

DRAMA-DC: Disaster Recovery Algorithm with Mitigation Awareness in Data Center EONs

Rujia Zou*, Shih-Chun Lin[†], Motoharu Matsuura[‡], Hiroshi Hasegawa[§], Suresh Subramaniam*

*The George Washington University, {rjzou,suresh}@gwu.edu

[†] North Carolina State University, slin23@ncsu.edu

[‡] University of Electro-Communications, m.matsuura@uec.ac.jp

[§] Nagoya University, hasegawa@nuee.nagoya-u.ac.jp

Abstract—Elastic optical networks (EONs) have become a highly effective solution to address the exploding traffic demand driven by 5G/6G and cloud applications. Leveraging advanced optical technologies, EONs offer high allocation efficiency and flexibility. With the network function virtualization paradigm, critical network functionalities can be instantiated as software-based virtual network functions (VNFs) deployed within the network's data centers. An end-to-end service can then be composed as an ordered set of VNFs called as service function chain (SFC). Developing effective strategies for recovering SFCs after a disaster is a crucial challenge. In this work, we propose a new disaster recovery algorithm for data center EONs called DRAMA-DC. An auxiliary graph framework is developed to recover the affected SFCs and minimize the blocking ratio. Simulation results demonstrate the superior performance of DRAMA-DC in comparison to a baseline recovery algorithm, and the performance tradeoffs for recovery parameters.

Index Terms—Elastic optical networks, data centers, disaster management

I. INTRODUCTION

Elastic optical networks (EONs) have become a promising solution for core networks due to their ability to flexibly allocate resources and assign spectrum [1]. EONs leverage advanced technologies and flexible grid spectrum allocation, enabling high-efficiency utilization. In EONs, the network lightpath is allocated bandwidth in terms of frequency slots (FSs), each of which is 12.5 GHz [2].

Virtual network functions (VNFs) are an efficient and cost-effective tool for network operators [3]. VNFs are software-based instances that replicate the functionalities traditionally performed by dedicated hardware devices, such as firewalls, flow controllers, and intrusion detection systems [4]. These VNFs are placed in the data centers (DCs) of the network [5]. By virtualizing these functions, VNFs offer network operators unparalleled flexibility and resource optimization. Based on the design of VNFs, Service Function Chains (SFCs) play a pivotal role in the context of network services [6]. Each SFC has a set of VNF requirements. The establishment of an end-to-end service between the source node and the destination node requires a specified set of VNFs to be executed in a predetermined sequence [5], [7]. Therefore, during the assignment of the SFC request, network operators need to find appropriate DCs to execute the VNFs. A lightpath is then established from the source node to the destination node of the SFC, through the selected DCs.

Survivability holds significant importance in EONs. Strategies for survivability can be categorized into protection and restoration/recovery approaches [8]. Disaster management is a special case of survivability. In this case, the robustness of the network faces significant challenges in the face of unexpected large events, such as natural disasters or large-scale network failures, which can disrupt network connectivity

severely and impact critical services. Protection is typically not suitable due to the need for a substantial amount of backup resources to guard against possible, but unlikely, large-scale disasters. Therefore, recovery is a more favorable solution. In [9], a network component recovery algorithm is proposed to maximize the routed traffic demand after the disaster. In [10], the investigation of traffic rerouting and flooding after a disaster in a WDM network has been proposed. In our previous work [11], an integer linear program (ILP) and a disaster recovery algorithm for translucent EONs are proposed. In [7], a disaster protection algorithm is presented.

This paper introduces a novel approach to disaster recovery for data center EONs, based on the concept of mitigation zone, first proposed in [11], [12]. The mitigation zone is an area of the network, usually proximal to the disaster, inside which degraded service recovery is tolerated more than outside of the zone. In this study, SFCs within the mitigation zone are recovered with a relaxed latency threshold to address the congestion problem caused by a large-scale network failure. The advantages of the mitigation zone have been validated in [11], [12], where we proposed a set of lightpath recovery algorithms in an EON with possibly degraded service. In these previous papers, the mitigation zone brings flexibility in the spectrum domain. In contrast to those papers, we consider disaster management in data center EONs and explore the benefit of mitigation in the latency domain in this paper. A recovery algorithm called DRAMA-DC is proposed using an auxiliary graph framework, where both the VNFs and the lightpath of the affected SFC are recovered. The latency constraint of an SFC that lies within the mitigation zone is relaxed to help with recovery. As far as we are aware, this is the first paper to address disaster recovery in the context of data center EONs.

