Degraded Service Provisioning in Mixed-Line-Rate WDM Backbone Networks Using Multipath Routing

Chaitanya S. K. Vadrevu, Rui Wang, Massimo Tornatore, *Member, IEEE*, Charles U. Martel, and Biswanath Mukherjee, *Fellow, IEEE*

Manuscript received March 20, 2013; accepted April 15, 2013; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor M. Reisslein. Date of publication May 13, 2013; date of current version June 12, 2014. Parts of this work were presented at the IEEE International Conference on Advanced Networks and Telecommunication Systems (ANTS), Bangalore, India, December 18–21, 2011, and the OSA Optical Fiber Communication Conference/National Fiber Optic Engineers Conference (OFC/NFOEC), Los Angeles, CA, USA, March 6–10, 2012.

C. S. K. Vadrevu is with Silver Spring Networks, Redwood City, CA 94063 USA, and also with the University of California, Davis, Davis, CA 95616 USA (e-mail: svadrevu@ucdavis.edu).

R. Wang, C. U. Martel, and B. Mukherjee are with the University of California, Davis, Davis, CA 95616 USA (e-mail: ruiwang@ucdavis.edu; martel@cs.ucdavis.edu; mukherje@cs.ucdavis.edu).

M. Tornatore is with the University of California, Davis, Davis, CA 95616 USA, and also with Politecnico di Milano, Milan 20133, Italy (e-mail: tornator@elet.polimi.it).

I. INTRODUCTION

TRAFFIC requests in optical backbone networks are becoming increasingly heterogeneous as they support emerging services such as cloud computing, VoD, e-science, etc. [1], [2]. Mixed-line-rate (MLR) networks offer a cost-effective solution to support heterogeneous traffic requests with various line rates, e.g., 10, 40, 100 Gb/s, etc. over different wavelength channels on the same fiber [3]. Higher bit rates offer high volume discounts that can be exploited in provisioning large traffic requests [3], [4]. The optical reach is lower for higher rates (for the same modulation format) due to linear and nonlinear physical impairments. Thus, MLR networks present design tradeoffs between volume discount and optical reach in selecting proper line rates to support heterogeneous traffic requests [3]–[5]. Cost can be minimized by multiplexing several low-bandwidth requests properly onto higher-bit-rate lightpaths, which need to be regenerated for large distances.

The Internet traffic is expected to grow significantly in future years with a higher proportion of bandwidth-intensive services [1]. Thus, cost-effective and survivable routing approaches exploiting the heterogeneity of mixed line rates in backbone networks are increasingly important.

Providing reliability for Internet services is crucial [6], and high reliability can be cost-effectively achieved by designing a survivable optical backbone network. The traditional design approach to backbone networks consists in provisioning full protection (where the primary capacity is 100% protected by a backup path), but this is expensive and consumes significant network resources. However, certain services may accept to be guaranteed only a part of the requested bandwidth under failure scenarios in exchange for paying a lower fee. These are known as *degraded services* (and the related protection approaches are known as *partial protection*) [7]–[9], [11]. For example, a user requiring file downloads might be willing to accept 60% of the requested bandwidth in the event of a failure in exchange for a lower fee. There are several studies on degraded service provisioning using multipath routing [8], [9], [12]. Disaster survivability is becoming a crucial topic with growing number of disasters and bandwidth resources become scarce during such circumstances. Partial protection is highly helpful during such situations and allows limiting the CAPEX needed to provide some minimal network connectivity under very large failure scenarios. However, all of these studies are confined to single-line-rate (SLR) networks, and this important topic has not been explored in MLR networks. Dedicated partial protection in SLR networks has been studied in [11]. Our present work studies degraded service provisioning in MLR networks by ensuring a user-specified fraction of requested bandwidth, even under failures, known as *partial-protection ratio*, and the entire bandwidth under normal operation. We particularly consider a minimum-cost network design problem to make decisions on which rate transponders and how many to use at each node. Dedicated protection and virtual topology design considering nonlinear interferences in MLR networks have been studied in [5] and [13], respectively. Reach-adaptive heuristics in MLR networks are discussed in [14].

Multipath routing is an efficient scheme for provisioning partial-protection services in which a service request is provisioned over multiple paths by routing part of the requested bandwidth on each path to efficiently utilize network resources [9], [10]. Thus, in the event of a failure, e.g., a link cut, only a part of the request bandwidth is affected, and the rest of the request can still be served over the other paths if they are link-disjoint. Multipath routing naturally offers protection from single-link failures. Multipath routing is much more challenging in MLR networks as several parameters need to be considered: 1) number of paths over which the request must be routed and the fraction of requested bandwidth on each path (known as partial-protection model); 2) survivability of routing to link failures, i.e., if a path fails, is there sufficient capacity on other paths to meet the desired bandwidth requirement?; and 3) cost-efficient rate assignment for the transponders. The heterogeneity in volume discount and optical reach of various line rates along with subwavelength grooming present further algorithmic challenges in rate assignment. Note that our approach can also be adopted to develop intelligent heuristics to support dynamic traffic requests [15].

The rest of this paper is structured as follows. In Section II, we describe the various partial-protection models. Section III provides a mathematical formulation and the theory for degraded service provisioning in MLR networks for a given set of static traffic demands. In Section IV, we present a computationally efficient heuristic for this problem. In Section V, we present illustrative results from typical backbone networks. Section VI concludes the paper.

II. Partial-Protection Models

We present two partial-protection models: dedicated partial-protection and multipath-based partial-protection. For better understating, let us consider a 100-Gb/s request with 0.6 partial-protection ratio (i.e., fraction of requested bandwidth that needs to be ensured even under failures is 0.6).

Dedicated Partial Protection (DPP): For each request, a primary circuit and a link-disjoint dedicated backup circuit are provisioned. This ensures survivability when a single physical link fails. The backup circuit supports backup traffic. Enough capacity in the backup circuit is reserved to meet the partial-protection requirement (ratio) of the primary circuit.

Fig. 1(a) shows a DPP solution with a primary and dedicated backup circuit of 100 and 60 Gb/s, respectively. The total capacity used is 160 Gb/s for supporting a 100-Gb/s request. This approach is expensive due to the high amount of idle backup capacity (60 Gb/s).

Multipath-Based Partial Protection (MPP): Now, the request is routed over multiple paths where each path acts as a primary circuit. The requested bandwidth is split over multiple paths so that a user-specified fraction (partial-protection ratio) of the requested bandwidth is always available even when one of the paths fails. There are two scenarios that arise.

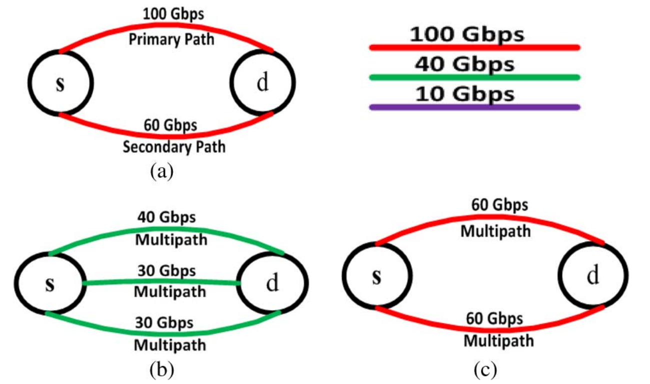

Fig. 1. Partial-protection models. (a) DPP. (b) MPP: Scenario I. (c) MPP: Scenario II.

