Availability-Aware Provisioning Strategies for Differentiated Protection Services in Wavelength-Convertible WDM Mesh Networks

Jing Zhang, *Student Member, IEEE*, Keyao Zhu, *Member, IEEE*, Hui Zang, *Member, IEEE*, Norman S. Matloff, and Biswanath Mukherjee, *Fellow, IEEE*

***Abstract*—In an optical WDM mesh network, different protection schemes (such as dedicated or shared protection) can be used to improve the service availability against network failures. However, in order to satisfy a connection's service-availability requirement in a cost-effective and resource-efficient manner, we need a systematic mechanism to select a proper protection scheme for each connection request while provisioning the connection. In this paper, we propose to use connection availability as a metric to provide *differentiated protection services* in a wavelength-convertible WDM mesh network.**

We develop a mathematical model to analyze the availabilities of connections with different protection modes (i.e., unprotected, dedicated protected, or shared protected). In the shared-protection case, we investigate how a connection's availability is affected by backup resource sharing. The sharing might cause backup resource contention between several connections when multiple simultaneous (or overlapping) failures occur in the network. Using a continuous-time Markov model, we derive the conditional probability for a connection to acquire backup resources in the presence of backup resource contention. Through this model, we show how the availability of a shared-protected connection can be quantitatively computed.

Based on the analytical model, we develop provisioning strategies for a given set of connection demands in which an appropriate, possibly different, level of protection is provided to each connection according to its predefined availability requirement, e.g., 0.999, 0.997. We propose integer linear programming (ILP) and heuristic approaches to provision the connections cost effectively while satisfying the connections' availability requirements. The effectiveness of our provisioning approaches is demonstrated through numerical examples. The proposed provisioning strategies inherently facilitate the service differentiation in optical WDM mesh networks.

***Index Terms*—Availability, connection provisioning, differentiated services, optical mesh network, protection, service reliability, WDM.**

Manuscript received August 16, 2003; revised September 3, 2004, and January 31, 2006; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor A. Fumagalli. This work was supported by the National Science Foundation (NSF) under Grant ANI-9805285, and by Sprint Advanced Technology Laboratories (ATL). A short, summarized version of this paper was presented at the IEEE International Conference on Communications (ICC), Anchorage, Alaska, May 2003.

J. Zhang is with Sun Microsystems, Menlo Park, CA 94025 USA (e-mail: j.zhang@sun.com).

K. Zhu is with Brion Technology Inc., Santa Clara, CA 95054 USA (e-mail: kzhu@briontech.com).

H. Zang is with Sprint Advanced Technology Laboratories, Burlingame, CA 94010 USA (e-mail: hui.zang@sprint.com).

N. S. Matloff and B. Mukherjee are with the Computer Science Department, University of California, Davis, CA 95616 USA (e-mail: matloff@cs.ucdavis.edu; mukherje@cs.ucdavis.edu).

Digital Object Identifier 10.1109/TNET.2007.896232

I. INTRODUCTION

WITH the maturing of wavelength-division multiplexing (WDM) technology, a single fiber link can carry a huge amount of data, which might be on the order of terabits per second. However, the failure of a network component (e.g., a fiber link, an optical crossconnect, an amplifier, a transceiver, etc.) can lead to a huge loss in data and revenue. *Protection*, a proactive procedure, is one of the important strategies to recover traffic when a failure occurs [1]–[8]. In protection, one path, referred to as primary path, is used to carry traffic during normal operation while extra backup resources are pre-reserved and they will be activated when the primary path fails. Protection schemes can be classified by the type of routing strategy as link-based versus path-based. In path-based protection, one (or multiple) link- or node-disjoint path(s) (referred to as backup path(s)) are pre-computed and the corresponding network resources are also reserved from the source node to the destination node to recover the traffic in case of a failure along the primary path. The resources on a backup path can be dedicated to one connection or shared among different connections as long as any two of these connections are not in the same shared-risk group (SRG). Since link failure is the dominant failure scenario, shared-risk link group (SRLG) is commonly used and will be referred to primarily in this paper.

Compared to a ring network, a WDM mesh network can provide a wide variety of protection schemes. What we lack, however, is a systematic methodology to efficiently select a cost-effective protection scheme for each connection while satisfying its quality-of-service (QoS) requirements. Usually, QoS can be measured in many different ways such as signal quality, service availability, service reliability, restoration time, service restorability, etc. Our interest is in the availabilities of service paths (i.e., connections) since availability is one of the key concerns of customers and usually defined in a Service-Level Agreement (SLA). The SLA is a contract between the network operator and a customer. A SLA violation may cause a certain amount of penalty to be paid by the network operator according to the contract, e.g., providing free services for one additional month. Thus, a cost-effective, availability-aware, connection-provisioning scheme is very desirable such that, for each customer's service request (static or dynamic), a proper protection scheme (dedicated, shared, or unprotected) is designed to guarantee the SLA-defined availability requirement and to reduce overall cost.

Connection availability is defined as the probability that the connection will be found in the operating state at a random time in the future [9]. It is defined only over a connection's lifetime and can be computed statistically based on the failure frequency and failure repair rate, reflecting the percentage of time a connection is "alive" or "up" during its entire service period. Although the problem of how the connection availability is affected by network failures is currently attracting more research interest [9]–[23], we still lack a systematic methodology to quantitatively estimate a connection's availability, especially when shared-protection schemes are applied.

It should be clear that a protection scheme will help improve a connection's availability since traffic on the failed primary path will be quickly switched to the backup path. For example, a path-protected connection will have 100% availability in the presence of any single failure if the contribution of the reconfiguration time from primary path to backup path towards unavailability is disregarded [since it is relatively small (usually on the order of a few tens of milliseconds) with respect to the failure repair time (on the order of hours) and the connection's holding time (on the order of weeks or months)].

Nevertheless, a more realistic failure scenario is multiple, simultaneous (or overlapping) failures where more than one failure occurs in the network and their failure states overlap in time. When the multiple-failure case is considered, a path-protected connection may become unavailable in some failure scenarios, e.g., when two concurrent failures occur, one on the backup path and the other on the primary path. Therefore, when considering multiple failures, connection availability depends intimately on the precise details of the failures (locations, repair times, etc.), how much backup resources are reserved (i.e., single backup route or multiple backup routes), and how the backup resources are allocated (i.e., dedicated or shared). Intuitively, the more backup resources (paths) there are, the higher is the connection availability, while more backup sharing leads to lower connection availability. Therefore, instead of simply stating that a connection has been protected, we need to quantitatively evaluate how well the connection is protected, i.e., we need to have a relatively accurate estimation of its availability so that the SLA can be satisfied.

Unlike a lot of previous work, in which single or double network failure scenario is assumed, we do not make any specific failure scenario assumption. Instead, the failure behavior of a network component will follow its physical characteristics. Therefore, the network may experience multiple network component failures concurrently. Consequently, a connection t will become unavailable in the following cases:

1) One failure occurs on primary path of t and a second failure occurs on backup path of t.
2) If t shares its backup wavelength with connection t_1 on one backup link, t will be unavailable if the primary paths of both t and t_1 fail but the shared backup wavelength is taken by t_1.

The failed connection will be in the "down" state until the failure on its primary path or backup path is repaired, or backup wavelengths are released by other connections.

Note that service availability is not the only QoS metric we need to consider to provide differentiated services in a WDM mesh network. For instance, two connections, A and B, may have the same availability during their entire service periods; however, A may experience fewer network failures with longer service downtime for each failure and while B may experience more network failures with shorter service downtime. Although A and B have the same service availability, they have different service disruption rates and failure-repair times, which may lead to different customer-perceived service qualities. In our current study, we focus on service availability and will incorporate other service-quality metrics, e.g., service disruption rate, in our future study.

The rest of the paper is organized as follows. Section II discusses related works and our contributions. Section III presents a mathematical availability-analysis model for connections with different protection schemes in WDM mesh networks. Section IV presents general provisioning strategies using the analytical model in which an appropriate level of protection is provided to each connection according to the customer's predefined (or desired) availability requirement. Both ILP and heuristic-based approaches are developed for static traffic where a given set of connection demands need to be provisioned. Illustrative numerical results are presented and analyzed in Section V. Section VI concludes the study.

II. Related Studies and Our Contributions

Availability analysis and the idea of providing differentiated reliability in SONET rings have been studied in the optical network literature [10]–[12], [14]. The authors in [10] have given an extensive review on availability in ring networks. The concept of differentiated reliability (DiR) has been proposed and studied in [12], [13] to provide multiple reliability degrees using a common protection mechanism in optical ring networks. The work in [14] analyses a number of long-haul network architectures from an unavailability point of view and shows that self-healing rings and dual fed systems offer the highest level of survivability, by eliminating service impacts caused by cable cuts and equipment failures.