II. SYSTEM MODEL AND PROBLEM STATEMENT

A. Network and service model

Our model of the data center EON is based on the model presented in [5], [7]. We are given a network $G(N, E)$, where N is the set of optical crossconnect nodes and E is the set of links, with each link consisting of a pair of oppositely-directed fibers. Some of the nodes are connected to a data center (DC) which hosts different types of VNFs. The set of DCs is denoted as B .

There is a set of VNFs F , each of which can be hosted by any of the DCs. A VNF service can be provided by a DC if it can host a VNF instance. In each DC, different VNFs have separate processors, and all the processors in different DCs have the same computational capacity limitation. The computational capacity is defined in units of Giga Operations Per Second (GOPS).

At the instant of the disaster, a set of active SFC requests S is given. Each SFC request is denoted as $s(src, des, r, V, V_{load})$, where src , des , and r represent the source node, destination nodes, and lightpath data rate requirement. V represents a sequence of VNFs, and V_{load} represents the computational loads of VNFs (in units of GOPS) in V respectively. The same VNF that is requested by different SFC requests may have different computational loads.

For each s , a lightpath (LP) is established from src to des through the DCs that have VNF instances in V successively. Examples are shown in Fig. 1. s_1 and s_2 are SFC requests. For s_1 , src and des are node 1 and node 6. V is (1, 3) which means s requests VNF_1 and VNF_3 (in that order) with computational loads of 2 and 4, respectively. s_1 can be regarded as a virtual path from its source node to the destination node, with VNFs executed in the middle of the path. The virtual path of s_1 is shown at the bottom of Fig. 1. In this case, the lightpath of the SFC originates from node 1. To fulfill the VNF request requirement, the lightpath is routed to DC A on node 2 and B on node 5 for VNF_1 and VNF_3 , respectively. The lightpath eventually terminates at node 6. For s_2 , there is only one VNF and it is processed in DC A on node 2.




Fig. 1. Examples of the SFC requests

During the LP routing of each SFC request, spectrum resources are assigned to the LP. We assume that each SFC has its own dedicated lightpath and spectrum resources [5], [7], [13]. There exist various modulation formats, each with a distinct physical distance limitation and spectrum efficiency. The lightpath of an SFC is allocated the highest modulation format feasible for its path length, and spectrum assignment follows continuity and contiguity constraints.

We assume that DCs have the capability to 3R regenerate optical signals, which means they can perform spectrum conversion and allow higher modulation formats for lightpaths [14], [15].

B. Disaster model

The *disaster zone* $D(C_d, R_d)$ is defined as the area within a circle with center C_d and radius R_d . We consider any node (including DCs) within the disaster zone and any link with either of its endpoints within the disaster zone as failed as a result of the disaster. Following the disaster, if there is no available path from an SFC's source node to its destination node, or if either the source or destination node fails, the SFC is defined as *unrecoverable*.

The *mitigation zone* $M(C_m, R_m)$ is described as the annulus bordered by the disaster zone circle and the circle centered at $C_m = C_d$ with a radius of $R_d + R_m$. The region outside both the mitigation zone and disaster zone is referred to as U .

Each SFC $s \in S$ is categorized as being in one of three zones - D , M , or U - based on the locations of their source and destination nodes. If either the source or destination node

is within D , then $s \in D$ and s is unrecoverable. For all recoverable SFCs, if the source and/or destination node is within M , then $s \in M$; otherwise, $s \in U$.

