Service Restoration in Multi-Entity Network-Cloud Ecosystems: How to Cooperate?

Subhadeep Sahoo, Sugang Xu, Sifat Ferdousi, Yusuke Hirota, Massimo Tornatore, Yoshinari Awaji, and
Biswanath Mukherjee

***Abstract*—In network-cloud ecosystems, cooperation among different entities, e.g., network carriers and datacenter providers (DCPs), is crucial to enhance resiliency, especially during large-scale failures or congestion. However, such cooperation is constrained by limited visibility of confidential information e.g., network topology, resource availability, etc., of different entities owing to proprietary and regulatory policies. To facilitate cooperation, we present and discuss the role of a third-party entity, called Provider Neutral Exchange (PNE), which acts as a broker/mediator and enables cooperation among multiple entities by sharing abstracted (instead of detailed) information of individual entities. We design novel cooperation strategies for post-disaster service restoration and categorize them as: (i) multi-carrier cooperation and (ii) DCP-carrier cooperation. Results under different failure scenarios show benefits of cooperation in terms of service-restoration efficiency, restoration time, and restoration cost.**

***Index Terms*—Carrier, datacenter provider, confidential information, disaster, recovery, service restoration, cooperation.**

I. INTRODUCTION

TODAY'S Internet services are mostly cloud-based (e.g., online storage and software/platform/infrastructure-as-a-service), and future network services (based on large-language models, generative artificial intelligence, etc.) will be more data-centric. These services are supported by network-cloud ecosystems, composed of Datacenters (DCs) and networking resources that offer DC Interconnections (DCIs) through carrier networks [1], [2]. Even though these network-cloud ecosystems, in some cases, are owned by a single entity (e.g., a large over-the-top service provider operating with its own transport network), often networking and computing infrastructures are owned by different entities, namely network carriers and DC providers (DCPs). As shown in Fig. 1, the first and second scenarios are referred to as integrated ecosystem (left) and non-integrated ecosystem (right), respectively. In integrated ecosystems, unified management of networking and computing resources enables efficient and flexible resource utilization. In non-integrated ecosystems, different entities manage their resources independently [3]; however, to ensure that the ecosystem can sustain evolving demands and unprecedented situations (e.g., congestion, failure, etc.), it is imperative that multiple entities cooperate with one another.

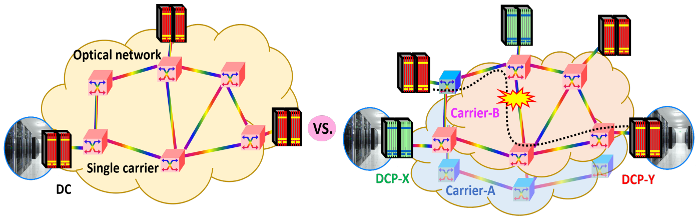

Fig. 1: Integrated vs. non-integrated ecosystem design with multiple entities.

We note that cooperation-based management is crucial, since such infrastructures can be vulnerable to large-scale failures (due to earthquakes, tsunamis, hurricanes, etc.), which may lead to severe resource crunch and service outages even if some baseline resiliency measures are in place [4], [5]. Now, affected services can be restored more effectively if different entities cooperate. For example, in Fig. 1 (right), if a link in Carrier-B fails, services of both Carrier-B and DCP-Y are disrupted (as DCP-Y connects its two DCs using Carrier-B's network). But, if Carrier-B can lease services from Carrier-A, while recovering/repairing its failed link, and DCP-Y cooperates by re-configuring its DCIs to adapt to the failure, affected services of DCP-Y can be restored sooner.

Considering the importance of cloud services today, affected services must be restored as soon as possible to minimize service downtime. In this regard, some proposals for post-disaster network and cloud recovery have been made [6]–[9]. However, these works typically propose recovery strategies for individual carriers/DCPs and do not consider the opportunity for distinct entities to cooperate and share resources during recovery. Some existing studies, such as [10], showed the benefits of coordinated recovery planning across network and DC domains to achieve higher service restoration efficiency compared to non-coordinated recovery. However, such studies consider exchange of detailed information (e.g., DC locations and content availability of DCPs, network topologies, resource availability, failure scenarios, etc.) among carriers and DCPs. In a non-integrated network-cloud ecosystem, sharing of detailed information may not be possible due to privacy and regulatory policies [11]. This limited visibility makes it challenging for the entities to engage in effective cooperation and achieve fast and efficient service restoration. To overcome these challenges, information from different entities can be abstracted and shared while preserving confidentiality and

---

Subhadeep Sahoo, Sifat Ferdousi, and Biswanath Mukherjee are with University of California, Davis, USA; Sugang Xu, Yusuke Hirota, and Yoshinari Awaji are with National Institute of Information and Communications Technology (NICT), Japan; Massimo Tornatore is with Politecnico di Milano, Italy; Biswanath Mukherjee is also with Soochow University, China.

