

University of Nebraska - Lincoln

DigitalCommons@University of Nebraska - Lincoln

CSE Journal Articles

Computer Science and Engineering, Department
of

8-2005

Shared Risk Link Group (SRLG)-Diverse Path Provisioning Under Hybrid Service Level Agreements in Wavelength-Routed Optical Mesh Networks

Lu Shen
IEEE

Xi Yang
IEEE

Byrav Ramamurthy
University of Nebraska-Lincoln, bramamurthy2@unl.edu

Follow this and additional works at: <https://digitalcommons.unl.edu/csearticles>



Part of the [Computer Sciences Commons](#)

Shen, Lu; Yang, Xi; and Ramamurthy, Byrav, "Shared Risk Link Group (SRLG)-Diverse Path Provisioning Under Hybrid Service Level Agreements in Wavelength-Routed Optical Mesh Networks" (2005). *CSE Journal Articles*. 78.

<https://digitalcommons.unl.edu/csearticles/78>

This Article is brought to you for free and open access by the Computer Science and Engineering, Department of at DigitalCommons@University of Nebraska - Lincoln. It has been accepted for inclusion in CSE Journal Articles by an authorized administrator of DigitalCommons@University of Nebraska - Lincoln.

Shared Risk Link Group (SRLG)-Diverse Path Provisioning Under Hybrid Service Level Agreements in Wavelength-Routed Optical Mesh Networks

Lu Shen, *Student Member, IEEE*, Xi Yang, *Member, IEEE*, and Byrav Ramamurthy, *Member, IEEE*

Abstract—The static provisioning problem in wavelength-routed optical networks has been studied for many years. However, service providers are still facing the challenges arising from the special requirements for provisioning services at the optical layer. In this paper, we incorporate some realistic constraints into the static provisioning problem, and formulate it under different network resource availability conditions. We consider three classes of shared risk link group (SRLG)-diverse path protection schemes: dedicated, shared, and unprotected. We associate with each connection request a lightpath length constraint and a revenue value. When the network resources are not sufficient to accommodate all the connection requests, the static provisioning problem is formulated as a revenue maximization problem, whose objective is maximizing the total revenue value. When the network has sufficient resources, the problem becomes a capacity minimization problem with the objective of minimizing the number of used wavelength-links. We provide integer linear programming (ILP) formulations for these problems. Because solving these ILP problems is extremely time consuming, we propose a tabu search heuristic to solve these problems within a reasonable amount of time. We also develop a rerouting optimization heuristic, which is based on previous work. Experimental results are presented to compare the solutions obtained by the tabu search heuristic and the rerouting optimization heuristic. For both problems, the tabu search heuristic outperforms the rerouting optimization heuristic.

Index Terms—Capacity minimization problem, dedicated-path protection, revenue maximization problem, shared-path protection, shared risk link group (SRLG), tabu search, wavelength-routed optical network.

I. INTRODUCTION

WDM optical transport networks are evolving from point-to-point connectivity to mesh networking. In such a network, optical-layer protection is of paramount importance due to the risk of losing large amount of data under a failure. However, provisioning wavelength-level services with the protection requirement involves complicated issues. On the one hand, service users are demanding various types

of services under diverse classes of protection. They expect service providers to deliver the right service at the right price. On the other hand, service providers are facing a lot of pressure on the capital investment return. Efficiently utilizing network resources is their top priority. In this situation, challenges exist in provisioning services in an efficient and economic manner. In this paper, we address the problem of static service provisioning with hybrid service requirements in wavelength-routed mesh networks.

Many optical-layer protection schemes for WDM networks have been proposed in the literature [1]–[4]. Generally, there are two categories of protection: path-based protection and link-based protection. In this paper, we consider path-based protection, which many previous studies considered to be easier to implement in the current phase [4], [5], [15]. In path-based protection, a protection path is used to prevent the services on the working path from disruption. The working path and protection path must be diverse to avoid any single point of failure, e.g., a fiber cut.

However, finding diverse protection paths at the optical layer is a special challenge. In a wavelength-routed network, a lightpath traverses optical cross-connects (OXC) that are connected by optical fibers. These fibers are placed into conduits, which are buried along the right of way (ROW). For economic reasons, service providers rent ROWs from third parties, such as the railroad companies. As a result, two diverse fiber links at the OXC layer may be placed into the same conduit at the conduit layer and are subject to a single point of failure. Such links cannot be regarded as diverse links when being used to compute working and protection path pairs. Shared risk link group (SRLG) was proposed to address this problem [6]. An SRLG is a group of links that are subject to a common risk, such as a conduit cut. Therefore, finding a pair of diverse paths at the optical layer involves computing a pair of SRLG-diverse paths. Each link at the OXC layer may be related to several SRLGs. Although the concept of SRLG was originally proposed to deal with conduit cuts, it can be extended to include general risks. For example, all the fiber links located in a geographic area may be assigned the same SRLG considering the risk of earthquakes. An SRLG risk in this paper represents a general risk.

Real-world service users may require different classes of protection. In this paper we consider three kinds of SRLG-diverse path protection schemes: dedicated-path protection, shared-path protection, and no protection. Dedicated-path protection assigns a pair of diverse working and protection paths for a connection at service provisioning time. The wavelengths on both paths

Manuscript received July 25, 2003; revised May 25, 2004; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor C. Qiao. This work was supported in part by the National Science Foundation under Grants ANI-0074121 and EPS-0091900.

L. Shen and B. Ramamurthy are with the Department of Computer Science and Engineering, University of Nebraska-Lincoln, Lincoln, NE 68588-0115 USA (e-mail: lshen@cse.unl.edu; byrav@cse.unl.edu).

X. Yang is with the University of Southern California, Information Science Institute (USC/ISI) East, Arlington, VA 22203 USA (e-mail: xyang@isi.edu).

Digital Object Identifier 10.1109/TNET.2005.852879

are reserved for this connection. If the working path is disrupted by a failure, the destination node will continue the disrupted services on the protection path. Dedicated-path protection provides high reliability and fast restoration speed, but consumes more resources. Shared-path protection is an alternative to dedicated-path protection. In shared-path protection, protection paths of different connections may share the same wavelength on a link as long as their working paths are not subject to a single point of failure. Shared-path protection is more efficient in terms of spare capacity utilization, but may not be as reliable as dedicated-path protection. Some users also request unprotected services because they may protect their services at upper layers, or simply decline protection to save cost.

The class of SRLG-diverse protection is specified in the service level agreement (SLA). An SLA is a contract between a service provider and a service user on the quantity, quality and price of a service. In addition to the class of protection, we further consider two other components in the SLA: the constraint on lightpath length and the price of a service. The lightpath length has significant impact on the quality of service for a connection. In this paper, we consider wavelength-routed all-optical networks, where no optical-electrical-optical (O-E-O) regeneration is available at the intermediate nodes of a lightpath. In such a network, the signal quality degrades with the increase of the lightpath length due to accumulated power losses, noises, and crosstalks [6], [7]. In addition, end-to-end delay is dominated by the propagation time, which is determined by the lightpath length. A service user may specify the maximum lightpath length in the SLA to limit the signal degradation and delay. The price of a service is also an important component of the SLA for both service users and providers. Different services may have different prices.

Before SLAs are signed, the service provider needs to examine the SLA requirements for all the connection requests. When the network resources are not sufficient to accommodate all the connection requests, the service provider is more inclined to provision those connection requests that result in the highest revenue. This is called the revenue maximization problem in [8]. When all the connection requests can be provisioned in the network, a different strategy is adopted. The objective becomes minimizing the network resource consumption, e.g., the number of used wavelength-links, so as to leave more resources to future services. This is called the capacity minimization problem in [8] (or the static lightpath establishment (SLE) problem in [9]). However, it is difficult to check whether a network can accommodate all the connection requests. Service providers can first solve the revenue maximization problem to accept as many connection requests as possible in terms of the total revenue value. Then, they can solve the capacity minimization problem, based on all the accepted connections, to further reduce network resource consumption.

The revenue maximization and capacity minimization problems studied in this paper are NP-hard because they contain several NP-complete sub-problems. SLE without protection has been proven to be NP-complete in [9]. Some recent studies have proved the NP-completeness of SRLG-diverse routing problems [10]–[13]. The routing problem under both wavelength capacity and path length constraints is also NP-complete [14]. In the lit-

erature, several integer linear programming (ILP) formulations for relevant survivable SLE problems have been given [1], [8], [15], [16]. Those ILP formulations that combine routing and wavelength assignment together have been shown to be prohibitively time-consuming [16]. The work in [1] and [8] used alternate working paths as input for the ILP formulation to reduce the complexity. The study in [15] applied Lagrangean relaxation technique to solve the ILP formulations. In [8], the ILP formulations of the capacity and revenue optimization problems were solved by using a multi-stage solution methodology. The authors considered three classes of protection at the OXC layer (dedicated, best-effort, and unprotected, which are different from our classification), but did not consider the SRLG constraints. Under the SRLG constraints, the ILP formulations become more complicated. In [16], the survivable provisioning problem under conduit-layer constraints was studied. The authors divided the combined ILP problem into separate routing and wavelength assignment problems and solved them by using both an ILP solver and heuristics. For solving the conduit-diverse routing sub-problem using such a divide-and-conquer approach, several heuristics have been proposed in [11], [16], [17], and [27].

In this paper, we formulate the revenue maximization and capacity minimization problems by incorporating some realistic constraints faced by service providers. We consider wavelength-level services with three classes of path protection under the SRLG constraints and lightpath length constraints. We provide ILP formulations for our problems. Based on these realistic constraints that were not addressed by previous work, our ILP formulations are an important step toward real-world applications. Heuristic approaches are considered important for dealing with the greatly increasing complexity in real-world, large-sized networks. In this paper, we design a rerouting optimization heuristic based on some optimization principles previously proposed in [15] and [16]. We propose a more sophisticated and a more powerful tabu search heuristic to solve these problems within reasonable times. In the context of hybrid SLA requirements for connection requests, the tabu search heuristic can solve both the revenue maximization and capacity minimization problems in a combined manner and yields better solutions than the rerouting optimization heuristic.

