A Novel Use Case of Future Multi-Carrier Interconnection—Swift and Low-Cost Disaster Recovery via Carrier Cooperation

Sugang Xu[1], Noboru Yoshikane[2], Xiaocheng Zhang[3], Subhadeep Sahoo[4], Sifat Ferdousi[4], Masaki Shiraiwa[1], Yusuke Hirota[1], Takehiro Tsuritani[2], Shigenari Suzuki[3], Massimo Tornatore[5], Yoshinari Awaji[1] and Biswanath Mukherjee[4, 6]

[1]National Institute of Information and Communications Technology (NICT), Tokyo, Japan,
[2]KDDI Research, Inc., Saitama, Japan,
[3]NTT Docomo Business, Tokyo, Japan,
[4]University of California, Davis, USA,
[5]Politecnico di Milano, Milan, Italy,
[6]Soochow University, Suzhou, China

Abstract— We present our latest development and experimental validation of carrier cooperative recovery for enhancing the resilience of optical packet transport networks. Experimental results prove that in case of resource crunch caused by, e.g., traffic congestion, failures, man-made/natural disasters, etc., swift and low-cost recovery can be achieved by exploiting the interconnection capability among carriers, which demonstrates a novel use case of multi-carrier interconnection technology.

1. INTRODUCTION

To meet the growing demand for cloud services, optical networks and data centers (DCs) together constitute network–cloud ecosystems (hereafter referred to as ecosystems). As these ecosystems evolve toward greater scale and openness, interconnection and cooperation among diverse entities, such as telecom carriers (carriers) and data center providers (DCPs), become both inevitable and essential. To promote innovation in large-scale network–cloud services and enrich these ecosystems, interconnection technologies are being actively explored in various open communities, including the Innovative Optical and Wireless Network (IOWN) Global Forum [1], Optical Internetworking Forum (OIF) [2], Telecom Infra Project (TIP) [3], and OpenROADM [4], etc. These initiatives are developing architectures and specifications to enable interoperability of both the Data Plane (D-Plane) and Control/Management Plane (C/M-Plane) in multi-domain, multi-carrier networks. In parallel, architectural design and field trials of a data center exchange (DCX) for interconnecting distributed DCs with carrier networks have also been undertaken [5]. These efforts collectively underscore the importance and feasibility of interconnection technologies.

As these ecosystems grow increasingly important, their resilience against resource shortages becomes critical to ensuring the reliability of essential services. In particular, the resilience of carriers is vital, as all entities depend on their communication services. To strengthen the ecosystem resilience, we have conducted modeling studies on confidentiality-preserving cooperative recovery planning among multiple entities, leveraging interconnection capabilities facilitated by a third-party entity known as the Provider-Neutral Exchange (PNE) [6–8]. Furthermore, in [9], we demonstrated a proof-of-concept prototype of a Multi-entity Cooperation Platform (MCP), developed using Distributed Ledger Technology (DLT). This prototype successfully validated both model-driven DCP-carrier cooperative recovery planning and cooperative recovery within the D-Plane, illustrating the potential for multi-entity cooperation to enable rapid ecosystem recovery.

In this paper, we present, for the first time, a proof-of-concept of carrier cooperative recovery, including experimental validation of both a DLT-based cooperative recovery planning mechanism and the corresponding D-Plane recovery. Specifically, we extend our previous work in [9] by introducing a policy-based control mechanism used in carrier recovery planning, that allows carriers to flexibly incorporate their individual concerns into (1) pricing and (2) resource allocation during cooperation. Furthermore, building on our earlier work on carrier recovery planning [6], which introduced a Cost-First (CF) strategy, we propose a novel Leadtime-First (LF) strategy to further reduce recovery time and cost in cooperative scenarios. Through this work, we successfully address the following key questions: (i) Can we experimentally demonstrate that carrier cooperative recovery is achievable without violating confidentiality? (ii) Can fast recovery be realized through cooperation among carriers using a model aligned with current direct interconnection practices? (iii) Can efficient cooperative recovery be achievable by applying appropriate policies and strategies in cooperative resource allocation?

