An online distributed protection algorithm for WDM networks

Xun Su, Ching-Fong Su

Abstract— **Failure protection and restoration is gaining significance as high capacity optical networks become increasingly ubiquitous. An important design issue that faces the network operator, lies in utilizing the wavelength sharing potential presented by non-concurrent failure events. In this study we propose a novel link metric and a set of routing algorithms that aim to maximally exploits these sharing opportunities. With the introduction of the light-weight aggregated link metrics termed "buckets", we provide guidance to construct the protection routes with minimal wavelength consumption. We show by simulations that our scheme significantly improves upon the baseline mechanism that lacks the proposed intelligence. We also extend the solution to a generic node-based approach, and show that further performance gain can be achieved by improving upon the way one applies our algorithms.**

I. INTRODUCTION

The popularity of the Internet has created a deluge of data traffic. It forces service providers to seriously consider new infrastructures that meet the explosive demand for bandwidth. Wavelength Division Multiplexing (WDM), which allows a single fiber to carry multiple signals simultaneously, is perceived to be a promising candidate to address the bandwidth shortage on the Internet. While the enormous amount of bandwidth provided by WDM may help alleviate the mounting pressure for higher access speed, it also makes protection/restoration a very important issue in network management. For example, current technology allows up to 128 wavelengths to be multiplexed in a single fiber, each with a data rate up to 10 Gbps. This roughly translates into millions of telephone calls on a single fiber. Hence it is easy to comprehend the catastrophic consequence a fiber cut may cause without appropriate protection mechanism in place.

Different types of protection schemes have been developed for optical networks[18]. Many existing transport networks use SONET (synchronous optical network) rings. Rings are simple topologies which contain two separate paths between any pairs of nodes, and they are resilient to any single link or node failures. Although simple and fast, the direct application of the ring architectures in WDM networks brings a number of problems. It is well known that ring-structured protection schemes typically rely on excessive capacity redundancy. By contrast, one can provide protection with substantially less spare capacity on mesh networks[7]. The protection schemes on mesh optical networks were intensively studied in the early 1990s[18], [7], [17], [15], [8]. Nevertheless, mesh-based SONET are not widely used due to certain inadequacies, notably the slow restoration process that sometimes takes more than 2 seconds[15]. Note that typical Digital Cross-Connect Systems (DCS) in transport networks have very limited functionality. Hence only simple restoration algorithms were developed for mesh networks in the aforementioned literature.

Recently the emerging Optical Cross Connect (OXC) on WDM networks are shedding new light on the survivability issues of mesh networks. These intelligent OXCs, unlike their predecessor DCS, function much more like ATM switches or IP routers. They offer dynamic configuration via light path switching and allow many management tasks to be carried out in a distributed manner. Because of the dominance of IP traffic, IP-oriented control plane are being considered for WDM-based optical networks in order to provide seamless data transport [10], [2], [9]. The goal is to provide integrated functionalities such as light path routing, signaling, and restoration. This brings forth a significant shift in the management paradigm from centralized control to distributed control[5].

This shift in management paradigm has significant impact on the design of protection solutions for WDM networks. Two major issues of network survivability—restoration time and resource efficiency, now can be addressed separately. In particular, the restoration is handled in two phases, planning and activation, each aiming at optimizing resource efficiency and restoration speed respectively. At the planning phase, protection light paths are pre-computed and stored in OXCs before failures occur. At the activation phase when actual failure occurs, OXCs switch to pre-determined protection light paths. Thus traffic can be re-routed promptly in real time. The planning phase of WDM restoration is the focus of this paper. There are much related work in the literature regarding the optimization of the resource utilization in survivable mesh networks, see [4], [11], [12], [7], [16], [14]. However, most of the work assumed the complete information about traffic demands a priori and the protection paths for these demands were computed in a *batch*, either in a centralized manner [7], [12], [16], [14] or through distributed algorithms [4], [11]. The batch computation may work well in conventional telecommunication networks where the traffic demand is relatively static, but they are not suitable in a dynamic, data-centric environment such as the bandwidth-on-demand paradigm now considered by Optical Domain Service Interconnect (ODSI) [6] and Internet Engineering Task Force (IETF) [5]. With batch computation, any incremental changes of traffic demand will cause every existing paths to be re-computed, which is not desirable.