Manuscript received Month xx, 2024; revised Month yy, 2024.

maintaining privacy [12]. Some technologies for inter-domain routing (e.g., BGP) allow to exploit available resources across multi-domain networks and route traffic based on path vectors. However, the problem investigated in this study, namely post-disaster service restoration based on cooperation among multiple entities, requires exchange of information, recovery planning, etc., besides traffic routing. Hence, existing inter-domain routing technologies are enablers for our proposed solutions, but are not sufficient.

To facilitate multi-entity cooperation within a network-cloud ecosystem, we consider a third-party entity, *Provider Neutral Exchange* (PNE), which acts as a broker/mediator and enables information/traffic exchange and resource sharing among different entities. PNE can be a consortium of Internet Exchange Points (IXPs) or carrier hotels operating in different cities/locations. PNE can build a PNE network (PNEN) for the area of interest (e.g., disaster-affected area), which serves as a public reference topology. PNEN comprises of several electronic switching nodes (e.g., IXPs) that interconnect different carrier nodes and DCs, located in close proximity (e.g., within same city), and facilitate traffic exchange [13]. Logical links between PNE nodes are chosen based on a general topology (e.g., a subset of *Japan photonic network* [14]). Now, PNE can broadcast the PNEN to all entities in the ecosystem. For cooperation, all entities can abstract their original topologies of the area of interest to PNEN (e.g., during post-disaster service restoration) and avoid leakage of detailed information. Thus, PNE provides a platform for different entities to cooperate.

This article first considers cooperation among carriers, and then among DCPs and carriers. For multi-carrier cooperation, we discuss how link resources (bandwidth) can be shared (based on some incentive) among the carriers for restoring disaster-affected connection services. We identify a set of overlapping and mutually-beneficial recovery tasks (e.g., link repairs) to minimize recovery effort and time. For DCP-carrier cooperation, we discuss how to jointly conduct DCI re-configuration (by DCP) and recovery planning (by carriers) to restore disaster-affected cloud services. Our contributions are summarized as follows:

- We present novel directions for post-disaster service restoration based on multi-entity cooperation.
- We provide insights regarding the information shared across different entities and its importance in facilitating service restoration; and study how the information can be abstracted to maintain confidentiality.
- We quantify, using realistic study cases, benefits of cooperation over non-cooperation in terms of restoration efficiency, restoration time, and restoration cost.

II. MULTI-ENTITY COOPERATION

Multi-entity cooperation for post-disaster service restoration can be categorized based on the cooperating entities, i.e., (i) multi-carrier cooperation and (ii) DCP-carrier cooperation. Both carriers and DCPs must enhance service restoration efficiency while reducing restoration time and restoration cost. We discuss network model for multi-entity cooperation below.

Before a disaster, carriers provide regular connection services to their customers (e.g., end users, DCPs, etc.) by allocating network resources. DCPs provide regular cloud services to the end users by constructing their own DCI topologies, which

