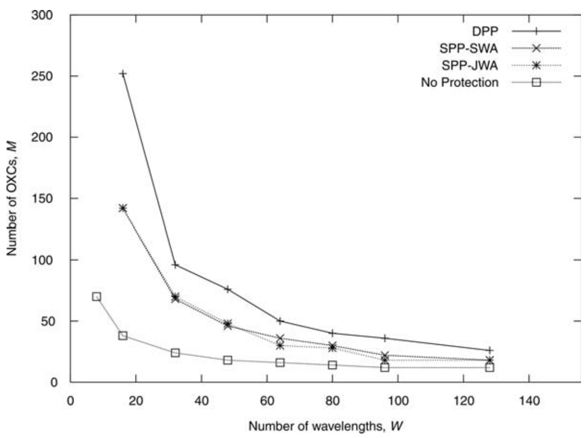

Fig. 5. OXC requirements of the various protection schemes ($N = 300, P = 64, \Delta = 8$).

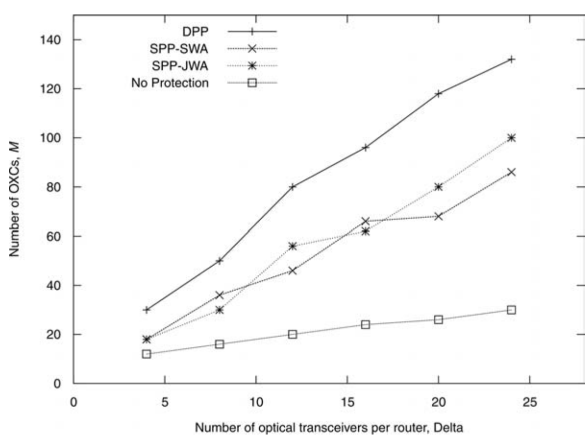

Fig. 6. OXC requirements of the various protection schemes ($N = 300, P = 64, W = 64$).

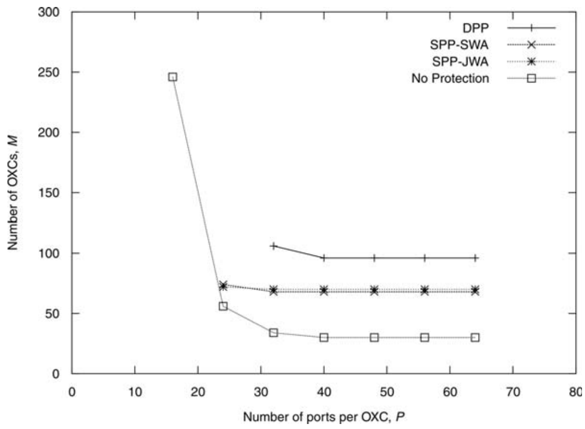

Fig. 7. OXC requirements as a function of the number of ports ($N = 300, P = 64, W = 32, \Delta = 8$).

three protection schemes; put another way, with limited port resources, a feasible solution may require a number of OXCs greater than the number of routers, which is clearly undesirable. But a solution with a reasonable number M of OXCs is possible for $P \geq 24$ (except for DPP, when $P = 24$). We observe that increasing the number of ports beyond 30 does not have a significant effect on the number of OXCs; this is mainly due to the fact that the number W of wavelengths is constant in this experiment. Note that, in general OXCs with many ports are more expensive than those with few ports. Therefore, our results indicate that optical OXC networks can be built cost-effectively with OXCs of medium size even when protection is required.

Finally, in Fig. 8 we plot the running time of our algorithms on a Sun Ultra workstation, as a function of the number N of routers; for these experiments we have let $M = N/20, P = 64$, and $\Delta = 4$. As expected, the no-protection algorithm takes the least amount of time, while the DPP algorithm takes by far the most amount of time. Among the shared path protection algorithms, SPP-SWA is more efficient than SPP-JWA, and its running time is close to that of the no-protection algorithm. In general, we have found that SPP-SWA typically takes one-half to one-third the computation time of SPP-JWA, making it a good choice for provisioning a network to survive single-link failures.

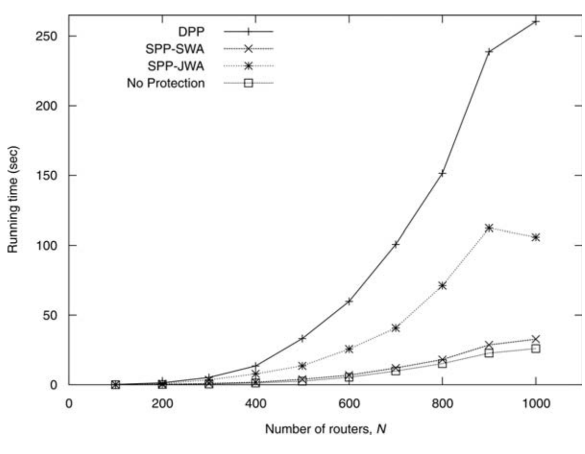

Fig. 8. Running time of the various protection schemes ($M = N/20, P = 64, W = 32, \Delta = 4$).

5 Concluding Remarks

We have described heuristic algorithms for dedicated path protection and shared path protection in optical networks. We use these algorithms within the framework of our previous work to design the physical and logical topology of large-scale WDM networks that can survive all single-link failures. Our results indicate that the additional resources required to guarantee network survivability depend on the protection scheme employed. In particular, shared path protection schemes are more cost-effective than dedicated path protection schemes, and require only a modest additional amount of resources over a network designed without protection in mind.

Acknowledgment

This work was supported by the NSF under grant ANI-9805016.

