

DRAMA-PRO: Disaster Recovery Algorithm with Mitigation Awareness for Protected Elastic Optical Networks

Rujia Zou*, Shih-Chun Lin†, Motoharu Matsuura‡, Hiroshi Hasegawa§, Gustavo B. Figueiredo¶,
Suresh Subramaniam*

*The George Washington University, †North Carolina State University,
‡University of Electro-Communications, §Nagoya University, ¶Federal University of Bahia,

Abstract—Elastic optical networks (EONs) have become a promising option to meet the substantial demand growth for 5G and cloud services. Due to the flexibility in resource assignment, EONs require advanced recovery mechanisms and disaster management to prevent connection loss during large-scale disasters. Traditionally, disaster recovery has focused on restoring the traffic without considering pre-disaster protection. In this paper, we consider multi-class traffic consisting of a mix of unprotected lightpaths and protected lightpaths. By utilizing the mitigation zone strategy proposed in our previous work, we present a disaster recovery algorithm with mitigation awareness for protected EONs called DRAMA-PRO. Simulation results highlight the effects of protection on disaster recovery performance and show that the proposed algorithm outperforms conventional disaster recovery algorithms.

Index Terms—Elastic optical networks, disaster management, protection

I. INTRODUCTION

Elastic optical networks (EONs) have emerged as a highly adaptable solution for core network infrastructures, marked by their ability to allocate resources and precisely assign spectrum dynamically. Unlike traditional fixed-grid networks, EONs adjust bandwidth and spectrum with a finer level of granularity [1]. This capability is achieved through the use of frequency slots (FSs), typically 12.5 GHz each, enabling precise and efficient spectrum utilization [2].

Survivability is an important aspect of EONs. Survivability strategies can be divided into protection and recovery [3]. Protection aims to reserve backup resources to prepare for network failures. As a pre-failure approach, it allows for immediate reaction to network failures with minimal service disruption. However, it requires maintaining redundant backup resources for all potential failures. Protection strategies are typically divided into dedicated backup protection (DPP) and shared backup protection (SBPP). There has been a large amount of work on protection in EONs (e.g., [4]–[7]). For dedicated backup protection, in [4] an integer linear programming(ILP) mode and a heuristic algorithm is proposed for routing and spectrum assignment with DPP. In [6], the authors proposed an SBPP algorithm with an ILP model to minimize the required spare capacity with consideration of transponder tunability and bandwidth-squeezed restoration. In [7], a hybrid protection algorithm is proposed for providing shared or dedicated backup path protection to minimize spectrum utilization. In [5], a p-cycle design is proposed for protection in EONs with spectrum sharing and defragmentation consideration. In our previous work [8], we proposed a p-cycle approach for translucent EONs to minimize the protection bandwidth by using regenerators.

Protection requires the provisioning of significant redundant resources and is usually designed for common failures such as

one or two nodes/links or shared risk groups [6]. Compared to protection, recovery or restoration is a post-failure strategy that aims to find available resources in the damaged network after the failure [9]–[11]. In [12], a novel data center provider services recovery strategy is designed with resource-driven demand matching, but it is not designed for EONs. In [13], the authors proposed a network recovery algorithm to maximize the routed traffic demand after the disaster. However, it mainly focuses on repairing the failed network components instead of service recovery. In [14], a heuristic traffic recovery algorithm is proposed for EONs, incorporating a genetic operator to optimize the service restoration sequence for failed services. In [15], the authors propose a capacity-constrained, maximally spatially-disjoint lightpath algorithm for EONs, specifically to address the provisioning and restoration of disrupted lightpaths in a post-disaster scenario. Still, the bandwidth recovery with degradation is not considered in the above two works.

Clearly, protection cannot be an efficient solution for disaster survivability because of a disaster's potential for extensive network damage and the need for an inordinate amount of redundant resources that have a low probability of ever being activated. Post-disaster recovery using the undamaged part of the network is naturally a better approach. Nevertheless, networks in practice would have to use a combination of protection and recovery – protection for common failures and recovery for unanticipated large-scale network damage, such as disasters. Despite the existence of an extensive amount of research separately on protection and recovery, we are not aware of any work that leverages the backup resources in protected EONs in disaster recovery.