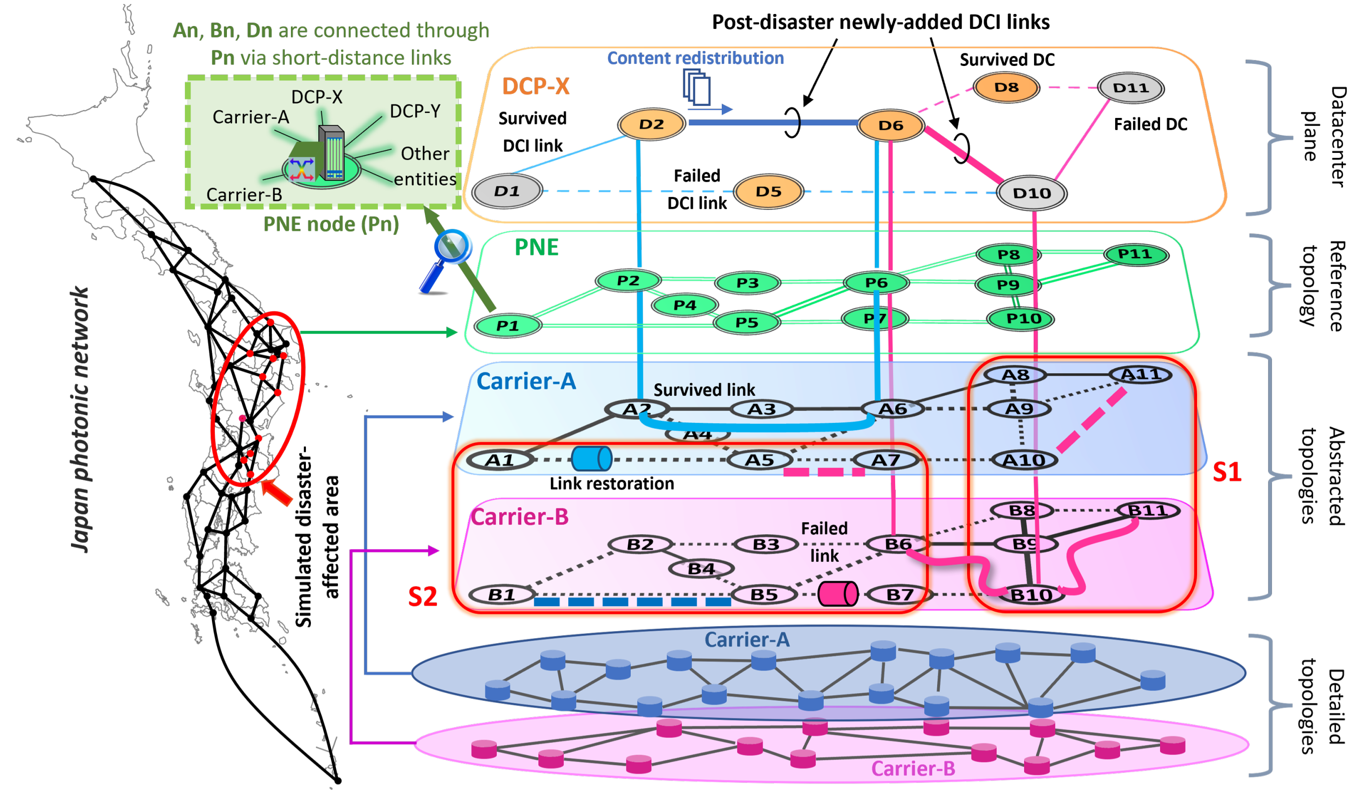

Fig. 2: Multi-entity network-cloud ecosystem.

can be built by leasing connection services (optical lightpaths or IP-over-WDM connections) from carriers. However, a large-scale disaster can cause widespread damage to the infrastructures of all entities and disrupt regular connection and cloud services. To restore affected services, cooperation among the multiple entities is required, for which the entities abstract their original topologies to the PNEN reference topology to avoid leakage of information of their infrastructures. Fig. 2 portrays an example of a network-cloud ecosystem, consisting of two carriers (A and B), one DCP (i.e., DCP-X), and one PNE. Let us consider a subset of *Japan photonic network* as a disaster-affected area (marked in red oval). PNE creates a public reference PNEN topology for that area, represented by *P1-P11*. We consider that each PNE node has enough capacity to handle traffic exchange among all the entities in the ecosystem for service restoration (a different business model for PNE may consider varying capacities depending on operational cost). The bottom two layers represent the topologies of Carrier-A and -B, which are abstracted to PNEN by checking the co-location of carrier nodes with PNE nodes; abstracted Carrier-A and -B topologies are represented by *A1-A11* and *B1-B11*, respectively. Inside abstracted area, failed and survived links of carriers are represented by dashed and solid lines, respectively.

In Fig. 2, for DCI topology, blue and pink lines represent the connection services leased by DCP-X from Carrier-A and -B, respectively. DCI topology is also affected by failures in carrier networks; blue and pink dashed lines represent failed DCI links. We consider that all DC nodes are co-located with PNE nodes and a few DC nodes (*D1, D10, D11*) are identified as failed due to disaster. The magnified image of a PNE node (*P1*) illustrates that it interconnects the co-located nodes (e.g., *A1*, *B1*, *D1*, etc.) of multiple entities through short-reach interfaces. Accordingly, we discuss a set of multi-entity cooperation problems.

A. Multi-Carrier Cooperation

Given disaster-affected carrier networks and disrupted connection services, the objective is to maximize service restoration while minimizing restoration time and cost for carriers by sharing network resources. This cooperation-based scheme enables, with aid of PNE, balanced recovery-task allocation and resource sharing among multiple carriers while ensuring information confidentiality through abstraction.

For example, consider ***disaster scenario 1*** (S1) in Fig. 2. Carrier-A needs to route traffic between *A10* and *A11* but these nodes are disconnected as no path can be established through Carrier-A's network between *A10* and *A11*. Thus, Carrier-A seeks cooperation from Carrier-B, as it has survived resources, and sends traffic through *A10-P10-B10*. Then, Carrier-B carries the traffic through its survived path *B10-B9-B11* and sends the traffic back to Carrier-A through *B11-P11-A11*. Now, consider ***disaster scenario 2*** (S2) in Fig. 2, where both carriers want to route traffic between *nodes 1* and *7*. However, both carriers do not have survived paths between these nodes. Now, PNE instantiates and schedules recovery tasks for both carriers. PNE determines commonly-desired connection requests from both carriers (i.e., between *nodes 1* and *7*), and assigns Carrier-A to recover the link between *A1* and *A5* and Carrier-B to recover the link between *B5* and *B7*, or vice versa. After recovering the link, Carrier-A carries traffic of both Carrier-A and -B from *A1* to *A5*, and Carrier-B carries traffic of both Carrier-A and -B from *B5* to *B7*. The traffic of two carriers is exchanged through PNE nodes (*P1*, *P5*, *P7*).

For recovery-time calculation, we consider that time required by a carrier to recover a link is one time unit (which can range from an hour to days depending on failure situation). Thus, in this example, recovery time is one time unit (for links recovery between *nodes 1* and *7*) which is less compared to standalone (i.e., non-cooperation) recovery, where both carriers need to recover two links (between nodes *1-5* and *5-7*) and hence recovery time would be two units. Therefore, restoration time of affected services can be reduced by reducing recovery time of the failed links and utilizing more survived resources. For both scenarios S1 and S2, carriers need to generate new connection requests to restore their affected services. Based on connection prices (for each PNE node pair) advertised by the carriers (see Sections II.C and III), connection requests can be generated such that total payment from a carrier to other carriers can be reduced. For example, in Fig. 2, connection price advertised between nodes *10* and *11* by Carrier-A is higher than Carrier-B, as there are no survived resources available in Carrier-A network between *A10* and *A11*. Thus, Carrier-A requests connection services from Carrier-B to accelerate service restoration.