Recently, increasing attention has been devoted to service availability and reliability in WDM mesh networks [4], [9], [15]–[23]. The work in [9] evaluates the restorability of span-restorable mesh networks when dual failures occur. The restorability of a network is defined as the average fraction of failed working capacity that can be restored within the spare capacity. This means that, when dual failures occur, a connection can be restored on the fly if both its primary path and pre-computed backup path get affected. It is reported that single-failure-designed mesh networks inherently have high levels of dual-failure restorability. The work in [4] examines the susceptibility of link-based and path-based protection schemes to multiple link failures. The susceptibility of a network is defined as the average fraction of failed connections during multiple link failures without allowing on-line restoration, i.e., a connection can only be carried by its primary path or pre-computed backup path. The results in [4] show that there is a trade-off between the capacity utilization and the susceptibility to multiple link failures, and shared-path protection is a little more susceptible to two-link failures than shared-link protection.

The authors in [17] extend the differentiated-reliability concept to shared-path protection in mesh networks with the assumption of single network failure. Their idea is to select some links along the primary path, and leave them unprotected so as to increase the backup resource sharability, but still guarantee the required maximum acceptable failure probability. The works in [19] and [20] consider the availability in multi-domain mesh networks, and they both show that partitioning a network into multiple domains increases the overall availability. In [21], an availability calculation model is studied to estimate both connection and system availability of different protection techniques such as 1:1, M:N, and mesh shared protection. Their model is close to our availability analysis model except that they introduce an approximation in analyzing the availability provided by shared protection. In [22], [23], the tradeoff between capacity requirement and service availability provided by reserved protection resources has been studied.

Unlike most previous work, we present a framework in this paper to provide differentiated protection services to meet customers' availability requirements cost effectively. We first develop an availability-analysis model for connections with different protection schemes (i.e., unprotected, dedicated protected, or shared protected). Through this model, we show how a connection's availability is affected by resource sharing.

Based on the analytical model, we then develop provisioning strategies (both integer linear program (ILP) and heuristic based) in which an appropriate level of protection is provided to each connection according to its predefined availability requirement. We consider full wavelength-conversion networks and static lightpath provisioning where a set of traffic demands is given in advance, each of which requires the full capacity of a wavelength channel, and the network operator needs to provision each connection with minimal network cost, and at the same time, meet the connections' availability requirements.

III. Availability Analysis in WDM Mesh Networks

We analyze the availability of a system (which could be a component, path, connection, etc.) in a mesh network with the following typical assumptions:

1) a system is either available (functional) or unavailable (experiencing failure);
2) different network components fail independently; and
3) for any component, the "up" times (or Mean Time To Failure, MTTF) and the repair times (or Mean Time To Repair, MTTR) are independent memoryless processes with known mean values.

The availability of a system is the fraction of time the system is "up" during the entire service time. If a connection t is carried by a single path, its availability (denoted by A_t) is equal to the path availability; if t is dedicated or shared protected, A_t will be determined by both primary and backup paths. Here, the contribution of the reconfiguration time for switching traffic from the primary path to the backup path (including signal propagation delay of control signals, processing time of control messages, and switching time at each node) towards unavailability is disregarded since it is relatively small, usually on the order of a few tens of milliseconds, compared to the failure-repair time (on the order of hours) and the connection's holding time (on the order of weeks or months).

TABLE I
Failure Rates and Repair Times (Bellcore) [14]

Metric	Bellcore Statistics
Equipment MTTR	2 hrs
Cable-Cut MTTR	12 hrs
Cable-Cut Rate	4.39/yr/1000 sheath miles
Tx failure rate in FIT	10867
Rx failure rate in FIT	4311

A. Methodology for Assessing Network-Component Availability

A network component's availability can be estimated based on its failure characteristics. Upon the failure of a component, it is repaired and restored to be "as good as new". This procedure is known as an alternating renewal process. Consequently, the availability of a network component j (denoted as a_j) can be calculated as follows [24]:

$$a_j = \frac{\text{MTTF}}{\text{MTTF} + \text{MTTR}}. \tag{1}$$

Component failure parameters usually can be obtained from the network operators. In particular, the MTTF of a fiber link is distance related and can be derived according to measured fiber-cut statistics. We also assume that the repair process of each link is independent of one another so two links will be repaired in parallel if their failure states overlap. Table I shows some typical data on failure rates and failure-repair times of network components (transmitters, receivers, fiber links, etc.) according to Bellcore (now Telcordia). In Table I, FIT (failure-in-time) denotes the average number of failures in 10^9 hours, Tx denotes optical transmitters, and Rx denotes optical receivers.

B. End-To-End Path Availability

Given the route of a path i, the availability of i (denoted as A_i) can be calculated based on the known availabilities of the network components along the route. Path i is available only when all the network components along its route are available. Let a_j denote the availability of network component j. Let G_i denote the set of network components used by path i. Then, A_i can be computed as follows:

$$A_i = \prod_{j \in G_i} a_j. \tag{2}$$

C. Availability for a Dedicated-Path-Protected Connection

In path protection, connection t is carried by one primary path p and protected by one backup path b that is link disjoint with p. By link disjoint, we mean that the backup path for a connection has no links in common with the primary path for that connection. Node failures can be also accommodated by making the primary and the backup paths node disjoint as well. However, one should also note that carrier-class optical crossconnects (OXCs) in network nodes must be $1 + 1$ (master/slave) protected in the hardware for both the OXC's switch fabric and its control unit. The OXC's port cards, however, do not have to

be $1+1$ protected since they take up the bulk of the space (perhaps over 80%) and cost of an OXC; also a port-card failure can be handled as link and/or wavelength channel failure(s). However, node failures are important to protect against in scenarios where an entire node (or a collection of nodes in a part of the network) may be taken down, possibly due to a natural disaster or by a malicious attacker. In this study, we require primary and backup of a connection to be link-disjoint and only consider link failures in the availability analysis. Extensions to include node failures when computing connection availability are open problems for future research.

If b's backup wavelengths are dedicated to connection t, then, when primary path p fails, traffic will be switched to b as long as b is available; otherwise, the connection becomes unavailable until the failed component is replaced or restored. t is "down" only when both paths p and b are unavailable, so A_t can be computed straightforwardly as follows:

$$A_t = 1 - (1 - A_p) \times (1 - A_b) = A_p + (1 - A_p) \times A_b \quad (3)$$

where A_p and A_b denote the availabilities of paths p and b, respectively. A connection may employ multiple backup paths to increase its availability. If all backup paths are disjoint and dedicated to this connection, the connection availability can be derived following the similar principle in (3).

D. Mathematical Model for Availability of a Shared-Path-Protected Connection

1) Issues Affecting Availability in Backup Sharing: In this section, we describe various issues or policies in backup sharing that will affect the availability of a shared-path-protected connection.

- **Share Per Single Backup Wavelength Versus Share Per Wavelength Pool** In shared-path protection, connection t is carried by primary path p, and protected by a link-disjoint backup path b; but the reserved wavelength on each link of b can be shared by other connections as long as SRLG constraints can be satisfied. Let S_t contain all the connections that share some backup wavelength on some link with t. We denote S_t as the sharing group of t. In the literature, backup sharing has been performed in two ways: share per single backup wavelength and share per wavelength pool [8], [25], [26]. In the first case, backup wavelength is fixed on every backup link of a connection while in the second case, on each backup link, a backup wavelength will be chosen when failure occurs from a pre-reserved backup wavelength pool. Connection availability will be calculated differently in the two schemes. We consider share per single backup wavelength in this study.
- **Reverting Versus Non-reverting** Connection t's traffic will be switched to b when a failure occurs on p. After the failure is repaired, connection t's traffic can be switched back to p, an approach which is called *reverting*; or it can stay on b for the remaining service time (or till b fails), an approach which is called *non-reverting*. Both the reverting and non-reverting strategies have their pros and cons. For example, traffic may be disturbed twice in the reverting strategy, which may be undesirable for some services. In the non-reverting strategy, the backup paths for the connections in S_t may need to be rearranged since some of the shared backup wavelengths on parts of their backup paths have been taken by t when t is switched to its backup path. These connections can become vulnerable during their backup-recomputation and backup-resource-reservation processes; and, furthermore, their successful backup rearrangement is not guaranteed; so, non-reverting may result in unpreferred service degradation. A network operator may choose policies based on operational cost and service characteristics. The reverting model may sometimes be preferable since it provides simplicity in network control and management. We assume a reverting model in our analysis.
  The concept of stub release refers to the release of capacity along the surviving upstream and downstream portions of a failed primary path, and making those capacity available for the restoration process. Since we only consider to restore a connection using the preplanned backup path (with static traffic demands in this paper) and assume a reverting model, stub release is not relevant for this modeling study. Stub release will become important for dynamic provisioning where connections come and go.
- **Active Recovery Versus Lazy Recovery** In the reverting model, after traffic is reverted back to p, the shared backup resources will be released. Similarly, when backup resources are fixed from a failure, they are also "up and free", which means that the backup resources are not in failing states (up) or being used by any connection (free). In both of the two cases, the fixed or released backup resources can be actively used to recover the connections in S_t that are experiencing failure and waiting for their backup resources to be fixed or released. We call this mechanism *active recovery*. On the contrary, if the backup resources wait to be activated when the next failure arrives, these currently failed connections cannot be recovered even though their backup is up and free now. This mechanism is called *lazy recovery*. In active recovery, the backup resources released by a connection may be able to recover more than one connection as b may traverse multiple links. Obviously, backup resources are utilized more intelligently in the active-recovery model so we assume an active-recovery system in our study.
  If active recovery is employed, another problem will arise, i.e., if there are multiple failed connections waiting for the backup resources, which connection should be chosen to recover next? Connections can be recovered in the exact order of their failure sequence, i.e., earliest failure recovered first. We call this a resource-locked system in the sense that a failed connection will "lock" all the up and free backup wavelengths it needs and wait for others to be fixed or released. And we further assume that the locked backup resources can only be released when the primary path of the failed connection is fixed. Fig. 1 shows an example, where t, t_1, and t_2 are three connections; t and t_1 share the same backup wavelength on link l_1; and t and t_2 share the same backup wavelength on link l_2. If the failure