In this paper, we introduce a novel recovery algorithm for protected and transparent EONs called DRAMA-PRO. Our recovery approach is based on the mitigation zone concept introduced in our previous work [16], [17]. The mitigation zone is an area immediately surrounding the disaster zone that aids in disaster recovery by allowing lightpaths within that zone to be recovered with degraded service in the network damaged by a disaster. The idea behind this concept is that it is more acceptable for a service closer to a disaster to be recovered in a degraded fashion rather than degrading a service far away from the disaster. The mitigation zone can be flexibly deployed by network operators to strike a balance between acceptable service degradation and recovery. The effectiveness of the mitigation zone has been demonstrated in [16], [17]. In this paper, we consider multi-class traffic that includes a mix of unprotected and protected connections. DRAMA-PRO aims to retain the protection for protected lightpaths in the damaged network, while also recovering the unprotected

lightpaths. This work is different from our previous work in [16]–[18]. In [16] we first proposed the mitigation zone concept and presented a disaster recovery algorithm for unprotected transparent EONs. In [17] we extended the previous work to translucent but unprotected EONs, and in [18] we proposed a recovery algorithm for service function chain requests in data center EONs. None of these papers considered protected EONs, which is the focus of this paper.

The rest of the paper is organized as follows. The problem statement, network model, and penalty model are presented in Section II. The DRAMA-PRO algorithm is presented in Section III. This is followed by simulation results in Section IV, and the paper is concluded in Section V.

II. NETWORK MODEL AND PROBLEM STATEMENT

A. Network and traffic model

The disaster recovery problem is defined within a network $G(N, E)$, where N represents the node set and E represents the set of links. Each link consists of a pair of fibers in opposite directions. At the time of the disaster, there is a set of active lightpaths denoted as T . A lightpath request is represented as $t(s, d, w, \theta)$, where s and d are the source and destination nodes respectively, and w indicates the data rate. θ is the class of the request. We assume that there are two types of lightpath requests: *first-class* and *normal-class*. A first-class request requires a working path and a protection path, while a normal-class request only needs a working path and is not required to be protected. We assume that first-class requests are protected using dedicated path protection and the working path and protection path do not share any common link or node. All working and protection paths are routed and assigned with spectrum continuity and contiguity requirements. Multiple modulation formats are available, each having different distance limitations and offering different levels of spectrum efficiency. All the lightpaths are assigned the highest feasible modulation format based on their physical distance.

During the recovery after a disaster, for a first-class lightpath request, the protection path of the request should also be recovered or re-accommodated. If the working path is affected but the protection path is not affected, the protection path will be directly used as the working path after the disaster without any data rate degradation. This protection deployment is done right after the disaster and before the recovery starts.

B. Disaster model and mitigation zone

We assume that the disaster zone $D(C_d, R_d)$ is defined as the circular area centered at C_d with a radius of R_d . Any node and any link with at least one endpoint within this area is failed by the disaster. If there is no viable path between a request's source and destination node, or if either the source or destination node fails, the lightpath is regarded as *unrecoverable*. The mitigation zone $M(C_d, R_m)$ is defined as the annulus between the outer edge of the disaster zone and an additional circle centered at the disaster center. These zones are depicted in Fig. 1. The area outside both the mitigation and disaster zones is designated as U .

Each lightpath request t from the set T is classified according to the zones of its source and destination nodes. If either source or destination node is within D , then t belongs to D and is

unrecoverable. If both nodes are outside D but one or both are within M , then t is placed in M . Otherwise, t belongs to U .¹

Based on the motivation of the mitigation zone, we assume that the unaffected lightpath request inside the mitigation zone could be re-accommodated with a degraded data rate to alleviate the bottleneck problem caused by the disaster. Service degradation is allowed for all lightpath requests during recovery, but re-routing is only allowed for a lightpath that is directly affected by the disaster or inside the mitigation zone. Service degradation leads to a penalty, and according to the motivation for the mitigation zone concept, the penalty inside the mitigation zone is kept lower than the penalty outside the mitigation zone to discourage the degradation of a service far away from the disaster. The next subsection presents more details.

Examples of lightpaths are shown in Fig 1. t_1 is an example of a first-class request and t_1^w and t_1^p are the working path and protection path, respectively, before the disaster. Since $t_1 \in U$ and its working path is affected by the disaster, the lightpath request will directly use the unaffected protection path t_1^p as the working path after the disaster. During the recovery, we will look for a new protection path to keep the request protected. The normal-class request t_2 is an example of a lightpath outside the mitigation zone but not affected by the disaster. In this case, t_2 is a candidate for degraded service with the same path as before the disaster. If the lightpath request is inside the mitigation zone, for both first-class request and normal-class request, no matter whether it is affected or not, it will be recovered or re-accommodated (e.g., t_3). If the lightpath request is unrecoverable, the lightpath is dropped (e.g., t_4).