The existing works on restoration path computation can be roughly categorized into either path-based or link-based approaches[4]. In the former case, upon the detection of the failure event by the destination node of the connection, a notification is sent to the traffic source where the backup path is activated. In the latter case a failure event is detected and dealt with locally, *i.e.*, a "detour" is set up around the failed link/node. On the

---

Xun Su is with ECE Department, University of Texas at Austin. This work was done when he worked at Fujitsu Labs of America as a summer intern in 2000 (email: xsu@ece.utexas.edu).

Ching-Fong Su is with Fujitsu Labs of America, Sunnyvale, CA 94085 (email: csu@fla.fujitsu.com).

one hand, path-based approach benefits from its ability to create the resource-efficient backup path, but it incurs longer response time. On the other hand, link-based approach may not be able to establish the "optimal" protection path, but the speed at which a backup is set up is much higher.

In this paper, we propose a novel link metric and a set of distributed routing algorithms that maximize the wavelength sharing among independent protection light paths. We formulate the problem in the link-based restoration context, then extend the proposed scheme into a generic node-based approach. The goal is to devise a generic algorithm that efficiently exploits the potential sharing opportunities among the protection paths. These algorithms support on-demand path computation, so information about the complete traffic demands is not required. In practice, the proposed algorithms can be easily implemented as an extension of the existing IP routing protocols, e.g., open shortest path first (OSPF). The protection paths are optimized to reduce the wavelength redundancy while working light paths are assumed to be routed using minimum-hop paths. This separate optimization of the working and protection path is a design choice we made. It is partly motivated by [14], where it was shown that joint working/protection path computation significantly complicates the path computations with only marginal gain in the performance. We leave to the future study the further investigation of the impact of this design choice.

In section II we formulate two integer programs which demonstrate the idea of wavelength sharing in protection path computation. Several observations are made to motivate the design of our distributed algorithms. In section III we introduce our "bucket-based" link metrics and the corresponding routing algorithms. A modification to the link-based restoration mechanism, termed "node"-based restoration, is also discussed. In section IV we present simulation results and evaluate the performance of the proposed solution. We conclude the paper with future directions in section V.

II. PROBLEM SETUP

Let us consider a WDM network $G(N, E)$, where N is the set of nodes and E is the set of links. Suppose there exists a set of demands U which requests light-paths to be established across the networks. These demands are protected by link-based restoration, *i.e.*, each link (i, j) on the working path of demand $u \in U$ is protected by an alternative path connecting i and j. Let $x^u_{(i,j)}$ denote the amount of reserved wavelengths on link (i, j) in order to carry the traffic of demand u. Similarly let $y^{(i,j,u)}_{(m,n)}$ be the wavelength reservation on link (m, n) for demand u in case link (i, j) fails. Hence $x^u_{(i,j)}$ and $y^{(i,j,u)}_{(m,n)}$ denote the routing of working and protection paths respectively. We make the following assumptions:

- The traffic demands are in the form of unit wavelength requests. *i.e.*, $x^u_{(i,j)} \in \{1, 0\}$ and $y^{(i,j,u)}_{(m,n)} \in \{1, 0\}$.
- The working path of u (i.e, $x^u_{(i,j)}$) is determined by the minimum-hop path.
- The number of wavelength on each link is unconstrained.
- Only single link failure may occur at any instance of time.