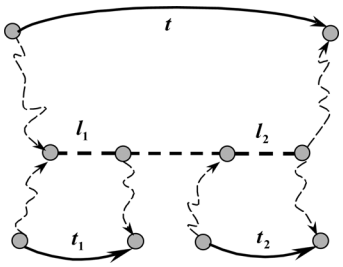

Fig. 1. A general backup sharing example.

sequence is t_1, t, and t_2, then t will lock the backup wavelength on link l_2 so it cannot be used by t_2 even though it is up and does not recover any connection when t_2 fails. A locked system can provide fairness in the context of a first-fail-first-served (FFFS) policy. Therefore, we assume a locked system in the following analysis.

2) Computation of the Conditional Probability That a Connection Succeeds in Backup-Resource Contention: The availability of connection $t(A_t)$ will be affected by the size of S_t and the availabilities of the connections in S_t. When one or more primary paths of connections in S_t fails together with t, either t or some of the failing connections in S_t can acquire the shared backup wavelengths. Hence, we need to compute the conditional probability (denoted as $\delta_t^{(k)}$) that t will successfully acquire the backup wavelengths when k connections' primary paths in S_t fail concurrently with t's primary path.

We employ a continuous-time Markov chain to derive $\delta_t^{(k)}$. Fig. 2(a) shows the corresponding state-transition diagram (when $k = 1$) for the Markov chain when an active-recovery, resource-locked system is applied. Let t_1 denote the other connection which shares backup resources with t. The label for each state in Fig. 2 is a 3-tuple (x, y, z), where x and y represent the status of the primary paths of connections t and t_1, respectively, and z represents which connection uses or locks the backup resources. Tuples x and y could be "Up" (U) or "Down" (D); and z could be "None" (0), "t", or "t_1". Note that we do not show the "Up/Down" state of shared backup resources or the "Up/Down" state of each connection in Fig. 2(a). $z = t$ (or $z = t_1$) does not mean that the traffic of t (or t_1) is being restored by backup resources, and only indicates that the shared backup resources are locked by t (or t_1) (which is the first failed connection, as we assume an active-recovery, resource-locked system). Actually, even though $z = t$ (or $z = t_1$) in some states, t (or t_1) will be down if the backup resources are down in these states.

Let $\mathrm{MTTF} = \frac{1}{\lambda}$ and $\mathrm{MTTR} = \frac{1}{\mu}(\mathrm{MTTF}_1 = \frac{1}{\lambda_1}$ and $\mathrm{MTTR}_1 = \frac{1}{\mu_1}$ to be the mean failure parameters for the primary path of connection $t(t_1)$. The state-transition probabilities can thus be represented by these parameters. Let π_i denote the long-run proportion of time the system is in state i. Again, $\delta_t^{(1)}$ is the conditional probability that t has the backup resources, given that both t and t_1 are down. After solving for π_i (the details of the solution are straightforward and not included here), $\delta_t^{(1)}$ can be computed as follows for Fig. 2(a) (as t and t_1 are

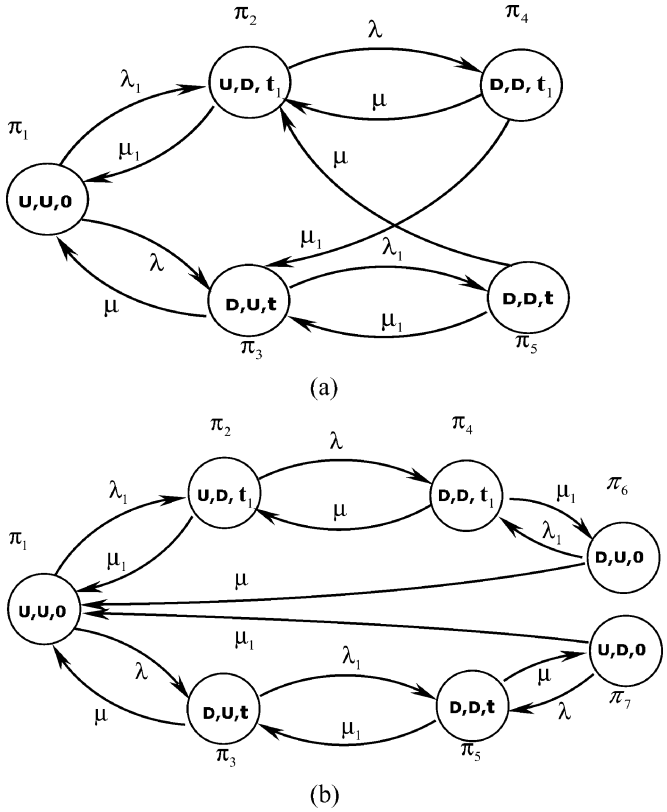

Fig. 2. State-transition diagram for computing $\delta_t^{(1)}$. (a) With active recovery. (b) With lazy recovery.

both down in states 4 and 5 but t has the backup resources only in state 5):

$$\delta_t^{(1)} = \frac{\pi_5}{\pi_4 + \pi_5} = \frac{\mu_1}{\mu + \mu_1}. \tag{4}$$

The solution implies solving the rate-based linear system according to the Kolmogoroff equations for the continuous-time Markov chain. Similarly, we can compute the conditional probability for one connection to acquire the backup wavelengths when $k (k \geq 2)$ primary paths in S_t are experiencing failures concurrently. Please see [27] for the state-transition diagram for $k = 2$ as an example. $\delta_t^{(2)}$ is shown as follows where $\mathrm{MTTR} = \frac{1}{\mu}(\mathrm{MTTR}_i = \frac{1}{\mu_i})$ denote the mean time to repair for connection t (t_i, where $i = 1$ or 2):

$$\delta_t^{(2)} = \frac{\mu_1 \mu_2}{\mu + \mu_1 + \mu_2} \times \left(\frac{1}{\mu + \mu_1} + \frac{1}{\mu + \mu_2} \right). \tag{5}$$

One may notice that the values of $\delta_t^{(1)}$ in (4) and $\delta_t^{(2)}$ in (5) are only determined by the repair rates of the concurrently failed primary paths (and not their failure rates)! Due to the complexity of deriving the repair rate of a path (which is related to the MTTF and MTTR of each individual link along the path), we use an approximation to simplify the value of $\delta_t^{(k)}, k = 1, 2$. We approximately assume that all the primary paths have the same repair rate, which is referred to as Approximation I in what follows. Then, we have $\delta_t^{(1)} = \frac{1}{2}$ from (4) and $\delta_t^{(2)} = \frac{1}{3}$ from (5). We further make a conjecture that $\delta_t^{(k)} = \frac{1}{k+1}$ for any $k \geq 3$ under

Approximation I. Intuitively, each one of the $k+1$ failed primary paths (including t) will approximately have equal chance to get the backup wavelengths if the conditional probability is only affected by the repair rates and under the approximation that all of them have the same repair rate. With Approximation I, the value of $\delta_t^{(k)}(k=1,2)$ is greatly simplified. Most importantly, without this approximation, computing $\delta_t^{(k)}$ when k is large is extremely complicated and time consuming as the size of the Markov chain will grow exponentially. We have conducted simulations to verify our model for a general backup sharing case and the results show that the error between the availability calculated using our model and that from the simulation is negligible (please see Section V-A for the results), which indicates that the error brought by Approximation I is negligible.

We can follow the same approach to derive $\delta_t^{(k)}$ for other recovery policies, e.g., lazy recovery, even though the Markov chain may be different. As an example, Fig. 2(b) shows the corresponding Markov chain to compute $\delta_t^{(1)}$ with the lazy-recovery policy.