Net service-restoration cost can be calculated by subtracting the payment earned from other carriers for resource sharing from the total link-recovery cost incurred by a carrier. Net service-restoration cost can be minimized by minimizing link-recovery tasks performed by carriers. This problem can be stated as follows.

- ***Inputs:*** set of failed links in each carrier, set of affected connection services of each carrier, and set of advertised connection prices for each PNE node pair by the carriers.
- ***Outputs:*** set of connection requests from a carrier to other carriers and set of recovery tasks allocated by PNE.
- ***Objective:*** maximize service restoration while minimizing restoration time and restoration cost of the carriers.
- ***Constraints:*** available network resources (bandwidth), recovery progress (e.g., one link can be recovered in one time unit), and recovery budget of each carrier.

Please see [12] for detailed mathematical model of the multi-carrier cooperation problem.

B. DCP-Carrier Cooperation

For service restoration of a DCP, given disaster-affected carrier networks and DCI topology, the objective is to maximize service restoration while minimizing restoration time for DCP and restoration cost for carriers. This can be achieved through reducing recovery tasks by utilizing survived resources of carriers. The cooperation-based scheme enables, with aid of PNE, resource-driven request generation from DCP to carriers while ensuring information confidentiality through abstraction.

For example, in Fig. 2, for DCP-X, DCs *D1*, *D10*, and *D11* are failed due to a disaster. Thus, content requests of damaged DCs need to be redirected to backup DCs. Let us consider that, for content *c*, *D10* is primary and *D2* is backup DC. Hence, all requests for *c* need to be redirected from *D10* to *D2*. However, redirecting all requests directly from *D10* to *D2* consumes more network resources (as backup DCs are usually placed at disaster-disjoint far-away locations) and might not be feasible due to post-disaster resource crunch in carrier networks (e.g., if links between *A1-A5*, *A5-A7*, *B1-B5*, and *B5-B7* are failed). Thus, DCP-X re-distributes (through replication) content *c* to DC *D6* which is near *D10* and redirects requests for *c* to *D6*. For such content redistribution and request redirection, DCP-X modifies its DCI topology and requests new connections from the carriers; carriers then provision these new requests accordingly, e.g., connection requests between *D2-D6* and *D6-D10* are provisioned by Carrier-A and -B, respectively. Similar to the previous problem, DCP generates new connection requests based on connection prices (for each PNE node pair) advertised by all carriers, to minimize total restoration cost. For example, in Fig. 2, connection price advertised between nodes *2* and *6* by Carrier-B is higher than that of Carrier-A, as there are link failures in Carrier-B network between *B2* and *B6*. Thus, DCP-X requests connection services from Carrier-A to accelerate service restoration. This problem can be stated as follows.

- ***Inputs:*** set of failed DCs, set of failed DCI links, set of affected DCP requests, and set of advertised connection prices for each PNE node pair by the carriers.
- ***Outputs:*** set of connection requests from DCP to carriers.
- ***Objective:*** maximize service restoration while minimizing restoration time of DCP and restoration cost of carriers.
- ***Constraints:*** available network resources (bandwidth), recovery progress (e.g., one link can be recovered in one time unit), and recovery budget of each carrier.

C. Information Sharing by Different Entities

Information sharing plays an important role in cooperation-based service restoration. When DCs and network resources are owned by a single entity, detailed information of network resource availability, damaged infrastructures, DC content location, etc., is available and can be analyzed to perform joint network and DC recovery. However, in multi-entity network-cloud ecosystems, confidential information of an entity should not be shared. Here, we summarize different types of information, ownership, level of confidentiality, importance for cooperation, and how information can be shared through abstraction.

(i) PNEN reference topology (public): PNE generates the reference topology (PNEN) which is broadcasted to all entities, which can create their abstracted topologies based on PNEN.

(ii) Carrier topology (private): Carriers must not share their own detailed topologies. They abstract their topologies to PNEN reference topology for traffic exchange with other entities through PNE nodes within proximity.

(iii) Failure data (private): Real failure footprint in carrier networks (i.e., which links are failed, which are survived) should not be shared. Since such information is crucial for efficient cooperation, details of failures are indirectly reflected in connection-service prices generated by carriers.

(iv) Resource availability (private): Resource availability (bandwidth) in carrier networks is private information but is essential in cooperation. Carriers abstract such information in connection-service prices.

(v) DCI topology (private): DCPs should not share their detailed DCI topologies. They abstract their DCI topologies to PNEN reference topology.

(vi) Content location and availability (private): Primary/backup DC location of a content and availability of a content (e.g., after DC failure) are private to the DCP.