Scenario I: Total bandwidth reserved over all the paths is equal to the requested bandwidth.

Scenario II: Total bandwidth reserved over all the paths is larger than the requested bandwidth.

Fig. 1(b) shows an example MPP Scenario-I solution with 40, 30, and 30 Gb/s of capacity reserved over three link-disjoint paths. The total capacity reserved is 100 Gb/s and is equal to the bandwidth requested. At least 60 Gb/s of capacity remains after any single-link failure, thereby meeting the partial-protection requirement. This approach is highly capacity-efficient, as rates of each of the paths could be assigned as 40 Gb/s.

Fig. 1(c) shows an example MPP Scenario-II solution. Now, the total bandwidth reserved on all the multiple paths is larger than the requested bandwidth to meet the partial-protection requirements. This approach is useful when the number of paths between a node pair is limited, and it is essential to overprovision to meet the partial-protection requirements. Two link-disjoint paths with 60 Gb/s are used to support the request in Fig. 1(c). The total reserved capacity of 120 Gb/s is larger than the requested bandwidth of 100 Gb/s. There is extra bandwidth reserved in this approach and rates of both paths are assigned as 100 Gb/s.

We note that rate assignment and partial-protection model are interrelated. For example, in Fig. 1(a), to support 60 Gb/s on backup circuit, we need either to assign a rate of 100 Gb/s to that path or use a combination of 40- and 10-Gb/s rates. Optical reaches of the rates also need to be considered.

III. Mathematical Formulation: DS-MLR MILP

The mathematical formulation of degraded services in a mixed-line-rate network (called DS-MLR) is formally described in the following as a mixed integer linear program (MILP).

The key aspects involve selection of a partial-protection model (from Section II), multipath routing of the paths, and cost-efficient rate assignment of transponders. The formulation explores various rate-assignment choices considering all possible grooming options that arise from each candidate rate assignment. The volume discount and optical reach of various line rates are also considered. We consider multiple lightpaths between each pair of nodes as an input in our approach and use m shortest paths for computing physical routing of these lightpaths. Note that the candidate paths need not be link-disjoint. Our approach also offers a framework to input an arbitrary subset of lightpaths. The formulation selects which lightpaths to route the requests. We formally state the problem as follows: Given a network topology, traffic matrix, partial-protection (degraded service) ratios (α_d) for the traffic demands, available line rates on each link (10/40/100 Gb/s), and costs of associated transponders. Assign lightpaths (including rates) to support all the traffic demands so that the overall transponder cost is minimized. The constraints are the following.

1) A lightpath can be assigned any rate provided that the physical distance of the lightpath is within the optical reach of desired line rate.
2) Number of lightpaths on a fiber link must be less than the number of wavelength channels on it.
3) Traffic demand d can be routed on a set of lightpaths provided that a fraction α_d of the requested bandwidth is always available under all single-link failures.
4) Lightpaths should be assigned the same wavelength on all the links along their paths (wavelength continuity).
5) Multiple subwavelength traffic demands can be aggregated on a lightpath (traffic grooming).

The parameters used in the MILP formulation follow.

Input Parameters:

- $G(N, E)$: Physical topology of the network with N nodes (optical switches) and E fiber links.
- r_k: Set of paths[1] where a path p is a sequence of links at a common wavelength.
- D: Set of demands.
- r_k: Set of available channel rates.
- c_k: Cost of a transponder with rate r_k.
- α_d: Partial-protection ratio for demand d.
- $\delta^+(n)$: Set of paths p in P initiated at node $n \in N$.
- $\delta^-(n)$: Set of paths p in P ending at node $n \in N$.

Notation:

- $s(d)$: Source node of demand d.
- $t(d)$: Destination (terminating) node of demand d.
- $h(d)$: Traffic in Gb/s of demand d.
- $s(p)$: Source node of path p.
- $t(p)$: Destination (terminating) node of path p.
- $w(k)$: Wavelength of path p.
- P_e: Set of paths in P passing over physical link e.

Variables:

- x_{dp}: Traffic of demand d that traverses path p.
- T_d: Total bandwidth reserved between $s(d)$ and $t(d)$ to support traffic demand d.
- $u_{pk\lambda}$: 1, if path p is active at rate r_k and wavelength λ.

The mathematical formulation is as follows:

$$\text{Minimize} \sum_{\lambda} \sum_{k} \sum_{p} u_{pk\lambda} c_k$$

subject to :

$$\sum_{p \in \delta^+(n)} x_{dp} - \sum_{p \in \delta^-(n)} x_{dp} = \begin{Bmatrix} T_d, & \text{if } s(d) = n \\ -T_d, & \text{if } t(d) = n \\ 0, & \text{otherwise} \end{Bmatrix} \quad \forall n \in N, \quad \forall d \tag{1}$$

$$T_d \geq h(d) \qquad \forall d \tag{2}$$

$$T_d - \sum_{p \in P_e} x_{dp} \geq \alpha_d h(d) \qquad \forall e, \forall d \tag{3}$$

$$\sum_{d} x_{dp} \leq \sum_{k} \sum_{\lambda} r_k u_{pk\lambda} \qquad \forall p \tag{4}$$

$$\sum_{k} \sum_{p \in P_e} u_{pk\lambda} \leq 1 \qquad \forall e, \forall \lambda. \tag{5}$$

The objective function minimizes the total transponder cost at different channel rates. The variable $u_{pk\lambda}$ is set to one, when the lightpath p on wavelength λ is lit at rate r_k. Constraint (1) is a flow-conservation constraint, which ensures that a bandwidth of T_d is reserved either over a single path or multiple paths for the traffic demand d. Constraint (1) is formulated on a network with nodes N the same as in $G(N, E)$ and links as precomputed paths p in P. Constraint (2) ensures that reserved bandwidth T_d is greater than or equal to the requested bandwidth $h(d)$. Bandwidth higher than the requested $h(d)$ needs to be reserved for some demands in order to meet with degraded QoS requirement, i.e., $\alpha_d h(d)$ of bandwidth needs to be always available even under any fiber link failure, and this technique is known as bandwidth overprovisioning. Constraint (3) computes the total bandwidth belonging to demand d that will become unavailable due to failure of a fiber link e. The bandwidth after failure of link e must be greater than or equal to $\alpha_d h(d)$ for all single-link failures in (3) and satisfy the degraded service requirement.

Constraint (4) ensures that traffic routed over a lightpath is assigned a rate r_k. This constraint determines the rate to be assigned to a lightpath and also takes care of grooming multiple demands over a lightpath. Constraint (5) ensures that every lightpath is assigned a unique channel rate. Solution of constraint (3) incorporates both partial-protection models, namely dedicated partial protection and multipath-based partial protection, i.e., the formulation can route a connection with MPP or DPP depending on which is effective

$$\sum_{p \in P_e} x_{dp} \leq 1 \qquad \forall d, \forall e. \tag{6}$$

Constraint (6) is an optional constraint, which ensures that all lightpaths assigned to support each demand d are link-disjoint. In the results section, we study MILP solutions with and without constraint (6) and call it DS-MLR-D (MILP).