Every SFC experiences an end-to-end latency, as described further below. The purpose of the mitigation zone is to allow post-disaster recovery with a relaxed latency for those SFCs inside the mitigation zone. For this purpose, we introduce a latency threshold, which is activated after a disaster happens. The latency threshold is considered as a constraint on the latency of an SFC after recovery. For an SFC outside the mitigation zone, i.e., $s \in U$, the latency threshold is set to the SFC latency before the disaster. In other words, there is no latency relaxation for SFCs outside the mitigation zone and they are expected to be recovered without degradation. For an SFC inside the mitigation zone, its latency threshold is set to a relaxation factor α ($\alpha \geq 1$) times its latency before the disaster. A higher α means more relaxation and leads to a larger room for recovery in terms of routing and DC selection. Note that $\alpha = 1$ here means that there is no relaxation, exactly as outside the mitigation zone. In this case, the mitigation zone can be considered to be non-existent.

Based on these three different zones, each SFC s is classified as one of the following four types (examples of SFCs are shown in Fig. 2):

- If $s \in M$, s is a recovery candidate with a new lightpath, DC, and relaxed latency threshold (e.g., s_2).
- If $s \in U$ and is affected by the disaster, s is a candidate for recovery without latency threshold relaxation (e.g., s_1).
- If $s \in U$ and is not affected, s will not be touched (e.g., s_4).
- If the SFC is unrecoverable, the SFC is dropped (e.g., s_3).




Fig. 2. Examples of SFC types and zones. The region defined by the red circle indicates the disaster zone, while the mitigation zone is the area enclosed by the green circle and excluding the disaster zone.

C. Latency model

The latency of an SFC is calculated as the total of the propagation latency of the SFC lightpath and the processing latency of all its VNFs. The propagation latency is defined as $l_{prop} = p/c'$, where p is the physical distance of the lightpath, and c' is the propagation speed of optical signals in fiber. We adopt the simple M/M/1 queuing model proposed in [5] for the processing latency of a VNF in a DC. This is given by:

$$l_{proc} = \frac{1}{\sigma - \tau}, \quad (1)$$

where σ is the computational capacity limitation of a VNF processor, and τ is the total load (used computational capacity) of the processor.

During recovery, an SFC can be blocked due to (a) spectrum unavailability, (b) insufficient computational capacity in the DC, or (c) inability to meet the latency constraint. We formulate the goal of disaster recovery in data center EONs as follows: Re-assign resources (lightpath, spectrum, DCs for VNFs) for the recoverable SFCs with the aim of minimizing the overall blocking ratio.

This problem is difficult since it includes the LP routing and spectrum assignment, which itself is an NP-hard problem in general. In [12], we presented a heuristic algorithm to recover LPs with a mitigation zone, but no SFCs or DCs were considered in that work, and the service degradation was in the form of lightpath bandwidth reduction. In this work, we introduce a novel algorithm for recovering SFCs with relaxed latencies for data center EONs.

III. DRAMA-DC

A. Auxiliary graph framework

We begin by presenting the auxiliary graph framework utilized to address the recovery problem.




Fig. 3. Example of the auxiliary graph.

For each SFC, we generate an auxiliary graph (Fig. 3). The source node and the destination node are created first. For each VNF requirement, we create a node for each DC that hosts this VNF. For example, suppose an SFC is requesting VNF_1 and VNF_2 successively, and VNF_1 is placed in DCs A and B while VNF_2 is placed in DCs B, C, and D. In this case, two candidate DC nodes DC_1^A and DC_1^B are created for VNF_1 , while three candidate DC nodes DC_2^B , DC_2^C , and DC_2^D are created for VNF_2 . These candidate DC nodes are used to represent DC selection options for this SFC since the lightpath of this SFC should be routed across one of them to meet VNF requirements. In order to find the appropriate DC for each VNF from its options, an attached dummy node is created for each candidate DC node. For instance, $DC_1^{A'}$ is created as the dummy node to DC_1^A .

An edge is created between each candidate DC node to its dummy node, shown as blue edges in the figure. This type of edge is named as *data center edge* and will be assigned a weight to measure the suitability of the DC to process a VNF.