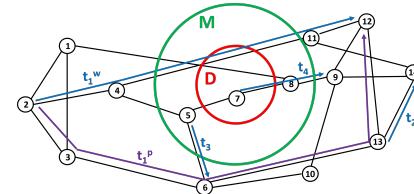


Fig. 1: Examples of mitigation zone and request types.

C. Penalty model

To evaluate recovery efficiency with data rate degradation on lightpaths, we use penalty functions to quantify the service level degradation. The penalty function design is proposed and validated in [16], [17]. In this paper, We assume that each lightpath contributes revenue to the network, which is assumed for simplicity to be equivalent to the data rate of the lightpath. For a first-class lightpath request, the working path and the protection path have the same revenue that equals the data rate of the lightpath request, while the revenue of a normal-class lightpath only comes from its working path. If the data rate of a lightpath is degraded after disaster recovery, there is an incurred penalty. The definition of the mitigation zone implies that the penalty for the same level of degradation should be higher outside the mitigation zone than inside the mitigation zone. While our algorithm can work with any penalty function, in this paper, we use the two penalty functions shown in Fig. 2 and defined in (1) and (2).

¹We assume circular disaster and mitigation zones for simplicity. In reality, the disaster zone can be of arbitrary shape, and the mitigation zone shape and size can be chosen by the network operator. The recovery algorithms we propose take these zones as inputs and are not restricted to circular zones.

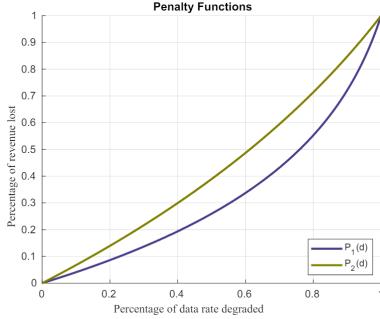


Fig. 2: Penalty functions.

$$P_1(d) = \frac{\log(1 - 0.9 \times d)}{\log(1 - 0.9 \times 1)} \quad (1)$$

$$P_2(d) = \frac{\log(1 - 0.6 \times d)}{\log(1 - 0.6 \times 1)}, \quad (2)$$

where d is the degradation factor, defined as the percentage of the data rate degraded (the reduction percentage of the data rate). There are two different penalty functions: $P_1(d)$ is for a lightpath inside the mitigation zone, and $P_2(d)$ is for a lightpath outside the mitigation zone. The penalty functions describe how the degradation factor correlates with the percentage of revenue lost. The absolute value of the penalty is the value of the penalty function times the total revenue of the lightpath.

In addition, we set up two balance factors: α and β to measure the penalty more reasonably. α is used to quantify the relative penalty for the protection path of the first-class request (relative to the working path of a first-class request), while β is used to quantify the relative penalty for the normal-class request. The actual penalty for these two types of lightpath equals the actual penalty given by the penalty function times the corresponding factor. In this paper, both α and β are set to 0.7. Again, this is a choice that can be set by the network operator, and the algorithm can work with any given choice.

The goal of the disaster recovery problem with protection is to recover the two different classes of affected traffic requests and re-accommodate the unaffected requests with appropriate routing, degradation, and protection, while minimizing the total penalty.

III. DRAMA-PRO

In this section, we present the DRAMA-PRO algorithm. Lightpaths are recovered one by one in a sequence in DRAMA-PRO. First, the recovery order of the recoverable lightpath requests is determined. Then we perform the recovery or re-accommodation with degradation factor selection for each working and protection path of the lightpath requests. The pseudocode of the recovery algorithm is presented in Algorithm 1.

In lines 1 to 2, we initialize an empty request set W and release all the spectrum in the damaged network.² From line 3 to line 9, if the working path is affected but the protection path is not affected, we directly use the protection path as the new working path after the disaster. From lines 11 to 14, we add all the recoverable requests to W .

²This step is done only for recovery path and data rate computation, and does not mean that all ongoing lightpaths are immediately disrupted. After the recovery algorithm completes, some lightpaths may be recovered with a reduced data rate and/or new paths in the damaged network, and that is when the actual rerouting and data rate selection is implemented.