IV. Heuristic: DS-MLR

The DS-MLR MILP solution is computationally expensive and may take unacceptably long times to solve for increasing traffic loads and large network topologies. Hence, we present a computationally efficient heuristic (DS-MLR) for provisioning degraded services in MLR networks that scales well to larger topologies and traffic loads and has efficient and practical computational times. We consider bundled traffic requests to achieve efficient network optimization. However, DS-MLR can also be adapted to dynamic traffic requests, in which case it can be used to provision individual traffic requests with different partial-protection ratios. Note that, for dynamic scenarios, it is a traffic engineering problem rather than a network design problem. In the following presentation of DS-MLR, we assume that candidate multipath routings are link-disjoint. However, DS-MLR can be easily adapted to support non-link-disjoint path also.

We will refer to the DS-MLR algorithm in Algorithm 1 during this discussion. We sort the requests in decreasing order of traffic demands, and provision requests one at a time (line 4 in Algorithm 1). Our approach examines p multipath routing solutions (i.e., 2-path solutions, 3-path solutions, etc.) for every traffic request (where n is the maximum number of disjoint

[1] A set of m shortest paths between each node pair is precomputed.

Algorithm 1: Heuristic DS-MLR: Degraded Services in Mixed-Line-Rate Networks

1: Let $\boldsymbol{T}$ be a sorted list of all demands in descending order
2: construct the auxiliary graph ***MLR-aux***
3: **while** there is a nonzero demand in $\boldsymbol{T}$ **do**
4: select the largest demand $\boldsymbol{D_{s,d}}$ in $\boldsymbol{T}$
5: choose the number of multipath routings to be explored $\boldsymbol{n}$
Route Assignment
6: **for** every multipath routing $\boldsymbol{i}$ ($\mathbf{2}$: $\mathbf{p}$ routings)
7: make a copy of current auxiliary graph in s
8: compute the ratios in which demand $\boldsymbol{D_{s,d}}$ is to be split over $\boldsymbol{i}$ paths $B=\left[\frac{\alpha}{i-1},\ldots\ldots(i-1 \text{ times}), \max\left\{(1-\alpha), \frac{\alpha}{i-1}\right\}\right]$
9: sort the list $\boldsymbol{B}$ in descending order
10: **for** each path $\boldsymbol{k}$ from $\mathbf{1}$: $\boldsymbol{i}$
11: select the top most ratio $\boldsymbol{b_k}$ from $\boldsymbol{B}$, where $\boldsymbol{b_k^* D_{(s,d)}}$ is the capacity to be routed over the $\boldsymbol{k}$th path
12: remove all links whose capacity is below $\boldsymbol{b_k^* D_{(s,d)}}$ from $\boldsymbol{MLR-aux_i}$
13: remove all those links from $\boldsymbol{MLR-aux_i}$ that are not link-disjoint to any provisioned paths (earlier $\boldsymbol{k}-\mathbf{1}$ paths)
14: run a shortest-path algorithm over the resultant $\boldsymbol{MLR-aux_i}$
15: **if** a path $\boldsymbol{P}$ is found **then**
16: remove any logical links from $\boldsymbol{P}$ and reduce capacity by $\boldsymbol{b_k^* D_{(s,d)}}$ on corresponding links in $\boldsymbol{MLR-aux_i}$
17: assign new lightpaths for any remaining segments in $\boldsymbol{P}$
Rate Assignment
18: **for** each new lightpath in **Step 17**
19: assign highest feasible rate ($\geq \boldsymbol{b_k^* D_{(s,d)}}$) without regeneration; **else**, assign the smallest rate ($\geq \boldsymbol{b_k^* D_{(s,d)}}$ units) with regeneration
20: route $\boldsymbol{b_k^* D_{(s,d)}}$ capacity on the lightpath
21: add a corresponding new logical link in $\boldsymbol{MLR-aux_i}$ and update residual capacity on it
22: **end for**
23: **else**
24: the $\boldsymbol{i}$th multipath routing is not feasible, go to **Step 6**
25: **end if**
26: **end for**
27: compute total capacity used on all newly setup lightpaths
28: **end for**
Request Provision
29: select the multipath routing $\boldsymbol{i}$ that minimizes additional capacity usage computed in **Step 27**
30: update auxiliary graph $\boldsymbol{MLR-aux}$ with $\boldsymbol{MLR-aux_i}$
31: remove demand $\boldsymbol{D_{s,d}}$ from $\boldsymbol{T}$
32: **end while**
Post-Processing
33: **for** every lightpath in $\boldsymbol{MLR-aux}$
34: **if** residual capacity on it $\geq$ ***threshold***
35: try to route on lower line rate or multiple lightpaths of lower rates to achieve transponder cost savings
36: **end if**
37: **end for**

paths between the request's source and destination) and selects the one that minimizes a cost metric (line 6 in DS-MLR). DS-MLR mainly uses four modules: *route assignment*, *rate assignment*, *request provision*, and *post-processing*. We choose the routing of the lightpaths in the route assignment phase. We assign line rates to the lightpaths in the rate assignment phase. For each of the p candidate multipath routing solutions, the heuristic runs *route assignment* and *rate assignment* modules. Then, the *request provision* module compares the costs of the n different candidate multipath routings and selects one to actually support the request. After provisioning all requests, the *post-processing* module checks whether any underutilized lightpath/transponders can use a lower line rate to reduce transponder cost without violating reach constraints. The flowchart in Fig. 2 shows a high-level flow of the heuristic algorithm. As shown in the flowchart, for each entry in the traffic matrix, we compute several multipath routing solutions in the first step (*route assignment*) followed by rate or transponder assignment to the lightpaths of each of the multipath solutions (*rate assignment*). Finally, we select one multipath routing solution based on a cost metric and provision it (*route provision*). *Update traffic matrix* module in Fig. 2 removes provisioned demands from the traffic matrix.

Route Assignment: We accomplish route assignment using an auxiliary-graph approach [5]. For each request, we construct an auxiliary graph called ***MLR-aux***, whose edge weights reflect

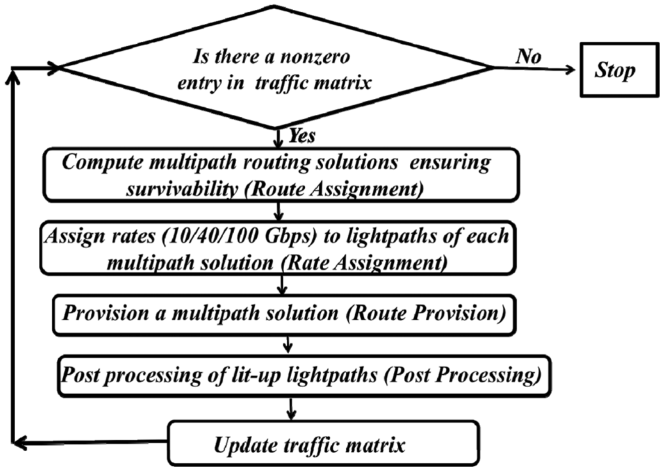

Fig. 2. Flowchart of heuristic approach.