The objective of a protection routing algorithm is to determine the protection paths for every u such that the total number of wavelength reserved for protection is minimized. It can be formulated as an optimization problem:

$$\min\{\sum_{(m,n)\in L} w_{(m,n)}\}$$

subject to:

$$\sum_n y^{(i,j,u)}_{(m,n)} - \sum_n y^{(i,j,u)}_{(n,m)} = \begin{cases} x^u_{(i,j)} & m = i \\ -x^u_{(i,j)} & m = j \\ 0 & m \neq i, m \neq j \end{cases} \tag{1}$$

$$w^{(i,j)}_{(m,n)} = \sum_u y^{(i,j,u)}_{(m,n)} \tag{2}$$

$$w_{(m,n)} \geq w^{(i,j)}_{(m,n)} \tag{3}$$

$$y^{(i,j,u)}_{(m,n)} \in \{0, 1\} \tag{4}$$

The constraint (1) is related to flow preservation on the protection path. In (2), $w^{(i,j)}_{(m,n)}$ is the amount of traffic on link (i, j) that will be moved to link (m, n) if link (i, j) fails. To take into account the reservation sharing among non-concurrent failures and to ensure enough wavelength is reserved for any link failure, $w_{(m,n)}$ needs to be reserved on link (m, n) , *i.e.,* (3) is to be satisfied.

This problem formulation fits well into a centralized management paradigm where the Network Management System (NMS) may optimally configures every protection paths, utilizing the complete knowledge of the demand set U. However, such an off-line algorithm is not desirable in an environment where the demands for light-paths arrive and depart dynamically. After all, it is costly to re-configure the whole networks whenever traffic demands change. Instead, an online protection routing algorithm is preferred in a dynamic environment.

An online algorithm determines the protection routing based on the *existing* network status. We do not assume that all future demands is known or the existing demands can be rerouted. Thus the objective of an online algorithm is to minimize the *marginal* wavelength requirement due to any newly arrived demand u^*. Suppose the working path $x^{u^*}_{(i,j)}$ has been determined by the minimum-hop path. Another optimization problem can be formulated to determine the protection path, i.e., $y^{(i,j,u^*)}_{(m,n)}$.

$$\min\{\sum_{(m,n)\in L} \Delta w^{(i,j)}_{(m,n)}\}$$

subject to:

$$\sum_n y_{(m,n)}^{(i,j,u^*)} - \sum_n y_{(n,m)}^{(i,j,u^*)} = \begin{cases} x_{(i,j)}^{u^*} & m = i \\ -x_{(i,j)}^{u^*} & m = j \\ 0 & m \neq i, m \neq j \end{cases} \tag{5}$$

$$b_{(m,n)}^{(i,j)} = \begin{cases} 0 & \text{if } w_{(m,n)}^{(i,j)} + 1 \leq w_{(m,n)} \\ 1 & \text{otherwise} \end{cases} \tag{6}$$

$$\Delta w_{(m,n)}^{(i,j)} \geq b_{(m,n)}^{(i,j)} \times y_{(m,n)}^{(i,j,u^*)} \tag{7}$$

$$y_{(m,n)}^{(i,j,u^*)} \in \{0,1\}, \tag{8}$$

In (6), $b_{(m,n)}^{(i,j)}$ is the *additional* wavelength requirement on link (m,n) if it is used by u^* to restore from the failure of link (i,j). It is determined based on the reservation sharing with other failures. For example, if $w_{(m,n)}^{(i,j)} + 1 \leq w_{(m,n)}$, no additional wavelength reservation is necessary since $w_{(m,n)}$ has been reserved on link (m,n).

From the problem formulation, we found that 1) determining $y_{(m,n)}^{(i,j,u^*)}$ is equivalent to finding the *minimum-cost* alternative path from i to j, and 2) the existing network status can be aggregated into $w_{(m,n)}^{(i,j)}$ and $w_{(m,n)}$. As a result, the protection routing problem in WDM networks can draw upon the conventional *shortest path* routing algorithms in data network. In the following we propose a novel link metric that provides the necessary network state information in an aggregated form and develop an online protection routing algorithm which fits into the framework of current Internet routing.