Algorithm 1 DRAMA-PRO Algorithm

Input: $G(N, E)$, T , $D(C_d, R_d)$, $M(C_d, R_m)$

Output: Lightpath requests recovery

```

1: Initialize two empty request sets  $W$ 
2: Release all the spectrum in  $G$ 
3: for each  $t \in T$  do
4:   if  $t$  is first-class then
5:     if  $t^w$  is affected and  $t^p$  is not affected then
6:       Assign  $t$  with the  $t^p$  as the working path
7:     end if
8:   end if
9: end for
10: for each  $t \in T$  do
11:   if  $t \notin D$  (i.e.,  $t$  is recoverable) then
12:     Add  $t$  to  $W$ 
13:   end if
14: end for
15: Calculate the shortest path for each  $t$  in  $W$ 
16: Sort all  $t \in W$  in non-increasing order of PRE
17: for each  $t \in W$  do
18:   if  $t^w$  is affected and  $t^p$  is not affected then
19:     Working path is the protection path before the disaster
20:   else
21:     if  $t \in M$  or  $t$  is affected then
22:       Calculate the working path with the CR algorithm
23:     else
24:       Select the same working path as before the disaster
25:     end if
26:     Calculate the degradation option for the working path
27:     Assign the working path with the selected path and degradation option
28:   end if
29:   if  $t$  is first-class then
30:     if  $t \in M$  or  $t$  is affected then
31:       Calculate the protection path with the CR algorithm
32:     else
33:       Select the same protection path as before the disaster
34:     end if
35:     Calculate the degradation option for the protection path
36:     Assign the protection path with the selected path and degradation option
37:   end if
38: end for

```

A. Order of recovery

In lines 15 to 16, we calculate the order of recovery. All the recoverable lightpath requests are sorted in terms of their Potential Revenue Efficiency (PRE) in non-increasing order. The PRE is defined as the ratio of the actual revenue of the lightpath to the number of FSs needed if the lightpath is routed with the shortest path. The PRE of a first-class lightpath request is the sum of the actual revenue of the working path and protection path to the sum of the number of FSs needed. The actual revenue is designed as the revenue factor of the lightpath request times the original revenue (equal to the data rate). The actual revenue of the lightpath shows how much revenue we can maintain. The number of FSs shows the general cost of this recovery or re-accommodation. Even though we may not use the shortest path to re-assign the lightpath eventually, the shortest path is an estimation of the spectrum cost.

B. Routing

Once we have the order, we start to recover the requests one by one. In lines 17 to 25 we calculate the working path. For each lightpath request inside the mitigation zone or affected by the disaster, we calculate the working path by using the Cost-Routing algorithm (CR) as follows. In the CR algorithm, we calculate the cost of each path in the K shortest path

pool and select the path with the lowest cost. The cost of a path is defined as $Cost = H \times M \times \gamma + \nu_F/\nu_A + \nu_F/\nu_L$, where H is the number of hops and M is the modulation factor of the path. The modulation factor for each lightpath is determined based on the physical distance of the path, with the highest feasible modulation format being selected. For modulation formats BPSK, QPSK, 8QAM, and 16QAM, the spectrum efficiencies are 1, 2, 3, and 4 bps/Hz, respectively. Consequently, the corresponding modulation factors are set at 1, 0.5, 0.34, and 0.25. γ is a large constant number used to make $H \times M$ dominate the value of cost. ν_F is the number of FSs needed without degradation and ν_A is the number of available FSs on the path with spectrum continuity constraint, while ν_L is the size of the largest spectrum fragment on the path. In case the value of ν_A is 0, the path is dropped from the pool. The value of $H \times M$ is used to measure the spectrum cost of the path, which dominates the routing selection, to save spectrum as much as possible. For two paths with the same spectrum cost, ν_F/ν_A and ν_F/ν_L are used to evaluate the load balance of the paths. The path that has more available FSs has a lower cost and is encouraged to be selected. If there is no working path available, we directly block the request.

C. Degradation selection

In line 26, once the working path is found, we select the degradation option for the working path. The modulation format is determined by the physical length of the lightpath, and we calculate the number of FSs required to maintain service without degradation. For each potential degradation option, we compute the Potential Penalty (PP), which is the sum of the Current Penalty (CP) and Future Penalty (FP). The degradation option with the lowest PP is then chosen for the lightpath.