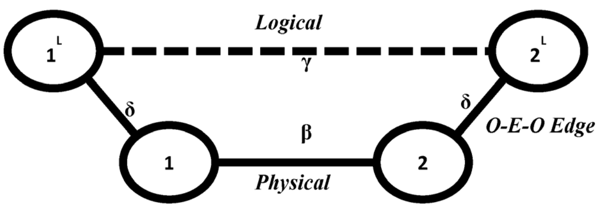

Fig. 3. Different edges in ***MLR-aux*** graph.

information about the current physical topology and logical topology (lightpath connectivity). It serves as a unified graph model for making routing decisions. ***MLR-aux*** has two layers: physical layer and lightpath layer, and three types of edges: physical edges, logical edges, and O-E-O edges. For every node N in the network's physical topology, a node N in the physical layer and another node N^{L} in the lightpath layer are created in the ***MLR-aux*** graph.

1) Physical Edges: Edges in the physical layer represent links from the physical topology and are assigned a weight $\boldsymbol{\beta}$, which represents the cost in routing new lightpaths over physical links.

2) Logical Edges: Edges in the lightpath layer represent lightpaths routed in physical topology and are assigned a weight $\boldsymbol{\gamma}$. They are constructed as follows: If a lightpath is routed between two nodes N_1 and N_2 in the physical topology, the nodes N_1^{L} and N_2^{L} in the lightpath layer in the ***MLR-aux*** graph are connected by a logical edge of weight $\boldsymbol{\gamma}$ and are indicated by a dashed line in Fig. 3.

3) O-E-O Edges: Represent O-E-O conversion. Every node N and its corresponding node N^{L} in the lightpath layer are connected by an O-E-O edge in ***MLR-aux*** and are assigned a weight $\boldsymbol{\delta}$. Fig. 3 shows part of ***MLR-aux*** with all three edges.

We use ***MLR-aux*** (lines 2 and 7 in DS-MLR) to make routing decisions, but not rate assignment. Hence, the weights of all *logical* edges are the same, and the weights of all *O-E-O* edges are the same and thus do not reflect the line rate. Weights of *O-E-O* and *logical edges* are chosen to be lower than *physical edges* to favor grooming using residual capacity on existing lightpaths to save the cost of lighting up new transponders (lines 2 and 7 in DS-MLR). Running Dijkstra on ***MLR-aux*** gives a solution with the following options (depending on the choice of edge weights): 1) grooming the requested bandwidth as much as possible on existing lightpaths (by preferring *logical edges* in the lightpath layer); 2) setting up new lightpaths (selecting *physical edges* in ***MLR-aux***) for the portion of the requested bandwidth that cannot be met on existing lightpaths.

For a multipath solution with $\boldsymbol{k}$ paths, we split the requested bandwidth as: $[(\alpha_{s,d}/k - 1)(k - 1 \text{ times})$ and $\max\{(1 - \alpha_{s,d}),(\alpha_{s,d}/k - 1)\}]$, where $\alpha_{s,d}$ is the *partial-protection* ratio on the $\boldsymbol{k}$ paths (line 8 in DS-MLR). For example, the requested bandwidth in a 2-path solution is split in the ratio $\{\alpha_{s,d},(1 - \alpha_{s,d})\}$ for $0 \leq \alpha_{s,d} \leq 0.5$ and $\{\alpha_{s,d},\alpha_{s,d}\}$ for $0.5 < \alpha_{s,d} \leq 1$ onto the two paths.

Link-disjoint routes for each of the $\boldsymbol{k}$ paths are computed sequentially using ***MLR-aux*** as follows: For path $\boldsymbol{i}$ from ***1 to k***: 1) remove those *logical edges* in ***MLR-aux*** whose residual capacity is less than the fractional requested bandwidth to be routed over the $\boldsymbol{i}$th path (line 12 in DS-MLR); 2) remove those *logical* and *physical* edges from ***MLR-aux*** that are not link-disjoint to earlier $\boldsymbol{i - 1}$ paths (line 13 in DS-MLR); and 3) run a shortest-path algorithm over the resulting ***MLR-aux*** (line 14 in DS-MLR). The computed routes will ensure grooming the request over residual capacity in existing lightpaths (*logical edges*) and any new lightpaths to be set up in a single step (lines 16 and 17 in DS-MLR).

Rate Assignment: We assign the line rate for new lightpaths setup in the *route assignment* phase. The bandwidth routed on each of the k paths is known from the *route assignment* phase. We assign the smallest line rate larger than the bandwidth routed over that lightpath, e.g., if the bandwidth routed over a lightpath is between 10–40 Gb/s, then 40 Gb/s line rate is assigned for that lightpath (lines 18–21 in DS-MLR). Consolidation of rate assignment to foster traffic grooming is performed in the Post-Processing phase. The lightpath is regenerated at appropriate intermediate nodes if the optical reach of the line rate is less than the physical distance to be traversed. We assume the regeneration can be accomplished by O-E-O conversion and requires additional transponder cost.

Request Provision: For each traffic request, we compute $\boldsymbol{n}$ different multipath routing solutions (one for each possible number of edge disjoint paths) during the *route assignment* phase. *Request provision* module compares the cost of each of the $\boldsymbol{n}$ candidate multipath solutions and selects one to support the request. The cost metric is an input to this module (lines 29–31 in DS-MLR). For example, transponder cost can be used as cost metric, i.e., we compare the cost of newly lit transponders for supporting the request for various multipath routing solutions and select the minimum. Another important metric can be residual capacity in newly lit lightpaths, i.e., for supporting other requests, we may need to set up new lightpaths. After provisioning the request, there may be unused capacity, and we wish to choose the solution that has large unused/residual capacity (line 29 in DS-MLR). Greater residual capacity can increase opportunities for grooming later traffic requests and lead to greater transponder savings. Our heuristic offers general framework to implement various *route assignment*, *rate assignment*, and *request provision* schemes.

Post Processing: We use this module after all the requests have been routed (lines 33–37 in DS-MLR). We perform the following steps for each lightpath provisioned.

1) We consider the line rate of each lightpath and the traffic routed on it.

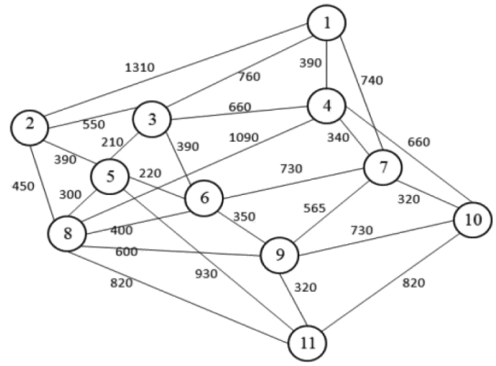

Fig. 4. Cost 239 network topology.

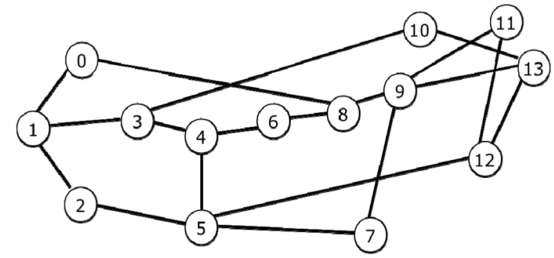

Fig. 5. 14-node NSFnet topology.