The rest of this paper is organized as follows. Section II presents the network model, problem definition, and ILP formulations for our problems. Section III introduces the rerouting optimization heuristic. Section IV describes the tabu search heuristic. Section V presents our experimental results. Section VI concludes this work.

II. NETWORK MODEL, PROBLEM DEFINITION, AND ILP FORMULATIONS

A. Network Model and Problem Definition

In this paper, we consider WDM wavelength-routed optical mesh networks. Such a network contains a set of OXCs interconnected by optical fiber links. We assume that each link has a single fiber in each direction, while each fiber has a fixed number of wavelengths. We assume that a lightpath is subject to the wavelength continuity constraint, i.e., a lightpath must use

the same wavelength to carry end-to-end user traffic. We also assume that an OXC node can add and drop any wavelength, i.e., an OXC node has sufficient number of transmitters and receivers.

We use a directed graph to represent the OXC layer network as follows:

- Each vertex of the graph represents an OXC node.
- Each link has a fixed number of wavelengths.
- Each link is associated with a number, representing the link length.
- Each link is associated with a set of numbers, representing a set of SRLG risks.

The SRLG risk set for a link consists of all SRLG risks that the link is subject to. Two links are SRLG-diverse if their risk sets contain no common risk. Two paths are said to be SRLG-diverse if no link on one path has a common risk with any link on the other path. A link is SRLG-diverse with a path if the risk set of this link contains no common risk with the risk set of any link on the path.

A connection request, which is a service demand for one wavelength, has the following attributes:

- It starts from a source OXC node and terminates at a destination OXC node.
- It requires one of the three classes of protection: dedicated-path protection, shared-path protection, and no protection.
- It has a constraint on the maximum path length.
- It generates a fixed revenue value, i.e., the price charged by the service provider. A given connection request has a fixed revenue value, while the revenue value may vary among different connection requests.

In this paper we address the following two problems:

1) *The Revenue Maximization Problem:*

Given: An OXC layer network and a set of connection requests with the above attributes.

Objective: Maximize the total revenue value, which is generated by all the successfully provisioned connection requests.

Constraints:

- Each provisioned connection is assigned a working path and a wavelength along the working path.
- Protection requirements of each provisioned connection are satisfied:
 - a) For the connection under dedicated-path protection:
 - Assign a protection path that is SRLG-diverse with its working path.
 - Assign an available wavelength along the protection path.
 - This wavelength cannot be shared.
 - b) For the connection under shared-path protection:
 - Assign a protection path that is SRLG-diverse with its working path.
 - Assign a wavelength along the protection path.
 - This wavelength can be shared by other shared-protection paths if the working path of the current connection

is SRLG-diverse with all other working paths protected by this wavelength.

- c) For the unprotected connection, no protection path is needed.
 - The path length of both working and protection paths is bounded by the maximum path length allowed by this connection.
 - The number of used wavelengths on a link is no more than the total number of wavelengths on the link.

2) *The Capacity Minimization Problem:*

Given: An OXC layer network and a set of connection requests.

Objective: Minimize the number of wavelength-links consumed by all the connection requests (That is the number of wavelengths used on all the links in the network).

Constraints:

- All the connection requests must be successfully provisioned.
- Other constraints remain the same as the revenue maximization problem.

B. ILP Formulations

In this subsection, we provide the ILP formulations for the revenue maximization and capacity minimization problems. The ILP formulations are extended from the work in [1] and [16], where [1] presented the ILP formulations for the SLE problem under different kinds of protections and [16] extended the formulations in [1] to address the duct-layer constraint. The contributions of our ILP formulations are threefold: 1) our formulations take into consideration the generic SRLG constraints (instead of conduit-layer constraints); 2) our formulations can handle the constraints imposed by hybrid SLAs, i.e., three types of protections, different path length constraints and revenue values; 3) the formulations can handle the revenue maximization problem, which has different constraints from previous work. In this paper we mainly present our contributed formulations with a few supporting constraints that are adopted from previous work.

1) *Notations:*

Input Variables:

$G(V, E)$	A directed graph, where V denotes the set of nodes and E denotes the set of links.
(i, j)	A link of graph $G(V, E)$, where $(i, j) \in E$.
$l_{i,j}$	Link length of the link (i, j) .
R	The set of all SRLG risks.
k	A number that represents a risk and $k \in R$.
$(i, j) \in k$	The link (i, j) is subject to the risk k .
W	The number of wavelengths on each link.
w	A wavelength number and $1 \leq w \leq W$.
X	The maximum number of connection requests between any pair of nodes.
C	The set of all connection requests.
(s, d, x)	The identifier of a connection request, where s denotes the source node, d denotes the destination node, and x is a number used to distinguish different connection requests between the same pair of nodes. Note that $1 \leq x \leq X$ and $(s, d, x) \in C$.

$u^{s,d,x}$	The revenue value generated by the connection request (s, d, x) .
$L^{s,d,x}$	The path length constraint for the connection request (s, d, x) .
$dd_{s,d,x}$	1, if there is a connection request (s, d, x) under dedicated-path protection; 0, otherwise.
$sd_{s,d,x}$	1, if there is a connection request (s, d, x) under shared-path protection; 0, otherwise.
$nd_{s,d,x}$	1, if there is an unprotected connection request (s, d, x) ; 0, otherwise.

ILP Variables:

$f_{i,j}$	The number of wavelengths used by working paths or by dedicated-protection paths on the link (i, j) .
$s_{i,j}$	The number of wavelengths used by shared-protection paths on the link (i, j) .
$F_{i,j,w}^{s,d,x,1}$	1, if the working path for the connection request (s, d, x) under dedicated-path protection uses the wavelength w on the link (i, j) ; 0, otherwise.
$F_{i,j,w}^{s,d,x,2}$	1, if the protection path for the connection request (s, d, x) under dedicated-path protection uses the wavelength w on the link (i, j) ; 0, otherwise.
$S_{i,j,w}^{s,d,x}$	1, if the working path for the connection request (s, d, x) under shared-path protection uses the wavelength w on the link (i, j) ; 0, otherwise.
$N_{i,j,w}^{s,d,x}$	1, if the working path for the unprotected connection request (s, d, x) uses the wavelength w on the link (i, j) ; 0, otherwise.
$\alpha_k^{s,d,x,1}$	1, if the working path for the connection request (s, d, x) under dedicated-path protection is subject to the risk k ; 0, otherwise.
$\alpha_k^{s,d,x,2}$	1, if the protection path for the connection request (s, d, x) under dedicated-path protection is subject to the risk k ; 0, otherwise.
$\lambda_{s,d,x,1}^w$	1, if the wavelength w is assigned to the working path of connection request (s, d, x) under dedicated-path protection; 0, otherwise.
$\lambda_{s,d,x,2}^w$	1, if the wavelength w is assigned to the protection path of connection request (s, d, x) under dedicated-path protection; 0, otherwise.
$\beta_{s,d,x}^w$	1, if the wavelength w is assigned to the working path of connection request (s, d, x) under shared-path protection; 0, otherwise.
$\chi_{s,d,x}^w$	1, if the wavelength w is assigned to the working path of unprotected connection request (s, d, x) ; 0, otherwise.
$m_{i,j}^w$	1, if the wavelength w on the link (i, j) is utilized by some shared-protection paths; 0, otherwise.
$\delta_{i,j,w,k}^{s,d,x}$	1, if the wavelength w on the link (i, j) is used by the shared-protection path for the connection request (s, d, x) , whose working path is subject to the risk k ; 0, otherwise.

2) Formulations:**Objective for the revenue maximization problem:**

$$\text{Maximize } \sum_{\forall (s,d,x) \in C} \sum_{\forall i:(s,i) \in E} \sum_{1 \leq w \leq W} \left(F_{s,i,w}^{s,d,x,1} + S_{s,i,w}^{s,d,x} + N_{s,i,w}^{s,d,x} \right) \times u^{s,d,x}.$$

Subject to the following constraints:

First, the number of established connections should not exceed the number of connection requests.

$$\sum_{1 \leq w \leq W} \lambda_{s,d,x,1}^w \leq dd_{s,d,x}, \quad \forall (s, d, x) \in C \quad (1)$$

$$\sum_{1 \leq w \leq W} \lambda_{s,d,x,2}^w = \sum_{1 \leq w \leq W} \lambda_{s,d,x,1}^w, \quad \forall (s, d, x) \in C \quad (2)$$

$$\sum_{1 \leq w \leq W} \beta_{s,d,x}^w \leq sd_{s,d,x}, \quad \forall (s, d, x) \in C \quad (3)$$

$$\sum_{1 \leq w \leq W} \chi_{s,d,x}^w \leq nd_{s,d,x}, \quad \forall (s, d, x) \in C. \quad (4)$$

Second, we use the flow-conservation constraints for end-to-end routing and wavelength assignment. We only present the formulations for the working paths under dedicated-path protection. The flow-conservation constraints on the protection paths under dedicated-path protection ($F_{i,j,w}^{s,d,x,2}$) and the flow-conservation constraints on the working paths under shared-path protection ($S_{i,j,w}^{s,d,x}$) and nonprotection ($N_{i,j,w}^{s,d,x}$) are similar to (5)–(7). We skip those formulations (The complete formulations can be found in [26]). Note that (6) and (7) are used to prevent the loop of a flow that originates from source and returns immediately. Such a loop is not a valid routing path, but may occur when maximizing the revenue value.

$$\begin{aligned} & \sum_{\forall i:(i,j) \in E} F_{i,j,w}^{s,d,x,1} - \sum_{\forall e:(j,e) \in E} F_{j,e,w}^{s,d,x,1} \\ &= \begin{cases} \lambda_{s,d,x,1}^w, & \text{if } j = d \\ -\lambda_{s,d,x,1}^w, & \text{if } j = s \\ 0, & \text{otherwise} \end{cases} \end{aligned} \quad \forall (s, d, x) \in C, j \in V, 1 \leq w \leq W \quad (5)$$

$$\sum_{\forall i:(i,s) \in E} F_{i,s,w}^{s,d,x,1} = 0, \quad \forall (s, d, x) \in C, \quad 1 \leq w \leq W \quad (6)$$

$$\sum_{\forall j:(d,j) \in E} F_{d,j,w}^{s,d,x,1} = 0, \quad \forall (s, d, x) \in C, \quad 1 \leq w \leq W. \quad (7)$$

Third, (8)–(12) and (13)–(23) are the SRLG diversity constraints on the working and protection paths of dedicated protection and shared protection respectively.