III. BUCKET-BASED LINK METRICS AND ALGORITHMS: MAXIMIZING THE SHARING

As motivated in the previous section, the problem we are faced with is to devise a distributed online restoration routing algorithm that maximizes the sharing among the protection paths. The key contribution we offer, lies in providing a protection scheme that is 1) bandwidth efficient; 2) computationally simple and conforming to the existing Internet routing framework; and 3) amenable to fast restoration. To achieve these goals, we propose a link-based (for goal 3), shortest-path (for goal 2) algorithm with a well-designed link metric (for goal 1). The link metric we devise is the central piece of the proposal.

The key observation is that the protection paths for different link failures can share the protection wavelengths since these paths need not to be activated at the same time. This suggests an efficient "bucket-based" link state representation. In the network $G(N,E)$, each link $l \in E$ maintains a set of "buckets", $h_l = (h_l^k, k \in E, k \neq l)$, as illustrated in Fig. 1. Each bucket h_l^k corresponds to a failure event k, and the "height" of the bucket, *i.e.*, the value of h_l^k, indicates the protection wavelengths that is reserved on the link l for the failure event k. In terms of the notation in the previous section, we have the correspondence $h_l^k = w_{(m,n)}^{(i,j)}$ for link $l = (m,n)$ and failure $k = (i,j)$. The key observation, is that the wavelengths that need to be actually reserved, equal to the maximum of the bucket heights, *i.e.*, $\max_k h_l^k$. Thus by maintaining a sequence of values indexed by the failure events, we captured the necessary information on the sharing potential offered by each link. Notice that this sharing potential is a function of the failure event. For example, in Fig. 1, link 4 maintains 3 buckets. Bucket h_4^2, corresponding to the failure of link 2, is the highest. This indicates that in order to protect an additional wavelength on link 2, link 4 has to reserve an extra wavelength if it is selected as part of the protection route. On the contrary, to protect an additional wavelength on link 2 or 3, no extra wavelength needs to be reserved.

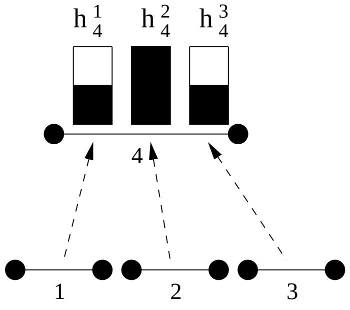

Fig. 1. Link metrics: buckets

Equipped with the proposed link metric representation, we now describe the computational procedure for the protection paths. It is a variant of the "shortest-widest" algorithm. we define the "width" $l\_width(l;k^*)$ of a link l with respect to a failure event k^*, as the normalized difference between the maximum bucket height, $\max_k h_l^k$ and the bucket corresponding to link failure k^*, $h_l^{k^*}$, *i.e.*, $l\_width(l;k^*) = 1 - \frac{h_l^{k^*}}{\max_k h_l^k}$ if $\max_k h_l^k > 0$, and $l\_width(l;k^*) = 0$ otherwise. Observe that the $l\_width(l,k^*)$ is between 0 and 1, and this value indicates the sharing capability that link l has to offer for the protection of the failure k^*.

We use a modified Bellman-Ford algorithm to identify the widest paths between the end nodes of the protected link, *i.e.*, the path that offer the most sharing. Here the width of the path p with respect to a link failure k^*, $p\_width(p,k^*)$, is defined to be the minimum of its link components, *i.e.*, $p\_width(p;k^*) := \min_{l \in p} l\_width(l;k^*)$. Observe that by this definition we implicitly establish that the marginal cost of traversing a path is dictated by that of the "narrowest" links along the path.

In the event that there are more than one such path candidates, and their widths are all 0 (*i.e.*, these paths all go through links with non-zero marginal wavelength consumption), we select the one that traverses the least number of "exhausted" links, *i.e.*, the links of width 0. On all other cases of tie breaking with positive path width, *i.e.*, the marginal costs are zero, we randomly select one widest path. Notice that this is obviously a locally optimal scheme. In other words, given only the current demand and without any knowledge about the future arrivals, the protection path we come up with is the most efficient in utilizing the sharing opportunities.