TABLE I
TRAFFIC MATRIX OF 350-Gb/s AGGREGATE CAPACITY FOR COST 239 NETWORK (EACH ENTRY IS IN UNITS OF Gb/s)

]	1	1	3	1	1	1	1	1	1	1
1	0	5	8	4	1	1	10	3	2	3
1	5	0	8	4	1	1	5	3	1	2
3	8	8	0	6	2	2	11	11	9	9
1	4	4	6	0	1	1	6	6	1	2
1	1	1	2	1	0	1	1	1	1	1
1	1	1	2	1	1	0	1	1	1	1
1	10	5	11	6	1	1	0	6	2	5
1	3	3	11	6	1	1	6	0	3	6
1	2	1	9	1	1	1	2	3	0	3
1	3	2	9	2	1	1	5	6	3	0

2) If the lightpath is underutilized, we reduce the lightpath to a lower line rate provided the traffic routed on it can be supported while ensuring reach constraints (e.g., the optical reach of 40 Gb/s is lower than 100 Gb/s in our studies). We regenerate if necessary. We also consider supporting the lightpath over multiple lightpaths of lower line rates when larger transponder cost savings is achieved.

As an example, consider a lightpath supporting 30 Gb/s of traffic that is assigned a 100-Gb/s transponder. We attempt to reduce this line rate to 40 Gb/s provided that its physical routing is within the optical reach of a 40-Gb/s line rate. Otherwise, the lightpath needs to be regenerated, and this is beneficial only when two 40-Gb/s transponders cost less than a single 100-Gb/s transponder. The above lightpath can also be supported on three 10-Gb/s lightpaths provided that this is a cost-efficient choice, according to the cost model used. Thus, *post-processing* enables us to achieve further transponder cost savings using better grooming techniques. We use first-fit wavelength assignment in the heuristic.

Computational Complexity: The complexity of the heuristic is dominated by the route assignment module, where link-disjoint routes have to be computed. The number of link-disjoint paths between any pair of nodes is bounded by the highest nodal degree in the topology, say D. Let the number of wavelength channels on each link be W. Hence, for C connection requests, the complexity of the heuristic is $O(|N|^2 DCW)$.

Transponder Upgrade: The *route assignment* module of the DS-MLR (Heuristic) can be modified to incorporate upgrading the transponders to higher line rates. In computing link-disjoint routes in the *route assignment* module, we ignore all those logical links in the ***MLR-aux*** graph that do not have sufficient capacity to route the requested capacity. However, the transponders of those logical links can be upgraded to higher line rates to generate sufficient residual capacity to support the new request. For example, a 40-Gb/s lightpath can be upgraded to a 100-Gb/s lightpath for the additional cost difference between the two transponders. This will enable more efficient grooming over existing lightpaths without having to set up new lightpaths. We denote the modified heuristic with *transponder upgrade* feature as DS-MLR (Heuristic) upgrade.

V. ILLUSTRATIVE NUMERICAL EXAMPLES

We present illustrative numerical results using the 11-node Cost 239 network and 14-node NSFnet topologies shown in Figs. 4 and 5, respectively, initially with eight wavelengths on each fiber link (and later with 80 wavelengths/fiber) [16]. Normalized transponder costs for 10/40/100 Gb/s are 1×, 3.3×, and 7×, respectively [17]. Different modulation schemes are used for different rates, namely, 50% RZ-DPSK, 50% RZ-DQPSK, and 50% RZ-DP-DQPSK with coherent receiver for 10, 40, and 100 Gb/s, respectively. Transmission reaches for 10/40/100 Gb/s are 5000, 2400, and 2700 km, respectively, which are calculated using 10G dispersion map only considering that legacy systems are 10 Gb/s; 40- and 100-Gb/s systems are likely to be installed on a 10-Gb/s system taking into account nonlinear interactions between adjacent wavelengths and other linear impairments [17]. The transmission reach of 100 Gb/s is larger than 40 Gb/s due to coherent reception and advanced modulation techniques. 100 Gb/s has good cost advantage over 10 × 10 Gb/s and makes it greatly favorable. However, the lower optical reach of 100 Gb/s allows choosing 10 Gb/s for long distances. Also, the MLR scenario allows us to avoid capacity overprovisioning when less capacity is needed through assigning 10 and 40 Gb/s for lesser bandwidth requirements. The traffic demands shown in Table I form an aggregate traffic of 350 Gb/s over the Cost 239 topology [5]. The weights of different edges in the ***MLR-aux*** graph are chosen as $\beta = 1$, $\gamma = 0.3$, and $\delta = 0.2$ to avoid traversing many logical hops and thus avoid large physical distances. Results with other weights are also presented.

A. DS-MLR (ILP) Versus DS-MLR (Heuristic) Versus Full Protection

Fig. 6 shows the normalized transponder cost obtained by our DS-MLR (MILP) partial-protection scheme compared to the cost of dedicated full protection for the Cost 239 topology. Dedicated full protection ($\alpha = 1$) is computed using the MLR dedicated protection strategies in [5] and compared to DS-MLR (MILP) with partial-protection ratios ranging from 0.1 to 1. The cost savings due to partial protection are evident as we notice growing transponder cost from partial-protection ratio of 0.1 to

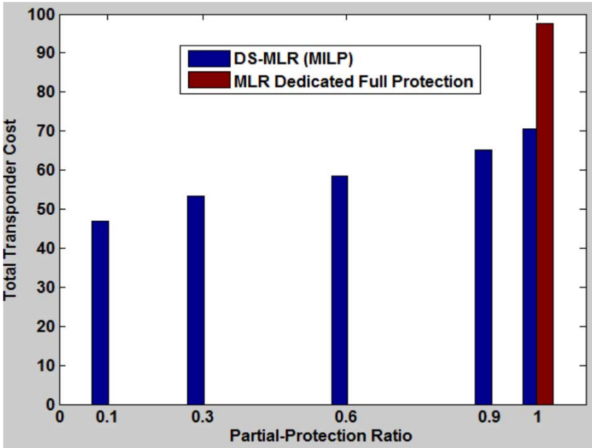

Fig. 6. Comparison of transponder cost of DS-MLR (MILP) approach (partial-protection) for various partial-protection ratios α to dedicated full protection in MLR networks on Cost 239 topology.

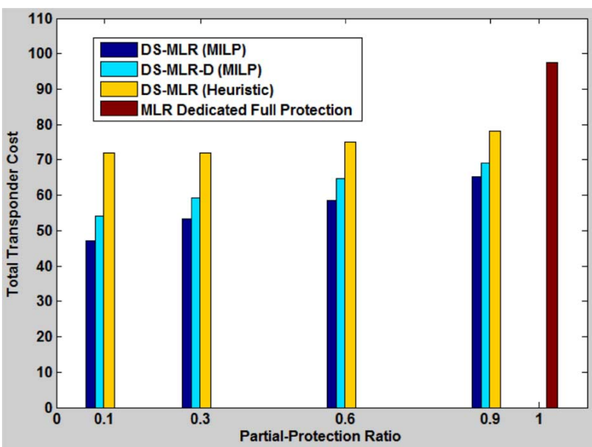

Fig. 7. Comparison of transponder cost for DS-MLR (MILP), DS-MLR-D (MILP), DS-MLR (Heuristic), and dedicated full protection in MLR networks for various partial-protection ratios α on Cost 239 topology.