References

[1] P. F. Fonseca, Pan-European multi-wavelength transport network network design, architecture, survivability and SDH networking, in: DRCN'98, (1998) p. P3.
[2] D. Zhou, S. Subramaniam, Survivability in optical networks, IEEE Network, vol. 14, no. 6, (November/December 2000), pp. 16–23.
[3] G. Mohan, C. Siva Ram Murthy, Ligthpath restoration in WDM optical networks, IEEE Network, vol. 14, no. 6, (November/December 2000), pp. 24–32.
[4] G. Ellinas, R. G. Hailemarian, T. E. Stern, Protection cycles in mesh WDM networks, IEEE Journal on Selected Areas in Communications, vol. 18, (October 2000), pp. 1924–1937.

[5] W. D. Grover, D. Stamatelakis, Cycle-oriented distributed preconfiguration: Ring-like speed with mesh-like capacity for self-planning network reconfiguration, in IEEE ICC'98, (1998) pp. 537–543.
[6] M. Medard, R. A. Barry, S. G. Finn, W. He, S. S. Lumetta, Generalized loop-back recovery in optical mesh networks, IEEE/ACM Transactions on Networking, vol. 10, no. 1, (February 2002) pp. 153–164.
[7] O. Gerstel, R. Ramaswami, Optical layer survivability—an implementation perspective, IEEE Journal on Selected Areas in Communications, vol. 18, no. 10, (October 2000), pp. 1885–1899.
[8] O. Crochat, J. L. Boudec, O. Gerstel, Protection interoperability for WDM optical networks, IEEE/ACM Transactions on Networking, vol. 8, no. 3, (June 2000), pp. 384–395.
[9] A. Fumagalli, L. Valcarenghi, IP restoration vs. WDM protection: Is there an optimal choice, IEEE Network Magazine, (November/December 2000), pp. 34–41.
[10] L. Sahasrabuddhe, S. Ramamurthy, B. Mukherjee, Fault management in IP-over-WDM networks: WDM protection vs. IP restoration, IEEE Journal on Selected Areas in Communications, vol. 20, no. 1, (January 2002), pp. 21–33.
[11] J. Doucette, W. D. Grover, Comparison of mesh protection and restoration schemes and the dependency on graph connectivity, in: DRCN'2001, (2001), pp. 1–9.
[12] V. Anand, C. Qiao, Static versus dynamic establishment of protection paths in WDM networks, Journal of High Speed Networks, vol. 10, no. 4, (2001), pp. 317–327.
[13] M. Lee, J. Yu, Y. Kim, Jinwoo Park, A restoration method independent of failure location in all-optical networks, Computer Communications, vol. 25, no. 10, (June 2002), pp. 915–921.
[14] S. Ramamurthy, L. Sahasrabuddhe, B. Mukherjee, Survivable WDM mesh networks, Journal of Lightwave Technology, vol. 21, no. 4, (April 2003), pp. 870–883.
[15] K. Zhu, L. Sahasrabuddhe, B. Mukherjee, Topology design and upgrade of an optical network by bottleneck-cut identification, Journal of High Speed Networks, vol. 10, no. 4, (2001), pp. 293–301.
[16] B. T. Doshi, S. Dravida, P. Harshavardhana, O. Hauser, Y. Wang, Optical network design and restoration, Bell Labs Technical Journal, (January–March 1999), pp. 58–84.
[17] R. Bhandari, Survivable networks: Algorithms for diverse routing, (Kluwer, USA, 1999).
[18] T. Eilam-Tzoreff, The disjoint shortest paths problem, Discrete Applied Mathematics, vol. 85, (October 1998), pp. 113–138.
[19] F. Guerriero, R. Musmanno, Label correction methods to solve multicriteria shortest path problems, Journal of Optimization Theory and Applications, vol. 111, no. 3, (December 2001), pp. 589–613.
[20] J. Vygen, NP-completeness of some edge-disjoint paths problems, Discrete Applied Mathematics, vol. 61, (1995), pp. 83–90.
[21] Y. Xin, G. N. Rouskas, H. G. Perros, Physical and logical topology design of large-scale optical networks, IEEE/OSA Journal of Lightwave Technology, vol. 21, no. 4, (April 2003), pp. 904–915.
[22] Z. Zhang, A. Acampora, A heuristic wavelength assignment algorithm for multihop WDM networks with wavelength routing and wavelength reuse, IEEE/ACM Transactions on Networking, vol. 3, no. 3, (June 1995), pp. 281–288.
[23] Y. Xin, Topology Design of Large-Scale Optical Networks, Ph.D. thesis, North Carolina State University, Raleigh, NC (August 2002).

George N. Rouskas 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, Atlanta, GA, in 1991 and 1994, respectively. He joined the Department of Computer Science, North Carolina State University, 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 2000 and December 2002 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 the 2003 NCSU Alumni Outstanding Research Award. He is a recipient of a 1997 NSF Faculty Early Career Development (CAREER) Award, and a co-author 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, North Carolina State University, 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*, Special Issue on Protocols and Architectures for Next Generation Optical WDM Networks, published in October, 2000, and is on the editorial boards of the *IEEE/ACM Transactions on Networking, Computer Networks*, and *Optical Networks*. He is a senior member of the IEEE, and a member of the ACM and of the Technical Chamber of Greece.

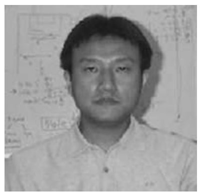

Yufeng Xin received his Ph.D. degree in Operations Research and Computer Science from North Carolina State University, Raleigh, NC. in 2002. He is currently a network research engineer at the Advanced Network Research Group at MCNC-RDI, RTP, NC. His research interests include optical networks, multicasting, and performance evaluation. He is a member of IEEE.