Differentiated Quality-of-Protection Provisioning in Optical/MPLS Networks*

Canhui (Sam) Ou and Biswanath Mukherjee

Department of Computer Science, University of California, Davis, CA 95616, USA
{ouc, mukherje}@cs.ucdavis.edu

Abstract. This paper investigates the problem of dynamic survivable lightpath provisioning against single node/link failures in optical mesh networks employing wavelength-division multiplexing (WDM). We present a new approach to provisioning lightpath requests according to their differentiated quality-of-protection (QoP) requirements. We focus on one of the most important QoP parameters, namely protection-switching time, since lightpath requests may have differentiated protection-switching-time requirements. For example, lightpaths carrying voice traffic may require 50-ms protection while lightpaths carrying data traffic may have a wide range of requirements. Numerical results show that, compared to shared-path protection, our approach achieves significant performance gain which leads to remarkable reduction in blocking probability. While our focus is on optical WDM network, the basic ideas of our approaches can be applied to multiprotocol label switching (MPLS) networks with appropriate variations, e.g., differentiated bandwidth granularities.

1 Introduction

In a wavelength-routed optical network, the failure of a network element can cause the failure of several lightpaths, thereby leading to large data and revenue loss. Protection, a proactive procedure in which spare capacity is reserved during lightpath setup [1–5], is essential for recovering from such failures in a short time period, e.g. 50 ms. Protection schemes can be classified by the type of routing used (link-based versus path-based) and by the type of resource sharing (dedicated versus shared). A path carrying traffic during normal operation is known as a *working* path[1]. When a working path fails, the lightpath is rerouted over a *backup* path. High bandwidth efficiency and short protection-switching time are two of the most important features of a protection scheme [6], where protection-switching time for a lightpath is the time period the lightpath takes to properly signal/configure the nodes along the backup path before switching traffic to the backup path after a failure occurs on the working path [5].

We consider the problem of dynamic survivable lightpath provisioning against single node (crossconnect) and single link (fiber) failures. Specifically, we focus on

---

[1] Working path is also referred to as primary path, active path, and service path.

shared protection (because of its desirable resource efficiency) with the assumptions that existing lightpaths cannot be disturbed and no knowledge of future arrivals is available at the time of provisioning the current lightpath request. While we consider full wavelength-convertible networks here, the extension to the wavelength-continuous case is straightforward.

Much work has been conducted on dynamic shared protection [7–10] in optical WDM networks and on dynamic routing of restorable bandwidth-guaranteed connections in MPLS networks [11–14]. A widely considered approach, called shared-path protection [5], is bandwidth efficient due to backup sharing. Consequently, how to increase backup sharing based on different cost models and route-computation techniques is of particular interest and has been reported in [12, 15–20]. The complexity of shared-path protection is high as shown in [9, 21] that it is NP-complete to find a working path and a backup path for a new lightpath request when backup sharing with existing backup paths is allowed. As a result, practical heuristics are usually employed.

One possible limitation of shared-path protection is that backup paths may sometimes become longer due to backup sharing [15]. Consequently, protection-switching time may increase because of longer backup paths. The relation between backup sharing and backup-path hop distance for path protection have been shown to be that one trades off another in [10, 15].

Furthermore, lightpath requests may have differentiated protection-switching-time requirements. For example, lightpaths carrying voice traffic may require 50 ms protection while lightpaths carrying data traffic may require a wide range of protection-switching-time requirements. Due to the path-wise node-/link- disjoint nature of path protection, shared-path protection may not provision lightpath requests according to their protection-switching-time requirements effectively in practical-sized networks [22, 23]. Clearly, proper mechanisms are needed to provision such lightpath requests in a resource-efficient manner.

Motivated by the above considerations, we present a new and effective approach to provisioning lightpath requests according to their protection-switching-time requirements while taking into account backup sharing. While our focus is on optical WDM network, in which the bandwidth requirement of a lightpath request is one wavelength, our approaches can also be directly applied to MPLS networks for provisioning restorable, bandwidth-guaranteed connections of differentiated bandwidth granularities with appropriate adjustments.

2 Provisioning Lightpaths of Differentiated Quality-of-Protection (QoP) Requirements

We focus on one of the most important QoP parameters, namely protection-switching time. The protection-switching time of a shared-path protected lightpath can be based on the hop count of the working/backup paths [5, 23, 24]. Therefore, we consider QoP in terms of hop count.