The remainder of this paper is organized as follows. Section 2 presents the system structure of carrier cooperative recovery planning. Section 3 introduces three policies and two strategies of carriers in cooperation. Section 4 demonstrates how is carrier cooperative recovery planning performed, and showcases the D-Plane recovery based on the recovery plan yielded from carrier cooperative recovery planning. Section 5 presents the numerical results of performance evaluation under different policies and strategies. Finally, Section 6 concludes the paper.

2. CARRIER COOPERATIVE RECOVERY PLANNING

2.1 System Structure for Interconnection and Cooperative Recovery

Figure 1(a) depicts a scenario of carrier cooperative recovery [6] in case of resource crunch. For regular services, carriers who have the overlapped coverage are interconnected via distributed PNE sites (P*x*) to enable the multi-carrier large-scale communication service. Carrier optical transport networks are multi-layer networks. The underlying optical networks offer lightpath services to upper-layer packet transport networks, e.g., Multiprotocol Label Switching (MPLS) networks. Figure 1(b) shows a direct interconnection model within a PNE site, which is close to the existing practice of interconnection between carriers. For example, nodes (packet routers/switches) of different carriers, DCPs, and users/enterprises in close proximity (e.g., in the same city) can be collocated in a single PNE site and directly interconnected to exchange traffic during regular service operation. Figure 1(c) illustrates a more futuristic indirect interconnection model, in which PNE nodes (packet routers/switches) can be introduced into PNE sites for facilitating indirect interconnection when it is required. In our current work, we assume the simple and direct interconnection model between carriers. The indirect interconnection model is envisioned as future work. As another novel use case of carrier interconnection beyond the large-scale multi-domain and multi-carrier regular communication service, in case of large-scale failures/disasters, carriers can employ these PNE interconnection sites to perform cooperative recovery by efficiently utilizing the sparsely distributed surviving resources in different carrier optical networks.

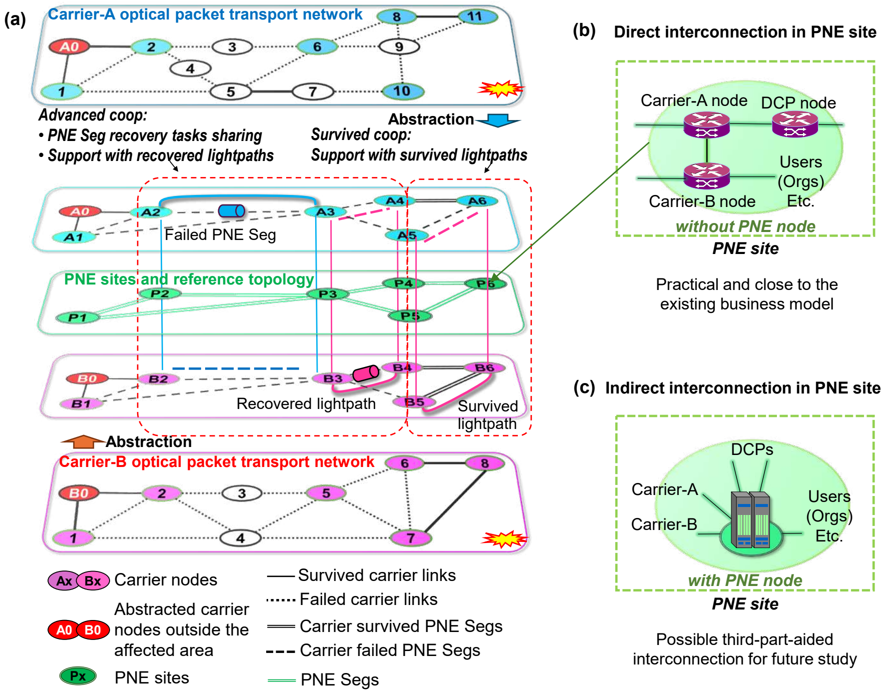

Figure 1: Concept of carrier cooperative recovery aided by PNE: (a) Shared view of carriers with abstraction; (b) Direct interconnection in PNE site; and (c) Indirect interconnection with PNE nodes.