Equations (8)–(12) are the SRLG diversity constraints on the working path and the dedicated protection paths, where (8)–(9) and (10)–(11) are used to define the variable $\alpha_k^{s,d,x,1}$ and $\alpha_k^{s,d,x,2}$, respectively. Equation (12) is the SRLG diversity constraint, i.e., a pair of working and dedicated-protection paths are not subject to a common risk.

$$\sum_{\forall (i,j) \in k} \sum_{1 \leq w \leq W} F_{i,j,w}^{s,d,x,1} \leq |E| \cdot \alpha_k^{s,d,x,1}, \quad \forall k \in R, \quad \forall (s, d, x) \in C \quad (8)$$

$$\sum_{\forall (i,j) \in k} \sum_{1 \leq w \leq W} F_{i,j,w}^{s,d,x,1} \geq \alpha_k^{s,d,x,1}, \quad \forall k \in R, \quad \forall (s, d, x) \in C \quad (9)$$

$$\sum_{\forall (i,j) \in k} \sum_{1 \leq w \leq W} F_{i,j,w}^{s,d,x,2} \leq |E| \cdot \alpha_k^{s,d,x,2}, \quad \forall k \in R, \quad \forall (s, d, x) \in C \quad (10)$$

$$\sum_{\forall(i,j) \in k} \sum_{1 \leq w \leq W} F_{i,j,w}^{s,d,x,2} \geq \alpha_k^{s,d,x,2}, \quad \forall k \in R, \quad \forall(s,d,x) \in C \quad (11)$$

$$\alpha_k^{s,d,x,1} + \alpha_k^{s,d,x,2} \leq 1, \quad \forall k \in R, \quad \forall(s,d,x) \in C. \quad (12)$$

Equations (13)–(21) are the SRLG diversity constraints on the working and protection paths of the shared-protection connection request (s, d, x) , where (13)–(20) define the flow-conservation constraints of the variable $\delta_{i,j,w,k}^{s,d,x}$ at the source and destination, and (21) defines the flow conservation constraint of variable $\delta_{i,j,w,k}^{s,d,x}$ at the intermediate nodes.

$$\sum_{1 \leq w \leq W} \sum_{\forall(i,j) \in k} S_{i,j,w}^{s,d,x} \leq |E| \cdot \sum_{1 \leq w \leq W} \sum_{\forall j:(s,j) \in E} \delta_{s,j,w,k}^{s,d,x} \quad \forall k \in R, \quad \forall(s,d,x) \in C \quad (13)$$

$$\sum_{1 \leq w \leq W} \sum_{\forall(i,j) \in k} S_{i,j,w}^{s,d,x} \geq \sum_{1 \leq w \leq W} \sum_{\forall j:(s,j) \in E} \delta_{s,j,w,k}^{s,d,x}, \quad \forall k \in R, \quad \forall s, d, x \in C \quad (14)$$

$$\sum_{1 \leq w \leq W} \sum_{\forall j:(s,j) \in E} \delta_{s,j,w,k}^{s,d,x} \leq 1, \quad \forall k \in R, \quad \forall(s,d,x) \in C \quad (15)$$

$$\sum_{1 \leq w \leq W} \sum_{\forall(i,j) \in k} S_{i,j,w}^{s,d,x} \leq |E| \cdot \sum_{1 \leq w \leq W} \sum_{\forall i:(i,d) \in E} \delta_{i,d,w,k}^{s,d,x}, \quad \forall k \in R, \quad \forall s, d, x \in C \quad (16)$$

$$\sum_{1 \leq w \leq W} \sum_{\forall(i,j) \in k} S_{i,j,w}^{s,d,x} \geq \sum_{1 \leq w \leq W} \sum_{\forall i:(i,d) \in E} \delta_{i,d,w,k}^{s,d,x}, \quad \forall k \in R, \quad \forall(s,d,x) \in C \quad (17)$$

$$\sum_{1 \leq w \leq W} \sum_{\forall i:(i,d) \in E} \delta_{i,d,w,k}^{s,d,x} \leq 1, \quad \forall k \in R, \quad \forall(s,d,x) \in V \quad (18)$$

$$\delta_{i,s,w,k}^{s,d,x} = 0, \quad \forall(s,d,x) \in C, \quad \forall(i,s) \in E, 1 \leq w \leq W, \quad \forall k \in R \quad (19)$$

$$\delta_{d,j,w,k}^{s,d,x} = 0, \quad \forall(s,d,x) \in C, \quad \forall(d,j) \in E, 1 \leq w \leq W, \quad \forall k \in R \quad (20)$$

$$\sum_{\forall i:(i,j) \in E} \delta_{i,j,w,k}^{s,d,x} - \sum_{\forall e:(j,e) \in E} \delta_{j,e,w,k}^{s,d,x} = 0, \quad \forall(s,d,x) \in C, \quad \forall j \in V : j \neq s \& j \neq d, \quad \forall k \in R, \quad 1 \leq w \leq W. \quad (21)$$

Equation (22) defines the constraint that the shared-protection path of the connection request (s, d, x) , whose working path is subject to the risk k , cannot traverse the link (i, j) if the link (i, j) is also subject to the risk k .

$$\delta_{i,j,w,k}^{s,d,x} = 0, \quad \forall k \in R, \quad \forall(i,j) \in E : (i,j) \in k, \quad \forall(s,d,x) \in C, \quad 1 \leq w \leq W. \quad (22)$$

Equation (23) defines the diversity constraint for connection requests under shared-path protection, i.e., two working paths for the connection requests under shared-path protection cannot

be subject to the same risk k , if they are protected by the same wavelength w on the link (i, j) .

$$\sum_{\forall(s,d,x) \in C} \delta_{i,j,w,k}^{s,d,x} \leq 1, \quad \forall k \in R, \quad \forall(i,j) \in E, 1 \leq w \leq W. \quad (23)$$

Fourth, Equations (24)–(26) are wavelength availability constraints.

Equations (24)–(25) define the variable $m_{i,j}^w$, which indicates whether the wavelength w on the link (i, j) is used by some shared protection paths.

$$m_{i,j}^w \leq \sum_{\forall(s,d,x) \in C} \sum_{\forall k \in R} \delta_{i,j,w,k}^{s,d,x}, \quad \forall(i,j) \in E, \quad 1 \leq w \leq W \quad (24)$$

$$|V|^2 \cdot X \cdot |R| \cdot m_{i,j}^w \geq \sum_{\forall(s,d,x) \in C} \sum_{\forall k \in R} \delta_{i,j,w,k}^{s,d,x}, \quad \forall(i,j) \in E, \quad 1 \leq w \leq W. \quad (25)$$

Equation (26) indicates that the wavelength w on the link (i, j) can only be used by either protection paths or a working path.

$$m_{i,j}^w + \sum_{\forall(s,d,x) \in V} \left(S_{i,j,w}^{s,d,x} + N_{i,j,w}^{s,d,x} + F_{i,j,w}^{s,d,x,1} + F_{i,j,w}^{s,d,x,2} \right) \leq 1, \quad \forall(i,j) \in E, \quad 1 \leq w \leq W. \quad (26)$$

Fifth, Equations (27)–(31) are the path length constraints for each connection request.

$$\sum_{1 \leq w \leq W} \sum_{\forall(i,j) \in E} F_{i,j,w}^{s,d,x,1} \cdot l_{i,j} \leq L^{s,d,x}, \quad \forall(s,d,x) \in C \quad (27)$$

$$\sum_{1 \leq w \leq W} \sum_{\forall(i,j) \in E} F_{i,j,w}^{s,d,x,2} \cdot l_{i,j} \leq L^{s,d,x}, \quad \forall(s,d,x) \in C \quad (28)$$

$$\sum_{1 \leq w \leq W} \sum_{\forall(i,j) \in E} S_{i,j,w}^{s,d,x} \cdot l_{i,j} \leq L^{s,d,x}, \quad \forall(s,d,x) \in C \quad (29)$$

$$\sum_{1 \leq w \leq W} \sum_{\forall(i,j) \in E} N_{i,j,w}^{s,d,x} \cdot l_{i,j} \leq L^{s,d,x}, \quad \forall(s,d,x) \in C \quad (30)$$

$$\sum_{1 \leq w \leq W} \sum_{\forall(i,j) \in E} \delta_{i,j,w,k}^{s,d,x} \cdot l_{i,j} \leq L^{s,d,x}, \quad \forall(s,d,x) \in C, \quad \forall k \in R. \quad (31)$$

We have presented the formulation for the revenue maximization problem. For the capacity minimization problem, we need to make the following adjustment: first, we need to switch their objective functions; second, we need to add two more constraints (32) and (33) on the wavelength usage on each link. third, we need to replace all the inequality symbols (\leq) in (1), (3), and (4) with equality symbols ($=$) and remove (6) and (7) which are special for the revenue maximization problem. Other

constraints remain the same.