0.9. The additional cost needed for partial-protection ratio 0.3 from 0.1 is 13%, 0.3 from 0.6 is 22.33%, 0.6 from 0.9 is 11.45%, and 0.3 to 0.9 is 22.33%. These results indicate the significant cost savings network operators can gain from partial protection. At partial-protection ratio of 1, the difference in transponder cost between DS-MLR (MILP) and MLR dedicated full protection is due to: 1) cost benefits of multipath routing in MLR as multipath routing offers an efficient mechanism for distribution of traffic over multiple lightpaths and does not need idle protection bandwidth as in dedicated full protection. The benefits of multipath routing in MLR networks are greater than partial protection. The total cost savings containing both partial-protection and multipath routing are in the range of 27%–52% over MLR dedicated full protection.

Fig. 7 compares the performance of DS-MLR (Heuristic), DS-MLR (MILP), DS-MLR-D (MILP), and MLR full-dedicated protection on the Cost 239 topology. DS-MLR (Heuristic) performs within 15%–20% versus DS-MLR (MILP) and is computationally efficient. DS-MLR-D (MILP) has around 5% larger cost over DS-MLR (MILP). This is due to using link-disjoint routings in DS-MLR-D (MILP). DS-MLR (Heuristic) reduces transponder cost significantly compared to MLR

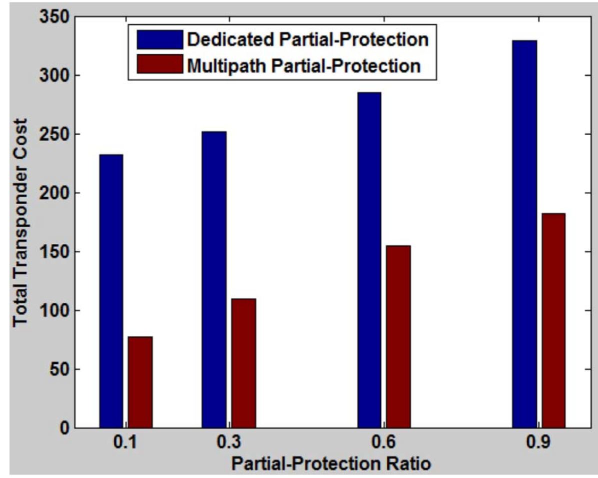

Fig. 8. Comparison of transponder cost for dedicated partial-protection and multipath partial-protection in MLR networks for various partial-protection ratios α on NSFnet topology.

TABLE II
COMPARISON OF TRANSPONDER COST OF DS-MLR (MILP) FOR 0.6 PARTIAL-PROTECTION RATIOS TO DEDICATED FULL PROTECTION IN MLR NETWORKS ON COST 239 TOPOLOGY

Traffic in Multiples of 350 Gbps	1	5	10
DS-MLR (MILP)	58.5	188.6	310.9
MLR Dedicated Full Protection	97.5	240	400

full-dedicated protection. This is primarily due to multipath routing, but is also due to partial protection. DS-MLR (Heuristic) also scales well with larger topologies such as the NSF network, US wide network, etc., and increasing traffic demands taking only a few seconds to minutes of computation time unlike the DS-MLR (MILP) solution. DS-MLR (MILP) takes about 2 days to run to completion using CPLEX optimization tool for the Cost 239 topology on a personal laptop. DS-MLR (MILP) has performance gaps to optimal on NSFnet even after 4–6 days of computation. Runtimes for larger topologies are expected to take significantly longer.

Table II shows transponder cost benefit of DS-MLR (MILP) over dedicated full protection for MLR networks with increasing traffic loads (where higher load means multiplying each element in Table I by a factor 5 or 10). We find that partial protection for DS-MLR (MILP) has decreasing (40%, 21.5%, and 22%) savings over dedicated protection for aggregate loads of 350 Gb/s, 1.75 Tb/s, and 3.5 Tb/s, respectively. This is due to the fact that, for higher traffic, grooming becomes less important, and so the amount of idle protection bandwidth in dedicated full protection decreases.

B. MPP Versus DPP

Fig. 8 shows the cost benefits of MPP versus DPP using traffic matrix in Table III [3]. In MPP, we consider both bandwidth overprovisioning and no bandwidth overprovisioning. We observe large cost benefits for MPP versus DPP, i.e., MPP only requires 44%–67% of cost of dedicated partial protection. The above benefits of MPP are due to splitting the request over multiple routes, which avoids the need to allocate extra bandwidth for backup paths. Multipath routing has inherent advantages in MLR as it allows exploiting volume discount of high-line-rate

TABLE III
TRAFFIC MATRIX OF 1 Tb/s AGGREGATE CAPACITY FOR 14-NODE NSFNETWORK (EACH ENTRY IS IN UNITS OF Gb/s)

0	2	1	1	1	4	1	1	2	1	1	1	1	1
2	0	2	1	8	2	1	5	3	5	1	5	1	4
1	2	0	2	3	2	11	20	5	2	1	1	1	2
1	22	2	0	1	1	2	1	2	2	1	2	1	2
1	8	3	1	0	3	3	7	3	1	1	5	2	5
4	2	2	1	3	0	2	1	2	2	1	1	1	2
1	1	11	2	3	2	0	9	4	20	1	8	1	4
1	5	20	1	7	1	9	0	27	7	2	3	2	4
2	3	5	2	3	2	4	27	0	75	2	9	3	1
1	5	2	2	3	2	20	7	75	0	1	1	2	1
1	1	1	1	1	1	1	2	2	1	0	2	1	61
1	5	1	2	5	1	8	3	9	1	2	0	1	81
1	1	1	1	2	1	1	2	3	2	1	1	0	2
1	4	2	2	5	2	4	4	1	1	61	81	2	0

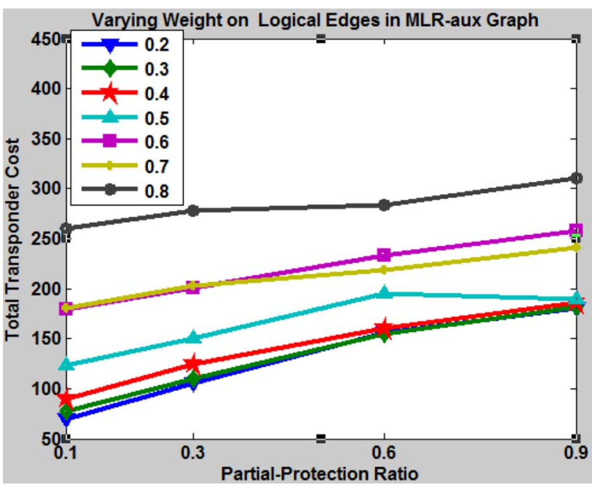

Fig. 9. Transponder costs for different *logical edge* weights in ***MLR-aux*** using DS-MLR (Heuristic) approach for various partial-protection ratios α on NSFnet topology.

transponders and selects appropriate line rate for cost-effective traffic grooming to maximize transponder reuse and thus minimize cost. We also note the cost benefits due to partial protection by comparing total transponder cost in MPP from partial-protection ratio 0.1 to 0.9. The additional cost needed for partial-protection ratio 0.3 from 0.1 is 42.21%, 0.6 from 0.3 is 41.28%, 0.9 from 0.6 is 17.32%, etc., showing the significant benefits of partial protection.