Then we add edge(s) from all the pair(s) of the dummy DC nodes of VNF_n to the candidate DC nodes of VNF_{n+1} . These edges represent the lightpaths between where VNF_n could potentially be executed for this SFC and where VNF_{n+1} could potentially be executed. An example of such an edge is the one between $DC_1^{A'}$ and DC_2^D . We also add edge(s) between the source node and all the first VNF's candidate DC nodes (e.g., Src to DC_1^A) and edge(s) between all the last VNF's dummy nodes to the destination nodes (e.g., $DC_2^{C'}$ to Des). This type of edge is named as *lightpath edge* because each of these edges constitutes a lightpath in the physical network. Examples of data center edges and lightpath edges are shown in Fig. 3.

Next, we add weights to the auxiliary graph edge. The weight of each edge is calculated as follows:

$$Weight = \tilde{C}_s + \frac{1}{\alpha} \times \tilde{C}_l, \quad (2)$$

where $\tilde{C}_s$ and $\tilde{C}_l$ are the spectrum cost and latency cost, normalized by the maximum actual spectrum cost and maximum actual latency cost over all the edges (explained further below), respectively. $\tilde{C}_s$ and $\tilde{C}_l$ are used to evaluate the spectrum and latency cost if the SFC is routed with this edge. Since the actual spectrum and latency cost are measured in different units, we use normalized values to balance the two types of cost. (Normalization makes the normalized costs fall between 0 and 1.) Here, the latency relaxation factor α is used as a multiplier of the latency cost $\tilde{C}_l$ in the edge weight. In other words, the larger α we have, the lower the latency cost's effect on the final edge weight. Larger α means that the latency threshold is larger, so the spectrum cost should be considered more than the latency cost.

Since $\tilde{C}_s$ and $\tilde{C}_l$ of each edge are normalized by the maximum actual spectrum cost and latency cost over all the edges, we have to calculate the actual spectrum cost C_s and actual latency cost C_l for each edge. We explain this calculation below.

The weights of *data center edges* are chosen to model the processing latency of the corresponding VNF within the DC. Accordingly, the actual spectrum cost C_s is set to 0, and the actual latency cost C_l is set to the processing latency of the corresponding VNF in this DC, calculated using eq. (1). If the VNF cannot be processed at this DC due to unavailable capacity, this *data center edge* will be removed from the auxiliary graph.

For all the *lightpath edges*, the actual spectrum cost and latency cost are determined by the selected path among the k-shortest paths (in the physical network) between the end nodes of the edge. For example, for the edge Src to DC_1^A in the auxiliary graph, we compute k-shortest paths between node Src and DC A in the physical network. The path with the lowest weight is selected as the corresponding lightpath of the edge in the auxiliary graph. The weight of each path is defined as follows:

$$Weight_p = \tilde{C}_{p,s} + \tilde{C}_{p,l}, \quad (3)$$

where $\tilde{C}_{p,s}$ and $\tilde{C}_{p,l}$ are the spectrum cost and latency cost of the path normalized by the maximum spectrum cost and maximum latency cost among all the k-shortest paths. For each path, the actual spectrum cost $C_{p,s}$ and actual latency cost $C_{p,l}$ are calculated as follows. The actual latency cost is defined as the propagation latency of the path. The actual spectrum cost is defined as follows:

$$C_{p,s} = \frac{H \times M}{L + F}, \quad (4)$$

where H is the number of hops on the path, and M is the modulation factor (reciprocal of the spectral efficiency) determined by the highest possible modulation format based on the physical distance of the segments. Considering that the corresponding spectrum efficiencies for BPSK, QPSK, 8QAM, and 16QAM are 1, 2, 3, and 4 bps/Hz, the modulation factors are accordingly set as 1, 0.5, 0.34, and 0.25, respectively. L is the length (in terms of frequency slots or FSs) of the longest spectral fragment on the path, and F is the total number of available FSs on the path with continuity and contiguity

constraints. If L (and F) is 0, this path is unavailable and we remove the path from this k-shortest path pool. If all the k-shortest paths are unavailable, we remove this lightpath edge from the auxiliary graph. The reasoning for using this cost function is the following. Lower hop length and modulation factors should be preferred for a path, and the cost should be lower if more spectrum is available.