A. Highlights of the proposal

In section IV we evaluate the effectiveness of our proposal via simulation. Let us summarize its desirable features at this point:

- This procedure *optimally* solves the on-demand protection problem we formulated in section II. In the sequel, its performance improvement over the baseline scheme is demonstrated in terms of the reduction in wavelength redundancy.
- This proposal conforms to the existing Internet routing framework. In particular, it is easy to modify the Dijkstra or Bellman-Ford algorithm to implement our shortest-widest algorithm. The amount of link states is proportional to the number of failures, not the number of light paths, and hence is manageable in practice. In particular, one can implement our bucket link metric via "opaque LSA (Link State Advertisement)", an OSPF option[3]. Moreover, it is worth noting that the link states are set up in a way that is flexible enough to accommodate the failure types other than link failure, *e.g.,* node failures.
- One inherent feature of this proposal is the association between the link state/cost and the failure event. For different failure events, one is presented with different network topology and link states. Our proposal exploits this feature to locate the protection path that incurs least marginal cost, for a particular link failure. It is easy to extend our scheme to encompass further consideration, *e.g.,* excluding paths that traverse too many hops, which may, incur lowest marginal cost, but takes too much time to activate when the failure occurs.

B. A generalization: node-based protection

In the previous section we introduce a "bucket"-based link metric and a corresponding shortest-widest routing algorithm. The study is set in the context of the link-based failure detection, *i.e.,* we assume the loss of the optical signal is caused by a certain link failure. The implication, in relation to our routing/protection design, is that we construct the protection paths that are constrained to go from one end of the protected link to the other end, as illustrated in Fig. 3-1.

We observe that this design possesses some obvious merits, especially the locality of the restoration operation. However, it is also worth pointing out that the restoration efficiency is negatively impacted by this choice. Being forced to go from one end of the failure link to the other end, even for the traversing demands that are destined to the different nodes, the group of restoration paths may unwittingly clog the "local area" and under-utilize the potential sharing capability in the network. Fig. 2 is an illustration of the problem, where the working path of 2 demands share link (a, b), but diverge to node c and node d respectively. With link-based restoration, the protection paths for both demands have to go from node a to node b. This may not be desirable. Alternatively, if we set up these protection paths such that they start from node a, but end at node c and node d, there is a better chance for the exploitation of the "network-wide" sharing potential. Fig. 3-(2) illustrates a different restoration mechanism, where the protection path for a given link ends at the node two hops away on the corresponding working path. We call this "node"-based restoration. For example, for the protection of link (a, b), instead of going from a to b, we construct a protection path from a to c. An special case is link (c, d), for which the node-based protection path has the same end nodes as its link-based counterpart, since d is the destination and there is no nodes further down the working path.

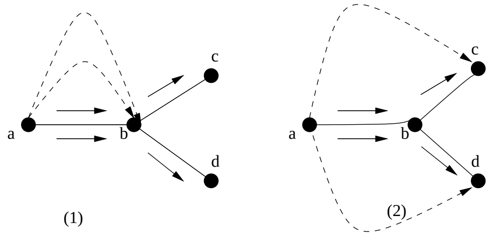

Fig. 2. (1) link-based computation; (2) node-based computation.

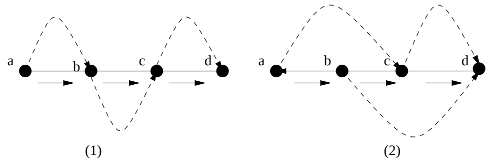

Fig. 3. (1) link-based computation; (2) node-based computation.