C. Sensitivity of the DS-MLR (Heuristic)

Fig. 9 compares the performance of DS-MLR (Heuristic) on NSFnet for increasing *logical edge* weights in ***MLR-aux***. We fix the weights of *physical edges* to 1 and *O-E-O edges* to 0.2 and vary the weights of *logical edges*. We note that as the *logical edge* weight increases, the preference for grooming traffic requests over existing lightpaths reduces (i.e., selection of *logical edges* in ***MLR-aux*** during *route assignment* reduces). Our DS-MLR (Heuristic) prefers to set up new lightpaths for increasing *logical edge* weights, and this is reflected in increasing total transponder cost in Fig. 9. Note that, for *logical edge* weights of 0.2, 0.3, and 0.4, the transponder cost is similar, showing that routing remains almost the same for these weights. As we increase the *logical edge* weights to 0.5, 0.6, 0.7, the transponder cost rises, showing that more new lightpaths are set up and grooming over existing lightpaths reduces. Note that for *logical edge* weight of 0.6 and 0.7, it is still economical to prefer one-logical-hop routing to setting up a new lightpath, but not if there are two or more hops. Fig. 9 also shows the abrupt increase in transponder cost when the *logical edge* weight is increased to 0.8. Now, the cost of logical edges is so high that they are not used by ***MLR-aux***, losing cost benefits due to grooming/multiplexing subwavelength services over high-capacity wavelength channels. Thus, it is extremely important to select appropriate weights for *logical*, *physical*, and *O-E-O edges*, which also depend on the network topology and traffic matrix being considered.

TABLE IV
TRANSPONDER COST FOR MULTIPATH PARTIAL PROTECTION IN SLR VERSUS MLR NETWORKS ON NSFNET TOPOLOGY FOR 3 Tb/s AGGREGATIVE CAPACITY AND VARIOUS PARTIAL-PROTECTION RATIOS α

α →	**0.5**	**0.6**	**0.7**	**0.8**	**0.9**
SLR 100G	487.5	513.75	528.75	540	566.25
SLR 40G	477.5	520	575	592.5	682.5
SLR 10G	429	484	591	655	735
MLR	387.75	437	496.75	554	550

D. MLR Versus SLR

Table IV shows the transponder cost of MLR networks versus SLR networks in case of multipath partial protection. DS-MLR (Heuristic) is modified to accommodate only one rate at a time for SLR networks. Note that as α increases, SLR 10G has the worst performance compared to SLR 40G, SLR 100G, and MLR. This is because as the traffic requests increase, the cost benefit obtained from volume discount in 100G and 40G is greater than benefits due to longer reach of 10G. We also notice the volume discount leveraged due to 100 Gb/s line rate from the cost savings of SLR 100G to SLR 10G and 40G for α of 0.7, 0.8, and 0.9. In MLR, the optimal rate is chosen from the set of all lines rates for every request, so it performs better than each of the three SLR networks for most scenarios. Cost savings of MLR over SLR 10G in our studies are on average 18.44% and as high as 33% for $\alpha = 0.9$. Cost savings of MLR over SLR 40G are on average 17.8% and over SLR 100G is 14.1%. Cost savings of MLR over 40G and 100G are more prominent for lower partial-protection ratios. The advantage of MLR networks is obvious, and significant savings are also experienced for growing traffic in backbone networks. These results also indicate inherent benefits of multipath routing in MLR over SLR due to the volume discount of higher line rates along with the ability in MLR to select and cost-effectively groom over appropriate line rate from the set of multiple line rates as per bandwidth needs of each request.

We also studied traffic requests with heterogeneous partial-protection ratios on NSFnet. We considered four scenarios shown in Table V, whose entries indicate the percentage of traffic requests with a particular partial-protection ratio. The traffic requests are randomly assigned to a partial-protection ratio. We note increasing transponder cost as the proportion of high partial-protection traffic increases—Transponder cost for Scenario I: 109.5; Scenario II: 115.3; Scenario III: 141.6; and Scenario IV: 169.1.

TABLE V
PERCENTAGE OF TRAFFIC REQUESTS BELONGING TO VARIOUS PARTIAL-PROTECTION RATIOS α ON NSFNET TOPOLOGY

α →	0.3	0.6	0.7	0.9	1
Scenario I	50%	20%	10%	5%	5%
Scenario II	30%	30%	20%	10%	10%
Scenario III	20%	20%	30%	20%	10%
Scenario IV	10%	10%	20%	30%	30%

TABLE VI
TRANSPONDER DISTRIBUTION FOR DS-MLR (HEURISTIC) VERSUS DS-MLR (HEURISTIC) UPGRADE ON NSFNET TOPOLOGY

Heuristic	α	Cost	No of Transponders		
			10G	40G	100G
DS-MLR	0.6	388.8	209	46	4
DS-MLR Upgrade	0.6	311.8	79	26	21
DS-MLR	0.9	518.1	216	47	21
DS-MLR Upgrade	0.9	458.6	92	22	42

E. Transponder Upgrade

Table VI shows transponder distribution for DS-MLR (Heuristic) versus DS-MLR (Heuristic) Upgrade on NSFnet by scaling the traffic matrix in Table III five times. We notice from Table VI that DS-MLR Upgrade achieves 10%–20% transponder cost savings over DS-MLR heuristic for partial-protection ratios of 0.6 and 0.9. DS-MLR Upgrade uses many more 100-Gb/s and much fewer 10-Gb/s transponders versus DS-MLR. This is due to transponder upgrade in DS-MLR Upgrade wherein several active lightpaths get upgraded to 100 Gb/s, avoiding the need to set up new lightpaths of lower rates, e.g., 10 Gb/s as shown in Table VI, and thus achieving significant grooming efficiency. The savings are not much for lower partial-protection ratios and traffic loads. In fact, for low partial-protection ratios (e.g., 0.1, 0.3), DS-MLR performs better versus DS-MLR Upgrade on NSFnet with 1-Tb/s aggregative traffic capacity (Table III). We observed similar trends on 24-node USnet topology [7]. This is because transponder upgrade to higher rates, e.g., 100 Gb/s, becomes cost-efficient when traffic loads are high. We also note from our experiments on the Cost 239 network that transponder upgrade is not merely dependent on traffic load, but also on the network topology due to different optical reaches of the various line rates and the associated regeneration cost. DS-MLR performs better than DS-MLR Upgrade on the Cost 239 network for aggregate traffic loads ranging from 350 Gb/s to 3.5 Tb/s. Thus, the network operator can consider both the heuristics and design his network based on whichever gives better savings.

The cost benefits of post-processing in the heuristic are observed up to 10%.

F. Results With 80 Wavelength Channels

Fig. 10 shows the cost benefits of MPP versus DPP with 80 wavelength channels over the NSFnet topology. We note that multipath partial protection requires 44%–60% of cost of dedicated partial protection. Our results are similar to our experimental results with eight wavelengths presented earlier. This shows that our approach achieves good optimization with growth in traffic and backbone capacity and scales well to traffic growth. Table VII shows the transponder cost benefit of MLR networks over SLR networks for multipath partial protection with 80 wavelengths. MLR performs better than each of the three SLR networks for all scenarios. Cost savings of MLR over SLR 10G, 40G, and 100G in our studies is on average 20%, 17%, and 15% as the optimal rate is chosen in MLR from the set of all lines rates for every request, in line with our experimental results with eight wavelengths.