After a disaster, in the affected area, the damaged fiber links of carrier networks are shown in the figure as dotted single lines. To enable carrier cooperation without violating confidentiality, PNE can declare a logical reference topology among the distributed PNE sites (e.g., the Japan photonic network model [10] as shown in Fig. 1(a), etc.). Each carrier can abstract its real network topology to this common public PNE reference topology for concealing their detailed network topology and damage information [6–8]. In these abstracted carrier topologies, the dashed lines are the carriers' abstracted damaged PNE segments (Segs, for short) (e.g., consisting of multiple damaged fiber links); the double lines are the survived PNE Segs (with survived reachability between PNE sites). Note that the red colored nodes A0 and B0 are the abstracted nodes representing the nodes outside the affected area for Carrier-A and -B, respectively. The abstracted topologies are shared among carriers forming a common view in the affected area to facilitate carrier cooperation for swift and low-cost recovery.

To achieve efficient recovery, carriers can offer each other lightpath supports with their survived resources ("*lightpath-SP-I*," for short) through PNE sites. For instance, Carrier-B offers a *lightpath-SP-I* between <B5, B6> to Carrier-A via PNE sites P5/P6. When the *lightpath-SP-I* is not applicable, carriers may need to restore the reachability between certain PNE sites. If this requirement is for the PNE Segs and carriers have mutually-desired PNE Segs, then carriers can share the recovery tasks of these mutually-desired PNE Segs. After the recovery of the PNE Segs, carriers can offer each other another type of lightpath support with the recovered resources ("*lightpath-SP-II*," for short). For instance, in the same PNE reference topology, Carrier-A and -B require the same damaged PNE Segs <A2, A3>, <A3, A4>, and <B2, B3>, <B3, B4>, respectively. They can undertake the recovery tasks of <A2, A3> and < B3, B4>, respectively, and offer each other *lightpath-SP-II* via PNE sites P2/P3/P4.

2.2 Framework of Cooperative Recovery Planning

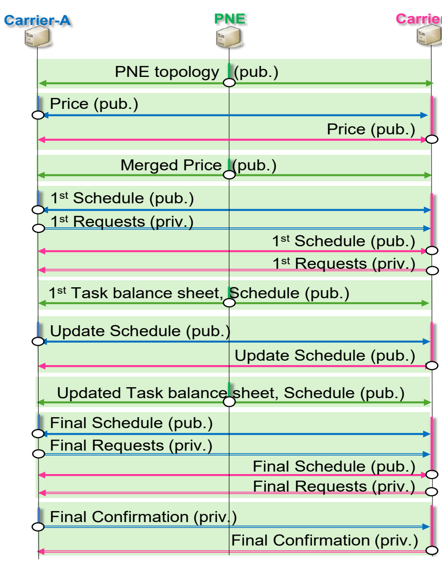

Ph.	Actions
Ph1	**[PNE]** PNE topology broadcast
Ph2	**[Carrier-A/B]** Standalone recovery planning, Price and Policy definition and broadcast
Ph3	**[PNE]** Merged price/policy broadcast
Ph4	**[Carrier-A/B]** Initial recovery task planning/scheduling and broadcast **[Carrier-A/B]** Analysis and requests of counterpart carriers' *lightpath-SP-I*
Ph5	**[PNE]** Initial recovery task matching/balancing for carriers
Ph6	**[Carrier-A/B]** Recovery task planning/scheduling refinement and broadcast
Ph7	**[PNE]** Recovery schedule adjustment and broadcast
Ph8	**[Carrier-A/B]** Final recovery task schedule confirmation and broadcast **[Carrier-A/B]** Final requests for counterpart carriers' *lightpath-SP-I*
Ph9	**[Carrier-A/B]** Analysis and final confirmation of counterpart carriers' requests for *lightpath-SP-I*

o Source of info sharing/exchange ↔ Public data sharing via broadcast → Private data exchange between the relevant carriers

Figure 2: Framework of multi-phase carrier cooperative recovery planning.

In cooperative recovery planning, since individual carriers may have their own concerns and policies, the planning process must be carried out in a distributed manner. Figure 2 illustrates the framework of carrier cooperative recovery planning used to compute the solution for actual cooperative recovery, based on our previous modeling work [6]. Within this framework, carriers and the PNE execute their respective planning subtasks (actions shown on the right-hand side) across nine distinct phases. Specifically, in phases Ph2, Ph4, Ph6, Ph8, and Ph9, carriers conduct their Carrier-Side Planning Tasks (CSPT) and associated scheduling (CSPT-Scheduling), which include: Ph2, standalone recovery planning and pricing; