Strategic Cooperation among Datacenter Providers and Optical-Network Carriers for Disaster Recovery

Subhadeep Sahoo[1], **Sugang Xu**[2], **Sifat Ferdousi**[1], **Yusuke Hirota**[2], **Massimo Tornatore**[3], **Yoshinari Awaji**[2], **Biswanath Mukherjee**[1, 4]
[1]University of California, Davis, USA, [2]National Institute of Information and Communications Technology (NICT), Japan, [3]Politecnico di Milano, Italy, [4]Soochow University, China
Email: {subsahoo, sferdousi, bmukherjee}@ucdavis.edu, {xsg, hirota.yusuke, yossy}@nict.go.jp, massimo.tornatore@polimi.it

***Abstract*— Cooperation among datacenter providers (DCPs) and network carriers is necessary to support today's ubiquitous cloud services. However, such cooperation can be constrained by limited visibility as confidential information, such as network topology, resource availability, etc., may not be disclosed among these entities due to regulatory policies. We study a DCP-carrier cooperation-based service restoration scheme during a disaster with the aid of a third-party mediator, namely a Provider Neutral Exchange (PNE). We propose a novel resource-driven demand-matching strategy to restore DCP services. When multiple DCPs compete for network resources (due to post-disaster resource crunch), resource balancing by PNE can achieve fair and efficient service restoration. To allow flexibility in demand-resource matching, DCPs generate multiple sets of connection requests and define varying priorities and bandwidth degradations for each request. Carriers evaluate the DCP requests and provide feedback (e.g., whether a request can be satisfied or not) based on their available resources. We present an eight-phase DCP-carrier cooperation framework, with each phase employing individual sub-tasks carried out by DCPs, carriers, and PNE. Results under different disaster scenarios show that our strategy significantly improves DCP service restoration, incurring less restoration time.**

***Keywords*— Carrier; datacenter provider; cooperation; confidentiality; demand-resource matching; disaster recovery.**

I. Introduction

Internet users are increasingly adopting cloud-based services to easily access computing and storage resources, with higher service availability and reliability [1, 2]. To accommodate the growing demand for cloud services [3], the underlying network infrastructures and datacenters (DCs) hosting these services are continuously evolving [4]. Contrary to the traditional monolithic cloud/network environment, where a single entity, such as a network carrier, owns both the network and DC infrastructures, in today's common scenario, the infrastructures are owned by separate entities, such as network carriers and DC providers (DCPs) [5]. To ensure efficient cloud-service provisioning, carriers and DCPs must perform joint/coordinated resource allocation [6]. This becomes especially challenging at times of resource crunch due to widespread failures, e.g., as a consequence of disasters (such as earthquake, tsunami, etc.).

After a disaster, network recovery and service restoration must be performed as soon as possible to minimize service downtime. Post-disaster service restoration is widely studied for both DCPs and carriers [7-10]. Ref. [10] presents the benefits of joint network and DC recovery after a large-scale disaster when DC and network infrastructures are owned by the same entity. This allows exchange of detailed information between DC and network operators which can achieve high restoration efficiency. However, in a distributed environment, where the network and DC infrastructures are owned by separate entities, confidential information such as detailed network topology and failure/ damage scenario of carriers, resource availability of carriers, content availability and locations of the DCPs, etc., may not be shared among the different entities due to regulatory policies [11]. Hence, service restoration with limited information visibility across different entities creates a challenge for the carriers and DCPs. In such cases, cooperation among the carriers and DCPs is crucial to accomplish fast and efficient service restoration. Such cooperation may include sharing of network resources, joint provisioning of user requests, sharing of recovery tasks (e.g., repair of failed/damaged links), etc., while ensuring mutual benefits in terms of recovery and cost.

In [12], we introduced a carrier-cooperation optimization model to perform post-disaster connectivity restoration where multiple carriers achieve fast and efficient recovery at lower cost through strategic cooperation (such as network resource sharing and identifying the overlapping and mutually beneficial recovery tasks). In [13], we developed heuristics for DCP-carrier cooperation for cloud service restoration considering a single DCP. However, in case of multiple DCPs, potential imbalance and inefficiency in service restoration across the various DCP are likely to arise, especially in a situation of resource crunch due to a disaster. Thus, how to best fit/match the *demands* (i.e., requests for network resources) from multiple customers (e.g., DCPs) to the *supplies* (i.e., available network resources) of the carriers, by exploiting more effective cooperation, needs to be investigated.