With the value of $\delta_t^{(k)}$, we can compute the availability of a shared-path-protected connection t now. Connection t will be available if: 1) path p is available; or 2) p is unavailable, b is available, and p can get the backup wavelengths when other primary paths of connections in sharing group S_t have also failed. Therefore, A_t can be computed as follows:

$$A_t = A_p + (1 - A_p) \times A_b \times \sum_{k=0}^{N} \delta_t^{(k)} \times p_k \tag{6}$$

where A_p and A_b denote the availabilities of p and b, respectively; N is the size of S_t; $\delta_t^{(k)}$ is the probability that t can get the backup resources when both p and other k primary paths in S_t fail; and p_k is the probability that exactly k primary paths in S_t are unavailable. We can enumerate all the possible k connection failures to compute p_k. Note that it may not be necessary for us to enumerate all the possible simultaneous failure cases (up to N) since the probability of k simultaneous failures decreases drastically as k increases. Hence, such failure scenarios will have little effect on the connection availability. In a practical network, instead of enumerating all possible failure scenarios, we may only consider up to B simultaneous connection failures, where B is known as the *approximation bound*. By properly choosing the value of B, we can get a very tight lower bound on the connection availability. The value of B depends on the network failure characteristics, i.e., the more fragile the network is, the larger the value of B should be, and vice verse. The computational complexity of (6) depends on: 1) the size of the sharing group (N); and 2) the approximation bound (B). We find that it will only take several seconds to compute (6) when several tens of connections are in the sharing group and B is around 10 using a computer with a 1.4-GHz Pentium processor and 512-Mbytes RAM; thus, the computation is feasible in a practical network.

Besides the availability analysis shown above, the mean down time, which a shared-path-protected connection t experiences after a failure of one of its primary links, can also be derived analytically. Please see [27] for this computation. The results can be used to assess the severity of the impact of network failures on connections.

IV. Proposed Connection-Provisioning Strategies

Based on the analytical model, we have developed connection-provisioning approaches in which differentiated protection services can be provided to each connection according to its predefined availability requirement. We first discuss how to compute the path with the highest availability between a node pair in the network, which is referred to as the *most-reliable path*. This idea will be frequently used in the following provisioning strategies. Then, we propose ILP and heuristic-based strategies to provision connections cost effectively while satisfying the connections' availability requirements by choosing appropriate protection schemes.

A. Techniques to Compute the Most-Reliable Path

Suppose a single path p is used to carry connection t. The availability of $p(A_p)$ is equal to the multiplication of the availabilities of components it traverses as we have discussed in Section III-B. In what follows, we consider links as the only network components used by a path but it is straightforward to incorporate other network components as well. Suppose path p traverses links $l_1, l_2, \ldots, l_n$. We call p to be a *reliable path* for connection t if and only if:

$$A_p = A_{l_1} \times A_{l_2} \times \cdots \times A_{l_n} \geq A'_t \tag{7}$$

where A_{l_i} is the availability of link $l_i, 1 \leq i \leq n$, and A'_t is the required availability of connection t. If we compute the logarithm of both sides of (7), we obtain:

$$\log A_p = \log A_{l_1} + \log A_{l_2} + \cdots + \log A_{l_n} \geq \log A'_t. \tag{8}$$

Since A_{l_i} and A'_t are between 0 and 1, $\log A_{l_i}$ and $\log A'_t$ have negative values. Multiplying both sides by -1, we get:

$$-\log A_p = -\log A_{l_1} - \log A_{l_2} - \cdots - \log A_{l_n} \leq -\log A'_t. \tag{9}$$

Now, we can observe that, if the cost of link $l_i(C_{l_i})$ is defined as a function of its availability (i.e., $C_{l_i} = -\log A_{l_i}$), the cost is additive and the path with minimum cost will be the path with maximum availability (i.e., the most-reliable path (MRP)). Through this *Multiplication-to-Summation* (MS) conversion technique, a standard shortest-path algorithm can be applied to compute the MRP.

Taking the logarithm is convenient but the derivation of the dynamic-programming algorithms works the same if multiplication is used instead of addition. An alternate way to compute the MRP with multiplication is as follows:

1) define link cost equal to link availability; and
2) modify the shortest-path algorithm with "multiplication" parameter to compute the MRP.

The standard shortest-path algorithm is computing the shortest path with the "addition" parameter. If the link cost is between 0 and 1, we can easily modify a standard shortest-path algorithm (such as Dijkstra's or Bellman-Ford algorithm) to compute the longest path with the length of the path defined as the multiplication of the cost of each link along the path.

If the availability of a MRP is smaller than A'_t, we know that protection is needed for connection t. Therefore, we can

categorize a connection as either a *one-path-satisfiable connection* whose availability requirement can be satisfied without using any backup path, or a *protection-sensitive connection*, otherwise.

In the remainder of this section, we present our availability-aware provisioning approaches for static traffic, including an ILP approach with dedicated-path protection and no protection as the candidate protection services and heuristics with dedicated-path protection, shared-path protection, and no protection as the candidate protection services. We are given the following inputs to the problem.

1) $G = (V, E, A, W)$, the physical network topology where V is the set of nodes, E is the set of unidirectional fiber links, A: $E \rightarrow (0,1)$ is the availability function for each link (where $(0,1)$ denotes the set of real numbers between 0 and 1), and W: $E \rightarrow Z^+$ specifies the number of free wavelengths on each link (where Z^+ denotes the set of positive integers).
2) $T = \{t = \langle s, d, A'_t \rangle\}$, a set of connection requests that need to be provisioned where s is the source, d is the destination, and A'_t is the availability requirement of request t.

Our goal is to determine the route for each request and protect them, if necessary, while minimizing the total network cost (wavelength links, particularly).

To optimize network-resource usage, we first classify the connection requests into two categories (by comparing availability of MRP with A'_t as described above): T_1, containing one-path-satisfiable connections, and T_2, containing protection-sensitive connections; and then, we provide different treatments to different connection sets, as follows.

1) For a connection in T_1, one path is needed to carry each of them. We use an ILP to find the routes that can satisfy the connections' availability requirements while minimizing the consumed resources (wavelength links). The ILP is given in Section IV-B.
2) Dedicated-path protection is considered to protect connections in T_2. The problem of providing dedicated-path protection while satisfying the connections' availability requirements is mathematically formulated in Section IV-C. We also discuss the nonlinearity of the formulations and propose two approximation schemes to solve them.

We then incorporate shared-path protection into the differentiated protection service model to further reduce network cost. Due to the complexity of availability analysis for a shared-path protected connection (see (6)), formulating the problem into a linear program would be extremely complicated and thus intractable mathematically. Therefore, we have to resort to heuristics when incorporating shared-path protection into the differentiated protection service model. Also, there are instances where the ILP approaches may have difficulty due to large network size and high volume of traffic demands even when shared-path protection is not considered. The heuristic algorithms are presented in Section IV-D.

B. ILP for One-Path-Satisfiable Connections

The MS conversion technique enables us to formulate the problem of provisioning connections in T_1 into an ILP since (nonlinear) multiplication has been converted into (linear) summation. We will use following notations in our mathematical formulations:

1) m and n denote end points of a physical fiber link; and
2) s and d denote source and destination of a given end-to-end connection request t.

The mathematical formulation for one-path-satisfiable connections is as follows.

- *Given:*
  - P_{mn}: Number of fiber links interconnecting node m and node n. $P_{mn} = P_{nm} = i, i > 0$, if and only if there exists i physical fiber links between nodes m and n; 0 otherwise.
  - W_{mn}: Number of wavelengths per fiber on link (m, n).
  - A_{mn}: Availability of link (m, n). If there are multiple fibers between a node pair, they have same availability if they traverse the same fiber bundles. (Note that fibers are usually laid in bundles.)
  - α_{mn}: Availability parameter of link (m, n) where $\alpha_{mn} = -\log A_{mn}$.
  - $T_1 = \{t = \langle s, d, \alpha_t \rangle\}$: Connection request set, where α_t is the minimum required availability parameter of connection t and defined as $\alpha_t = -\log A'_t$.
- *Variables:*
  - P^t_{mn}: $P^t_{mn} = 1$ if request t is routed through fiber link (m, n); otherwise, $P^t_{mn} = 0$.
- *Objective:* Minimize the total wavelength links used[1]:

$$\text{Minimize: } \sum_t \sum_{m,n} P^t_{mn}. \tag{10}$$

- *Constraints:*
  - On physical route flow-conservation constraints:

$$\sum_m P^t_{mk} = \sum_n P^t_{kn} \text{ if } k \neq s, d \quad \forall t, k \tag{11}$$

$$\sum_n P^t_{sn} = \sum_m P^t_{md} = 1 \quad \forall t. \tag{12}$$

  - On link-capacity constraints:

$$\sum_t P^t_{mn} \leq P_{mn} \times W_{mn} \quad \forall m, n. \tag{13}$$

  - On connection-availability constraints:

$$\sum_{m,n} P^t_{mn} \times \alpha_{mn} \leq \alpha_t \quad \forall t. \tag{14}$$

Note that P_{mn} and α_{mn} are given (constants), so (10)(14) are linear.

In (13), we assume that all of the wavelengths on link (m, n) (i.e., $P_{mn} \times W_{mn}$) can be utilized for the provisioning. However, for a general static connection-provisioning problem, the number of wavelengths on a link that can be used may also need to be optimized to avoid over-utilizing or congesting links. In Section V-B, we propose a simple approach to first determine the minimal number of wavelengths (denoted as W) through

[1]If the cost of each link is given, we can also incorporate the cost into the objective by minimizing the total cost: $\text{Minimize} : \sum_t \sum_{m,n} c_{mn} P^t_{mn}$ where c_{mn} is the cost of link (m, n).

which all the connection requests can be carried. Then we optimize the total number of consumed wavelength-fiber links given that the number of wavelengths on each link is constrained by W. Please see Section V-B for detailed approach.

C. Mathematical Formulation for Protection-Sensitive Connections

For connections in T_2, we provide dedicated-path protection to them. The problem to be solved now is to route each connection t in T_2 using two link-disjoint paths while satisfying A'_t and minimizing the resources used. The problem is mathematically formulated as follows (using the same notations as in the formulations in the previous section):

- *Variables:*
  - P^{tp}_{mn}: $P^{tp}_{mn} = 1$ if primary path of connection t is routed through fiber link (m, n); otherwise, $P^{tp}_{mn} = 0$.
  - P^{tb}_{mn}: $P^{tb}_{mn} = 1$ if backup path of connection t is routed through fiber link (m, n); otherwise, $P^{tb}_{mn} = 0$.
- *Objective A:* Minimize the total wavelength links used:

$$\text{Minimize: } \sum_t \sum_{m,n} (P^{tp}_{mn} + P^{tb}_{mn}). \tag{15}$$

- *Constraints:*
  - On physical route flow-conservation constraints: They are similar to (11)–(12) except that separate constraints are needed for both primary (P^{tp}_{mn}) and backup (P^{tb}_{mn}) paths.
  - On link-disjoint constraints:

$$P^{tp}_{mn} + P^{tb}_{mn} \le 1 \quad \forall m, n, t. \tag{16}$$

  - On link-capacity constraints:

$$\sum_t (P^{tp}_{mn} + P^{tb}_{mn}) \le P_{mn} \times W_{mn} \quad \forall m, n. \tag{17}$$

  - On connection-availability constraints:

$$\text{Define } x = \sum_{mn} P^{tp}_{mn} \times \alpha_{mn} \tag{18}$$

$$y = \sum_{mn} P^{tb}_{mn} \times \alpha_{mn} \tag{19}$$

$$1 - (1 - e^{-x}) \times (1 - e^{-y}) \ge A'_t \quad \forall t. \tag{20}$$

Please note that availability of the primary path is e^{-x} and that of the backup path is e^{-y} in (20) as $\alpha_{mn} = -\log A_{mn}$. Due to the nonlinearity of (20), the problem cannot be solved as an ILP. One approximation approach is to solve the formulation without the constraints in (20) as an ILP, i.e., optimize network resources to provide dedicated-path protection for connections of T_2 without considering the availability constraints. Since the dedicated-protection scheme may significantly improve the connections' availabilities, it is expected that the availabilities of most of the connections in T_2 can be satisfied using this approximation.

Another solution is to solve the formulation without the constraints in (20) and modify the objective A in (15) as follows.

- *Objective B:*

$$\text{Minimize: } \sum_t \sum_{m,n} \left(P^{tp}_{mn} \times \alpha_{mn} + P^{tb}_{mn} \times \epsilon\right). \tag{21}$$

This objective tries to maximize the availabilities of the primary paths, and at the same time, minimizes the total wavelength links used by the backup paths. ϵ is a positive number which is assigned a small value such that maximizing the availabilities of the primary paths is of higher priority.

D. Heuristic Algorithms

As we have mentioned in Section IV-A, we resort to heuristics when incorporating shared-path protection into the differentiated protection service model due to the complexity of availability analysis for a shared-path-protected connection (see (6)). We start by investigating several heuristics to provision connections with unprotected or dedicated-path-protection services. Then, we downgrade a dedicated-path-protected connection's protection service to shared-path protection as long as the connection's availability requirement can still be met. The heuristics are *fixed-alternate-routing based* [28], i.e., for each node pair, M candidate routes or link-disjoint route-pairs are pre-computed, and availability of each route is calculated. Therefore, a request $t = \langle s, d \rangle$ can pick routes (or route-pairs) that satisfy its requirement from the M candidate routes from s to d.

In the numerical examples shown in Section V-C, we precompute $M = 9$ candidate routes for each node pair, among which four are single paths and five are link-disjoint path-pairs. The ways to compute the candidate routes are described here to facilitate reproduction of our results by others. Route 1 is the shortest path (SP) by hop distance. Route 2 is the SP by hop distance after removing the link with lowest availability on route 1. Route 3 is the MRP. Route 4 is the MRP after removing the link with highest availability on route 3. Route 5 is the shortest path-pair by hop distance computed using the two-step approach[2] [29]. Route 6 is the shortest path-pair by hop distance computed using Suurballe algorithm[3] [30]. Route 7 is the shortest path-pair computed using the two-step approach where the cost of link l_i is defined as a function of its availability, i.e., $C_{l_i} = -\log A_{l_i}$. Route 8 is the shortest path-pair computed using Suurballe algorithm where the cost of link l_i is defined as a function of its availability, i.e., $C_{l_i} = -\log A_{l_i}$. In route 9, first path is the MRP and second path is the SP by hop distance after removing the links along the MRP.

The main concern for computing candidate routes is the trade-off between resource utilization and availability. For example, if resource (i.e., hop distance) is used as the only metric to compute routes, we cannot control the availability of each route; if availability is the only metric, we may end up with extensively utilizing the links with high availabilities, which will create congested links in the network. We incorporate these concerns when computing the candidate routes. Note that it is not guaranteed that each node-pair will have M distinct routes,

[2]In the two-step approach, the first path is the shortest path and the second path is the shortest path after removing the first path.

[3]In Suurballe algorithm, the two paths are jointly computed such that the total cost of the two paths is minimum among all such link-disjoint path pairs.

e.g., the SP by hop distance and the MRP for the same node-pair may follow the same route. So the number of candidate routes for each node pair is equal to or smaller than M in our study. One can also apply other algorithms for computing M shortest paths or path-pairs [31] and study their performance, but we feel that our route choices are a bit customized for the current problem.

Let Q_t denote a set containing all routes or route-pairs among the M candidates that can satisfy the availability requirement of request t. Let $R_{\rm best}(\rm RP_{best})$ denote the route (route-pair) with the highest availability in Q_t. Each request t can select its route using one of the following approaches.

- *Iteratively-select*: Randomly pick one request t, and randomly pick one route or route-pair r from Q_t. Use r to carry t if replacing current route of t by r could reduce total cost (wavelength-links); otherwise, keep current route. Repeat above steps until no route replacement occurs in a large number of continuous iterations (10^5 in our numerical simulations).
- *Most-reliable*: If $R_{\rm best}$ can satisfy the availability requirement of request t, use $R_{\rm best}$; use $\rm RP_{best}$ otherwise.
- *Just-above-threshold*: Choose the route or route-pair with minimal availability in Q_t to carry request t.
- *Minimal-cost*: Choose the route or route-pair with minimal cost in Q_t to carry request t.

After route selection, a connection can be either unprotected or dedicated-protected. In order to further reduce network cost without sacrificing service availability, we can downgrade a dedicated-path-protected connection's protection service to shared-path protection as long as the availability requirements of this connection and of all the connections in its sharing group can still be met. Algorithm 1 describes how to assign wavelengths to connection t's backup links (after the route is picked) while the sharability is optimized without downgrading t's availability below the required value. After the backup wavelength is fixed and sharing group is identified, we can compute the connection's availability according to (6).

An important property of the shared-path protection scheme used in Algorithm 1 is that backup sharing is allowed only when the service availabilities of connections that participate in the sharing can still be met. Define *sharing degree* of a connection as number of connections that share backup resources with this connection. Using Algorithm 1, we consciously control the sharing degree of each connection so that network resources are utilized more intelligently but connections still meet their availability requirements. We call this sharing SLA-constrained sharing. The relationship between sharing degree and service availability provided by shared-path protection has been studied in [23]. The authors find that dual-failure restorability of shared-path protection is affected by the sharing degree so they provide methods to optimize the capacity requirements of shared-path protection with explicit limits on the sharing degree. However, in our approach (i.e., Algorithm 1), we do not place explicit limits on the sharing degree. Instead, the degree is automatically controlled by the availability requirements and the availabilities of connections in the sharing group, which provides more flexibility.

Algorithm 1: SLA-Constrained Sharing Algorithm (SCSA)

1) For each backup link l_i of t, check every existing backup wavelengths w_j on l_i for the following two conditions (S_t is empty initially):

a) Sharing possibility: Let $U(w_j, l_i)$ contain all the connections that have been protected by w_j on link l_i. Check whether t can share w_j with connections in $U(w_j, l_i)$ under SRLG constraint;

b) Availability constraints: Re-compute the availabilities of t and the connections in $U(w_j, l_i)$. Check whether their availability requirements can still be met.