Objective for the capacity minimization problem:

$$\text{Minimize } \sum_{\forall (i,j) \in E} (f_{i,j} + s_{i,j}).$$

Equations (32)–(33) define the number of wavelengths used by working paths or by dedicated-protection paths on the link (i, j) and the number of wavelengths used by shared-protection paths on the link (i, j) .

$$f_{i,j} = \sum_{\forall (s,d,x) \in C} \sum_{1 \leq w \leq W} \left(S_{i,j,w}^{s,d,x} + N_{i,j,w}^{s,d,x} + F_{i,j,w}^{s,d,x,1} + F_{i,j,w}^{s,d,x,2} \right) \quad \forall (i,j) \in E \quad (32)$$

$$s_{i,j} = \sum_{1 \leq w \leq W} m_{i,j}^w, \quad \forall (i,j) \in E. \quad (33)$$

III. REROUTING OPTIMIZATION HEURISTIC

The ILP formulations presented in Section II-B are very complex. Although solving the ILP problems is still possible for some small-sized networks, it is not practical for the problems in real-world, large-sized networks. Previous work has proposed some optimization heuristic approaches for the SLE problem [15], [16]. The main principle of these heuristics is as follows: first, they find an initial solution using a divide-and-conquer approach and greedy techniques; and then they improve the initial solution by rerouting the existing connections. We develop a heuristic based on the same principle for both revenue maximization and capacity minimization problems. Hereafter, we refer to such heuristics as rerouting optimization heuristics. Because such a rerouting optimization heuristic stops at local optima, we propose a tabu search heuristic to further improve the solution of these problems. In this section, we describe the rerouting optimization heuristic. We present our tabu search heuristic in the next section.

A. Finding k -Shortest Paths

The rerouting optimization heuristic operates on a set of k -shortest paths for each connection request. These k -shortest paths are subject to the following constraints:

- Looplessness, i.e., repeated nodes are not allowed on each path.
- Path length constraint, i.e., the length of each path is no longer than the maximum path length allowed by each connection request.

The k -shortest paths for a connection request are used as candidates for the working path of a connection. We use Yen's k -shortest paths algorithm in our heuristic [19]. We make an adjustment to stop the algorithm when the length of a selected path exceeds the path length constraint required by a connection request. k is a pre-defined number. Note that the k -shortest paths represent as many paths as possible, but up to k . It may consist of fewer than k paths if the constraints are not satisfied on k paths.

The k -shortest paths of each connection request provide a reasonable search space for our heuristics (both rerouting optimization

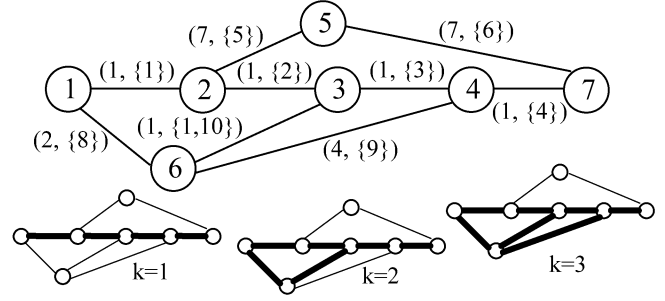


Fig. 1. Illustration of trap paths and how k -shortest paths help avoid trap paths.

tion and tabu search). Another key point of using the k -shortest paths is to avoid trap paths. The concept of trap paths comes from a simple method of finding SRLG-diverse paths. In such a method, a shortest path is computed by using any standard shortest path algorithm. After removing the links on this path as well as any other links that are not SRLG-diverse with this path, a second shortest path is selected. These two paths are a pair of SRLG-diverse paths. However, in some situations, even if a pair of diverse paths does exist, removing the first shortest path and its risk-related links will disconnect the source and the destination so that this method fails. Under such a situation, the first shortest path is called a trap path [20]–[22].

Fig. 1 gives an example of trap paths. Each link in the graph is associated with a length and a set of SRLG risks, which are represented in the form of $(l_{i,j}, R_{i,j})$. There exists a pair of SRLG-diverse paths between node 1 and node 7 in the graph, i.e., path (1-6-4-7) and (1-2-5-7). The path (1-2-3-4-7) is the shortest path between node 1 and node 7. It is also a trap path because removing it will disconnect node 1 and node 7. Another trap path is (1-6-3-4-7). Finding k -shortest candidate paths may avoid trap paths. As k increases, the probability of having all the k -shortest paths to be trap paths reduces. Fig. 1 also gives an example on how the trap paths are avoided by k -shortest paths. The bold lines in the mini-graphs represent the k -shortest paths. When k is one or two, all the k -shortest paths are trap paths. After k increases to three, the path (1-6-4-7) in the k -shortest paths is not a trap path.

After finding the set of k -shortest paths for each connection request, we test whether each path has a corresponding SRLG-diverse path. The test is conducted as follows:

1. Eliminate all the links on the path being tested as well as all the links that are not SRLG-diverse with the path.
2. Check whether a path exists from the source to the destination after eliminating these links.

If no path is found in step 2, the path being tested is removed from the k -shortest paths of this connection because it is a trap path for this connection. As a result, each path in the k -shortest paths has at least one SRLG-diverse path.

B. Provisioning Procedure

The provisioning procedure is a building block of our heuristics (for both rerouting optimization and tabu search). It implements routing and wavelength assignment (RWA) for a connection request. It has three input parameters: a network topology with the current wavelength availability information, a working

path for the connection request, and the SLA information of the connection request. If the provisioning procedure succeeds, it returns the following: a wavelength for the working path, a protection path (if required by the connection request), and a wavelength for the protection path. We use the first-fit method to assign a wavelength for the working path. If the first-fit method fails to find a wavelength for the working path, the provisioning procedure fails. If a protection path is needed by the connection request, it is calculated as follows:

1. Eliminate all the links on the working path and all the links that are not SRLG-diverse with the working path.
2. Find k -shortest paths from the source to the destination on the residual graph (obtained in Step 1), satisfying the path length constraint for this connection request. If no satisfactory path is found, the provisioning procedure fails. Note that the k -shortest paths found in this step serve as the protection path candidates. They are different from the k -shortest paths used as the working path candidates.
3. If the connection request is for dedicated-path protection, use last-fit method to find a wavelength for each of the k -shortest paths (Obtained in step 2) and select a least congested path as the protection path (the paths that failed in the last-fit wavelength assignment are not considered). The least-congested path has the least value of $\sum_{(i,j) \in P} \text{weight}(i,j)$, where P denotes the protection path and $\text{weight}(i,j)$ is the congestion weight for the link

$$(i,j) \cdot \text{weight}(i,j) = \begin{cases} |V|, & \text{if } aw_{i,j} = 1 \\ 1/(aw_{i,j} - 1) & \text{if } aw_{i,j} > 1 \end{cases}$$

where $aw_{i,j}$ denotes the number of available wavelengths on the link (i,j) . This least-congested path selection scheme is similar to the one introduced in [23].

4. If the connection request is for shared-path protection, a wavelength with the least number of unshared links is selected for each path in the k -shortest path set (obtained in step 2). The number of unshared links for a connection on a wavelength is the number of links on which this connection does not share this wavelength with other connections. The path with the least number of unshared links will be selected as the protection path (we break ties using path length).

Note that we also use this provisioning procedure to provision connection requests in the tabu search heuristic.

C. Rerouting Optimization for Revenue Maximization

The rerouting optimization heuristic for the revenue maximization problem establishes a connection based on the k -shortest paths for the connection request. A solution can be found by provisioning one connection request at a time for all connection requests. To establish a connection, each path in the k -shortest paths is tried as its working path according to the nonincreasing order of the path length. The path that succeeds in the provisioning procedure with the least-congested weight (congestion weight is defined in Step 3 of the provisioning procedure) is selected as its working path. And, the corresponding wavelengths and protection path (if any) found in the provisioning procedure are selected for this connection.

```

Set repeat ← 1
while repeat = 1
  repeat ← 0
  for each established connection
    ▪ release the wavelength used by its working and
      protection paths (if any).
    ▪ re-provision the connection based on its k-shortest
      path set using the provisioning procedure. The path
      that causes least wavelength-links is used.
    ▪ if the total number of wavelength-links is reduced,
      set repeat ← 1.
  end for
end while

```

Fig. 2. Optimization heuristic for capacity minimization.

The rerouting optimization heuristic first finds a solution by provisioning connection requests according to the non-increasing order of their revenue values. Then, it runs the same procedure many times with random orders of connection requests for finding a better solution. Finally, the solution with maximum revenue is chosen. This procedure can be stopped if all the connection requests are provisioned (i.e., an optimal solution is found).

This rerouting optimization heuristic for revenue maximization is different from the previous approaches that assumed all the connection requests can be provisioned. We consider the case in which network resources are not sufficient to accommodate all the connection requests. Our rerouting optimization heuristic operates on many restarts with different orders with the hope of finding a better solution. A more sophisticated tabu search heuristic (see Section IV) can improve solution quality dramatically.

D. Rerouting Optimization for Capacity Minimization

The rerouting optimization heuristic for capacity minimization operates on all the existing connections. We use the solution for the revenue optimization problem as an initial solution for the capacity minimization. The optimization heuristic improves the initial solution by iteratively re-provisioning each existing (established) connection. It operates until the number of used wavelength-links cannot be reduced. Fig. 2 presents the procedure of rerouting optimization for capacity minimization. Note that step 3 of the provisioning procedure (in Section III-B) needs to be changed here so that the shortest path instead of the least-congested path is selected as the protection path to minimize the number of used wavelength-links. Similar optimization heuristic approaches were also used in [15] and [16].

IV. TABU SEARCH HEURISTIC

The optimization heuristics introduced in Section III stop at local optima. To further improve the solution quality, we develop a tabu search heuristic to solve both the revenue maximization and capacity minimization problems. Tabu search is a meta-heuristic for solving hard combinatorial optimization problems [18]. It provides a set of principles that guide heuristics exploring possible solution space, leading to high quality solutions. In this section, we first give a general description of tabu search. Then, we present the design of the tabu search heuristic for solving our problems.

Without loss of generality, we describe the general tabu search procedure for solving a minimization problem. Given its optimization objective function $f(x)$, the tabu search looks for an optimal or close-to-optimal solution i^* in the solution space S . S represents the set of all possible solutions for this problem. A search operation from one solution to another is called a move. The tabu search is carried out through a number of such moves. If a move results in a better solution than the current best solution, it is called an improving move. Otherwise, it is a nonimproving move. A general tabu search procedure consists of the following steps:

- Step 1. Choose an initial solution i in S and set $i^* = i$, where i^* denotes the current best solution.
- Step 2. Find the best solution j in $N^*(i)$, such that $f(j) \leq f(j^*)$, $\forall j^* \in N^*(i)$. $N^*(i)$ is a subset of the neighborhood of the current solution i . The neighborhood of the solution i , denoted by $N(i)$, is a set of solutions obtained by going one move from i .
- Step 3. Set $i = j$. If $f(j) < f(i^*)$, set $i^* = j$.
- Step 4. If a stopping condition is met, then stop. Otherwise, go to Step 2.

The stopping condition may be defined in the following ways:

a) a fixed number of iterations of nonimproving moves have been executed since the last improving move; b) a predefined number of iterations have been executed; c) $N^*(i)$ is empty. In the implementation, a tabu list is used to record the moves that have recently been executed. Each move in the tabu list is assigned a tabu tenure. Until the tabu tenure of a move expires, the move is invalid to be executed. By adding a recent move into the tabu list and giving it a tabu tenure, the tabu search procedure may be prevented from entering into loops. However, the rule of tabu tenure is not always strict. If a move in the tabu list can generate a high quality solution (e.g., an improving move), it may become valid to be executed. This is called the aspiration condition.

Another key technique for tabu search is diversification. Diversification refers to the methods that lead the tabu search into new regions of solution space. Two popular methods for the diversification are the move-frequency-penalty for objective function and multi-start. The move-frequency-penalty for objective function is used to direct the search into a new region based on the frequency of previous moves [24], while multi-start restarts multiple rounds of search procedure by creating new initial solutions.

In this section, we first give the definitions for describing our tabu search heuristics. We then present the tabu search heuristic for the revenue maximization problem, followed by a description on the adjustment needed for the capacity minimization problem. We close this section by analyzing the time complexity of our tabu search heuristics.

A. Definitions for the Heuristic Design

The tabu search heuristic operates on a set of connection requests $\{n | 1 \leq n \leq M\}$, where M is the maximum number of connection requests. Other definitions are:

- u_n is the revenue value of the n th connection request.

- $P_n = \{p_n^i | 0 \leq i \leq K\}$ is a set of k -shortest paths from the source to the destination of the n th connection request and an empty path p_n^0 .
- $\text{move}(p_n^g, p_n^h)$ denotes a move that changes the current working path from p_n^g to p_n^h for the n th connection request ($g \neq h$).
- WL is the total number of used wavelength-links.
- WL_n^g is the number of wavelength-links consumed by the n th connection request when using the path p_n^g as its working path.

B. Tabu Search Iteration Procedure

The tabu search iteration procedure operates on an initial solution. The initial solution is computed by provisioning one connection request at a time for all connection requests, according to the nonincreasing order of their revenue values. It is the same as the solution found by rerouting optimization heuristic without random ordering (see Section III-C).

Our tabu search heuristic transforms one solution into another by moving the current working path p_n^g to another path p_n^h among the k -shortest paths of the n th connection request. First, the wavelengths used by the current working path p_n^g and its protection path (if any) are released from the network. Then, the provisioning procedure (described in Section III-B) is called using p_n^h as the working path. We define $\text{move}(p_n^g, p_n^h)$, ($g \neq h$), to be a move that changes the current working path from p_n^g to p_n^h for the n th connection request. If the provisioning procedure fails for a move, the move is invalid to be executed. $\text{move}(p_n^g, p_n^0)$ (from a nonempty working path to an empty path) means to release all the wavelengths that are used by the n th connection. The best move on the current solution is the move that generates the greatest move value among all the connection requests (i.e., $\forall n, 1 \leq n \leq M$). The move value of a connection request is defined as

$$\text{move-value}(p_n^g, p_n^h) = \begin{cases} u_n, & \text{if } g = 0 \text{ and } h \neq 0 \\ \left(WL_n^g - WL_n^h \right) / WL - \alpha \cdot \text{freq}(p_n^h), & \text{if } g \neq 0 \text{ and } h \neq 0 \\ -u_n - \alpha \cdot \text{freq}(p_n^h), & \text{if } g \neq 0 \text{ and } h = 0 \end{cases}$$

where $\text{freq}(p_n^h)$ denotes the frequency (the number of times) that the path p_n^h has been used as the working path of the n th connection and α is the move-frequency-penalty coefficient which is a predefined positive constant.

The three expressions in the above move value function can be explained as follows: if $g = 0$ and $h \neq 0$, the move value is u_n , which represents the revenue gain when a new connection is established; if $g \neq 0$ and $h \neq 0$, the working path of the n th connection is switched between two nonempty paths, i.e., the total revenue value remains the same. $(WL_n^g - WL_n^h) / WL$ represents the number of reduced wavelength-links normalized to the total number of wavelength links used in the network. This factor is used to favor the move that saves more wavelength-links when the total revenue value does not change; if $g \neq 0$ and $h = 0$, the working path of the n th connection request is changed from a nonempty path to an empty path,

which means a decrease in the total revenue value. $-u_n$ represents the decreased revenue value. This factor is used to favor the move with the least revenue loss. We maintain the revenue value of each connection to be greater than 1 in the input so that $-u_n < (WL_n^g - WL_n^h)/WL < u_n$. Thus, moves that tend to increase more revenue (or decrease less revenue) are granted higher priority of being selected. $\alpha \cdot \text{freq}(p_n^h)$ is the move-frequency-penalty for the move value function and is used to facilitate the diversification. By adding a frequency-based penalty for the moves that do not cause any new connection to be provisioned, the search tries to use the move that has been used less frequently and jump into a new search region when no move can improve the current solution.

We compute the move value of every possible move for all the connection requests based on the current solution. Then, the best move with the largest move value, say $\text{move}(p_n^g, p_n^h)$, is used to change the current solution. After the move, p_n^h is assigned a tabu tenure t and is inserted into the tabu list. This prevents p_n^h from moving back before its tabu tenure expires, unless it is an improving move (i.e., the move improves the current best solution).

The tabu search iteration procedure stops if one of the following conditions is satisfied:

- There is no valid move for the current solution.
- The number of iterations without an improving move has exceeded a predefined value since the last improving move is found.
- An optimal solution is found, e.g., the revenue value of the best solution equals the sum of the revenue values of all the connection requests.
- The heuristic has been running for a fixed amount of time.

C. Multi-Start

In addition to the move-frequency-penalty, multi-start is used to further facilitate the diversification by searching new regions of the solution space. Multi-start restarts the tabu search procedure with new initial solutions after the stopping condition of tabu search is met.

In each restart, we re-generate k -shortest paths for each connection request based on their trap path sets. The trap path set for a connection request is established when the k -shortest paths are computed. If a path in the k -shortest paths is a trap path (detected by the method introduced in Section III-A), it is inserted into the trap path set. The paths in the trap path set may contain some critical links (e.g., the link (6, 3) in Fig. 1) that may be a root cause of trap paths. By eliminating such a critical link in the graph, we may avoid obtaining some trap paths when computing the k -shortest paths.

The restart procedure operates as follows: first, a link is removed from the graph; then, the k -shortest paths for this connection request are re-generated based on the residual graph. (Note that the removed link needs to be restored in the graph when computing k -shortest paths for the next connection request). The link to be removed is selected according to the following rules: if the trap path set of the connection request is not empty, among all links on the trap paths, a link with the maximum cardinality of its risk set is selected; if the trap path set of

the connection request is empty, a link on the path that causes the provisioning procedure to fail for the most number of times is selected using the same rule (i.e., maximum cardinality of the risk set). After re-assigning the k -shortest paths for every connection request, the method for finding a new initial solution remains the same. Then, the tabu search iteration procedure is restarted based on the new initial solution. Our rule of removing a link to avoid trap paths is similar to the one introduced in [13].

D. Solving the Capacity Minimization Problem

The above description of the tabu search heuristic is for the revenue maximization problem. Some changes are needed for solving the capacity minimization problem. When minimizing the capacity, i.e., the total number of used wavelength-links, all the connection requests must be provisioned in the initial solution. Therefore, no empty path should exist in the set of k -shortest paths for each connection request. We use the solution of the revenue maximization problem as an initial solution for the capacity minimization problem. We need to change the move value function as follows:

$$\text{move-value}(p_n^g, p_n^h) = (WL_n^g - WL_n^h) - \alpha \cdot \text{freq}(p_n^h)$$

where

$$\alpha = \begin{cases} c & \text{if } (WL_n^g - WL_n^h) \leq 0 \\ 0 & \text{if } (WL_n^g - WL_n^h) > 0 \end{cases}$$

c is a positive constant, and $(WL_n^g - WL_n^h)$ is the difference in the wavelength-links consumed by the n th connection when changing its working path from the path P_n^g to P_n^h . The move that produces the greatest move value is executed in each tabu search iteration procedure. Due to the change in the objective function, we also need to change step 3 of the provisioning procedure (see Section III-B). In step 3 of the provisioning procedure, after using last-fit to find a wavelength for each protection path candidate, the shortest path instead of the least-congested path should be selected as the protection path to minimize the number of used wavelength-links.

Jointly using the heuristics for revenue maximization and capacity minimization can bring an extra benefit. The solutions obtained by revenue maximization may not be capacity-efficient because its objective only takes into consideration the revenue value. To further reduce the number of used wavelength-links, we can run the heuristic for capacity minimization, confining its operation on those connections that have been successfully provisioned by the revenue maximization heuristic. These extra wavelength links (saved by the capacity minimization) may be used to accommodate more future connections.

E. Time Complexity of the Tabu Search Heuristic

Yen's k -shortest path algorithm is based on the shortest path algorithm. If Dijkstra's shortest path algorithm is used, Yen's k -shortest path algorithm has a time complexity of $O(k \cdot |V|^3)$ [19], where k is the number of k -shortest paths, and $|V|$ is the number of network nodes. By M we denote the total number of connection requests. Thus, the procedure of finding k -shortest paths for all connection requests has a time complexity of $O(M \cdot$

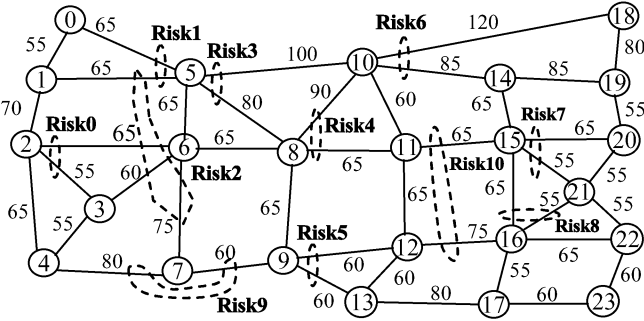


Fig. 3. US-NET (24-node and 86 Links).

$k \cdot |V|^3$). The running time for finding an initial solution is $k \cdot M$ times as the running time of the provisioning procedure, whose four steps (described in Section III-B) have the time complexity $O(|R|^2 \cdot |E|)$, $O(k \cdot |V|^3)$, $O(W \cdot k \cdot |V|)$, and $O(k \cdot W \cdot |R|^2 \cdot M \cdot |V|)$ respectively, where $|E|$ is the total number of links in the network, W is the number of wavelengths on each link, and $|R|$ is the size of the SRLG risk set. Note that only one of steps 3 and 4 in the provisioning procedure can be executed for each connection request. Therefore, the time for finding the initial solution is $O(k \cdot M \cdot (|R|^2 \cdot |E| + k \cdot |V|^3 + k \cdot W \cdot |R|^2 \cdot M \cdot |V|))$ in the worst case. It is equal to $O(k^2 \cdot M^2 \cdot |R|^2 \cdot W \cdot |V|)$, if $|R| \geq |V|$ and $|E| \leq |V|^2 \leq M \cdot |V|$. One iteration of the tabu search heuristic has the same complexity as finding the initial solution. If the number of iterations to be executed is I , the time complexity of our tabu search heuristic is $O(k^2 \cdot M^2 \cdot |R|^2 \cdot W \cdot |V| \cdot I)$.

V. EXPERIMENTAL RESULTS

In this section, we present numerical results for our solution approaches through experiments. We use the CPLEX software to solve the ILP formulations on a SUN Ultra-60 workstation with a 450 MHz UltraSparc II processor. The ILP problem for the 10-node, 1-wavelength network under 6 connection requests has 89 215 variables and 103 064 constraints. Solving the ILP problem becomes extremely time-consuming when we increase the size of the problem. We tried to solve an ILP problem for a 10-node, 4-wavelength network with 45 connection requests, but failed to obtain a feasible integer solution after running CPLEX for 40 hours (a feasible integer solution does exist for this case and can be found by the tabu search heuristic). Previous work also indicated that the ILP approach is not capable of handling large-scale problems in practice. In our experiments, we focus on the study of the heuristic approaches for realistic network settings.

We conduct our experiments on two example networks: US-NET which has 24 nodes and 86 unidirectional links (shown in Fig. 3) and the Italian network which has 21 nodes and 72 unidirectional links (shown in Fig. 4). Each link in these networks represents two unidirectional fiber links in opposite directions. A link length is labeled on each link in kilometers. We assume that the two unidirectional links between two adjacent nodes are subject to a common SRLG risk. In addition, a dashed circle in the figures represents an SRLG risk, and all the fiber links covered by a dashed circle are subject to that risk. Note that these SRLG risks are set to

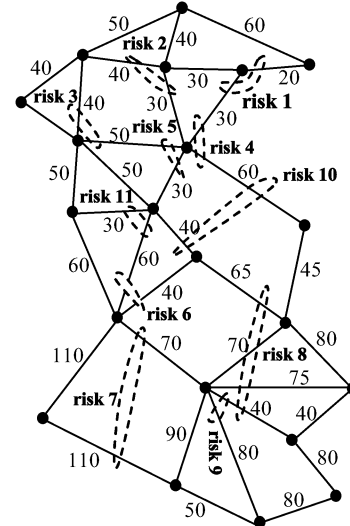


Fig. 4. Italian network (21 node and 72 links).

reflect not only conduit risks (e.g., risk 1, risk 6, and risk 7 in the US-NET), they also contain some logical risks (e.g., risk 2 and risk 10 in the US-NET). In the experiments, connection requests in each class of protection schemes are uniformly distributed among all node pairs. The number of connection requests (demands) under each protection class is manually specified. The revenue value of connection requests under each protection class is randomly generated in a range: from 1.5 to 3.5 for an unprotected connection request, from 4 to 8 for a connection request under shared-path protection, and from 6 to 10 for a connection request under dedicated-path protection. The path length constraint for each connection request is set to 600 km, which is considered as a typical reach distance of all-optical signals using today's amplification and dispersion compensation technologies [25]. This length is also acceptable for most types of services in terms of propagation delay.

A. Comparison of Tabu Search and Rerouting Optimization

In this subsection, we compare the tabu search heuristic and the rerouting optimization heuristic in the US-NET and Italian network with 4, 8, 16 wavelengths and up to 180 connection requests. These network settings are close to the real-world networks. To run the tabu search heuristic, the move-frequency-penalty coefficient α is set to 1 and the number of k -shortest paths (k) is set to 15 for the cases with 4 and 8 wavelengths, and 10 for the cases with 16 wavelengths. We set the tabu tenure to 5 for the cases with the number of connection requests less than 100, and to 10 for those cases with more than 100 connection requests. We set a 7200-second or 2-hour time limit for the running time of both heuristics. This time is long enough for obtaining good solutions and short enough for practical use since the static provisioning problems is normally solved in an off-line manner. The stopping condition of the tabu search heuristic is set as follows: (1) the number of consecutive nonimproving moves has exceeded k times of the total number of connection requests (M), or (2) the neighborhood of the current solution is empty.

Tables I and II show the solution results of tabu search and rerouting optimization for the revenue maximization problem

TABLE I
RESULTS OF TABU SEARCH AND REROUTING OPTIMIZATION HEURISTICS FOR SOLVING THE REVENUE MAXIMIZATION PROBLEM IN THE US-NET

Case #	W	Demands total (n-s-d)	Revenue Value		Improvement	Running Time (seconds)	
			Tabu	Rerouting		Tabu	Rerouting
1	4	30 (10-10-15)	168	161	4.4%	218	1228
2	4	40 (10-20-10)	195.5	188.5	3.7%	1537	1744
3	4	50 (15-20-15)	227	217.5	4.4%	5338	5481
4	8	70 (20-30-20)	391.5	368.5	6.2%	3644	484
5	8	80 (25-30-25)	420.5	389.5	8.0%	3702	534
6	8	90 (30-30-30)	485.5	456.5	6.4%	1295	605
7	16	140 (40-60-40)	751	722	4.0%	5437	556
8	16	160 (50-60-50)	921.5	880.5	4.7%	2063	1676
9	16	180 (50-80-50)	1009	963	4.8%	7197	421

TABLE II
RESULTS OF TABU SEARCH AND REROUTING OPTIMIZATION HEURISTICS FOR SOLVING THE REVENUE MAXIMIZATION PROBLEM IN THE ITALIAN NETWORK

Case #	W	Demands total (n-s-d)	Revenue Value		Improvement	Running Time (seconds)	
			Tabu	Rerouting		Tabu	Rerouting
1	4	30 (10-10-15)	161	152	5.9%	92	4853
2	4	40 (10-20-10)	184.5	175.5	5.1%	5536	4405
3	4	50 (15-20-15)	217	205	5.9%	3502	5290
4	8	70 (20-30-20)	377.5	358	5.4%	5216	6653
5	8	80 (25-30-25)	386	366.5	5.3%	1322	3231
6	8	90 (30-30-30)	434.5	411.5	5.6%	6625	2943
7	16	140 (40-60-40)	711	678	4.9%	2948	6970
8	16	160 (50-60-50)	826.5	792	4.4%	3355	142
9	16	180 (50-80-50)	875	837	4.5%	2686	159

in the US-NET and Italian network under 9 representative network settings respectively. The first column of the table shows the case number. The second column (W) gives the number of wavelengths on each link in the network. The third column gives the number of connection requests (demands) used in each case in the form of: **total** number of connection requests (**no** protection, **shared-path** protection, **dedicated-path** protection). The following two columns give solutions found by the tabu search and rerouting optimization. The improvement column shows the improvement of tabu search over rerouting optimization. The running time is the time when the final (best) solution for each case was found. Tables I and II show that there is a constant increase in revenue value obtained by tabu search in each case. The improvement of tabu search over rerouting optimization varies from 4% to 8%. The results also show that tabu search needs a little bit more time to obtain better solutions for most cases.

Tables III and IV show the results obtained from solving the capacity minimization problems for 12 cases in the US-NET and Italian network respectively. The case number marked with an asterisk denotes that all the connection requests are established and the capacity minimization is based on all of these connections. Other cases are based on the final solutions of revenue maximization found by the tabu search heuristic. Compared to rerouting optimization, tabu search saves wavelength-links in each case, ranging from 3.1% to 9.2%. Because the rerouting optimization stops at local optima, it has a shorter running time (less than 3 minutes for all the cases). The tabu search obtains better solutions at the expense of more time for exploring the solution spaces. However, this time is acceptable and worthwhile.