Note that if two consecutive VNFs are processed in the same DC, then the corresponding *lightpath edge* will have a weight of 0 because there is no physical lightpath and hence there is no spectrum or latency cost. An example of such a case is shown as the purple edge in the auxiliary graph of Fig. 3. Here VNF_1 and VNF_2 are both processed at DC B , and so the edge between $DC_1^{B'}$ and DC_2^B is assigned a weight of 0.

Now, it is straightforward to see that a path from Src to Des in the auxiliary graph for an SFC request corresponds to a selection of DCs to process the SFC's VNFs in order, and the routes for the lightpaths connecting those DCs.

B. Algorithm

The pseudocode of the proposed DRAMA-DC algorithm is presented in Algorithm 1.

Algorithm 1 DRAMA-DC Algorithm

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

Output: SFC recovery

```

1: Initialize two empty SFC sets  $P$  and  $W$ 
2: for each  $s \in S$  do
3:   if  $s \in D$  (i.e.,  $s$  is unrecoverable) then
4:     Release the spectrum and VNFs of  $s$ 
5:   else if  $s \in M$  or  $s$  is affected by disaster then
6:     Release the spectrum and VNFs of  $s$ , add  $s$  to  $W$ 
7:   else
8:     Add  $s$  to  $P$ 
9:   end if
10: end for
11: Calculate the shortest path for each  $s$  in  $W$ 
12: Sort all  $s \in W$  in increasing order of total modulation factors of the shortest path
13: for each  $s \in W$  do
14:   Generate the auxiliary graph of  $s$ 
15:   Calculate K-shortest paths in the auxiliary graph
16:   Sort K-shortest paths in increasing order of its weight
17:   for each path in K-shortest paths pool do
18:     if path is available then
19:       Assign  $s$  with the corresponding path and DC of this path with First Fit spectrum assignment
20:       Add  $s$  to  $P$ 
21:       break
22:     end if
23:   end for
24:   Remove  $s$  from  $W$ 
25: end for

```

In lines 1-10, the spectrum and VNFs allocated to unrecoverable SFCs, and SFCs that are waiting for recovery or re-assignment, are released. The SFCs requiring recovery and re-assignment are added to the waiting set W . SFCs that were neither affected by the disaster nor require re-assignment will be added to the post-disaster set P .

In line 11, for each SFC, we first calculate the shortest path (SP) in the physical network in terms of number of hops from the source to the destination going through each of the

intermediate DCs for the VNFs in the SFC. This shortest path is cut into segments by DCs because of their 3R regeneration capability. The total modulation factor is defined as the sum of the modulation factors of the segments. If the SFCs have the same total modulation factors, the SFC that has a shorter SP distance in the number of hops will be recovered first. SFCs are considered to be recovered in this order because these SFCs tend to have the highest spectrum efficiency. Note that this process is used only for determining the order of recovery of the SFCs, and may not be the final paths chosen for recovering the SFC. In line 12, all the SFCs in W are sorted in increasing order of total modulation factors.

In lines 13-25, we perform recovery and re-assignment. The auxiliary graph is created for each s . Then the k-shortest paths are calculated in the auxiliary graph. Since we may remove edge(s) while adding the weights, it is possible that the path does not exist from the source node to the destination node. In this case, the SFC is blocked. If at least one k-shortest path(s) exists, for each path, according to its corresponding lightpath and DC, we check whether the path is *available* as follows:

- If there are enough FSs on the lightpath;
- If the latency of the SFC is lower than its latency threshold;
- If the latencies of SFCs already placed in the network (i.e., the set P) are still lower than their thresholds. This check is to ensure that the latency threshold is not violated due to the potential increase in the VNF load of a DC as an additional SFC is recovered.