CP is the actual penalty of the current degradation option, calculated with the penalty function. FP is calculated based on the consequence if the LP is assigned with a degradation option: $FP = \nu_W \times \nu'_F/\nu_A \times p'/p$, where ν_W and ν'_F are the amount of maintained data rate and FSs of the given degradation option, ν_A is the number of available FSs as used in the CR algorithm, and p' and p are the number of requests that have already been recovered and the number of total recoverable requests. ν_W is the amount of data rate that the algorithm gives to this request. ν'_F/ν_A is designed as the factor to measure the spectrum utilization; a larger value means the more FSs are used and the higher FP we will have. p'/p is used to account for the progress in the recovery procedure. As the algorithm proceeds towards the end of the recovery process, we may have a limited number of available FSs, and these must be judiciously spread over the remaining lightpaths yet to be recovered. The FP is accordingly scaled up so that it becomes more dominant than the CP.

In line 27, once the degradation of the working path is selected, we assign the working path first. From lines 29 to 37, if the request is a first-class request, we then find the dedicated protection path and its degradation. We first remove the path that has a joint node or link with the working path from the K shortest path pool, and then we keep using the CR algorithm to find the dedicated protection path. The degradation of the protection path is also selected with PP, CP, and FP, similar to the working path. Note that during the degradation selection of the protection path, the recovered protection bandwidth should

not be larger than the recovered working bandwidth, but it could end up being smaller.

The time complexity of the CR routing algorithm is $O(K \cdot F \cdot N)$, as the number of hops in a path is N . After that, for the degradation selection, the time complexity is $O(F \cdot N)$. Therefore, the time complexity of DRAMA-PRO is $O(K \cdot F \cdot N)$.

IV. SIMULATION RESULTS

We now present some simulation results. The NSF network (14 nodes and 22 links, shown in Fig. 3) topology is used for testing. Each fiber has 352 FSs. To thoroughly evaluate the performance of DRAMA-PRO, two typical disasters are tested, shown as red circles in Fig. 3. D_1 happens at the center of the network and D_2 happens near the edge of the network. In addition, we also show the average performance for random disasters. The examples of the mitigation zones of D_1 are shown in Fig. 3 as M_s , M_m , and M_L for the small, medium, and large mitigation zones. We also test recovery with the *no mitigation zone* and the *all mitigation zone*, i.e., when the whole network outside D is the mitigation zone.

For the pre-disaster scenario, a set of lightpath requests is generated with source and destination nodes selected uniformly at random. The data rates for these lightpaths are set at 40 Gbps, 100 Gbps, and 400 Gbps, assigned with probabilities of 0.2, 0.5, and 0.3, respectively. We utilize four modulation formats: BPSK, QPSK, 8-QAM, and 16-QAM. The physical distance limitations and the number of required FSs for each data rate adhere to the parameters previously established in [17].

Before the disaster happens, 1000 lightpath requests are generated. These lightpaths are assigned in random order with the shortest path in the number of hops and first fit spectrum assignment. A lightpath is blocked if the selected path does not have available slots. For a first-class request, if a protection path cannot be found due to insufficient slots, the lightpath request is not established (both the working path and the protection path will not be assigned).

In the simulation, 50 different lightpath request sets are generated and average results are presented. As there are no existing algorithms in the literature for this problem, we evaluate the performance of DRAMA-PRO with algorithms in which one or more of the three performance-optimizing features of DRAMA-PRO is disabled. The three features are (a) order of recovery using PRE sorting, (b) CR routing, and (c) degradation selection. The baseline algorithms use either sorted recovery (SR) order or random recovery (RR) order, shortest path (SP) or CR routing (CR), and degradation (D) or no degradation (ND). For example, RR-SP-ND is the most basic baseline algorithm in which the requests are recovered with RR order, SP routing, and no degradation. As another example, SR-SP-ND uses SR recovery but with SP routing and no degradation. All the shown penalty values consider the protection lightpath factor α and normal-class lightpath factor β , both of which are set to 0.7.

First, we show the performance results for different mitigation zones of D_1 with the percentage of first-class requests set to 50%. In Fig. 4, we can see that the total penalty of DRAMA-PRO decreases as the size of the mitigation zone increases, which shows that the mitigation zone can benefit recovery and solve the bottleneck problem. We can also observe that the

TABLE I: Recovery of normal and first-class requests for D_1 .