In this work, we study how to maximize service restoration and minimize restoration time for multiple DCPs when there is resource crunch in the network by considering different damage scenarios in carrier networks due to disasters. Below, we portray an example of a DCP-carrier cooperation framework. Fig. 1 represents the abstracted network model considered in our work, which shows that DCPs (e.g., DCP-X and DCP-Y) build their DC interconnect (DCI) networks among the geographically-distributed DCs by leasing connection services from several

978-1-6654-3540-6/22/$31.00 ©2022 IEEE

underlying carriers (e.g., Carrier-A and Carrier-B). A third-party mediator entity, called PNE [12], provides the interconnection points (based on optical-electrical-optical, OEO, conversion) among the carrier network nodes and the DCs. PNE also acts as a facilitator for cooperation between the carriers and the DCPs.

The major contributions in this work are summarized below.

- We design a cooperation-based post-disaster service restoration framework in which different entities employ individual cooperation sub-tasks in a distributed manner.
- We identify the information that is crucial to ensure efficient service restoration and that can be shared by the different entities without violating confidentiality.
- We numerically illustrate realistic scenarios where DCP-carrier cooperation brings significant improvement in terms of fast and efficient cloud service restoration.

The rest of the study is organized as follows. Section II provides the problem statement and describes the network model. Section III presents the proposed cooperation-based strategy. Section IV evaluates the proposed strategy using illustrative results. Finally, section V concludes the study.

II. Problem Statement and Network Model

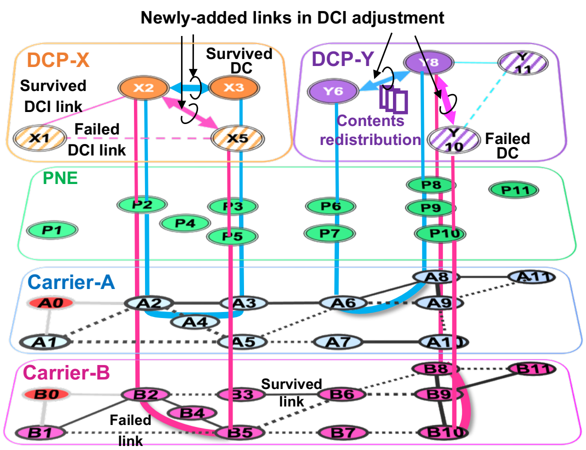

Fig. 1: Network model in DCP-carrier cooperation framework.

We study the problem of maximizing service restoration and minimizing restoration time for DCPs after a disaster. As shown in Fig. 1, before a disaster, DCPs create their own DCI network (interconnecting the geographically-distributed DCs) by leasing connection services, e.g., optical lightpaths or IP-over-WDM connection, from the underlying carriers based on connection prices advertised by the carriers. The solid blue and pink lines illustrate that DCP-X and DCP-Y lease connection services from Carrier-A and Carrier-B, respectively. To avoid leakage of confidential damage information after a disaster, the carriers abstract their detailed topologies of the disaster area to the public reference topology presented by PNE network (PNEN) [12, 13]. The PNEN topology consists of several electronic switching nodes (*P1-P11*) to interconnect the carrier nodes (consisting of the underlying optical nodes and upper-layer IP nodes) and the DCP nodes in the same proximity as the PNEN nodes. As shown in Fig. 1, nodes 1 to 11 (e.g., *A1-A11* and *B1-B11*, of Carrier-A and Carrier-B, respectively) represent the abstracted affected area after a disaster, and the unaffected area is represented as node 0. The damaged and survived links in the carrier networks are represented by dashed and solid lines, respectively. After the disaster, the DCI networks of both DCP-X and DCP-Y also get affected. For simplicity, we discuss the details of DCP-Y only.

The dashed blue and pink lines represent the damaged DCI links, and we consider that a few DC nodes (*Y11, Y10*) are also damaged. So, content requests of the damaged DCs need to be redirected to other DCs, where the backups of the contents are located. For example, in Fig. 1, *Y10* and *Y6* are the primary and backup DC of a content *c,* respectively, so requests for content *c* need to be redirected from damaged DC *Y10* to backup DC *Y6*. However, redirecting all requests from *Y10* to *Y6* will consume high network resources. So, DCP-Y replicates/redistributes content *c* from *Y6* to *Y8*, which is adjacent to *Y10*, and redirects requests for content *c* to *Y8*. For such content redistribution and request redirection, the DCP needs to modify its DCI topology by requesting new connection services from the underlying carriers. Based on the connection prices (for each PNEN node pairs) advertised by the carriers (see Section III.B), DCP tries to request for new connections by aiming to minimize its total payments towards the carriers. We name this strategy as *joint content redistribution and DCI adjustment*. As shown in Fig. 1, connection requests between *Y6-to-Y8* and *Y10*-to-*Y8* are provisioned by Carrier-A, and Carrier-B, respectively.