Below, we argue that new mechanisms are needed to provision lightpaths of differentiated QoP requirements and present a new approach to achieve the QoP.

2.1   Motivation

Lightpath requests may have differentiated protection-switching-time requirements. For example, lightpaths carrying voice traffic may require 50 ms protection while lightpaths carrying data traffic may need a wide range of protection-switching-time requirements. While some mission-critical lightpath requests can be dedicate protected, it is not economically viable to provide dedicated protection to each lightpath request due to its excessive resource requirement.

Below, we show that shared-path protection cannot provide the desired level of protection-switching time either. Let us consider a simple case in which the backup-path hop count of any lightpath cannot exceed a constant H_b (ignoring the constraint on working path for now). A lightpath request will be blocked if the computed backup path is longer than H_b hops. We simulate a dynamic network environment with the assumptions that the lightpath-arrival process is Poisson and the lightpath-holding time follows a negative exponential distribution. In every experiment, 10^6 lightpath requests are simulated; they are uniformly distributed among all node pairs; average lightpath-holding time is normalized to unity; the cost of any link is unity; and our example network topology with 16 wavelengths per fiber is shown in Fig. 1.

A commonly used route-computation heuristic for shared-path protection is as follows: compute a least-cost path as the working path and then compute as the backup path a link (or node) disjoint path of least additional cost [15]. In general, we can compute K candidate working paths, repeat the above heuristic for each candidate working path, and select the working and backup path pair of minimum cost.

Figure 2(a) plots the blocking probability of shared-path protection for $H_b = 6$ for the network shown in Fig. 1 with different values of ϵ. Figure 2(a) confirms the conclusion in [15, 10] that a larger value of ϵ leads to shorter backup path but decreased backup sharing, and a smaller value of ϵ leads to increased backup sharing but longer backup path. Please note that $H_b = 6$ for this network is reasonably large since the average hop distance is 2.99. However, regardless of the values of ϵ, the blocking probability in Fig. 2(a) is quite high. The main reason is that some lightpath requests are blocked because their backup paths span more than $H_b = 6$ hops. Figure 2(b) shows the impact of H_b on shared-path protection. While the blocking probability drops significantly as H_b increases, the blocking for $H_b = 7$, which is quite large, is still unacceptable.

As network size grows, it is clear that shared-path protection cannot achieve reasonable blocking for practical values of H_b due to its fundamental limitation: the backup path has to be end-to-end node-/link- disjoint to the working path.

Consequently, new mechanisms, which can relax the end-to-end node-/link disjointness, are needed. One such mechanism is segment protection, as described in the next section.

2.2   Segment Protection (SP)

Various forms of segment protection (SP) have been reported in [25–28]. The approaches proposed in [25] addressed single-link failures by dividing a work-
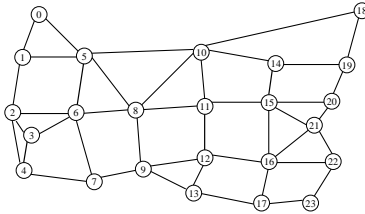
Fig. 1. A representative topology whose average hop distance is 2.99 and average nodal degree is 3.58.

(a) $K = 2$ and $H_b = 6$. (b) $K = 2$ and $\epsilon = 0.99$.
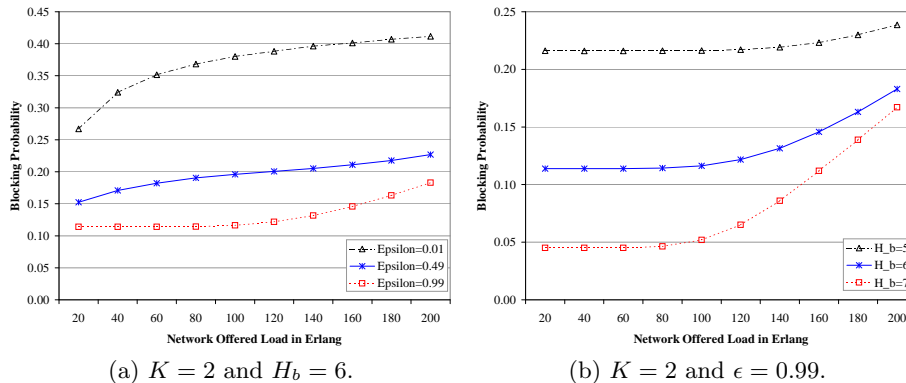
Fig. 2. Blocking probability of shared-path protection.

ing path into a sequence of non-overlapping segments and protecting each such segment individually. Nodes were assumed to be robust in these approaches as consecutive non-overlapping segments share the same node failure.