| | $ T_f $ | $ T_n $ | $ R_f $ | $ R_n $ | BW_f^w | BW_f^p | BW_n | P_f^w | P_f^p | P_n | P_{total} |
|-----------|------------------|------------------|---------------|---------------|----------|----------|--------|---------|---------|--------|-------------|
| RR-SP-ND | 173.7 (158.9) | 364.7 (314.6) | 150.7 (94.8%) | 264.1 (83.9%) | 132.1 | 86.5 | 125.9 | 979.3 | 5608.7 | 4615.6 | 11203.7 |
| SR-SP-ND | | | 152.9 (96.2%) | 263.7 (83.8%) | 133.1 | 93.8 | 124.2 | 860.2 | 4973.9 | 5117.4 | 10951.6 |
| RR-CR-ND | | | 153.5 (96.5%) | 269.4 (85.6%) | 133.8 | 87.5 | 125.9 | 722.4 | 5425.8 | 4334.7 | 10482.9 |
| RR-SP-D | | | 153.7 (96.6%) | 275.3 (87.4%) | 133.1 | 87.1 | 125.6 | 711.9 | 5406.2 | 4029.6 | 10147.8 |
| DRAMA-PRO | | | 157.8 (99.3%) | 289.2 (91.9%) | 135.1 | 93.4 | 126.9 | 396.5 | 4859.9 | 4120.2 | 9376.7 |

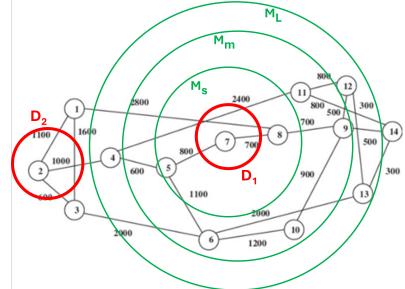
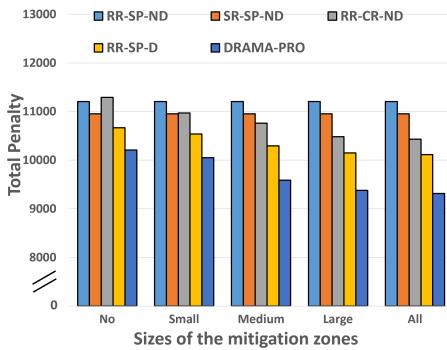
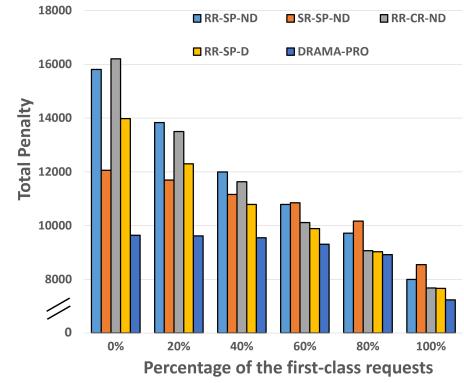


Fig. 3: 14-node NSF network.

Fig. 4: Performance results for D_1 .

combination of all three optimization features of DRAMA-PRO is useful, as DRAMA-PRO is better than all the baseline algorithms. When the mitigation zone extends to the whole network, the penalty of DRAMA-PRO is reduced by 16%, 14%, 10%, and 8% over that of RR-SP-ND, SR-SP-ND, RR-CR-ND, and RR-SP-D. We also note that SR-SP-ND, RR-CR-ND, and RR-SP-D have lower penalties compared to the most basic recovery algorithm RR-SP-ND. Among the baseline algorithms, RR-SP-D has the lowest penalty indicating that recovery with degraded service is very helpful in decreasing the penalty. The penalties of RR-CR-ND and RR-SP-D decrease as the size of the mitigation zone increases because they benefit from the mitigation zone while the penalties of RR-SP-ND and SR-SP-ND are fixed for different mitigation zone sizes because they do not have any re-routing or degradation selection advantage.

In Table I, we show more detailed results when the size of the mitigation zone is *large*. $|T_f|$ and $|T_n|$ are the number of first-class and normal-class requests before the disaster, followed by the number of recoverable first-class and normal-class requests in parentheses. $|R_f|$ and $|R_n|$ are the number of recovered first-class and normal-class requests. The percentage numbers are the recovery ratio, defined as the ratio of the number of recovered requests to the number of recoverable requests. BW_f^w , BW_f^p , and BW_n are the average recovered bandwidth (in

Fig. 5: Performance results for different class ratios for D_1 with *large* mitigation zone.

Gbps) of the first-class requests' working path and protection path, and the normal request lightpath, respectively. P_f^w , P_f^p , and P_n are the average penalty for the three different lightpath types. P_{total} is the total penalty. We can see that DRAMA-PRO can recover more requests and the average bandwidth is generally higher than for baseline algorithms. We note that the penalty improvement for DRAMA-PRO is much more than the bandwidth improvement, because the penalty is weighted differently for different types of lightpaths and DRAMA-PRO is able to take this into account during recovery.