If one of these checks is negative, the path is regarded as *unavailable*, and we move to the next path. If all the path(s) are *unavailable*, the SFC is blocked. Once the shortest available path is selected, we assign s with the corresponding path and DC of the shortest available path with first fit spectrum assignment.

The time complexity of setting up the auxiliary graph is $O(K \cdot N^3 \cdot F \cdot |V| \cdot |B|)$, where F is the number of slots on each fiber. After that, for calculating the K-shortest paths in the auxiliary graph, the time complexity is $O(K \cdot (|V| \cdot |B|)^3)$. Therefore the time complexity of DRAMA-DC is $O(K \cdot N^3 \cdot F \cdot |V| \cdot |B| + K \cdot (|V| \cdot |B|)^3)$.

IV. SIMULATION RESULTS

A. Simulation setting

For evaluating DRAMA-DC's performance, we use the NSF network (14 nodes and 22 links, shown in Fig. 4) topology. Each fiber has 352 FSs. We evaluate the performance for three specific disaster scenarios, shown as red circles in Fig. 4. These representative disasters represent incidents occurring at the center of the network (D_1), near the edge of the network (D_2), and a disaster that fails a DC (D_3). We assume that all the nodes, links, and DCs that lie within the disaster zone are disabled by the disaster. For instance, for disaster zone D_2 , node 2 and links 1-2, 1-3, 2-3, and 2-4 are disabled. For disaster zone D_3 , the DC that is located on node 6 is disabled as well. Besides these specific scenarios, we also evaluate the average performance for random disasters.

We assume that three DCs are placed in the network and here is how we determine the locations of the DCs. First, all the nodes in the network are sorted by their betweenness centrality (BC) value [11], [16]. BC measures the interaction degree between a node and other nodes, so it can be used to evaluate how a node can serve as a bridge from one part of a network to another. Since SFCs are routed through the DCs, placing DCs




Fig. 4. 14-node NSF network. D_1 , D_2 , and D_3 are three disaster zones. DCs are placed on nodes 4, 6, and 9, shown by green circles.

according to nodes' BCs makes them geographically close to the two ends of the SFCs. Therefore, we place the three DCs at the nodes with the highest BC (nodes 4, 6, and 9 - green circles in Fig. 4).

We assume there are 10 VNFs, i.e., $|F| = 10$, and each VNF is hosted in various DCs. The number of copies of each VNF is randomly selected between 2 and 3. In each DC, each VNF has its own separate processor. We assume that the computational capacity of the processor is 150 GOPS and all the processors for all the VNFs in different DCs have the same computational capacity limitation.

A set of SFC requests is generated for the pre-disaster scenario with uniformly randomly selected source and destination nodes. The required data rate of the SFC is 40 Gbps, 100 Gbps, or 400 Gbps with probability 0.2, 0.5, and 0.3, respectively [5], [7]. We assume that four modulation formats are utilized: BPSK, QPSK, 8-QAM, and 16-QAM. The transparent physical distance limitations for each modulation format are provided in parentheses in Table I (we assume that BPSK has no physical distance limitation). The number of required FSs of each SFC is determined by the selected modulation format and its data rate. Table I shows the number of FSs corresponding to different data rates and different modulation formats.

TABLE I
Required FSs and distance limitations [17].

Modulation	Data Rate		
	40	100	400
16QAM (500 km, 50 Gbps)	1	2	8
8QAM (1000 km, 37.5 Gbps)	2	3	11
QPSK (2000 km, 25 Gbps)	2	4	16
BPSK (>2000 km, 12.5 Gbps)	4	8	32

The necessary number of FSs for an SFC with a given modulation format is computed as $\lceil r/\eta_m \rceil$, where r is the SFC request's data rate, and η_m is the spectrum efficiency of the selected modulation format m (defined as data rate per FS, shown in Table I following the modulation format). Since we assume that regenerators are placed in every DC, the η_m of the SFC request is determined by the physical distance of the longest transparent segment of the end-to-end lightpath.