Even though node failures are not as often as link failures, they need to be carefully treated because the impact of node failures is much more disastrous than that of link failures. The work in [26–28] addressed single-node/link failures by dividing a working path into a sequence of overlapping segments and protecting each such segment separately. As shown in Fig. 3, the lightpath from node s to node d is partitioned into two overlapping segments: one with working segment $\langle s, i, j, u \rangle$ and backup segment $\langle s, u \rangle$; another with working segment $\langle j, u, v, d \rangle$ and backup segment $\langle j, d \rangle$. (There are other nodes along the two backup segments $\langle s, u \rangle$ and $\langle j, d \rangle$. Those nodes are not shown in Fig. 3 to preserve clarity.)

The advantages of SP are numerous. The end-to-end protection entity is *a segment* in SP as opposed to *a path* in path protection. When a failure occurs along a working path (segment), the source node of that path (segment) switches to its backup. Since a segment is typically shorter than a path in terms of hop count, SP is expected to have shorter protection-switching time.

Meanwhile, two segments (or two lightpaths in path protection) can share backup wavelength links as long as their working segments (or working paths in path protection) do not share the same node/link failure. Since, in general, a

segment is shorter than a path, the probability of two working segments sharing the same risk is typically lower than the probability of two working paths sharing the same risk. As a result, SP can have better backup sharing compared to shared-path protection.

Furthermore, SP has more flexibility in routing compared to path protection since path protection is a special case of SP in which every lightpath has exactly one segment. Apart from these differences, it is clear that the longer the working path is, the more difficult it is to find a node-disjoint backup path [29].

Based on SP, we present below an approach, called SP_QoP, which: (1) dynamically divides a judiciously selected working path into multiple overlapped working segments, (2) computes a backup segment for each working segment, (3) accommodates backup sharing, and (3) ensures the hop-count constraint on working/backup segments.
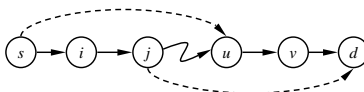
Fig. 3. Segment protection: node j and node u can be adjacent or separated by other nodes. (The solid lines from node s to node d represent the working path, and the dashed lines represent the backup segments. While only two segments are shown in these illustrations, in general, a path may employ many segments.)

2.3 SP_QoP Heuristic

Upon the arrival of a new lightpath request, the network management system needs to compute a working path l_w and a list of backup segments $\{l_b^i\}$, which divide the working path into overlapping segments $\{l_w^i\}$ such that l_w^i and l_b^i are node-/link- disjoint. New backup segments $\{l_b^i\}$ can share wavelength links with existing backup segments as well as among themselves. Unfortunately, it is NP-hard to determine if there exists an eligible solution as we have proved the NP-completeness of the existence version of shared-path-protection problem, which is a special case of SP with the number of segments being one in [9]. As a result, we resort to a heuristic.

A network is represented as a weighted, directed graph $G = (V, E, C, \lambda)$, where V is the set of nodes, E is the set of unidirectional fibers (referred to as links), $C : E \to R^+$ is the cost function for each link (where R^+ denotes the set of positive real numbers), and $\lambda : E \to Z^+$ specifies the number of wavelengths on each link (where Z^+ denotes the set of positive integers).

A conflict set is associated with a link to identify the sharing potential between backup segments[2]. The conflict set ν_e for link e defines the set of nodes

---

[2] The conflict set is similar to the conflict vector in [8], the aggregated square matrix in [17], and the "bucket" link metric in [19].

traversed by such working segments whose backup segments utilize wavelengths on link e. The conflict set ν_e for link e can be represented as an integer set, $\{\nu_e^u \mid \forall u \in V, 0 \leq \nu_e^u \leq \lambda(e)\}$, where ν_e^u specifies the number of working segments which traverse node u and are protected by link e (or, in other words, their corresponding backup segments traverse link e). The number of wavelengths reserved for backup segments on link e is thus $\nu_e^* = \max_{\forall u}\{\nu_e^u\}$. Clearly, the union of the conflict sets for all the links aggregates the per-segment-based information, and the size of the conflict set depends only on the number of nodes, not on the number of segments.

For a candidate working path l_w^k, our heuristic, called SP_QoP, performs the following recursive procedure to compute a list of eligible backup segments.