To facilitate fast and efficient service restoration, identifying the connection requests (demand) from the DCPs best fit to the available network resources (supply) of the carriers is required. Information regarding the requests, such as connection priority, bandwidth degradation tolerance as per SLA, etc., are crucial, since the carriers need to evaluate the DCP requests and offer feedback (e.g., providing a score based on whether the request can be satisfied or not, possible bandwidth degradation in the satisfied request, possible delay in request provisioning, etc.) based on the available resources and the recovery planning/ scheduling of the damaged links. Also, when multiple DCPs compete for network resources, resource balancing by PNE is required to ensure fair service restoration.

The problem described above can be summarized as how to maximize DC-service restoration (while minimizing restoration time) in a disaster scenario where multiple optical-network carriers and multiple DCPs are affected. Particularly, given a set of connection requests from multiple DCPs to multiple underlying optical-network carriers, the problem of cooperation-based service restoration identifies, with the aid of a PNE, a balanced resource-driven matching of DCP requests over the available physical network resources of the carriers under resource crunch, while ensuring information confidentiality (of damaged links, connection priority, etc.) among the different entities, such that DCP service restoration is maximized using less restoration time.

III. Proposed Framework and Strategy

In our cooperation-based service restoration framework, we organize the cooperation tasks into eight phases; and in each phase, DCPs, carriers, and PNE perform different sub-tasks in parallel. From DCP's perspective, the objective is to maximize service restoration in less time by generating flexible connection requests (demands) to the carriers. From carriers' perspective, the objective is to maximize the accepted DCP requests and simultaneously minimize their recovery tasks and their total recovery time (assuming that at most L number of damaged fiber links can be simultaneously repaired in each unit of time). From PNE's perspective, as PNE has visibility of requests from multiple DCPs and responses to the requests from multiple carriers, the objective is to perform demand-supply matching between DCPs and carriers (e.g., resource balancing, identifying DCP requests that are most suitable based on carrier resources, discovering common recovery tasks of the carriers, etc.). The eight phases of the cooperation-based restoration strategy is shown in Fig. 2, and details of each phase are explained below.

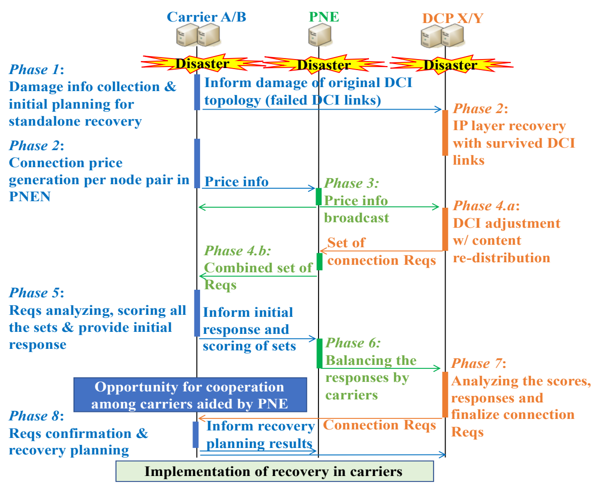

Fig. 2: Distributed eight-phase planning for DCP-carrier cooperation.

A. *Phase 1 (Damage Evaluation)*

In Phase 1, carriers evaluate the post-disaster damage in their network and conduct initial planning of the standalone recovery tasks for the original DCI networks of their customers (DCPs) as a basic recovery strategy. Then, they analyze the DCI link failures of DCPs, informing them about the failed DCI links.

B. *Phase 2 (Price Evaluation & Initial DCP Recovery)*

After receiving the failure information from the carriers, DCPs perform the rerouting in the IP layer by using the survived DCI links. In case of DC failures, content requests are redirected from the failed primary DCs to the backup DCs. However, only a few requests can be restored using re-routing at the IP layer over the survived DCI links [13]. On the other hand, the carriers generate the connection price for each PNEN node pair based on the damage in their networks. The exact damage situation is abstracted in the price to avoid leakage of confidential information. For a connection that requires physical recovery of the damaged links (which incurs additional delay), extra dummy price is added [13], and this guides the DCPs to avoid requesting such high-priced connections. Consequently, this reduces the recovery burden for the carriers immediately after the disaster.

C. *Phase 3 (Connection Service Price Declaration)*

Carriers declare their connection service price for each PNEN node pair; and the PNE broadcasts the price information to all other carriers and DCPs.