Before the disaster happens, 1000 SFC requests are generated. Each SFC request has 1 to 3 required VNF(s). These SFCs are assigned with the shortest path in the number of hops and first fit spectrum assignment. The DC selections of VNFs of the SFC are also determined by this shortest path. An SFC is blocked if the selected path does not have available slots or the selected DC does not have enough capacity. In this case, the SFC is not established.

In the simulation, for each disaster, 20 different SFC sets and 5 different VNF placement schemes are generated. The results are the average among these 100 combinations. Average

results are shown with 95% confidence interval. As there are no existing algorithms in the literature for this problem, we compare DRAMA-DC with a baseline shortest-path-first-fit (SPFF) algorithm in terms of the SFC blocking ratio and bandwidth blocking of the recovery (shown as a curve in Fig. 5 and Fig. 6). In SPFF, the mitigation zone is activated but the routing and DC selections for an SFC are based on the shortest path in the damaged network. Different sizes of the mitigation zone are tested with different values of the relaxation factor α . The size of the mitigation zone here is characterized by the difference in radius between the mitigation and disaster zone circles.

B. Results

As we can see from Fig. 5, DRAMA-DC is better than SPFF in terms of the SFC blocking ratio and bandwidth blocking ratio when no DC is damaged. For disaster D_1 , when the mitigation zone is 600 km, DRAMA-DC's SFC blocking ratio is reduced by 85% to 95% over SPFF for different α values, and its bandwidth blocking ratio is reduced by 68% to 89%. We also observe that both the blocking ratios of DRAMA-DC decrease as the mitigation zone expands and the value of α increases. For the largest mitigation zone cases, blocking ratios are very close to 0 and all the recoverable SFC requests can be recovered. Therefore, DRAMA-DC can take advantage of the additional flexibility due to a larger mitigation zone and relaxation factor α . The blocking ratio of DRAMA-DC with different values of α is always better than SPFF with $\alpha = 1.4$ (which is the best case for SPFF, as larger α leads to better recovery performance). When the size of the mitigation zone is larger than 2000 km, the recovery improvement with the increase of α is lower than the cases when the size of the mitigation zone is 500 km. The reason is that the flexibility given by the mitigation zone is larger than that provided by α . Compared to relaxing the latency threshold, the re-assignments of SFCs inside the mitigation zone help significantly more during the recovery, and the improvement space of larger α is limited.

Now we test DRAMA-DC in the case of a single DC failure (D_3). When the DC fails, all the SFCs that visit that DC are affected, and need to be recovered. In this case, the problem is very challenging since the bottleneck issue becomes severe as many SFCs are waiting to be recovered and the failed DC cannot be used during the recovery. In Fig. 6(a) and 6(b), we again note that DRAMA-DC recovery algorithm is always better than SPFF. Larger values of α also bring lower SFC blocking ratios. For example, when α is set to 2.6, the blocking of DRAMA-DC is reduced by 33% to 91% over SPFF for different sizes of the mitigation zone. For the DRAMA-DC itself, in the case of the mitigation zone being 2600 km, DRAMA-DC's blocking is reduced by 43% to 91% over SPFF for five different values of α . The gap between DRAMA-DC and SPFF also increases as the mitigation zone size increases.

We can also observe that the blocking ratios of D_2 and D_1 are lower than that of D_3 because D_3 fails a DC and the SFC requests have to be recovered by the two surviving DCs. Also, the blocking ratios of D_2 are larger than D_1 's, because D_2 is closer to a DC so the routing of the recoverable SFCs after disaster is more challenging.

We now present results for random disasters in Fig. 6(c) and 6(d). In this scenario, we select a random node as the center of the disaster zone, and the zone's radius is uniformly distributed in the range [100, 400] km. We obtain average results for 10 different disasters. Once again, we see that DRAMA-DC




Fig. 5. Performance results of recovery: (a) SFC blocking ratio of D_1 , (b) Bandwidth blocking ratio of D_1 , (c) SFC blocking ratio of D_2 , (d) Bandwidth blocking ratio of D_2 .