III. COOPERATION STRATEGY

Figure 3 illustrates various phases (Ph.) of the underlying cooperations. Both cooperation strategies are composed of multiple sub-tasks conducted by different entities and, after a disaster event, PNE initiates cooperation by broadcasting PNEN reference topology to all entities at phase 0.

A. Multi-Carrier Cooperation

- ***Phase 1:*** Carriers analyze post-disaster failure situation in their network infrastructures and perform initial planning for standalone recovery based on survived resources.
- ***Phase 2:*** Carriers generate price of connection service for each PNE node pair based on failure evaluation (failure information is abstracted in price). Carriers add an extra price for connections that are routed over failed links (which need to be repaired) to minimize workload on link recovery and maximize utilization of survived links.
- ***Phase 3:*** After gathering connection prices from each carrier, PNE broadcasts the prices to all carriers.
- ***Phase 4:*** Carriers first calculate the recovery cost and recovery time of failed links without considering cooperation from other carriers. Then, carriers calculate the recovery cost by considering connection support from other carriers and evaluate cost benefits. If cooperation can reduce recovery cost and maintain recovery budget, carriers generate connection requests for other carriers and send them to the PNE.
- ***Phase 5:*** Upon receiving the connection requests from all carriers, PNE identifies the set of requests that are desired by multiple carriers. To accommodate these requests efficiently, PNE assigns recovery tasks among carriers to ensure a balanced workload; this allows multiple carriers to perform recovery simultaneously, thereby minimizing recovery time.
- ***Phase 6:*** Upon receiving connection requests from PNE, carriers analyze them and inform the PNE whether a request is: (a) immediately acceptable, (b) acceptable but requires restoration time of *k* units, or (c) rejected. PNE broadcasts the information to all carriers. Then, carriers determine whether the restoration time is acceptable or not, and finalize the connection requests. Phases 4-6 are repeated till all services are restored.

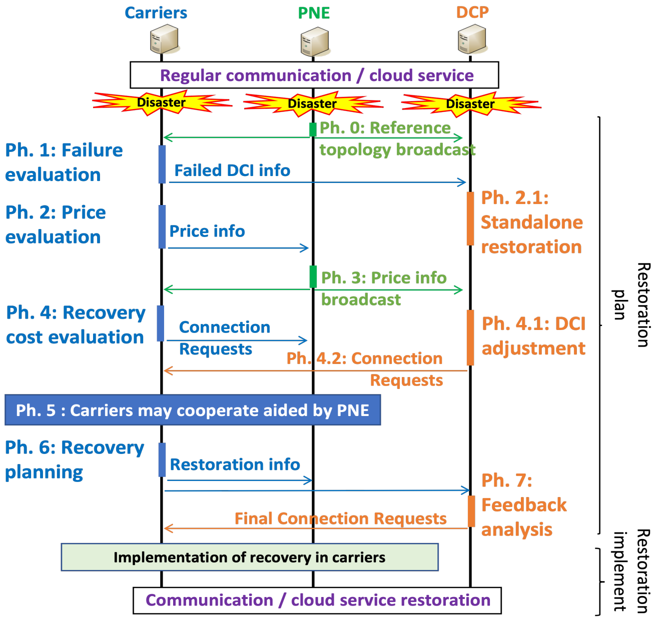

Fig. 3: Different phases of multi-entity cooperation.

B. DCP-Carrier Cooperation

- ***Phase 1:*** Similar to Phase 1 in multi-carrier cooperation, carriers analyze post-disaster failure situation in their network infrastructures and perform initial standalone recovery based on survived resources for the original DCI topology. Then, they inform DCP about the failed DCI links that could not be recovered by standalone recovery.
- ***Phase 2, 2.1, 3:*** Phase 2 is similar to Phase 2 in multi-carrier cooperation. Also, in Phase 2.1, after receiving the DCI failure information from the carriers, DCP performs service restoration by utilizing available bandwidth of survived DCI links. In case of DC failures, content requests are redirected from the failed primary DCs to backup DCs. However, only a few requests can be restored using the survived DCI links. In Phase 3, PNE broadcasts the price information to all carriers and DCP.
- ***Phase 4.1:*** Upon receiving the price information, DCP performs content redistribution to replicate content among the survived DCs and to accommodate post-disaster traffic. Then, affected users of failed primary DCs are redirected to new backup DCs. For this joint content redistribution and user redirection, DCP modifies its original DCI topology and generates connection requests for carriers to establish new DCI links while avoiding high-priced connections.
- ***Phase 4.2:*** Upon receiving all connection requests from DCP, PNE forwards them to the respective carriers for evaluation.
- ***Phase 6:*** Upon receiving all connection requests, carriers try to maximize the provisioning of connection requests in a best-effort manner while minimizing recovery tasks for their own networks and maintaining the recovery budget. After evaluation, carriers send feedback to PNE on whether a connection request is: (a) immediately acceptable, (b) acceptable but requires restoration time of k units, or (c) rejected.
- ***Phase 7:*** After receiving feedback for all the connection requests from carriers, PNE sends the feedback to DCP, which analyzes the feedback of each request, finalizes the connection requests, and sends it to the corresponding carriers. Then, carriers implement their recovery planning and provision the DCP requests.