1. Starting from node s, compute a least-cost path to all the other nodes along l_w^k, where the cost function is C_1 defined in Algorithm 1.
2. Starting from node d and following the reverse direction of l_w^k, find the first node v which satisfies the constraint that the least-cost path from node s to node v is of at most H_b hops.
3. If node v is the destination node d, the heuristic succeeds and terminates; otherwise, starting from all the nodes between node s and node v (excluding nodes s and v) along l_w^k, recursively apply the above procedure. (If there is no node between node s and node v, the heuristic fails.)

Our SP_QoP heuristic is specified in Algorithm 1. For a node $u \in V$, $PC(u)$ denotes the cost of the least-cost path destined to node u; $HC(u)$ represents the hop count of the least-cost path; and $PH(u)$ records the previous hop along the least-cost path. For a path l_w, $Head(l_w)$ returns the first node along l_w.

We make the following remarks. 1) In Algorithm 1, the candidate working path l_w^k is given. This is just for the purpose of simplifying the presentation; in our implementation, we dynamically compute K candidate working paths based on Yen's K shortest loopless paths algorithm [30], execute Algorithm 1 for each candidate working path, and select the working path and the list of backup segments of minimal cost.

2) Backup sharing in this case is tricky as the situation shown in Fig. 4 can arise. In the presence of H_b, the path $\langle s, x, p, q, y, d \rangle$ is not valid when $H_b = 4$. However, the two segments $\langle s, x, p, q, u \rangle$ and $\langle j, p, q, y, d \rangle$ so formed are still valid. Our heuristic accommodates this type of backup sharing in Steps 3 and 7 of Algorithm 1 since the freshly reserved backup wavelengths for a newly computed backup segment is used for computing later backup segments for the same lightpath request.

3) Sometimes, it may be desirable that the hop count of any working segment plus the hop count of its backup segment is no more than some constant H. We can modify Step 5 to satisfy this constraint as follows. For any node v along path l_w, denote as $l_b^{i,v}$ the least-cost path destined to node v and denote as h_w^v the number of hops from $Head(l_b^{i,v})$ to node v along l_w. Starting from node d and following the reverse direction of l_w, find the first node v which satisfies the constraint $HC(v) + h_w^v \leq H$. Other constraints based on combinations of working and backup segment hop count also can be easily incorporated.

Algorithm 1 SP_QoP

*Input*: $G = (V, E, C, \lambda)$, $\nu = \{\nu_e \mid e \in E\}$, s, d, a candidate working path l_w^k
Output: a list of backup segments $\{l_b^i\}$, each of which spans no more than H_b hops and they collectively partition l_w^k into overlapped segments $\{l_w^{k,i}\}$ such that $l_w^{k,i}$ and l_b^i are node-/link- disjoint; otherwise, NULL if no such list is found.

1. $S \leftarrow \{s\}, L_b \leftarrow \phi, l_w \leftarrow l_w^k, i \leftarrow 0$
2. $V' \leftarrow V$; $\forall u \in S, PC(u) \leftarrow 0, HC(u) \leftarrow 0, PH(u) \leftarrow NULL$; $\forall u \in V \wedge u \notin S, PC(u) \leftarrow \infty, HC(u) \leftarrow \infty, PH(u) \leftarrow NULL$; $\quad i \leftarrow i + 1$
3. define link-cost function $C_1(e)$, $e \in E$, with respect to l_w:

$$C_1(e) := \begin{cases} +\infty & \text{if } l_w \text{ traverses link } e, \text{ or } \nu_e^* \text{ is equal to } \nu_e^u \text{ for some node } u \text{ along} \\ & l_w \text{ and link } e \text{ does not have any free wavelength} \\ \epsilon \times C(e) & \text{if for any node } u \ (u \neq s, d) \text{ traversed by } l_w, \nu_e^u < \nu_e^* \\ C(e) & \text{otherwise} \end{cases}$$

4. while $(V' \neq \phi)$ do {
    $u \leftarrow \arg \min_{u \in V'}\{PC(u)\}$, $V' \leftarrow V' - \{u\}$
    if $(u = Head(l_w))$ or $(l_w$ does not traverse $u)$ {
        $\forall v \in V'$, s.t. $\langle u, v \rangle \in E$
        if $PC(v) > PC(u) + C_1(\langle u, v \rangle)$ then {
            $PC(v) \leftarrow PC(u) + C_1(\langle u, v \rangle)$
            $HC(v) \leftarrow HC(u) + 1$
            $PH(v) \leftarrow u$
        } // if
    } // if
} // while
5. starting from node d and following the reverse direction of l_w, find the first node v which satisfies $HC(v) \leq H_b$
6. retrieve the least-cost path destined to node v by following $PH(v)$ and denote the path as l_b^i
7. allocate backup wavelengths along l_b^i: let l_w^i be the working segment starting from $Head(l_w)$ and ending at v along l_w (inclusively); for any link e that l_b^i traverses and for any node u along l_w^i (excluding the source and the destination nodes of l_w^i), $\nu_e^u \leftarrow \nu_e^u + 1$; if $\nu_e^u > \nu_e^*$, then reserve one more wavelength on link e and let $\nu_e^* \leftarrow \nu_e^u$
8. if v is d, then return $\{l_b^i\}$
9. $S \leftarrow$ all the nodes between $Head(l_w)$ and node v along l_w, excluding $Head(l_w)$ and node v; if S is empty, then undo any changes made to G in Step 7 and return NULL
10. $l_w \leftarrow$ the path starting from node v to node d along l_w
11. go to Step 2

Computational Complexity: The computational complexity of Algorithm 1 is $O(|V|^3 + |E|)$. In particular, the computational complexities for Steps 1-11 are $O(1)$, $O(|V|)$, $O(|E|)$, $O(|V|^3 + |E|)$, $O(|V|)$, $O(|V|)$, $O(|V|)$, $O(|E|)$, $O(|V|)$, $O(|V|)$, and $O(1)$, respectively. If we compute K candidate working paths and
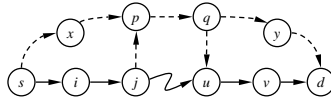
Fig. 4. Two backup segments, $\langle s, x, p, q, u \rangle$ and $\langle j, p, q, y, d \rangle$, of the same lightpath share the same wavelength link on link $\langle p, q \rangle$, assuming $H_b = 4$.

execute Algorithm 1 for each candidate working path, then the computational complexity is $O(K \cdot (|V|^3 + |E|))$.

3 Illustrative Numerical Results

We now quantitatively compare SP_QoP to shared-path protection under the simulation configuration as described in Section 2.1. For the illustrative results shown here, we use $K = 2$ as we found the performance improvement is marginal if we increase K to any larger value.

3.1 Blocking probability under different values of ϵ

Figure 5(a) plots the blocking performance for $H_b = 6$ under $\epsilon = 0.01, 0.49$, and 0.99. We observe that our SP_QoP approach has significantly lower blocking probability than shared-path protection under the same ϵ. We further observe that large values of ϵ, e.g., $\epsilon = 0.49$ or $\epsilon = 0.99$, are preferable as both SP_QoP and shared-path protection have significantly lower blocking when ϵ has a large value. Later, we shall use large values of ϵ.

   Figure 5(b) shows the *performance gain*, defined as the percentage of light-path requests which are blocked in shared-path protection but can be accepted by SP_QoP. Performance gain can be calculated as follows. Whenever shared-path protection needs to block a lightpath request, we apply SP_QoP to check whether the same lightpath request can be provisioned under the same network state (but we do not set up the lightpath request even if it can be provisioned). We observe that SP_QoP has a remarkable performance gain (over 70% across all load regions). The huge performance gain results from the fact that SP_QoP relaxes the path-wise node-/link- disjointness to segment-wise node-/link- disjointness and computes segments with respect to H_b.

3.2 Blocking probability under different values of H_b

Figure 6(a) examines the impact of H_b on both SP_QoP and shared-path protection with $\epsilon = 0.99$. We observe that: (a) SP_QoP has much lower blocking probability when load is not very high. (b) When H_b increases from 5 to 6, SP_QoP has noticeable reduction in blocking probability while the reduction is marginal when H_b further increases to 7. (c) As H_b increases, the blocking probability of shared-path protection drops significantly. However, shared-path
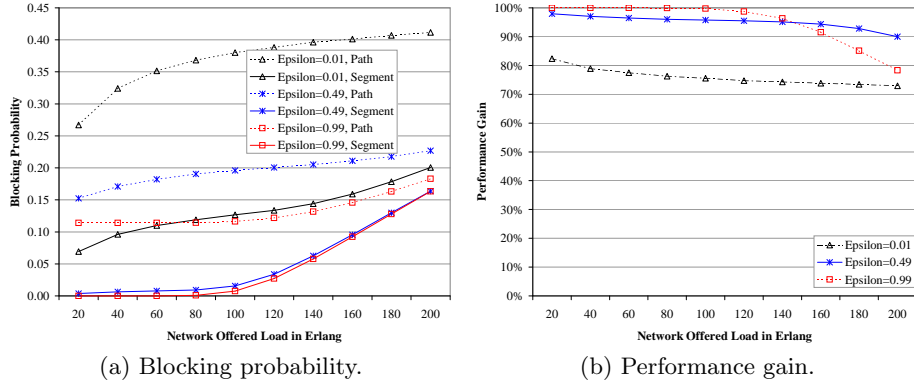
(a) Blocking probability.

(b) Performance gain.

Fig. 5. Impact of ϵ ($H_b = 6$).

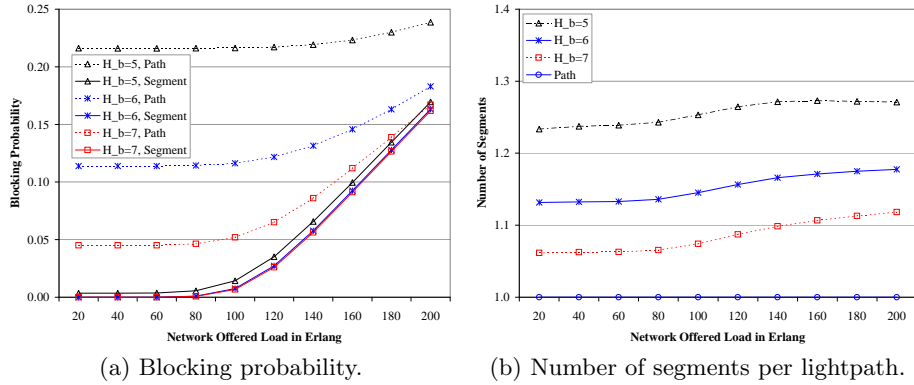(a) Blocking probability.

(b) Number of segments per lightpath.

Fig. 6. Impact of H_b ($\epsilon = 0.99$).

protection still has remarkable blocking (above 4%) even when the network offered load is low, e.g., 20 Erlangs which translates to about 8.5% average link utilization. This is due to the path-wise end-to-end node-/link- disjoint nature of shared-path protection.

Figure 6(b) shows that more segments are needed for smaller value of H_b. However, the average number of segments per lightpath is still quite low, e.g., less than 1.3, even for $H_b = 5$. This implies that the control and management overhead due to segmentation is not very significant.

3.3 Blocking probability for lightpath requests with differentiated QoP requirements

Different lightpath requests may have differentiated QoP requirements, as discussed earlier in Section 2.1. Figure 7 compares the performance of SP_QoP to shared-path protection under two types of traffic. The QoP of the lightpath re-

(a) Blocking probability.　　　　　　(b) Perforance gain.
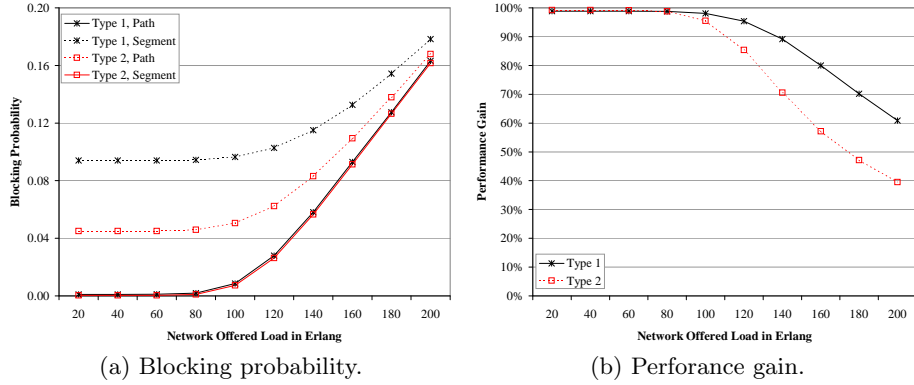
Fig. 7. Impact of differentiated QoP requirements ($\epsilon = 0.99$). In Type 1, H_b follows $5 : 6 : 7 : \infty = 30 : 20 : 10 : 40$; In Type 2, H_b follows $5 : 6 : 7 : \infty = 10 : 20 : 20 : 50$.