Fig. 6. Performance results of recovery: (a) SFC blocking ratio of D_3 , (b) Bandwidth blocking ratio of D_3 , (c) SFC blocking ratio of random disasters, (d) Bandwidth blocking ratio of random disasters

has a better performance in both the SFC blocking ratio and bandwidth blocking ratio. For instance, when the α is set to 2.6, the SFC blocking ratio of DRAMA-DC is reduced by 61% to 97% over SFPP for different sizes of the mitigation zone.

V. CONCLUSION

Disaster management is an important issue in elastic optical networks. In this paper, we proposed a new disaster recovery algorithm with mitigation awareness for EONs with data centers. Based on the concept of the mitigation zone, SFCs are recovered or re-assigned with a latency relaxation factor. Both the lightpath and VNF of the SFCs are re-accommodated to minimize the blocking ratio. The results show that DRAMA-DC performs better than a baseline algorithm in several scenarios, as well as the effectiveness of the mitigation zone in disaster recovery.

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. Matsuo, "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] S. Sharma, A. Engelmann, A. Jukan, and A. Gumaste, "VNF availability and SFC sizing model for service provider networks," *IEEE Access*, vol. 8, pp. 119 768–119 784, 2020.
- [4] L. Liu, S. Guo, G. Liu, and Y. Yang, "Joint dynamical VNF placement and SFC routing in NFV-enabled SDNs," *IEEE Transactions on Network and Service Management*, vol. 18, no. 4, pp. 4263–4276, 2021.
- [5] M. Zhu, J. Gu, T. Shen, J. Zhang, and P. Gu, "Delay-aware and resource-efficient service function chain mapping in inter-datacenter elastic optical networks," *Journal of Optical Communications and Networking*, vol. 14, no. 10, pp. 757–770, 2022.
- [6] A. Khatiri, G. Mirjalily, and Z.-Q. Luo, "Balanced resource allocation for VNF service chain provisioning in inter-datacenter elastic optical networks," *Computer Networks*, vol. 203, p. 108717, 2022.
- [7] Y. Liu, F. Zhou, T. Shang, and J.-M. Torres-Moreno, "Disaster protection for service function chain provisioning in EO-DCNs," *IEEE Transactions on Network and Service Management*, 2023.
- [8] 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.
- [9] H. Yu and C. Yang, "Partial network recovery to maximize traffic demand," *IEEE communications letters*, vol. 15, no. 12, pp. 1388–1390, 2011.
- [10] H. Nasralla, T. El-Gorashi, M. Musa, and J. Elmigehani, "Routing post-disaster traffic floods in optical core networks," in *2016 International Conference on Optical Network Design and Modeling (ONDM)*, 2016, pp. 1–5.
- [11] 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.
- [12] 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.
- [13] D. Dietrich, A. Abujoda, and P. Papadimitriou, "Network service embedding across multiple providers with nestor," in *2015 IFIP Networking Conference (IFIP Networking)*. IEEE, 2015, pp. 1–9.
- [14] M. Aibin and K. Walkowiak, "Adaptive modulation and regenerator-aware dynamic routing algorithm in elastic optical networks," in *2015 IEEE International Conference on Communications (ICC)*. IEEE, 2015, pp. 5138–5143.
- [15] A. Lozada, F. Calderón, J. N. Kasaneva, D. Bórquez-Paredes, R. Olivares, A. Beghelli, N. Jara, A. Leiva, and G. Saavedra, "Impact of amplification and regeneration schemes on the blocking performance and energy consumption of wide-area elastic optical networks," *IEEE Access*, vol. 9, pp. 134 355–134 368, 2021.
- [16] U. Brandes, "A faster algorithm for betweenness centrality," *Journal of mathematical sociology*, vol. 25, no. 2, pp. 163–177, 2001.
- [17] C. Wang, G. Shen, and S. K. Bose, "Distance adaptive dynamic routing and spectrum allocation in elastic optical networks with shared backup path protection," *Journal of Lightwave Technology*, vol. 33, no. 14, pp. 2955–2964, 2015.