IV. EVALUATION OF COOPERATION STRATEGY

A. Evaluation Setup

We demonstrate the effectiveness of the cooperation strategies by evaluating a network-cloud ecosystem with two carriers (Carrier-A/B), one DCP (DCP-X), and one PNE (see Fig. 2). Abstracted topologies of both carriers have 11 nodes and 15 links; DCI topology has 7 DC nodes and 11 DCI links. Failure patterns are generated using a correlation value of 0.8 among link failures in both carrier networks, i.e., if a link fails in Carrier-A network, the co-located link in Carrier-B also fails with probability 0.8. We consider three failure patterns: (i) **H**eavy damage (10:10), where 10 links are failed in both carriers; (ii) **M**ixed damage (10:5), where 10 links are failed in Carrier-A and 5 links are failed in Carrier-B or vice-versa; (iii) **L**ight damage (5:5), where 5 links are failed in both carriers. For simplicity, we assume that a carrier can recover one link in one time unit. Recovery cost of a link (i.e., cost of physical repair) is a random value uniformly selected in [1,10] units of price. For each failure pattern, 150 instances are generated randomly to observe consistency in numerical analysis. We assume that four lightpath channels (each with 100 Gbps capacity) are available in each optical fiber link to accommodate the post-disaster congestion in carrier networks. We consider that 3 to 4 DCs are failed in different instances of disaster. We assume that about 10% of existing user requests are affected by a disaster, i.e., 20,000 requests of each carrier for *multi-carrier* and 15,000 requests of DCP for *DCP-carrier* cooperation. Bandwidth requirement of each user request is uniformly distributed between 70 to 100 Mbps in steps of 1 Mbps. We set the price as 1 unit for 100 Gbps if a connection can be established by the survived links. An extra dummy price of 50 units is added (to abstract failure information) if a carrier needs to recover a failed link before establishing a connection. However, final price of a connection with recovered links during payment is set as 4 units. Average computational time of these simulations (using IBM CPLEX and Python on a server equipped with Xeon Gold5115 2.4-GHz 20-core CPU, 128 GB memory) is in the order of minutes for each instance.

B. Illustrative Results

To evaluate the advantages of cooperation, we focus on two metrics: *service restoration efficiency*, i.e., number of restored services over total number of affected services; and *net restoration cost*, which is calculated by subtracting the payment earned from other carriers for resource sharing from the total link-recovery cost incurred by a carrier.

*1) **Multi-Carrier Cooperation**:* Fig. 4(a) shows that cooperation among multiple carriers achieves higher restoration efficiency with less restoration time compared to baseline scheme (i.e., non-cooperation). For heavy damage, both strategies

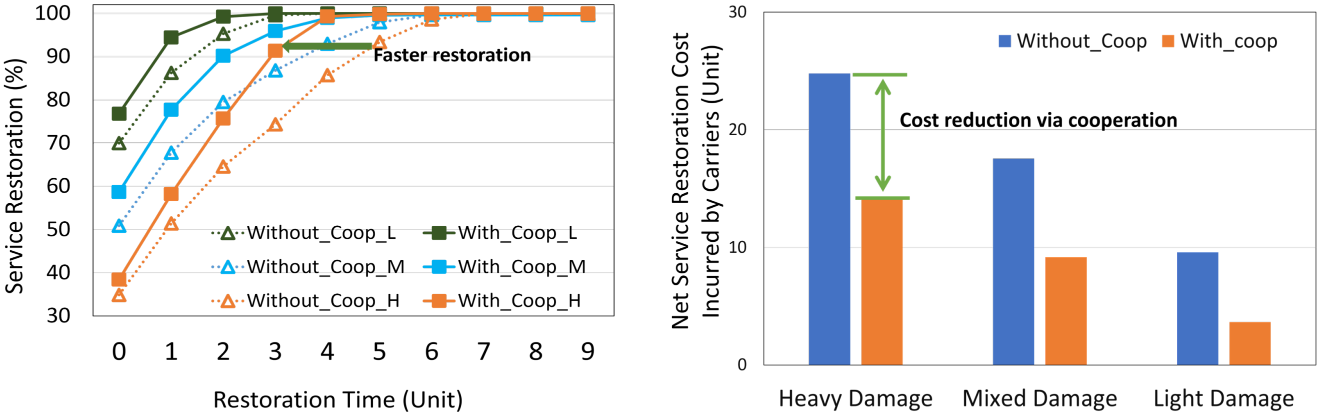

Fig. 4: (a) Service restoration efficiency of carriers, (b) Net cost incurred by carriers to achieve full service restoration.