2) Assign the lowest-numbered wavelength (say w_x) to connection t for link l_i if both of the two conditions can be satisfied; then, update $S_t(S_t = S_t \bigcup U(w_x, l_i))$ and for each connection in $U(w_x, l_i)$, put t into its sharing group; assign a new wavelength to t for link l_i if none of the existing backup wavelengths is qualified.

V. Illustrative Numerical Results

A. Verification of Availability Analysis

We verify the availability analysis for a shared-path-protected connection through simulations. The US nationwide network shown in Fig. 3 is used as a sample topology in our study. It has 26 nodes and 80 unidirectional links. Each edge in Fig. 3 is composed of two unidirectional links, one in each direction. The number next to each edge shows the lengths for the links in both directions. The number next to each node is the node id. The average failure rate is normalized in the unit of FIT. For illustration purposes, we assume that the fiber-failure rate depends on the fiber length in this verification simulation. Failures occur independently on each fiber link following a Poisson process. Failure repair time (or holding time) follows a negative exponential distribution with a mean value of 12 hours (see Table I). We assume that the failure repair-time distribution is universal for each link.

The connection request set T has 1000 connections, which are randomly generated and uniformly distributed among all node pairs, and each of them requires full capacity of a wavelength channel. As an example, each connection t is assumed to have infinite holding time, and its routing is fixed as the shortest path-pair by hop distance computed using Suurballe algorithm, where the shorter one is used as the primary path and the other one is the backup path. All connections are shared-path protected. The wavelength assignment is first fit for both primary and backup paths. In this configuration, the average size of the sharing group for a connection is 7 and, in the maximal case, 31 connections share backup resources with one connection. When a backup wavelength is released by a connection, it will be used in a FFFS manner to recover other failed connections that share this backup wavelength, i.e., this is a reverting, active-recovery, resource-locked system.

TABLE II
DIFFERENCE (ERROR%) BETWEEN SIMULATED AND THEORETICALLY-COMPUTED CONNECTION AVAILABILITIES FOR DIFFERENT FIT VALUES

FIT	400	800	2000	4000	6000	8000
TA-C	0.999988	0.999951	0.999702	0.998847	0.997485	0.995665
Error%	0.00026	0.00109	0.00645	0.02493	0.05326	0.09096
95% Confidence Interval	[0.00024, 0.00029]	[0.00100, 0.00118]	[0.00593, 0.00697]	[0.02299, 0.02687]	[0.04923, 0.05728]	[0.08430, 0.09761]

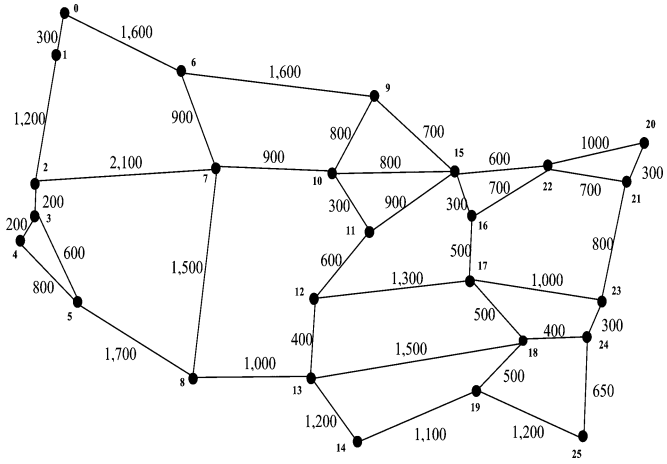

Fig. 3. A sample network topology.

The simulated availabilities and the theoretically-computed availabilities for primary paths (connections) are denoted as SA-P and TA-P (SA-C and TA-C), respectively. TA-C is computed according to (6) with $B = 10$. We find that SA-P perfectly match TA-P in our results. Table II shows the difference between SA-C and TA-C (denoted as Error% and computed as $|\mathrm{SA-C-TA-C}|/\mathrm{SA-C}$) for different FIT values, averaged over all connections. The FIT value is chosen such that TA-C is in the range from 0.999988 to 0.995665, which covers the availabilities most customers are interested in. The last row in Table II is the 95% confidence interval for Error%. We observe that the Error% is small when FIT is small but increases a little when FIT increases. However, it is only 0.09096% when TA-C is as low as 99.5665%, which indicates very good accuracy of our analytical model.

B. Results From Provisioning Strategies—ILP Approaches

The network shown in Fig. 3 is used as a sample topology. To incorporate the different rates of fiber cuts (e.g., due to different fiber types, construction areas, etc.), a more realistic model to estimate link availability needs to be obtained from the network operator based on their network-outage statistics. For illustration purposes, in what follows, the availability of each link is a pre-assigned value which could be 0.99, 0.999, or 0.9999 with equal probability. To make results reproducible, the exact values of the link availabilities have been given in [27]. The same traffic demand set T is used as in Section V-A. The availability requirements of the requests are uniformly distributed among five classes: 0.98, 0.99, 0.995, 0.997, or 0.999[4], which are referred to as Class C_1 to Class C_5, respectively.

TABLE III
RESULTS FROM ILP APPROACHES FOR FIVE PROVISIONING SCHEMES

	W	ASR	W-Links
Scheme I	87	30.2%	3361
Scheme II	174	98.8%	8442
Scheme III	163	98.6%	6707
Scheme IV	163	99.9%	7048
Scheme V	163	100%	6720

Table III compares the performance of different ILP-based provisioning schemes in terms of the number of wavelengths needed (W), connection availability satisfaction rate (ASR), and total wavelength links (W-Links). ASR represents the fraction of connections whose availability requirements have been satisfied through different schemes. As we have mentioned in Section IV-B, to avoid over-utilizing or congesting a link, the number of wavelengths on a link needs to be constrained. In each scheme, we first determine the minimal number of wavelengths (denoted as W) through which all the connection requests can be carried. We achieve this by simply setting $W_{mn}(\forall m, n)$ to be an initial (high) value, and then reducing the value by one if all the connections can be set up in the optimization. We repeat this process until some connections cannot be set up. Then, W is fixed as the value in previous loop. In this case, all the connections can be carried and there is no blocking. We calculate the total number of consumed wavelength-fiber links (denoted as W-Links) in each scheme given that the number of wavelengths on each link is constrained by W. Thus, in our work, the performance of a scheme is demonstrated by jointly considering both W and W-Links. We can also simply fix W as a large value so that all the connections can be set up in all schemes, and then compare their performance by only analyzing W-Links. However, for such a static connection-provisioning problem, the number of wavelengths on a link should also be optimized to avoid over-utilizing or congesting links.

We compare the performance of five different ILP-based provisioning schemes. They are described as follows.

- In Scheme I, all connections are provisioned without any protection, and network resources are optimized without any connection-availability consideration.
- In Scheme II, all connections are provisioned with dedicated-path protection, and the network resources are optimized without any availability consideration.

[4]The numbers shown here are for illustration purposes. When requests arrive with very high availability requirements, e.g., 0.9999 or 0.99999, the provisioning approaches may need to be extended to provide multiple backup paths to connections.

TABLE IV
STATISTICAL RESULTS FOR SCHEME III

	# of Conn.	W	# of Unsatisfied Conn.	W-Links
T_1	453	73	0	1455
T_2	547	90	14	5252

- In Scheme III, connections are classified into T_1 and T_2. Connections of T_1 are first provisioned using the ILP approach in Section IV-B. Connections of T_2 are then provisioned using the ILP approximation (in Section IV-C) with Objective A (i.e., (15)) by taking into account the connections in T_1 as existing connections. (Note that, for a general provisioning problem, solving which ILP first may affect the result but the difference is negligible in our numerical results.)
- Scheme IV is similar to Scheme III except that, for connections in T_2, Objective B (i.e., (21)) is used.
- Scheme V is a variation of Schemes III and IV, and it will be explained below.

We observe from Table III that Scheme I consumes the least amount of resources compared with the other schemes. But, in Scheme I, only 30.2% of the connections can meet their required availabilities. By providing dedicated-path protection to all connections, Scheme II can significantly improve the connection availability satisfaction rate (ASR); however, it also consumes a large amount of resources. One can also observe that there are still some connections whose availability requirements are not satisfied in Scheme II even though dedicated-path protection is provided to every connection. This is because, for these connections, the primary and the backup paths are the most resource-efficient path pair but they may not be reliable enough.

Through connection classification and traffic optimization, both Schemes III and IV jointly optimize ASR and resource usage. Schemes III and IV use less wavelength channels and around 20% less W-Links compared with Scheme II. Table IV shows statistical information on connection classification, resource usage, and service satisfaction for Scheme III. We observe that all connections in T_1 and most of the connections in T_2 (except 14 connections) receive the required services, which leads to Scheme III's 98.6% ASR, shown in Table III.

Compared with Scheme III, Scheme IV further improves ASR by consuming a little more network resources (i.e., W-Links). Based on this observation, we develop another approach, Scheme V, which can be viewed as a joint procedure of Schemes III and IV. In Scheme V, instead of applying the optimization objective B ((21)) to all the connections in T_2, we only apply it to the 14 availability-unsatisfied connections in Scheme III. By consuming 13 more W-Links (but no more wavelength channels), Scheme V can achieve 100% ASR compared to Scheme III.

However, although we have proposed the algorithm for finding the most-reliable path (MRP) for a connection, the problem of finding the pair of link-disjoint paths for a connection with the highest overall availability is expected to be NP-complete; and the mathematical formulation for the optimization problem with availability constraints is shown to be nonlinear. Therefore, the proposed schemes—Schemes III, IV, and V—are approximation schemes. However, we can expect that they can provide high ASR as protection-sensitive connections are all dedicated protected in the proposed schemes. We have tried other network topologies and different traffic demands, and Schemes III, IV, and V constantly demonstrate better performance in both W and W-Links compared to Scheme II. However, it is hard to predict which scheme (among Schemes II, III, IV, and V) will perform best in terms of ASR as none of them has availability constraints. Even though the ILP approaches studied here cannot provide an optimal solution, they can help us understand the property of the problem and they can be used to effectively provision a set of given traffic demands.

TABLE V
RESULTS FROM HEURISTIC ALGORITHMS WITHOUT SHARING

	W	ASR	W-Links
Iteratively-select	205	100%	6594
Most-reliable	339	100%	7289
Just-above-threshold	216	100%	7132
Minimal-cost	212	100%	6589

In the following section, we first show the performance of the heuristics without allowing backup resource sharing to compare to the ILP approaches. Then, the SLA-constrained sharing is incorporated into the differentiated service model and its performance is compared to a general shared-path-protection scheme without availability constraints.

C. Results From Provisioning Strategies—Heuristics

1) Without Allowing Backup Resource Sharing: Table V shows the performance (W, W-Link, and ASR) of the heuristics. In each heuristic, W is also equal to the minimal number of wavelength channels through which all the requests can be carried, as described in Section V-B. Then, we obtain the minimal W-Link used in each scheme given that the number of wavelengths on each link is constrained by W. Again, the performance of each heuristic is demonstrated by jointly considering both W and W-Links. One may notice that the order of routing connections will affect W-Link in heuristics *Most-reliable*, *Just-above-threshold*, and *Minimal-cost* as demands are routed sequentially. So, we tried a large number of different sequences and picked the solution with minimal W-Link. For each demand t, Q_t contains the candidate routes that satisfy A'_t and have available resources on links. Then, the best route in Q_t is chosen according to the policy in each heuristic.

We observe that all heuristics can provide 100% ASR because the route for request t is selected from Q_t, in which all routes can satisfy A'_t. We also observe that *Iteratively-select* demonstrates good performance if jointly considering both W and W-Links compared with other heuristics, and its performance is comparable to that of Scheme III in the ILP approaches. This is because the *Iteratively-select* algorithm employs a simple but effective back-tracking property. Please note that other sophisticated approximation algorithms, e.g., simulated annealing, genetic algorithm, etc., may also be used to further improve the

TABLE VI
COMPARING SLA-CONSTRAINED SHARING TO GENERAL SHARED-PATH PROTECTION (WITH ITERATIVELY-SELECT ALGORITHM)

	W	ASR	W-Links
S_1	205	100%	6594
S_2	194	100%	5476
S_3	194	93.1%	5461
S_4	197	92.7%	6109

TABLE VII
PERCENTAGE OF CONNECTIONS USING EACH PROTECTION SCHEME FOR EACH SERVICE CLASS IN SLA-CONSTRAINED SHARING SCHEME (S_2)

	Overall	Class C_1	Class C_2	Class C_3	Class C_4	Class C_5
Unprotected	44.9%	95.18%	51.21%	38.5%	32.39%	9.84%
Dedicated-path protected	3.3%	0.54%	0.97%	6%	4.70%	4.15%
Shared-path protected	51.8%	4.28%	47.82%	55.5%	62.91%	86.01%

overall performance. It is straightforward to see that choosing the routes or route-pairs with less cost would help reduce the overall cost; hence, *Minimal-cost* consumes less resources than *Most-reliable* and *Just-above-threshold* since it always chooses the candidate with minimal cost from Q_t. Again, heuristics are fixed-alternate-routing based so their performances are not as good as those of the ILP approaches where routing is not limited by candidate routes.

2) With Backup Resource Sharing: Tables VI and VII show the overall performance when we incorporate backup resource sharing into the service model. These results are shown here only for the *Iteratively-select* algorithm. Other heuristics show similar performance trends; hence, they are not included here.

What we support is SLA-constrained sharing where backup sharing is allowed only when connections' service availabilities can still be met. In the general version of shared-path protection, availability is not a concern and backup sharing is allowed as long as SRLG constraints are satisfied. It is possible that the service availabilities may also be satisfied without consuming too much resources if we incorporate general shared-path protection into the service model instead of SLA-constrained sharing. Thus, we compare SLA-constrained sharing to general shared-path protection in Table VI.

Table VI compares the performance of four different schemes. In Scheme S_1, differentiated services are provided (according to their availability requirements) without allowing any backup sharing, i.e., connections are either unprotected or dedicated-path protected. In Scheme S_2, differentiated services are provided with SLA-constrained sharing, i.e., a connection can be shared-path protected through Algorithm 1 (SCSA). Scheme S_3 is similar to S_2 except that general shared-path protection is offered instead of SLA-constrained shared-path protection. This means that sharing is allowed without checking the availability constraints defined in Algorithm 1, i.e., backup resource sharing is allowed as long as SRLG constraints are met. In Scheme S_4, uniform protection service (general shared-path protection) is offered to all connections without considering availability requirements.

As we expect, the network performance can be significantly improved after incorporating shared-path protection (either SLA-constrained sharing or general sharing). We can also observe that providing uniform sharing service consumes more network resources than the differentiated service schemes with either SLA-constrained sharing or general sharing. This is because, in the differentiated-service schemes, protection services are only provided on an as-needed basis according to service-availability requirements.

Comparing the sharing schemes S_2 and S_3, it is clear that, by employing a little more resources, the SLA-constrained sharing scheme (S_2) can significantly improve the ASR, from 93.1% to 100% in this case. The results from Schemes S_3 and S_4 indicate that general sharing, which is unaware of connection availability, could lead to a certain amount of service-quality degradation because the sharing degree is not carefully controlled. To conclude, we find that providing SLA-constrained shared-path protection in the service model can cost-effectively provide availability guarantee.

Table VII shows how protection services are differentiated for connections in each service class in the SLA-constrained sharing scheme (S_2). We observe that, overall, 44.9% connections are unprotected, 51.8% connections are shared-path protected, and only 3.3% connections are dedicated-path protected. We also find that more connections are protected by dedicated-path or shared-path protection with the service availability requirement becoming more and more stringent, e.g., only 4.28% of Class C_1 connections are shared-path protected while this percentage increases to 86.01% for Class C_5 connections. This shows that the proposed framework on differentiated services can provide an appropriate protection service to a request according to its service requirement.

VI. CONCLUSION

We presented a novel connection-provisioning framework which can cost-effectively provide differentiated protection services according to customers' availability requirements. The framework consisted of two parts: 1) theoretical availability analysis for a WDM mesh network under different protection schemes; and 2) ILP and heuristic-based connection-provisioning approaches.

We proposed a new sharing concept—SLA-constrained sharing where backup sharing is allowed only when connections' service availabilities can still be met. Through numerical examples, we found that, by employing a little more resources, the SLA-constrained sharing scheme can significantly improve the availability satisfaction rate (ASR). Our results also indicated that general shared-path protection, which is unaware of connection availability, could lead to a certain amount of service-quality degradation because the sharing degree is not carefully controlled.

ACKNOWLEDGMENT

The authors gratefully acknowledge the helpful comments from the editor and the reviewers, which significantly improved the paper.

REFERENCES

[1] B. Mukherjee, *Optical WDM Networks*. New York: Springer, 2006.

[2] B. Mukherjee, "WDM optical networks: Progress and challenges," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 10, pp. 1810–1824, Oct. 2000.

[3] S. Ramamurthy and B. Mukherjee, "Survivable WDM mesh networks, part I—Protection," in *Proc. IEEE INFOCOM'99*, New York, NY, Mar. 1999, vol. 2, pp. 744–751.

[4] S. Ramamurthy, L. Sahasrabuddhe, and B. Mukherjee, "Survivable WDM mesh networks," *J. Lightwave Technol.*, vol. 21, no. 4, pp. 870–883, Apr. 2003.

[5] O. Gerstel and R. Ramaswami, "Optical layer survivability: A services perspective," *IEEE Commun. Mag.*, vol. 38, no. 3, pp. 104–113, Mar. 2000.

[6] W. Wen, B. Mukherjee, and S. J. B. Yoo, "QoS based protection in MPLS controlled WDM mesh networks," in *Photon. Network Commun.*, Jul. 2002, vol. 4, pp. 297–320.

[7] H. Zang and B. Mukherjee, "Connection management for survivable wavelength-routed WDM mesh networks," *SPIE Opt. Networks Mag.*, vol. 2, no. 4, pp. 17–28, Jul. 2001.

[8] G. Mohan, C. S. R. Murthy, and A. K. Somani, "Efficient algorithms for routing dependable connections in WDM optical networks," *IEEE/ACM Trans. Networking*, vol. 9, no. 5, pp. 553–566, Oct. 2001.

[9] M. Clouqueur and W. D. Grover, "Availability analysis of span-restorable mesh networks," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 4, pp. 810–821, May 2002.

[10] W. D. Grover, "High availability path design in ring-based optical networks," *IEEE/ACM Trans. Networking*, vol. 7, no. 4, pp. 558–574, Aug. 1999.

[11] D. A. Schupke, "Reliability models of WDM self-healing rings," in *Proc. Design of Reliable Communication Networks (DRCN) 2000*, Apr. 2000.

[12] A. Fumagalli and M. Tacca, "Optimal design of optical ring networks with differentiated reliability (DIR)," in *Proc. Int. Workshop on QoS in Multiservice IP Networks*, Jan. 2001, pp. 299–313.

[13] A. Fumagalli and M. Tacca, "Differentiated reliability (DIR) in WDM ring without wavelength converters," in *Proc. ICC'2001*, Jun. 2001, pp. 2887–2891.

[14] M. To and P. Neusy, "Unavailability analysis of long-haul networks," *IEEE J. Select. Areas Commun.*, vol. 12, no. 1, pp. 100–109, Jan. 1994.

[15] M. Clouqueur and W. D. Grover, "Computational and design studies on the unavailability of mesh-restorable networks," in *Proc. Design of Reliable Communication Networks (DRCN) 2000*, Apr. 2000, pp. 181–186.

[16] D. Zhou and S. Subramaniam, "Survivability in optical networks," *IEEE Network*, vol. 14, no. 6, pp. 16–23, Nov./Dec. 2000.

[17] A. Fumagalli, A. Paradisi, S. M. Rossi, and M. Tacca, "Differentiated reliability (DIR) in mesh networks with shared path protection: Theoretical and experimental results," in *Proc. OFC'2002*, Mar. 2002, pp. 490–492.

[18] M. Tacca, A. Fumagalli, A. Paradisi, F. Unghvary, K. Gadhiraju, S. Lakshmanan, S. M. Rossi, A. de Campos Sachs, and D. S. Shah, "Differentiated reliability in optical networks: Theoretical and practical results," *J. Lightwave Technol.*, vol. 21, no. 11, pp. 2576–2586, Nov. 2003.

[19] A. Hac, "Improving reliability through architecture partitioning in telecommunication networks," *IEEE J. Sel. Areas Commun.*, vol. 12, no. 1, pp. 193–204, Jan. 1994.

[20] A. A. Akyamac, S. Sengupta, J. Labourdette, S. Chaudhuri, and S. French, "Reliability in single domain vs. multi domain optical mesh networks," in *Proc. National Fiber Optic Engineers Conf.*, Sep. 2002, pp. 240–249.

[21] D. Arci, G. Maier, A. Pattavina, D. Petecchi, and M. Tornatore, "Availability models for protection techniques in WDM networks," in *Proc. Design of Reliable Communication Networks (DRCN) 2003*, Oct. 2003, pp. 158–166.

[22] G. Willems, P. Arijs, W. V. Parys, and P. Demeester, "Capacity vs. availability trade-offs in mesh-restorable WDM networks," in *Proc. Design of Reliable Communication Networks (DRCN) 2001*, Oct. 2001.

[23] J. Doucette, M. Clouqueur, and W. D. Grover, "On the availability and capacity requirements of shared backup path-protected mesh networks," *SPIE Optical Networks Mag.*, vol. 4, no. 6, pp. 29–44, Nov. 2003.

[24] K. S. Trivedi, *Probability and Statistics with Reliability, Queuing, and Computer Science Applications*. Englewood Cliffs, NJ: Prentice-Hall, 1982.

[25] Y. Liu, D. Tipper, and P. Siripongwutikorn, "Approximating optimal spare capacity allocation by successive survivable routing," in *Proc. IEEE INFOCOM'2001*, Apr. 2001, vol. 2, pp. 699–708.

[26] X. Su and C.-F. Su, "An online distributed protection algorithm in WDM networks," in *Proc. ICC'2001*, Jun. 2001, vol. 5, pp. 1571–1575.

[27] J. Zhang, "Architectures and algorithms for fault management in optical WDM networks," Ph.D. dissertation, Univ. of California, Davis, CA, 2005.

[28] S. Ramamurthy and B. Mukherjee, "Fixed-alternate routing and wavelength conversion in wavelength-routed optical networks," *IEEE/ACM Trans. Networking*, vol. 10, no. 3, pp. 351–367, Jun. 2002.

[29] R. Bhandari, *Survivable Networks: Algorithms for Diverse Routing*. Boston, MA: Kluwer Academic, 1999.

[30] J. W. Suurballe and R. E. Tarjan, "A quick method for finding shortest pairs of disjoint paths," *Networks*, vol. 14, pp. 325–336, 1984.

[31] J. Y. Yen, "Finding the K shortest loopless paths in a network," *Manage. Sci.*, vol. 11, pp. 712–716, Jul. 1971.

Jing Zhang (S'02) received the B.S. degree from Peking University, Beijing, China, in 1998, and the M.S. and Ph.D. degrees from the University of California, Davis, in December 2001 and January 2005, respectively.

Her research interests include fault management, algorithm design, performance evaluation, and reliability analysis in communication networks. Currently, she is a Performance Engineer at Sun Microsystems, Inc., Menlo Park, CA.

Keyao Zhu (S'98–M'05) received the B.S. degree from Peking University, Beijing, China, in 1998, and the M.S. and Ph.D. degrees from the University of California, Davis, in July 2000 and September 2003, respectively.

From August 2003 to September 2004, he was with Research and Innovation, Alcatel Shanghai Bell. Currently, he is a Software Engineer with Brion Technology Inc., Santa Clara, CA.

Dr. Zhu has served as a Technical Committee Member of ICC'04 and ICC'05. In June 2004, he received the Zuhair A. Munir Award for the Best Doctoral Dissertation in the College of Engineering, University of California, Davis, for his research on WDM optical networks.

Hui Zang (S'97–M'02) received the B.S. degree in computer science from Tsinghua University, Beijing, China, and the M.S. and Ph.D. degrees in computer science from the University of California, Davis.

In 2000, she joined Sprint Advanced Technology Laboratories, Burlingame, CA, where she is a Research Scientist. She is the author of the book *WDM Mesh Networks: Management and Survivability* (Kluwer Academic, 2002). She has published over 30 conference papers and journal articles, and currently has one U.S. patent granted and five pending in the field of networking and communications. Her research interests include performance and security issues in wireless, IP, and optical networks.

Dr. Zang serves or has served as technical committee members of a number of conferences. She was one of the guest editors of *IEEE Network* special issue on "Traffic Engineering in Optical Networks."

Norman S. Matloff received the Ph.D. degree in theoretical mathematics from the University of California, Los Angeles, in 1975.

He is a Professor of computer science at the University of California, Davis, and was formerly a Professor of mathematics and statistics at the same university. His current research interests are performance modeling and data mining. He is a former appointed member of IFIP Working Group 11.3, an international committee concerned with database software security. His work on optical multiprocessor computers was awarded a U.S. patent. He has been a member of the ACM since 1982.

Biswanath Mukherjee (S'82–M'87–F'07) received the B.Tech. (Hons) degree from the Indian Institute of Technology, Kharagpur, India, in 1980, and the Ph.D. degree from the University of Washington, Seattle, in June 1987. At Washington, he held a GTE Teaching Fellowship and a General Electric Foundation Fellowship.

In July 1987, he joined the University of California, Davis, where he has been a Professor of computer science since July 1995, and Chairman of Computer Science since September 1997. His research interests include lightwave networks, network security, and wireless networks.

Prof. Mukherjee is co-winner of paper awards presented at the 1991 and the 1994 National Computer Security Conferences. He serves on the editorial boards of the IEEE/ACM TRANSACTIONS ON NETWORKING, *IEEE Network*, *ACM/Baltzer Wireless Information Networks (WINET)*, *Journal of High-Speed Networks*, *Photonic Network Communications*, and *Optical Network Magazine*. He also served as Editor-at-Large for optical networking and communications for the IEEE Communications Society. He served as the Technical Program Chair of the IEEE INFOCOM'96 conference. He is the author of the textbook *Optical Communication Networks* (McGraw-Hill, 1997), a book which received the Association of American Publishers, Inc.'s 1997 Honorable Mention in Computer Science.