TABLE VII
TRANSPONDER COST FOR MULTIPATH PARTIAL PROTECTION IN SLR VERSUS MLR NETWORKS ON NSFNET TOPOLOGY WITH 80 WAVELENGTHS PER LINK

α →	**0.5**	**0.6**	**0.7**	**0.8**	**0.9**
SLR 100G	602	733.2	833.16	854	868
SLR 40G	633.6	702.9	811.8	884.4	904.2
SLR 10G	647	733	868	924	1026
MLR	537.4	611	694.3	737.8	787.8

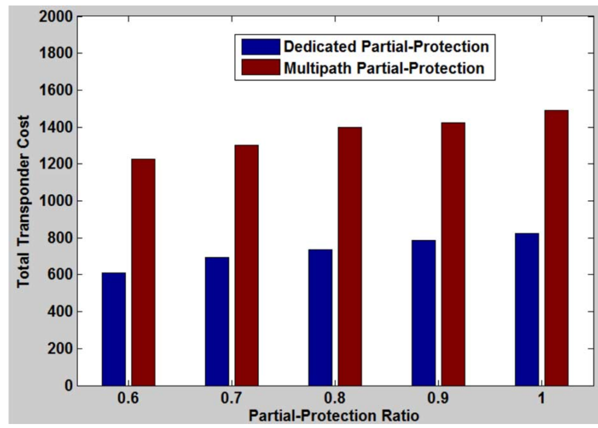

Fig. 10. Performance of dedicated partial-protection and multipath partial-protection with 80 wavelengths per fiber link on NSFnet.

We also considered another cost model from a recent study [18]. The model assumes on–off keying (OOK) modulation for 10 Gb/s, Differential phase-shift keying (DPSK) for 40 Gb/s, and polarization-division-multiplexed quadrature phase-shift keying (PDM-QPSK) for 100 Gb/s. The relative costs of 10/40/100-Gb/s transponders are assumed to be 1/3/6, and transparent reaches are assumed to be 3000 km for 10 Gb/s, 1600 for 40 Gb/s, and 800 for 100 Gb/s. We found that our approach gives good cost savings and shows the promise of our approach with various cost models.

VI. CONCLUSION

We studied the problem of provisioning degraded services in MLR networks, where a service accepts some reduction in bandwidth during link failures for lower cost, called *partial protection*. We used multipath routing to achieve degraded services/partial protection in MLR networks. Multipath routing in MLR networks is challenging compared to SLR networks due to the heterogeneities introduced by multiple rates, and it is a new topic. We developed a mathematical formulation (MILP) for provisioning degraded services in MLR networks using multipath routing and a computationally efficient heuristic considering various partial-protection models. Illustrative results show that significant cost savings can be achieved due to partial protection in MLR networks versus traditional full protection considered in prior works [5]. The cost savings are noted to be in the range 27%–52% in our experiments on the Cost 239 topology and include savings due to both partial protection and multipath routing. The savings due to multipath routing are observed to be greater than due to partial protection, but the savings of partial protection are themselves significant for network operators. Similar results are observed with our experiments on the NSF topology. We note that MLR networks have special advantages with multipath routing due to the volume discount of higher line rates along with the ability in MLR to select and cost-effectively groom over appropriate line rate from the set of multiple line rates. The cost savings are observed in the range of 15%–34% versus SLR. The heuristic has good performance compared to the MILP solution considering its computational efficiency and is highly scalable with large topologies.

References

[1] A. Chiu *et al.*, "Network design and architectures for highly dynamic next-generation IP-over-optical long distance networks," *J. Lightw. Technol.*, vol. 27, no. 12, pp. 1878–1890, Jun. 2009.

[2] A. Mittal, A. Pande, and P. Kumar, "Content-based network resource allocation for real time remote laboratory applications," *Signal, Image, Video Process.*, vol. 4, no. 2, pp. 263–272, 2010.

[3] A. Nag and M. Tornatore, "Optical network design with mixed line rates," *Opt. Switching Netw.*, vol. 6, pp. 227–234, 2009.

[4] C. Meusburger, D. A. Schupke, and A. Lord, "Optimizing the migration of channels with higher bitrates," *J. Lightw. Technol.*, vol. 28, no. 4, pp. 608–615, Feb. 2010.

[5] M. Liu, M. Tornatore, and B. Mukherjee, "New strategies for connection protection in mixed-line-rate optical WDM networks," *J. Opt. Commun. Netw.*, vol. 3, no. 9, pp. 641–650, Sep. 2011.

[6] C. V. Saradhi, M. Gurusamy, and L. Zhou, "Differentiated QoS for survivable WDM optical networks," *IEEE Commun. Mag.*, vol. 45, no. 2, pp. S8–10, Feb. 2004.

[7] R. Roy and B. Mukherjee, "Degraded-service-aware multipath provisioning in telecom mesh networks," in *Proc. OFC/NFOEC*, San Diego, CA, USA, 2008, pp. 1–3.

[8] G. Kuperman, E. Modiano, and A. Narula-Tam, "Analysis and algorithms for partial protection in mesh networks," in *Proc. IEEE INFOCOM Mini-Conf.*, Shanghai, China, 2011, pp. 516–520.

[9] A. Das, C. Martel, and B. Mukherjee, "A partial protection approach using multipath provisioning," in *Proc. IEEE Int. Conf. Commun.*, Dresden, Germany, 2009, pp. 1–5.

[10] S. Misra, G. Xue, and D. Yang, "Polynomial time approximations for multi-path routing with bandwidth and delay constraints," in *Proc. IEEE INFOCOM*, Rio de Janeiro, Brazil, 2009, pp. 558–566.

[11] M. Mahesh, J. Fang, K. M. Sivalingam, and A. K. Somani, "Design and analysis of partial protection mechanisms in groomed optical WDM mesh networks," *J. Opt. Netw.*, vol. 7, no. 6, pp. 617–634, 2008.

[12] P. H. Ho, J. Tapolcai, A. Haque, S. Shen, T. Cinkler, and M. Desroches, "A novel dynamic availability-aware survivability routing architecture with partial restorability," in *Proc. 23rd Bienn. Symp. Commun.*, 2006, pp. 360–363.

[13] R. Aparicio-Pardo, P. Pavon-Marino, and S. Zsigmond, "Mixed line rate virtual topology design considering nonlinear interferences between amplitude and phase modulated channels," *Photon. Netw. Commun.*, vol. 22, no. 11, pp. 230–239, 2011.

[14] K. Christodoulopoulos, K. Manousakis, and E. Varvarigos, "Reach adapting algorithms for mixed line rate WDM networks," *J. Lightw. Technol.*, vol. 29, no. 21, pp. 3350–3363, Nov. 2011.

[15] Y. Chen, N. Hua, X. Zheng, and H. Zhang, "Dynamic connection provisioning in mixed-line-rate optical networks," in *Proc. ACP*, Shanghai, China, 2010, pp. 262–263.

[16] J. Simmons, *Optical Network Design and Planning*. New York, NY, USA: Springer, 2008.

[17] P. Chowdhury, M. Tornatore, A. Nag, E. IP, T. Wang, and B. Mukherjee, "On the design of energy-efficient mixed-line-rate (MLR) optical networks," *J. Lightw. Technol.*, vol. 30, no. 1, pp. 130–139, Jan. 2012.

[18] O. Rival and A. Morea, "Cost-efficiency of mixed 10–40–100 Gbps networks and elastic networks," in *Proc. OFC/NFOEC*, Los Angeles, CA, USA, 2011, pp. 1–3.