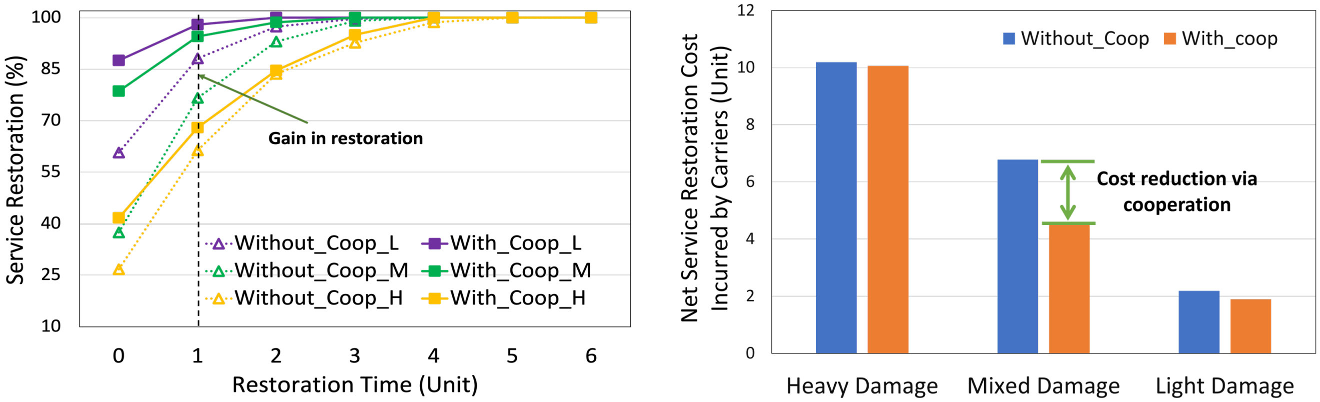

Fig. 5: (a) Service restoration efficiency of DCP, (b) Net cost incurred by carriers to achieve full service restoration of DCP.

perform similarly at the beginning (at restoration time 0) due to resource crunch in both carriers. However, with progression of restoration time, cooperation outperforms the baseline and achieves higher service restoration efficiency in less restoration time. For example, with cooperation, restoration of 90% of the services is achieved in 40% less time than baseline because (1) by sharing survived resources among each other, carriers restore their connection service and minimize the link-recovery tasks; and (2) by sharing the recovery tasks prompted by PNE, better restoration planning is achieved, leading to further minimization of link-recovery tasks. Similarly, for mixed and light damage, cooperation achieves higher restoration efficiency by at least 15%. By varying link-recovery costs, similar trends are observed.

Fig. 4(b) illustrates net restoration cost incurred by carriers to achieve full service restoration. Net cost is higher for heavy damage and lower for mixed and light damage, since more links of both carriers need to be recovered in case of heavy damage. With cooperation, net service restoration cost is reduced by 42%, 48%, and 62% compared to baseline, for heavy, mixed, and light damage, respectively. With cooperation, recovery tasks are shared (and hence minimized) which results in lower recovery cost. Moreover, by sharing resources, carriers achieve income from each other and minimize net restoration cost.

*2) **DCP-Carrier Cooperation**:* Fig. 5(a) reports service restoration efficiency of a DCP w.r.t. restoration time and compares the gain with baseline (non-cooperation), in which DCP restores its affected services using survived resources in DCI network without modifying the original DCI topology. For all damage scenarios, at restoration time 0, cooperation between DCP and carriers increases the restoration efficiency by 50% compared to baseline. We observe that, with progression of restoration time, cooperation strategy restores all affected services faster than baseline. The reason is that, with cooperation, DCP modifies its original topology by adapting the failure pattern. Thus, carriers can offer their survived resources to support the modified DCI topology and recover a few links also, if required. However, with baseline, DCP cannot utilize survived resources in carrier networks and need to wait for carriers to recover the original failed DCI links, which leads to higher restoration time. Fig. 5(a) illustrates that, among all damage scenarios, cooperation achieves higher gain for mixed damage. This is because more connection service requests from DCP could be satisfied immediately by one of the carriers that has comparatively more survived resources.

Fig. 5(b) illustrates the net restoration cost incurred by carriers for full DCP service restoration. In heavy damage, net cost is higher because it is difficult for DCP to avoid requesting new connection services from carriers that need to be recovered first, as both carriers lack survived resources. However, in mixed damage, with cooperation, net cost is reduced by 30% compared to baseline because, with cooperation, DCP modifies its original topology by adapting to the failure pattern. Thus, DCP tries to maximize the utilization of survived resources in both carriers and avoids connection requests that require carriers to recover first. This leads to minimizing the recovery cost for carriers. In light damage, DCP has limited requirements for connection service from carriers as most of the affected DCP services can be recovered by the survived DCI links. Hence, net cost is lower but not significantly reduced compared to baseline.

Due to post-disaster resource crunch in carrier networks, service restoration becomes more challenging when multiple DCPs are introduced in the ecosystem competing for the same connection services from the carriers. Readers can refer to [15] which presents how PNE adjusts the imbalance in resource allocation and achieves efficient service restoration by matching the demand of connection services from multiple DCPs with the supply of bandwidth resources from carriers.

V. CONCLUSION AND OPEN PROBLEMS

We presented multi-carrier and DCP-carrier cooperation strategies for post-disaster service restoration. We investigated the role of a third-party neutral entity (PNE), which acts as a broker/mediator and enables cooperation among multiple entities. We analyzed information to be exchanged across different entities (carriers and DCPs) and presented how the information can be abstracted to maintain confidentiality. Numerical results show that cooperation strategies can achieve significant improvement in service restoration and reduction in restoration time and restoration cost.