From a more practical standpoint, we notice that the failed optical cross-connects, depicted in Fig. 3 as the black nodes, may also cause disrupted service, in which case all the links adjacent to the failed node will "fail" simultaneously. Under such condition, the construction of the protection path for a particular link should consciously exclude the links that are experiencing the problem at the same time. To differentiate the link or nodal failure, it often takes time and causes undesirable delay in service recovery. In certain cases it is more advantageous to be conservative, *i.e.,* to use the nodal failure as the general model of the failure event, and treat the single link failure as a special case. The "jump-ahead" operation proposed above suits well to implement such a strategy.

IV. PERFORMANCE EVALUATION

In order to evaluate the performance of our proposal, we performed a number of simulations. Due to space limitation, in the following we present our results associated with network topology shown in Fig. 4 [13], though we conducted simulations and drew similar conclusions for a number of other network topologies, *e.g.,* NSFnet[1]. Unless stated specifically otherwise, the source and the destination node of the traffic demand are distributed uniformly across all the nodes in the network. The demands arrive in sequence and are routed one at a time. Each demand requests one unit of wavelength. We measure the "redundancy", *i.e.,* the ratio between the number of protection wavelengths and that of the working wavelengths, and plot it against an increasing number of demands in the network (on the log-scaled x-axis). As a reference, we solved the integer program formulated in section II. We are able to obtain the solution for 100 demands. For 15-node network, the redundancy is 61%. We

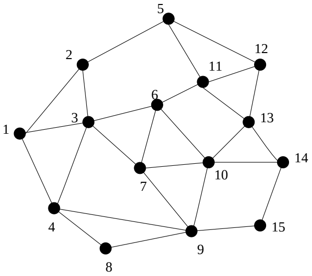

Fig. 4. 15-node network

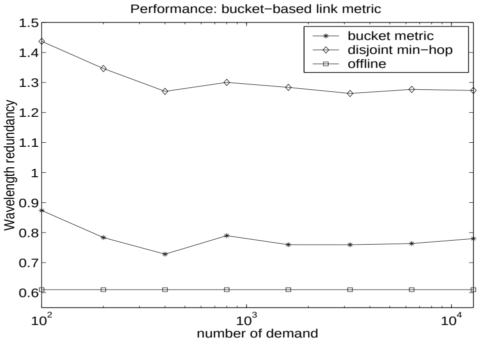

Fig. 5. Performance evaluation: 15-node net

are not able to solve for bigger number of demands due to the inability of GAMS/BonsaiG to deal with larger-sized problem. Note that if the further incoming demands are merely repetition of the first 100 demands, then the solution for the integer program will stay the same, *i.e.,* 61%. As it stands we think it is sensible to extrapolate this value as an approximation for the ideal off-line performance with larger number of demands. As a reminder, we note that this is a solution that assumes the complete knowledge of all the demands so it only serves as a lower bound for the distributed online solution approaches. We compare our proposal and a baseline online algorithm, where disjoint minimum hop-count paths is established to protect each link along the working path. We will see in the sequel that as we shift the problem domain into online context, the baseline protection scheme exhibits significant redundancy increase. Our proposal is an worthwhile approach to push the redundancy value towards that of the ideal off-line solution.

A. Performance improvement over the baseline

We first evaluate the performance of our proposal against the baseline algorithm. The baseline algorithm operates without utilizing the aggregate bucket information as our scheme does, hence provides a legitimate referencing system to 1) demonstrate the dramatic increase in wavelength redundancy when we shift the problem domain from off-line computation to online computation, and 2) evaluate the impact of the additional bucket state information on the overall wavelength utilization. We observe from Fig. 5 that significant reduction in redundancy is obtained due to the introduction and the use of the bucket link states. In addition, we see a near-constant capability of our approach to realize the network's sharing potential, when it is sufficiently loaded, *i.e.,* around 75% for 15-node network.

B. Gauge the room for improvement: iteration

An interesting indicator of the quality of the our proposal, is the degree of the sharing potential that has *not* been utilized after we configure the protection paths. By allowing iteration, or re-routing of the protection paths, we can gauge the amount of unrealized sharing opportunities after initial run. The "iteration" here simply means we allow the re-computation of the restoration paths with renewed aggregate network information, *i.e.,* bucket link states. It is possible to find a better protection path than the current choice. The degree of the improvement indicates how close to the optimality the initial configuration is. The smaller the improvement, the higher the quality of the first run/initial configuration.

demands #	working wavelengths	ite-0	ite-1	ite-2	ite-3
100	222	202	194	194	194
1000	2174	1687	1665	1665	1665
12800	28220	21588	21531	21531	21531

TABLE I
ITERATION: 15-NODE NETWORK

Table I demonstrates that we obtain a marginal amount of reduction in wavelength redundancy with one iteration, and the iterations after that do not improve the performance. For example, in the case of 1000 demands, the first iteration reduces the protection wavelength consumption by 22 units, a mere 1.3% decrease. And the further iterations make no additional improvement. Thus it is justified to state that our routing scheme is able to realize almost all of the sharing potential during its first run.

C. Performance evaluation: node-based vs. link-based algorithms

In Fig. 6 we compare the performance of the link-based and the node-based scheme. The reduction in the redundancy is evident when we apply the node-based protection mechanism. In particular, a roughly 7% reduction in wavelength redundancy is achieved by simply switching from link-based to node-based scheme when using our proposed link metrics, as can be observed from the lowest two curves in Fig. 6. This validates our intuition that a relaxation on the destination nodes allows for more diversified search for the best protection paths, which leads to better utilization of the "network-wide" sharing potentials.

V. CONCLUSIONS

In this study we explored a novel design of routing algorithms that aims to provide efficient failure protection in the WDM networks. The core idea of the "bucket"-based link metrics can be applied to a range of restoration mechanism, including link-based, node-based, or path-based approaches. We evaluated the effectiveness of this proposal via simulations and the results are

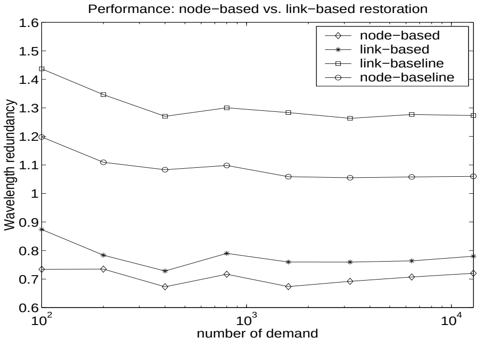

Fig. 6. Link-based vs. node-based, 15-node net

```
function routing-protection($G(N, E), s, t, \mathbf{w}(E)$)
  inputs:
    $G = (N, E)$: a network;
    $N$: the set of nodes of $G$;
    $E$: the set of links of $G$, also the set of failure events;
    $s$: source node in $N$;
    $t$: destination node in $N$;
    $\mathbf{h}_e$: a vector link metric associated with link $e \in E$,
      $\mathbf{h}_l = (h_l^1, h_l^2, \cdots h_l^k, \cdots h_l^{|E|}, k \in E, k \neq l)$;
    $h_l^k$: the wavelengths reserved on link $l$, for the protection of failure $k$;
  returns:
    working path $r(s,t)$ and associated protection paths $p[k^*]$ for all links
    $k^* \in r(s,t)$.

$r(s,t) = compute\_working\_path(s,t)$.
for each link $k^* \in r(s,t)$
  $s\_node = source(l), d\_node = end(l)$.
  $s\_node$ performs the following:
    for each link $l \in E$
      $l\_width(l; k^*) = 1 - \frac{h_l^{k^*}}{\max_k h_l^k}$
    endfor
    $l\_width(k^*; k^*) = -1$;      // "remove" link $k^*$
    $(widest\_paths, widest\_width)$
      $= compute\_widest\_paths(s\_node, d\_node)$;
    if $||widest\_paths|| > 1 \quad and \quad widest\_width == 0$ then
      for each $p \in widest\_paths$
        $c[p] = count\_saturated(p)$;
      endfor
      $p[k^*] = \min_p c[p]$;
    else if $||widest\_paths|| > 1 \quad and \quad widest\_width > 0$ then
    $p[k^*] = random\_select(widest\_paths)$;
    else
    $p[k^*] = widest\_paths$;
    endif
endfor
```