D. *Phase 4.a (Content Redistribution and DCI Adjustment)*

First, DCPs perform content redistribution to replicate contents among the survived backup DC and to accommodate the post-disaster traffic pattern. For redirecting the user requests of the damaged DCs to the backup DCs, DCPs then create new DCI topologies (DCI adjustment). Upon receiving the carriers' connection price information from PNE (in Phase 3), DCPs generate their connection requests to establish the new DCI topologies, while avoiding the selection of the high-priced connections. This strategy is called joint content redistribution and DCI adjustment (see [13]). A brief summary of connection request generation is given below.

First, DCP aggregates the bandwidth requirements for content c hosted in a failed primary DC d^p and finds the backup DC d^b of c. Then, DCP checks the advertised connection prices (from all carriers) in two ways: (a) direct connection price, $P_{d^p-d^b}$, between d^p and d^b; and (b) connection price, $P_{d^p-d^i}$, between an adjacent survived DC d^i of d^p and price, $P_{d^b-d^i}$, for replicating content c from d^b to d^i. If, $P_{d^p-d^b} > (P_{d^p-d^i} + P_{d^b-d^i})$, connection requests are created between $d^p - d^i$ and $d^i - d^b$; else, a connection request created between $d^p - d^b$.

E. *Phase 4.b (Candidate Set Generation)*

Since selecting appropriate connection requests plays an important role in service restoration for DCPs (especially during resource crunch in carrier networks), DCPs introduce two features: 1) DCPs allow flexibility in terms of bandwidth degradation, i.e., provide an acceptable range of bandwidth, between minimum bandwidth requirement, $BW_{\min}$, and maximum bandwidth requirement, $BW_{\max}$, for individual connection request; and 2) instead of providing a single set of connection request to the carriers, DCPs generate multiple sets of connection requests as candidate sets, out of which the best-fit candidate can be selected based on the available resources from the carriers. Thus, DCPs can offer great flexibility in their demands which enables better demand-supply matching for carriers, especially in case of resource crunch.

We consider that DCPs can create multiple candidate sets of connections requests. In each candidate set, each connection request is denoted by a list of parameters: **a**) source node, **b**) destination node, **c**) minimum bandwidth requirement, $BW_{\min}$ (calculated based on degradation tolerance level of individual user requests of the DCP), **d**) maximum bandwidth requirement, $BW_{\max}$, (i.e., required aggregated bandwidths to be restored of a DCP) **e**) priority grade, and **f**) target carrier's ID. The priority V of a connection request is calculated by DCPs using:

$$V = \alpha * n + \sum_i (w_i * k_i) \quad (1)$$

where n is the total number of user requests associated with the connection and α is a weight parameter which denotes the significance of the connection in terms of volume. $i \in \{high, mid, low\}$ is the set of priority levels for the user requests. k_i is the number of user requests with priority level i and w_i is the weight parameter for the priority level i. α and w_i $(\forall i)$ sum to 1. To conceal the confidential DCP information regarding the actual priority of their connection requests, the values of parameter V (from Eq. (1)) are abstracted as *priority grades*, e.g., connections with highest value of V are encoded as *Grade 1*, connections with the second highest value of V are encoded as *Grade 2*, and so on.

Fig. 3 shows an example of candidate set generation and selection. Here, DCP-X generates candidate sets *X.1* and *X.2*, and DCP-Y generates candidate sets *Y.1* and *Y.2*. Candidate set *X.1* of DCP-X has two connection requests *X.1.1* and *X.1.2*. Based on the parameters (**a-f**) mentioned above, *X.1.1* contains {2, 5, 80, 100, 2, A} and *X.1.2* contains {8, 10, 70, 110, 1, B}. Similarly, *Y.1.1* contains {8, 10, 60, 90, 1, B} and *Y.1.2* contains {6, 8, 50, 70, 2, A}. These candidate sets are then delivered to the PNE for further processing which involves combining the requests from different DCPs and forwarding them to the appropriate carriers. PNE tasks in this phase are as follows.

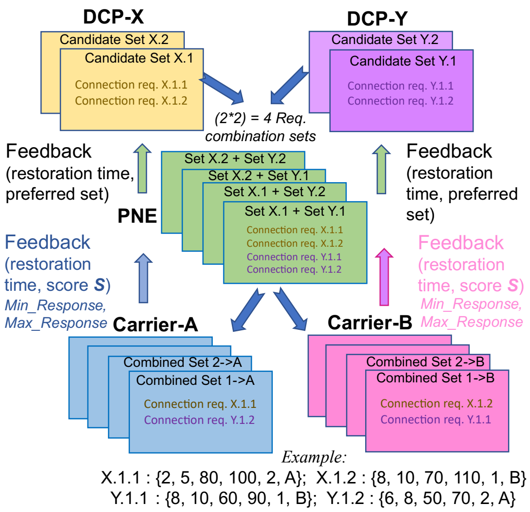

Fig. 3: Information exchange flow during candidate set selection.

After receiving all the candidate sets of connection requests from the DCPs (e.g., *X.1*, *X.2*, *Y.1*, and *Y.2*), PNE first generates the combination sets (to allow flexible matching of DCP requests with carrier resources). For example, as shown in Fig. 3, if DCP-X creates $x = 2$ sets and DCP-Y creates $y = 2$ sets, then total possible combinations will be ($x*y = 4$) candidate sets; combination set 1 contains {*X.1*, *Y.1*}, combination set 2 contains {*X.1*, *Y.2*}, and so on. From each combination set, PNE segregates the connection requests for individual target carriers and delivers the $x*y$ sets of connection requests to the respective target carriers (e.g., from combination set 1, *X.1.1* and *Y.1.2* are forwarded to Carrier-A).

F. Phase 5 (Carrier-side Planning Task)

After receiving all the $x*y$ sets of connection requests from PNE, for each connection request set, each carrier evaluates the requests by performing a so-called "carrier-side planning task" (CSPT) in its network using the following steps:

Step 1: Based on $BW_{\min}$ of individual connection request, the carrier performs the initial recovery planning by maximizing the number of acceptable high-priority connection requests of the DCPs and minimizing the recovery tasks of the entire network. The mathematical model in [12] is employed here.

Step 2: The carrier identifies the recovery tasks from Step 1 and performs link recovery time evaluation by minimizing the total restoration time for the satisfied connection requests. The response to the connection requests based on $BW_{\min}$ is denoted as *Min_response*, which includes restoration time and a score *S 1*, if $BW_{\min}$ is satisfied; or *0*, otherwise.

Step 3: For all satisfied requests which have been solved in Step 1, the carrier adds the corresponding $BW_{\min}$ and $BW_{\max}$ as the new lower and upper bound constraints, respectively. Then it performs the recovery planning by maximizing the satisfiable bandwidth (BW_{sat}) for these connection requests (by improving $BW_{\min}$ in a best-effort manner) and minimizing the recovery tasks of the entire network. Note that this is an extended version of the mathematical model in [12].

Step 4: For each satisfied connection request with potentially-improvable bandwidth allocation, the carrier calculates a score *S* based on Eq. (2). The score of a connection request is 100 if $BW_{\max}$ is satisfied; otherwise, range of the score is between 1 and 100, based on BW_{sat}, $BW_{\min}$, and $BW_{\max}$.

$$S = \begin{cases} 100, & \mathit{if}\ BW_{\text{sat}} = BW_{\max} \\ \dfrac{BW_{\text{sat}} - BW_{\min}}{BW_{\max} - BW_{\min}} * 100, & \mathit{otherwise} \end{cases} \tag{2}$$

Step 5: The carrier identifies the recovery tasks from Step 3 and performs the link recovery time evaluation by minimizing the total recovery time for all the connection requests. The response to the connection requests with BW_{sat} is denoted as *Max_response*, which includes restoration time and a score *S*.

After evaluating all the connection request sets, each carrier provides feedback containing $x*y$ sets of evaluation responses to PNE including *Min_response* and *Max_response*. For example, the *Min_response* and *Max_response* for *Y.1.1* by Carrier-B are illustrated as {restoration time: 0, score: 1} and {restoration time: 2, score: 50}, respectively. Here, we assume that, for *Y.1.1*, 60 Gb/s can be immediately satisfied, but to satisfy 75 Gbps, the restoration time is 2.

G. Phase 6 (Resource-Driven Demand Matching by PNE)

Balancing resource allocation is important to ensure that all customers have equal opportunity to restore their services. Thus, after receiving the initial responses for all the combination sets from carriers, PNE starts balancing the BW_{sat} for the DCP connection requests which are overlapping (i.e., same source node, destination node, and corresponding carrier). Finally, PNE performs the demand-resource matching between multiple DCPs and carriers. The steps are as follows:

Step 1: For each combination set, PNE selects the best between *Min_response* and *Max_response*, according to total score and restoration time. To minimize the recovery tasks for carriers and restoration time of the DCP requests, total restoration time gets more importance than total score while selecting between *Min_response* and *Max_response*.

Step 2: For an overlapping connection request, if (BW_{sat} of DCP-X / BW_{sat} of DCP-Y) is not equal to ($BW_{\max}$ of DCP-X / $BW_{\max}$ of DCP-Y), then PNE balances BW_{sat} for both DCP-X and DCP-Y without exceeding sum of satisfied bandwidths.

Step 3: PNE performs resource-driven demand matching by finding the best choice among all possible combination sets based on the information shared by DCPs and feedback from carriers (e.g., priorities, restoration times, and scores of the connection requests presented in a set). Then, PNE informs the DCPs of the best choice among all candidate sets.

For example, *Min_response* is selected as the best response for *Y.1.1* because *Y.1.1* can be immediately restored with 60 Gb/s. Thus, BW_{sat} for *Y.1.1* is 60 Gb/s without any improvement. *Y.1.1* and *X.1.2* are the overlapping connection requests to Carrier-B with same source and destination nodes. *X.1.2* requests 70 Gb/s and 110 Gb/s as $BW_{\min}$ and $BW_{\max}$, and it can be immediately restored with 100 Gb/s of BW_{sat}. As BW_{sat} is not fair and balanced as per $BW_{\max}$ of DCP-Y and DCP-X, PNE balances BW_{sat} for both DCPs, by increasing it to 72 Gb/s for *Y.1.1* and decreasing it to 88 Gb/s for *X.1.2*.

H. Phase 7 and 8 (Confirmation of DCP Requests)

Each DCP checks the responses, finalizes a single set of connection requests, and sends it to the corresponding carriers. Then, in Phase 8, carriers can perform the carrier-side subtasks based on carrier-cooperation scheme [12], to further minimize the recovery tasks, recovery time, and costs. Note that, in this study, we focus on the DCP-carrier cooperation; the carrier-cooperation is treated as an optional phase in this framework.

IV. Simulation Environment and Numerical Analysis

A. Simulation Environment

To evaluate our proposal, we consider a network model with two carriers (i.e., Carrier-A and Carrier-B), two DCPs (i.e., DCP-X and DCP-Y), and a PNE, as shown in Fig. 1. The PNEN is modeled as a subset of *Japan network* [14] with 11 nodes. We consider that both carrier networks are abstracted to this common PNEN reference topology, and all the carrier nodes are co-located with the PNE nodes. In addition, 15 co-located links are considered for both the carrier networks. Similarly, DCP-X and DCP-Y are modeled as subsets of the PNEN reference topology, and all DC nodes are co-located with PNEN nodes. For DCP-X and DCP-Y, the pre-disaster DCI topologies have 7 and 8 DC nodes, and 11 and 12 DCI links, respectively.

To generate the damage patterns, we consider a strong correlation value of 0.8 among the link failures in both carrier networks (e.g., if a link fails in Carrier-A network, the collocated link of Carrier-B will also fail with probability 0.8). In our work, three different damage scenarios are generated in the carrier networks. (i) *Heavy damage* (10:10): in both carrier networks, 10 links are damaged; (ii) *Mixed damage* (10:5): in Carrier-A network, 10 links are damaged and in Carrier-B network, 5 links are damaged; (iii) *Light damage* (5:5): in both carriers, 5 links are damaged. We assume that the carriers can perform one fiber link recovery per unit time. For each damage scenario, 30 instances are generated randomly to increase statistical confidence in our numerical analysis. To simulate the post-disaster congestion in the carrier networks, we consider that only 4 lightpath channels (each with 100-Gbps capacity) are available in each optical fiber link. For both the DCPs, we assume that 3 to 4 DCs are damaged due to the disaster, and a total of 15000 requests, previously allocated to the damaged DCs, are required to be restored. We consider that the bandwidth requirement of each user request is uniformly distributed between 70 to 100 Mbps in steps of 1 Mbps. We assume that there are three priority levels for the user requests, namely, high, mid, and low; and each level has 5000 users. For high-priority users, the bandwidth degradation tolerance is 10%, while for mid- and low-priority users, it is 20% and 30%, respectively. To calculate the priority V of a connection request, we set α as 0.5, w_i as 0.25, 0.15, and 0.1 for priority levels, $i \in \{high, mid, low\}$, respectively. We set the connection price as 1 unit for 1 Gbps of bandwidth (IP-over-WDM) for survived links, and an additional dummy price of 100 units is added if a carrier needs to recover the links for abstracting the damage information (as mentioned in Section III.B). However, the final payments from the DCPs to the carriers are linear with regular service price (i.e., 1 unit for 1 Gbps of bandwidth). Also, for simplicity, we consider that both DCPs generate two sets (x=y=2) of connection requests, which means four (x*y) sets of combination can be created by the PNE.

B. Numerical Analysis

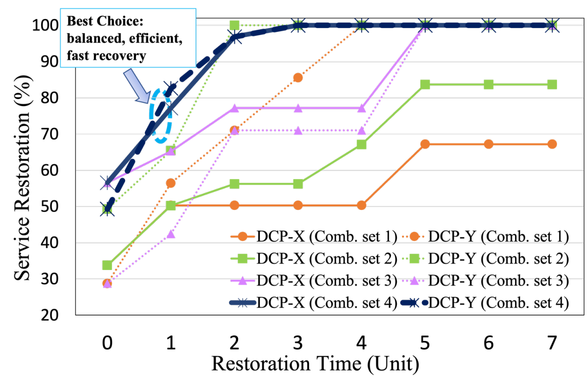

Fig. 4: Service restoration with different combination sets.

Let us now start with analysis of multiple combination set generation by the PNE. Fig. 4 illustrates the service restoration efficiency for all possible combination sets for both the DCPs (considering heavy damage scenario only). We observe that, with combination sets 1, 2, and 3, DCP-X and DCP-Y experience an unbalanced amount of service restorations, e.g., for combination set 1, DCP-Y is provisioned with most of the available resources from the carriers causing much lower service restoration for DCP-X. We also see that, for combination sets 1 and 2, the service restoration efficiency does not reach to 100% even with higher restoration time. Thus, we can say that not all sets of connection requests can be matched efficiently with the available resources in carrier networks. However, with combination set 4, service restoration is higher and balanced for both the DCPs using lower restoration time. Because combination set 4 perfectly matches the available resources in the carrier networks (hence we call this the preferred set), PNE could balance the resource allocation while reducing the competition for resources by multiple DCPs. This is achieved by identifying the sets with less overlapping connection requests from different DCPs. However, based on the damage patterns in the carrier networks, the preferred set may vary.

Let us now compare our proposed strategy with a baseline scheme presented in [13] in terms of service restoration percentage (defined as number of restored user requests over the total number of affected user requests of a DCP) and restoration time. As a baseline scheme, we propose a DCP-carrier cooperation scheme, where each DCP generates only one set of connection request, in which each request has equal priority and a fixed bandwidth requirement.

We evaluate the service restoration efficiency with respect to the service restoration time of our proposed scheme, namely,

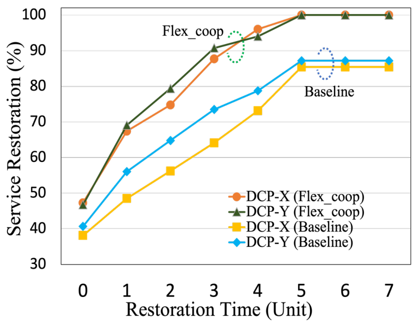

Fig. 5: Service restoration in heavy damage.

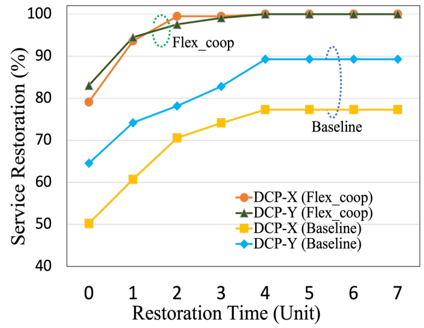

Fig. 6: Service restoration in mixed damage.

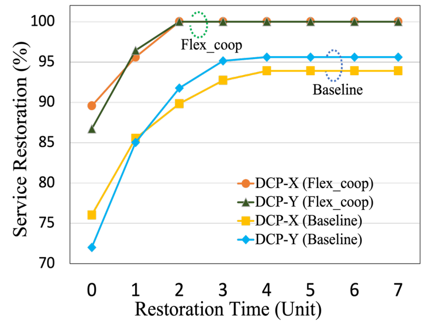

Fig. 7: Service restoration in light damage.

Flex_coop, against the baseline scheme. We analyze for heavy, mixed, and light damage as shown in Figs. 5 to 7 (considering only the preferred set). We see that, for both DCP-X and DCP-Y, *Flex_coop* performs better with about 20% higher efficiency in terms of service restoration with at least 50% lower restoration time compared to the baseline scheme. For example, in Fig. 5, we see that, with *Flex_coop*, 100% service restoration could be achieved for both the DCPs, but with the baseline scheme, the maximum service restoration that could be achieved is just about 87%. Also, we see that *Flex_coop* could achieve 80% service restoration for both the DCPs in less than three units of restoration time, while the baseline scheme requires at least five units of restoration time. This is because: first, with joint content redistribution and DCI adjustment, DCP avoids the selection of high-price connections; and second, bandwidth degradation flexibility in each DCP connection request helps to allocate the available resources in the carrier networks during resource crunch more efficiently, hence it avoids unnecessary rejection of connection requests. As shown in Figs. 5 to 7, with the baseline scheme, for both DCP-X and DCP-Y, some requests are rejected by the carriers (e.g., service restoration never reaches to 100% in any damage scenario). Especially, in case of light and mixed damage scenarios, in which resource availability is higher than that for the heavy damage scenario, full service restoration could not be achieved even with higher restoration time. Also, due to efficient demand and resource matching in *Flex_coop*, the preferred combination set best fits the available carrier resources and hence yields to higher service restoration. PNE plays an important role in balancing the allocation of overlapping connection requests from both the DCPs. This is evident when balanced service restoration is observed for both DCPs with *Flex_coop* in Figs. 5 to 7. Finally, as DCPs emphasize the priorities of each request in a set, the carriers try to satisfy the higher-priority requests first. This helps the carriers to minimize their recovery tasks and link recovery time. Analysis of recovery costs of carriers and DCP payments to carriers are not reported here due to limited space.

V. Conclusion

In this study, we presented a novel DCP-carrier cooperation framework for post-disaster service restoration. In this scheme, connection requests from DCPs are matched with the available network resources of the carriers, under resource crunch due to disaster, so that service restoration is maximized for the DCPs. To allow greater flexibility in demand allocation, DCPs generate candidate sets of connection requests of varying priority, bandwidth degradation tolerance, etc., which gives the carriers the opportunity to evaluate the candidate DCP requests and provide feedback accordingly to PNE based on their resource availability and recovery planning. PNE then performs balanced and efficient resource-driven demand matching between DCPs and carriers, while maintaining information confidentiality. Numerical results show that our proposal can achieve significant improvement in service restoration for DCPs using less restoration time.

Acknowledgment

This work is supported by US-Japan JUNO2 project: NSF Grant no. 1818972.

References

[1] Forbes Technology Council, "Why Migrate To The Cloud: The Basics, Benefits And Real-Life Examples," Mar. 2021.
[2] T. Fukai, et al., "Live migration in bare-metal clouds," IEEE Transactions on Cloud Computing, vol. 9, no. 1, Mar. 2021.
[3] Cisco, "Press Release: Global Cloud Index Projects Cloud Traffic to Represent 95 Percent of Total Data Center Traffic by 2021," Feb. 2018.
[4] S. J. B. Yoo, et al., "Intra and inter datacenter networking: the role of optical packet switching and flexible bandwidth optical networking," Proc. 16th International Conference on Optical Network Design and Modelling (ONDM2012), Colchester, UK, Apr. 2012.
[5] B. Shariati, et al., "Demonstration of federated learning over edge-computing enabled metro optical networks," Proc. European Conference on Optical Communications (ECOC), Brussels, Belgium, Dec. 2020.
[6] C. Li, et al., "High-precision edge-cloud collaboration with federated learning in edge optical network," Proc. Optical Fiber Communications Conference (OFC), San Francisco, CA, USA, June 2021.
[7] Y. Liu, et al., "Disaster protection in inter-datacenter networks leveraging cooperative storage," IEEE Transactions on Network and Service Management, vol. 18, no. 3, Sept. 2021.
[8] R. Xu, et al., "Disaster survivability in elastic optical datacenter networks," Proc. IEEE Optoelectronics Global Conference (OGC), Shenzhen, China, Sept. 2016.
[9] G. Le, et al., "Reliable provisioning for dynamic content requests in optical metro networks," Proc. Optical Fiber Communications Conference (OFC), San Francisco, CA, USA, June 2021.
[10] S. Ferdousi, et al., "Joint progressive network and datacenter recovery after large-scale disasters," IEEE Transactions on Network and Service Management, vol. 17, no. 3, Sept. 2020.
[11] P. Safari, et al., "Privacy-preserving distributed learning framework for 6G telecom ecosystems," arXiv preprint arXiv:2008.07225, 2020.
[12] S. Xu, et al., "A novel carrier-cooperation scheme with an incentive to offer emergency lightpath support during disaster recovery," Photonic Network Communications, vol. 40, no. 3, July 2020.
[13] S. Sahoo, et al., "Datacenter-carrier cooperation over optical networks during disaster recovery," Proc. Optical Fiber Communications Conference (OFC), San Diego, CA, USA, Mar. 2022.
[14] T. Sakano, et al., "A study on a photonic network model based on the regional characteristic of Japan," Proc. IEICE Technical Report PN2013-01, Jan. 2013.