The problem can be extended by introducing additional entities, such as Internet Service Providers and enterprises into the ecosystem and developing the business model of the PNE. Furthermore, enhanced cooperation strategies (e.g., using game theory) can be investigated to ensure fairness among the entities. This involves considering various revenues, such as income from serving own customers and sharing resources with competitors, while aligning incentives appropriately.

ACKNOWLEDGMENTS

This work is supported by US-Japan JUNO3 project: NSF Grant No. 2210384.

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] B. Shariati et al., "Demonstration of federated learning over edge-computing enabled metro optical networks," *Proc. European Conference on Optical Communications*, Brussels, Belgium, Dec. 2020.

[4] Y. Liu et al., "Disaster protection in inter-datacenter networks leveraging cooperative storage," *IEEE Transactions on Network and Service Management*, vol. 18, no. 3, Sep. 2021.

[5] N. H. Bao et al., "Adaptive path splitting based survivable virtual network embedding in elastic optical networks," *Optical Fiber Technology*, vol. 54, Jan. 2020.

[6] A. Agrawal et al., "Network and risk modeling for disaster survivability analysis of backbone optical communication networks," *Journal of Lightwave Technology*, vol. 37, no. 10, May 2019.

[7] K. Hazra et al., "A novel network architecture for resource-constrained post-disaster environments," *Proc. International Conference on Communication Systems & Networks*, Bengaluru, India, Jan. 2019.

[8] R. Zou et al., "DeepDRAMA: deep reinforcement learning-based disaster recovery with mitigation awareness in EONs," *Proc. IEEE Global Communications Conference*, Madrid, Spain, Dec. 2021.

[9] L. Tomás et al., "Disaster recovery layer for distributed openstack deployments," *IEEE Transactions on Cloud Computing*, vol. 8, no. 1, Aug. 2017.

[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, Sep. 2020.

[11] B. Shariati et al., "Inter-Operator Machine Learning Model Trading over Acumos AI Federated Marketplace," *Proc. Optical Fiber Communication Conference*, San Francisco, CA, USA, June 2021.

[12] S. Xu et al., "A Novel strategy of carrier cooperation with coordinated scheduling for swift failure/disaster recovery," *Proc. Optical Network Design and Modeling Conference*, Coimbra, Portugal, May 2023.

[13] Light Reading Inc., "The Divide: Connected Nation's Brent Legg on the need to build more carrier-neutral IXPs," Mar. 2023.

[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.

[15] S. Sahoo et al., "Strategic cooperation among datacenter providers and optical-network carriers for disaster recovery," *Proc. IEEE Global Communications Conference*, Rio de Janeiro, Brazil, Dec. 2022.

VI. BIOGRAPHIES

Subhadeep Sahoo received B.Tech. and M.S. degrees from West Bengal University of Technology, India, and Chongqing University of Posts and Telecommunications, China, in 2015 and 2020, respectively. He is currently pursuing Ph.D. in Computer Science at University of California, Davis.

Sugang Xu joined National Institute of Information and Communications Technology (NICT), Tokyo, Japan, in 2005. He has been engaged in research on new-generation network architecture and resilience of photonic networks. He is a member of IEEE and IEICE.

Sifat Ferdousi received B.S. degree in Electronics and Telecommunication Engineering from North South University, Bangladesh, in 2007, and M.S. and Ph.D. degrees in Electrical and Computer Engineering from University of California, Davis, in 2011 and 2017, respectively, where she is currently a Postdoctoral Researcher. Her research interests include optical network design and cloud network resiliency.

Yusuke Hirota (M'08) received Ph.D. degree from Osaka University, Japan, in 2008. He joined Osaka University, Japan, in 2008, as an Assistant Professor. Since 2017, he has been with NICT, Tokyo, Japan. He is a member of Optica, IEEE, and IEICE.

Massimo Tornatore is a Professor at Politecnico di Milano. He has also held appointments as Adjunct Professor at University of California, Davis, and as visiting professor at University of Waterloo, Canada. His research interests include optimization and design of communication networks, network reliability, and machine learning application for network management. Tornatore is an IEEE Fellow.

Yoshinari Awaji received B.E. degree in electronic physics in 1991 and M.E. and D.Eng. degrees in 1993 and 1996, respectively, from Tokyo Institute of Technology, Japan. In 1996, he joined Communications Research Laboratory (CRL: currently NICT). He is a Director at NICT, Japan. He was engaged in information security strategies at the Cabinet Secretariat from 2004 to 2006. His research interests include resilient optical network technologies. He is a member of IEEE/Photonics Society and IEICE.

Biswanath Mukherjee is Distinguished Professor Emeritus at University of California, Davis. He received B.Tech. degree from Indian Institute of Technology, Kharagpur (1980) and Ph.D. from University of Washington, Seattle (1987). He was General Co-Chair of OFC 2011, TPC Co-Chair of OFC '09, and Technical Program Chair of IEEE INFOCOM '96. He is also President and Founder of Ennetix, which specializes in AIOps solutions for application performance and security analytics. He is an IEEE Fellow.