In Fig. 5, we show results for different ratios of first-class to normal-class requests. We note that DRAMA-PRO continues to perform better than the other algorithms for all ratios, but the overall penalty sharply decreases as the ratio of first-class requests becomes high. The reason is that the total number of requests before the disaster is lower when the percentage of first-class requests is high and the total revenue is also lower. In the 20% case, the penalty of DRAMA-PRO is reduced by 30%, 18%, 28%, and 21% over RR-SP-ND, SR-SP-ND, RR-CR-ND, and RR-SP-D. We can see that the penalty reduction between DRAMA-PRO and baseline algorithms decreases for high percentages of first-class requests. This is because much of the spectrum is occupied after protection is "activated" before the recovery, so the flexibility due to the mitigation zone is limited.

In Fig. 6, we show performance results when the disaster is D_2 . We again see that DRAMA-PRO always has the lowest penalty; for instance, when the entire network is the mitigation zone, the penalty is reduced by 30%, 27%, 28%, and 15% over RR-SP-ND, SR-SP-ND, RR-CR-ND, and RR-SP-D. Interestingly, compared to the D_1 disaster scenario, the total penalty values are lower because the bottleneck problem is more severe in D_1 because D_2 occurs at the edge of the network. Both D_1 and D_2 cause one failed node, so the number of recoverable requests is the same on average. However, the penalty decrease percentage of DRAMA-PRO over the baseline algorithms is

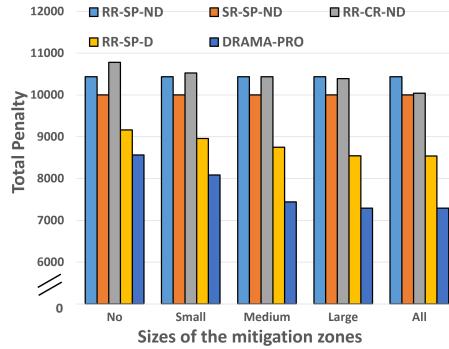
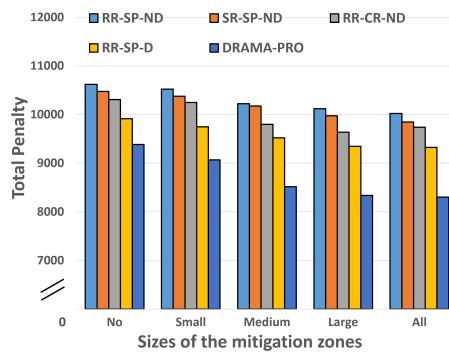
Fig. 6: Performance results for D_2 .

Fig. 7: Performance results for random disasters.

larger with disaster D_2 . This is because the flexibility of degradation selection and rerouting during recovery and re-accommodation is limited by the bottleneck problem in D_1 .

We show performance results for random disasters in Fig. 6. In this case, one node is randomly selected as the center of the disaster, while the size of the disaster is uniformly distributed in the range [100, 400] km. Results are the average of 10 distinct disasters. We can see that DRAMA-PRO has the lowest penalty and the penalty decreases as the mitigation zone size increases. When the size of the mitigation zone is *all*, the penalty of DRAMA-PRO is reduced by 17%, 15%, 14%, and 11% over baseline algorithms.

V. CONCLUSION

Disaster management is an important problem in EONs. EONs typically have a mix of protected and unprotected lightpaths, but protection is designed for common failures such as one or two links/nodes. In this paper, we introduce DRAMA-PRO, which leverages the backup resources of the protected lightpaths in disaster recovery. DRAMA-PRO also uses the concept of the mitigation zone to recover or re-assign multi-class lightpath requests with degraded service. Our results demonstrate that DRAMA-PRO outperforms baseline algorithms across various scenarios and highlights the effectiveness of the mitigation zone in enhancing disaster recovery efforts.

ACKNOWLEDGEMENT

This work was supported in part by NSF grants CNS-1818858 and CNS-2210343.