(a) Blocking probability ($\epsilon = 0.99$).　　(b) Performance gain ($\epsilon = 0.99$).
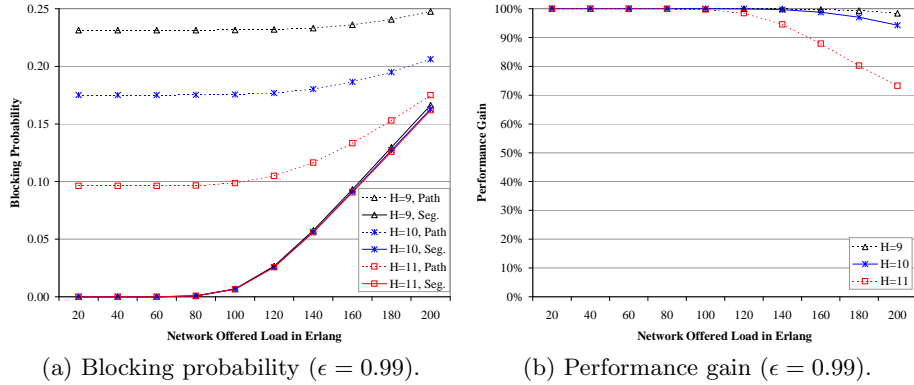
Fig. 8. Impact of H ($\epsilon = 0.99$).

quests in terms of H_b follows the distribution $5 : 6 : 7 : \infty = 30 : 20 : 10 : 40$ in Type 1 and $5 : 6 : 7 : \infty = 10 : 20 : 20 : 50$ in Type 2.

SP_QoP has much lower blocking probability than shared-path protection, as shown in Fig. 7(a). For shared-path protection, the large difference between the blocking probability for the two types of traffic implies that shared-path protection cannot effectively provision lightpath requests based on their differentiated QoP requirements. However, the difference between the blocking probability for the two types of traffic in SP_QoP is very small. This indicates that SP_QoP can properly provision lightpath requests according to their differentiated QoP requirements.

As shown in Fig. 7(b), when load is modest or low, SP_QoP achieves close to 100% performance gain; even when load is high, SP_QoP still achieves more than 35% performance gain.

3.4 Blocking probability for different values of H

Figures 8(a) and 8(b) examine the impact of H on SP_QoP and shared-path protection. The curves in Fig. 8 have similar trend to the ones in Fig. 6 and can be explained similarly. Meanwhile, since H applies to both working and backup segments, as opposed to H_b which applies only to backup segments, the H constraint is more stringent than the H_b constraint. As a result, the performance gain for different values of H is even higher, above 70% across all load regions, as shown in Fig. 8(b).

4 Conclusion

This paper considered the problem of dynamic survivable lightpath provisioning against single node/link failures in optical mesh networks. We presented a new approach to provisioning lightpath requests according to their differentiated quality-of-protection (QoP) requirements with focus on protection-switching time since lightpath requests may have differentiated protection-switching-time requirements. Numerical results showed that, compared to shared-path protection, our approach achieves significant performance gain which leads to remarkable reduction in blocking probability.

While our focus is on optical network, our approaches can also be applied to MPLS networks with appropriate variations, e.g., differentiated bandwidth granularities.

References

1. Doshi, B.T., Dravida, S., Harshavardhana, P., Hauser, O., Wang, Y.: Optical network design and restoration. Bell Labs Technical Journal **4** (1999) 58–84
2. Fumagalli, A., Tacca, M., Unghvary, F., Farago, A.: Shared path protection with differentiated reliability. In: Proc. IEEE ICC. (2002) 2157–2161
3. Lumetta, S.S., Medard, M., Tseng, Y.C.: Capacity versus robustness: a tradeoff for link restoration in mesh networks. IEEE J. Lightwave Technology **18** (2000) 1765–1775
4. Modiano, E., Narula-Tam, A.: Survivable lightpath routing: a new approach to the design of WDM-based networks. IEEE J. Selected Areas in Communications **20** (2002) 800–809
5. Ramamurthy, S., Sahasrabuddhe, L., Mukherjee, B.: Survivable WDM mesh networks. IEEE J. Lightwave Technology **21** (2003) 870–883
6. Koo, S., Subramaniam, S.: Trade-offs between speed, capacity, and restorability in optical mesh network restoration. In: Proc. OFC. (2002) 487–489
7. Hauser, O., Kodialam, M., Lakshman, T.V.: Capacity design of fast path restorable optical networks. In: Proc. IEEE INFOCOM. Volume 2. (2002) 817–826
8. Mohan, G., Murthy, C.S.R., Somani, A.K.: Efficient algorithms for routing dependable connections in WDM optical networks. IEEE/ACM Trans. Networking **9** (2001) 553–566
9. Ou, C., Zhang, J., Zang, H., Sahasrabuddhe, L., Mukherjee, B.: Near-optimal approaches for shared-path protection in WDM mesh networks. In: Proc. IEEE ICC. (2003) 1320–1324

10. Xiong, Y., Xu, D., Qiao, C.: Achieving fast and bandwidth-efficient shared-path protection. IEEE J. Lightwave Technology **21** (2003) 365–371
11. Kar, K., Kodialam, M., Lakshman, T.V.: Routing restorable bandwidth guaranteed connections using maximum 2-route flows. In: Proc. IEEE INFOCOM. (2002) 113–121
12. Kodialam, M., Lakshman, T.V.: Dynamic routing of bandwidth guaranteed tunnels with restoration. In: Proc. IEEE INFOCOM. Volume 2. (2000) 902–911
13. Kodialam, M., Lakshman, T.V.: Dynamic routing of locally restorable bandwidth guaranteed tunnels using aggregated link usage information. In: Proc. IEEE INFOCOM. Volume 1. (2001) 376–385
14. Qiao, C., Xu, D.: Distributed partial information management (DPIM) schemes for survivable networks – Part I. In: Proc. IEEE INFOCOM. (2002) 302–311
15. Bouillet, E., Labourdette, J.F., Ramamurthy, R., Chaudhuri, S.: Enhanced algorithm cost model to control tradeoffs in provisioning shared mesh restored lightpaths. In: Proc. OFC. (2002) ThW2
16. Li, G., Wang, D., Kalmanek, C., Doverspike, R.: Efficient distributed path selection for shared restoration connections. In: Proc. IEEE INFOCOM. (2002) 140–149
17. Liu, Y., Tipper, D., Siripongwutikorn, P.: Approximating optimal spare capacity allocation by successive survivable routing. In: Proc. IEEE INFOCOM. Volume 2. (2001) 699–708
18. Su, C., Su, X.: Protection path routing on WDM networks. In: Proc. OFC. Volume 2. (2001) TuO2–T1–3
19. Su, X., Su, C.: An online distributed protection algorithm in WDM networks. In: Proc. IEEE ICC. Volume 5. (2001) 1571–1575
20. Wang, H., Modiano, E., Medard, M.: Partial path protection for WDM networks: end-to-end recovery using local failure information. In: Seventh International Symposium on Computers and Communications (ISCC 2002). (2002) 719–725
21. Ellinas, G., Bouillet, E., Ramamurthy, R., Labourdette, J., Chaudhuri, S., Bala, K.: Routing and restoration architectures in mesh optical networks. SPIE Optical Networks Magazine **4** (2003) 91–106
22. Li, J., Park, H., Lee, H.: Shared sub-path protection with overlapped protection areas in WDM networks. In: Proc. OFC. (2003) FQ7
23. Ou, C., Zang, H., Mukherjee, B.: Sub-path protection for scalability and fast recovery in optical WDM mesh networks. In: Proc. OFC. (2002) 495–496
24. Assi, C., Ye, Y., Shami, A., Dixit, S., Ali, M.: Efficient path selection and fast restoration algorithms for shared restorable optical networks. In: Proc. IEEE ICC. (2003) 1412–1416
25. Anand, V., Chauhan, S., Qiao, C.: Sub-path protection: A new framework for optical layer survivability and its quantitative evaluation (Dept. of CSE, State University of New York at Buffalo, Tech. Report 2002-01, Jan. 2002)
26. Gummadi, K., Pradeep, M., Murthy, C.: An efficient primary-segmented backup scheme for dependable real-time communication in multihop networks. IEEE/ACM Trans. Networking **11** (2003) 81–94
27. Ho, P.H., Mouftah, H.: A framework for service-guaranteed shared protection in WDM mesh networks. IEEE Communication Magazine **40** (2002) 97–103
28. Ho, P.H., Mouftah, H.T.: A novel strategy for protection domain allocation in dynamic mesh WDM networks. In: Proc. IEEE ICC. (2003) CQ22–1
29. Liu, Y., Tipper, D.: Successive survivable routing for node failures. In: Proc. IEEE Globecom. (2001) 2093–2097
30. Yen, J.Y.: Finding the K shortest loopless paths in a network. Management Science (1971) 712–716