TABLE III
RESULTS OF TABU SEARCH AND REROUTING OPTIMIZATION HEURISTICS FOR SOLVING THE CAPACITY MINIMIZATION PROBLEM IN THE US-NET

Case #	W	Demands total (n-s-d)	Wavelength-Links		Improvement	Running Time (seconds)	
			Tabu	Rerouting		Tabu	Rerouting
1	4	30 (10-10-15)	194	206	5.8%	114	15
2	4	40 (10-20-10)	205	226	9.3%	3951	43
3	4	50 (15-20-15)	217	230	5.7%	166	42
4	8	70 (20-30-20)	405	427	5.2%	6778	87
5	8	80 (25-30-25)	439	439	4.4%	540	94
6	8	90 (30-30-30)	471	471	5.4%	3444	104
7	16	140 (40-60-40)	829	866	4.3%	307	136
8	16	160 (50-60-50)	845	882	4.2%	6674	90
9	16	180 (50-80-50)	935	986	5.2%	6745	139
10*	16	50 (15-20-15)	282	298	5.4%	3062	18
11*	16	60 (20-20-20)	335	358	6.4%	1673	21
12*	16	70 (20-30-20)	389	412	5.6%	4816	34

TABLE IV
RESULTS OF TABU SEARCH AND REROUTING OPTIMIZATION HEURISTICS FOR SOLVING THE CAPACITY MINIMIZATION PROBLEM IN THE ITALIAN NETWORK

Case #	W	Demands total (n-s-d)	Wavelength-Links		Improvement	Running Time (seconds)	
			Tabu	Rerouting		Tabu	Rerouting
1	4	30 (10-10-15)	175	182	3.8%	1100	9
2	4	40 (10-20-10)	196	208	5.8%	1707	11
3	4	50 (15-20-15)	214	222	3.6%	2013	15
4	8	70 (20-30-20)	354	369	4.1%	3326	30
5	8	80 (25-30-25)	369	395	6.6%	1972	26
6	8	90 (30-30-30)	407	430	5.3%	1935	27
7	16	140 (40-60-40)	728	728	3.1%	5464	36
8	16	160 (50-60-50)	761	761	3.2%	3853	51
9	16	180 (50-80-50)	817	817	4.2%	1563	51
10*	16	50 (15-20-15)	282	293	3.8%	3062	18
11*	16	60 (20-20-20)	335	348	3.7%	1673	21
12*	16	70 (20-30-20)	389	402	3.2%	4816	33

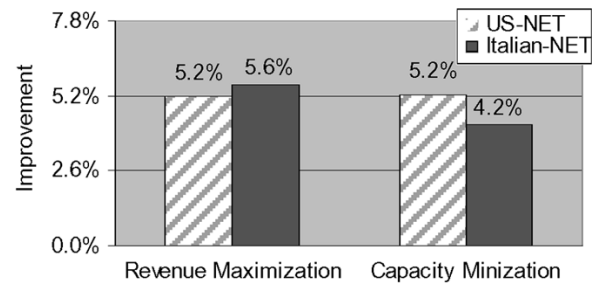


Fig. 5. Improvement of tabu search over rerouting optimization.

For both revenue maximization and capacity minimization, the tabu search heuristic obtains better solutions than rerouting optimization in all of these cases. Fig. 5 shows the average improvement of tabu search over rerouting optimization for revenue maximization and capacity minimization in these two sample networks respectively. On average, tabu search generates 5.2% and 5.6% more revenue than rerouting optimization in the US-NET and Italian networks respectively. It saves an average of 5.2% and 4.2% more wavelength-links in these two networks, respectively.

Using the tabu search heuristic on revenue maximization and capacity minimization as a combined tool can bring an extra benefit by further reducing the number of used wavelength-links after maximizing the revenue value. Tables V and VI show

TABLE V
WAVELENGTH-LINKS BEFORE AND AFTER SOLVING THE CAPACITY MINIMIZATION PROBLEM BASED ON THE SOLUTIONS OBTAINED BY THE TABU SEARCH HEURISTICS FOR REVENUE MAXIMIZATION IN THE US-NET

Case #	W	Demands total (n-s-d)	Wavelength-Links		Improvement
			before capacity minimization	after capacity minimization	
1	4	30 (10-10-15)	213	194	8.9%
2	4	40 (10-20-10)	231	205	11.3%
3	4	50 (15-20-15)	236	217	8.1%
4	8	70 (20-30-20)	472	405	14.2%
5	8	80 (25-30-25)	467	439	6.0%
6	8	90 (30-30-30)	522	471	9.8%
7	16	140 (40-60-40)	930	829	10.9%
8	16	160 (50-60-50)	980	845	13.8%
9	16	180 (50-80-50)	1064	935	12.1%

TABLE VI
WAVELENGTH-LINKS BEFORE AND AFTER SOLVING THE CAPACITY MINIMIZATION PROBLEM BASED ON THE SOLUTIONS OBTAINED BY THE TABU SEARCH HEURISTICS FOR REVENUE MAXIMIZATION IN THE ITALIAN NETWORK

Case #	W	Demands total (n-s-d)	Wavelength-Links		Improvement
			before capacity minimization	after capacity minimization	
1	4	30 (10-10-15)	189	175	7.4%
2	4	40 (10-20-10)	211	196	7.1%
3	4	50 (15-20-15)	222	214	3.6%
4	8	70 (20-30-20)	398	354	11.1%
5	8	80 (25-30-25)	418	369	11.7%
6	8	90 (30-30-30)	445	407	8.5%
7	16	140 (40-60-40)	840	728	13.3%
8	16	160 (50-60-50)	868	761	12.3%
9	16	180 (50-80-50)	903	817	9.5%

the number of used wavelength-links for the solutions of revenue maximization before and after capacity minimization. The results show that a significant number of wavelength-links is saved by solving the capacity minimization problem after revenue maximization in all these cases.

B. Testing the Tabu Search Heuristic

The better performance of the tabu search heuristic over the rerouting optimization heuristic comes from its capability of exploring promising solution space. We used some sophisticated techniques in the design of our tabu search, e.g., move-frequency-penalty and multi-start. This subsection provides experimental data to examine and validate our design of the tabu search heuristic.

Fig. 6 plots the tabu search solution trace of total revenue value versus running time during tabu search procedure of its second multi-start for the experiment case 6 in Table I. After the solution reaches its local optima, a move during the search may temporarily cause a worse solution. However, the long-term trend is to jump out of the local optima and then move to better solutions. Fig. 7 depicts the best solution found during the tabu search procedure of each multi-start for the same experiment case. We have two observations: 1) not every multi-start generates better solution than the previous ones; 2) multi-start may lead to a better solution, e.g., the second restart. The solution traces for other cases have similar patterns as those in Figs. 6 and 7.

Tables VII and VIII show the results for revenue maximization and capacity minimization with and without move-fre-

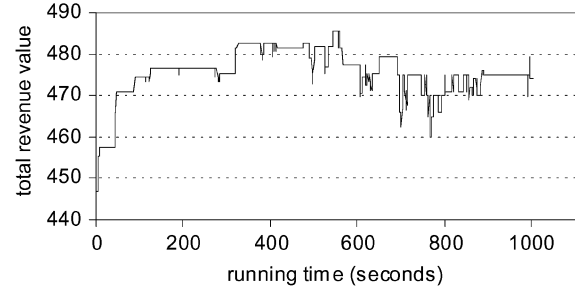


Fig. 6. Tabu search solution trace of total revenue value versus running time during its search procedure of the second multi-start in the case with 90 connection requests and 8 wavelengths in the US-NET.

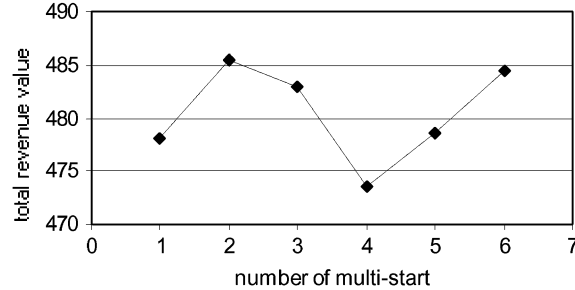


Fig. 7. Tabu search solution trace of revenue value versus the number of multi-start for the case with 90 connection requests and 8 wavelengths in the US-NET.

TABLE VII
RESULTS OF REVENUE MAXIMIZATION BY TABU SEARCH WITH AND WITHOUT MFP IN US-NET

Case #	W	Demands total (n-s-d)	Revenue Value		Improvement
			with MFP	without MFP	
1	4	30 (10-10-15)	195.5	188.5	3.7%
2	8	40 (10-20-10)	420.5	408	3.1%
3	16	50 (15-20-15)	751	738	1.8%

TABLE VIII
RESULTS OF CAPACITY MINIMIZATION BY TABU SEARCH WITH AND WITHOUT MFP IN US-NET

Case #	W	Demands total (n-s-d)	Used Wavelength-Links		Improvement
			with MFP	without MFP	
1	16	50 (15-20-15)	282	294	4.1%
2	16	60 (20-20-20)	335	348	3.7%
3	16	70 (20-30-20)	389	389	4.2%

quency-penalty (MFP) in the US-NET. We show three cases in each table (results for other cases are similar). We observe that tabu search with MFP results in better solutions for both revenue maximization and capacity minimization. These results show that such a diversification technique in the tabu search heuristic is especially important for solving our problems.

Table IX reports the number of trap paths found in each case in the US-NET and Italian network. The third column (K) is the number of *k*-shortest paths. Table IX shows both the number of trap paths and its percentage in terms of the total number of paths found. The results reveal that a minimum of 3.3% and maximum of 11.4% trap paths exist in the US-NET, while the percentage of trap paths is close to 20% for most cases in the Italian network. These results further prove the effectiveness of our tabu search heuristic under different network settings because tabu search obtains a constant improvement over rerouting optimization in both networks. Although many trap paths exist,

TABLE IX
NUMBER OF TRAP PATHS FOR ALL THE CONNECTION REQUESTS

case #	Demands	K	US-NET		Italian Network	
			# of trap paths	% of total paths	# of trap paths	% of total path
1	30	15	51	11%	79	18%
2	40	15	79	13%	123	21%
3	50	15	106	14%	126	17%
4	60	10	20	3%	80	13%
5	70	15	86	8%	196	19%
6	80	15	137	11%	232	19%
7	90	15	127	9%	254	19%
8	140	10	85	6%	262	19%
9	160	10	129	8%	269	17%
10	180	10	132	7%	334	19%

the speed of running our tabu search heuristic is not affected by these trap paths. This is because we insert all of these trap paths into the trap path sets at the beginning of each multi-start and then the tabu search heuristic operates on the other non-trap paths.