REFERENCES

- [1] M. Jinno, H. Takara, B. Kozicki, Y. Tsukishima, Y. Sone, and S. Matsuoka, "Spectrum-efficient and scalable elastic optical path network: architecture, benefits, and enabling technologies," *IEEE communications magazine*, vol. 47, no. 11, pp. 66–73, 2009.
- [2] Y. Hirota, H. Tode, and K. Murakami, "Multi-fiber based dynamic spectrum resource allocation for multi-domain elastic optical networks," in *2013 18th OptoElectronics and Communications Conference held jointly with 2013 International Conference on Photonics in Switching (OECC/PS)*, June 2013, pp. 1–2.
- [3] G. Shen, H. Guo, and S. K. Bose, "Survivable elastic optical networks: survey and perspective," *Photonic Network Communications*, vol. 31, no. 1, pp. 71–87, 2016.
- [4] M. Klinkowski and K. Walkowiak, "Offline RSA algorithms for elastic optical networks with dedicated path protection consideration," in *2012 IV International Congress on Ultra Modern Telecommunications and Control Systems*. IEEE, 2012, pp. 670–676.
- [5] M. Ju, F. Zhou, S. Xiao, and H. Wu, "Leveraging spectrum sharing and defragmentation to p-Cycle design in elastic optical networks," *IEEE Communications Letters*, vol. 21, no. 3, pp. 508–511, 2016.
- [6] G. Shen, Y. Wei, and S. K. Bose, "Optimal design for shared backup path protected elastic optical networks under single-link failure," *Journal of Optical Communications and Networking*, vol. 6, no. 7, pp. 649–659, 2014.
- [7] N. G. Anoh, M. Babri, A. D. Kora, R. M. Faye, B. Aka, and C. Lishou, "An efficient hybrid protection scheme with shared/dedicated backup paths on elastic optical networks," *Digital Communications and Networks*, vol. 3, no. 1, pp. 11–18, 2017.
- [8] R. Zou, S. Subramaniam, H. Hasegawa, and M. Jinno, "P-cycle design for translucent elastic optical networks," in *2019 IEEE Global Communications Conference (GLOBECOM)*. IEEE, 2019, pp. 1–6.
- [9] Z. Li, M. Ibrahim, Y. Zhao, B. Mukherjee, J. Zhang, and M. Tornatore, "Content evacuation in inter-dc optical networks under post-disaster cascading failures," in *2024 International Conference on Optical Network Design and Modeling (ONDM)*, 2024, pp. 1–5.
- [10] Y. Liu, F. Zhou, C. Chen, Z. Zhu, T. Shang, and J.-M. Torres-Moreno, "Disaster protection in inter-datacenter networks leveraging cooperative storage," *IEEE Transactions on Network and Service Management*, vol. 18, no. 3, pp. 2598–2611, 2021.
- [11] S. Ferdousi, M. Tornatore, F. Dikbiyik, C. U. Martel, S. Xu, Y. Hirota, Y. Awaji, and B. Mukherjee, "Joint progressive network and datacenter recovery after large-scale disasters," *IEEE Transactions on Network and Service Management*, vol. 17, no. 3, pp. 1501–1514, 2020.
- [12] S. Sahoo, S. Xu, S. Ferdousi, Y. Hirota, M. Tornatore, Y. Awaji, and B. Mukherjee, "Strategic cooperation among datacenter providers and optical-network carriers for disaster recovery," in *GLOBECOM 2022 - 2022 IEEE Global Communications Conference*, 2022, pp. 2151–2156.
- [13] H. Yu and C. Yang, "Partial network recovery to maximize traffic demand," *IEEE communications letters*, vol. 15, no. 12, pp. 1388–1390, 2011.
- [14] S. Li, R. Gu, G. Zhang, Y. Wang, Y. Wang, and Y. Ji, "Order aware service recovery algorithm in elastic optical network with multiple failures," in *2019 International Conference on Networking and Network Applications (NaNA)*, 2019, pp. 135–141.
- [15] M. Ashraf, S. Idrus, R. Butt, and F. Iqbal, "Post-disaster least loaded lightpath routing in elastic optical networks," *International Journal of Communication Systems*, vol. 32, no. 8, p. 3920, 2019.
- [16] R. Zou, H. Hasegawa, and S. Subramaniam, "DRAMA: disaster management algorithm with mitigation awareness for elastic optical networks," in *2021 17th International Conference on the Design of Reliable Communication Networks (DRCN 2021)*, Apr. 2021.
- [17] R. Zou, H. Hasegawa, M. Jinno, and S. Subramaniam, "DRAMA+: Disaster management with mitigation awareness for translucent elastic optical networks," *IEEE Transactions on Network and Service Management*, vol. 19, no. 3, pp. 2587–2599, 2022.
- [18] R. Zou, S.-C. Lin, M. Matsuura, H. Hasegawa, and S. Subramaniam, "DRAMA-DC: disaster recovery algorithm with mitigation awareness in data center EONs," in *2024 International Conference on Optical Network Design and Modeling (ONDM)*, Madrid, Spain, May 2024.