Fig. 7. routing-protection: a shortest-widest algorithm

promising. In the future we will look into the various issues that may impact the performance of our scheme. In particular, we would like to explore the possible ways to improving the coordination of the protection paths computation. The tradeoff between link-based and path-based protection schemes is theoretically interesting and practically useful. The benefit of joint working/protection path design, or the lack of which, is not studied in this paper and warrants further investigation.

APPENDIX

I. A PROTECTION PATH COMPUTATION ALGORITHM

See Fig. 7 for the shortest-widest algorithm we use to compute the protection paths.

REFERENCES

[1] G. Apostolopoulos, R. Guérin, and S. Tripathi. Quality of service based routing: A performance perspective. In *Proc. ACM Sigcomm*, 1998.
[2] P. Bonenfant and A. Rodriguez-Moral. Optical data networking. *IEEE Comm. Magazine*, Vol. 38 No. 3:63–70, 2000.
[3] R. Coltun. The OSPF opaque LSA option. *RFC 2370*, 1998.
[4] B. T. Doshi, S. Dravida, P. Harshavardhana, O. Hauser, and Y. Wang. Optical network design and restoration. *Bell Labs Technical Journal*, pages 58–84, Jan. - March 1999.
[5] D. Awduche et al. Multi-protocol lambda switching: Combining MPLS traffic engineering control with optical crossconnects (draft-awduche-mpls-te-optical-02.txt). *work in progress, Internet Draft*, July 2000.
[6] G. Berstein et al. Optical domain service interconnect (ODSI) functional specification (version 1.4). *work in progress, ODSI Draft*, August, 2000.
[7] T. Chujo et al. The design and simulation of an intelligent transport network with distributed control. *Network Operations Management Symposium*, 1990.
[8] W.D. Grover et al. Development and performance verification of a distributed asynchronous protocol for real-time network restoration. *IEEE JSAC*, 9(1):112–125, 1991.
[9] N. Ghani, S. Dixit, and T.S. Wang. On IP-over-WDM integration. *IEEE Comm. Magazine*, Vol. 38 No. 3:72–84, 2000.
[10] A. Greenberg, G. Hjalmtysson, and J. Yates. Smart routers-simple optics: A network architecture for ip over wdm. *Optical Fiber Conference*, 2000.
[11] R. R. Iraschko and W. D. Grover. A highly efficient path-restoration protocol for management of optical network transport integrity. *IEEE JSAC*, V. 18. No. 5:779–794, 2000.
[12] R.R. Iraschko, M.H. MacGregor, and W.D. Grover. Optimal capacity placement for path restoration in mesh surviviable networks. *IEEE ICC*, V. 18. No. 5, 1996.
[13] Murali Kodialam and T.V. Lakshman. Dynamic routing of restorable bandwidth guaranteed tunnels using aggregated network resource usage information. In *Proc. IEEE Infocom*, 2000.
[14] K. Murakami and H.S. Kim. Optimal capacity and flow assignment for self-healing ATM networks based on line and end-to-end restoration. *IEEE/ACM Transactions on Networking*, 6(2):207–221, 1998.
[15] Bellcore Special Report. Digital cross-connect systems in transport network survivability. *SR-NWT-002514*, issue 1, 1993.
[16] H. Sakauchi, Y. Nishimura, and S. Hasegawa. A self-healing network with an economical spare-channel assignment. In *Proc. IEEE Globecom*, pages 438–443, 1990.
[17] T.H. Wu. A passive protected self-healing mesh network architecture and applications. *IEEE/ACM Trans. Networking*, 2(1):40–52, 1994.
[18] T.H. Wu. Emerging technologies for fiber network survivability. *IEEE Comm. Magazine*, pages 58–74, Feb.,1995.