VI. CONCLUSION

In this paper, we addressed the problem of SRLG-diverse path provisioning under hybrid service level agreements. We considered some realistic constraints faced by service providers, including multiple classes of protection schemes and the light-path length constraints. This problem was formulated into two sub-problems: the revenue maximization problem and the capacity minimization problem. We provided ILP formulations for both problems. Although solving these formulations are extremely time-consuming for real world problems, they can serve as the basis for the solution approaches that start from relaxing the ILP formulations.

For the problems studied in this paper, heuristic approaches are essential for dealing with their complexity in real-world networks. We designed a rerouting optimization heuristic to encompass some existing optimization principles in previous work. We then proposed a tabu search heuristic. We conducted experiments to compare the tabu search heuristic with the rerouting optimization heuristic in two example networks. Our experimental results showed that our proposed tabu search heuristic can solve the large-scale problems within a reasonable amount of time, while obtaining a better solution than the rerouting optimization heuristic. We also showed that using the tabu search heuristic to jointly solve the revenue maximization and capacity minimization problems can reduce the network resource consumption significantly. Additional experiments illustrated how our design of the tabu search heuristic facilitates an efficient search procedure.

REFERENCES

- [1] S. Ramamurthy and B. Mukherjee, "Survivable WDM mesh networks, part I—Protection," in *Proc. IEEE INFOCOM*, Mar. 1999, pp. 744–751.
- [2] O. Gerstel and R. Ramaswami, "Optical layer survivability: A services perspective," *IEEE Commun. Mag.*, vol. 38, no. 3, pp. 104–113, Mar. 2000.
- [3] —, "Optical layer survivability—An implementation perspective," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 10, pp. 1885–1889, Oct. 2000.

- [4] R. Doverspike and J. Yates, "Challenges for MPLS in optical network restoration," *IEEE Commun. Mag.*, vol. 39, no. 2, pp. 89–96, Feb. 2001.
- [5] M. Goyal, G. Li, and J. Yates, "Shared mesh restoration: A simulation study," in *OFC 2002 Tech. Dig.*, Mar. 2002, pp. 489–490.
- [6] J. Strand, A. Chiu, and R. Tkach, "Issues for routing in the optical layer," *IEEE Commun. Mag.*, vol. 39, no. 2, pp. 81–87, Feb. 2001.
- [7] B. Ramamurthy, S. Yaragorla, and X. Yang, "Translucent optical WDM networks for the next-generation backbone networks," in *Proc. IEEE GLOBECOM*, Nov. 2001, pp. 60–64.
- [8] M. Sridharan, A. Somani, and M. Salapaka, "Approaches for capacity and revenue optimization in survivable WDM networks," *J. High Speed Netw.*, vol. 10, no. 2, pp. 109–125, Aug. 2001.
- [9] I. Chlamtac, A. Ganz, and G. Karmi, "Purely optical networks for terabit communication," in *Proc. IEEE INFOCOM*, vol. 3, Apr. 1989, pp. 887–896.
- [10] J. Q. Hu, "Diverse routing in mesh optical networks," *IEEE Trans. Commun.*, vol. 51, no. 3, pp. 489–494, Mar. 2003.
- [11] G. Li, B. Doverspike, and C. Kalmanek, "Fiber span failure protection in mesh optical networks," *Proc. Optical Networking and Communications (OptiComm)*, pp. 130–141, Aug. 2001.
- [12] E. Modiano and A. Narula-Tam, "Survivable lightpath routing: A new approach to the design of WDM-based networks," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 4, pp. 800–809, May 2002.
- [13] G. Ellinas *et al.*, "Routing and restoration architectures in mesh optical networks," *Opt. Netw. Mag.*, vol. 4, no. 1, pp. 91–106, Jan. 2003.
- [14] M. Garey and D. Johnson, *Computers and Intractability: A Guide to the Theory of NP Completeness*. San Francisco, CA: Freeman, 1979.
- [15] B. Doshi *et al.*, "Optical network design and restoration," *Bell Labs Tech. J.*, vol. 4, no. 1, pp. 58–84, Jan. 1999.
- [16] H. Zang and B. Mukherjee, "Path-protection routing and wavelength-assignment (RWA) in WDM mesh networks under duct-layer constraints," *IEEE/ACM Trans. Netw.*, vol. 11, no. 2, pp. 248–258, Apr. 2003.
- [17] R. Bhandari, *Survivable Networks: Algorithms for Diverse Routing*. Boston, MA: Kluwer, 1999.
- [18] F. Glover and M. Laguna, *Tabu Search*. Boston, MA: Kluwer, 1997.
- [19] J. Yen, "Finding the K shortest loopless paths in a network," *Manag. Sci.*, vol. 17, pp. 712–716, 1977.
- [20] D. Dunn, W. Grover, and M. MacGregor, "Comparison of k-shortest paths and maximum flow routing for network facility restoration," *IEEE J. Sel. Areas Commun.*, vol. 12, no. 1, pp. 88–89, Jan. 1994.
- [21] D. Xu, Y. Xiong, and C. Qiao, "Novel algorithms for shared segment protection," *IEEE J. Sel. Areas Commun.*, to be published.
- [22] D. Xu, Y. Xiong, C. Qiao, and G. Li, "Trap avoidance and protection schemes in networks with shared risk link groups," *IEEE J. Lightw. Technol.*, vol. 21, no. 11, pp. 2683–2693, Nov. 2003.
- [23] C. Assi, A. Shami, M. Ali, R. Kurtz, and D. Guo, "Optical networking and real-time provisioning: An integrated vision for the next-generation Internet," *IEEE Network*, vol. 15, no. 4, pp. 36–45, Jul./Aug. 2001.
- [24] M. Laguna and F. Glover, "Bandwidth packing: A tabu search approach," *Manag. Sci.*, vol. 39, Apr. 1993.
- [25] R. Wagner, L. Nederlof, and S. Maeschalck, "The potential of optical layer networks," in *OFC 2001 Tech. Dig.*, Mar. 2001, pp. TuT3.1–TuT3.3.
- [26] L. Shen, X. Yang, and B. Ramamurthy, "Shared risk link group (SRLG)-diverse path provisioning under hybrid service level agreements in wavelength-routed optical mesh networks: Formulation and solution approaches," in *Proc. OptiComm*, Dallas, TX, Oct. 2003, pp. 126–138.
- [27] P. Ho and H. Mouftah, "Reconfiguration of spare capacity for MPLS-based recovery for the internet backbone networks," *IEEE/ACM Trans. Netw.*, vol. 12, no. 1, pp. 73–84, Feb. 2004.



Lu Shen (S'02) received the M.S. degree in computer science from the University of Nebraska-Lincoln in 2002. He is currently pursuing the Ph.D. degree in the Department of Computer Science and Engineering, University of Nebraska-Lincoln (UNL).

Before he joined UNL, he was with Nortel Networks and Motorola Software Center as a Software Engineer. His research interests include survivable network design, control and management in WDM optical networks, IP over WDM, and MPLS/GMPLS protocol design.



Xi Yang (S'01–M'05) received the B.S. degree in communication engineering in 1997 and the M.S. degree in communication and information systems in 2000 from the University of Electronic Science and Technology, China, and the Ph.D. degree in computer science from the University of Nebraska-Lincoln in 2004.

He is currently a Research Associate with the University of Southern California, Information Science Institute (USC/ISI) East, Arlington, VA. He worked with Lucent Technologies for one year as a Member of Technical Staff in Bell Labs, Beijing, China. His research interests include design of WDM optical networks, wavelength routing and switching, optical network survivability, and generic network design, control and management.



Byrav Ramamurthy (S'96–M'98) received the B.Tech. degree in computer science and engineering from the Indian Institute of Technology, Madras, India, in 1993, and the M.S. and Ph.D. degrees in computer science from the University of California at Davis (UCDavis), in 1995 and 1998, respectively.

During 1998–2003, he was an Assistant Professor and since August 2003, he has been an Associate Professor in the Department of Computer Science and Engineering, University of Nebraska-Lincoln (UNL). He is the co-director of the UNL Academic

Program Priority Initiative in the areas of Simulation and Computing Engineering (SCE) and Information Technology and Telecommunications (ITT). He is the founding co-director of the Advanced Networking and Distributed Experimental Systems (ANDES) Laboratory at UNL. He is author of the textbook *Design of Optical WDM Networks—LAN, MAN, and WAN Architectures* (Boston, MA: Kluwer Academic Publishers, 2000) and coauthor of the textbook *Secure Group Communications Over Data Networks* (Boston, MA: Springer 2004). His research areas include optical networks, wireless networks, network security, distributed computing, and telecommunications.

Prof. Ramamurthy serves as the Technical Program Committee Co-chair for the Optical Systems and Networks Symposium to be held at the ICC 2006 Conference. From 2001–2003, he served as the founding secretary of the IEEE ComSoc Optical Networking Technical Committee (ONTC), for which he currently serves as the Online Content Chair. He was a recipient of the Indian National Talent Search scholarship and was a fellow of the Professors for the Future program at UCDavis. He is a recipient of the UNL Research Council Grant-in-Aid award for 1999, the College of Engineering and Technology Faculty Research Award for 2000 and the UNL CSE Dept. Students Choice Award for the Best Graduate Professor for 2002–2003. His research has been supported by the U.S. National Science Foundation, Agilent Tech., and OPNET Inc.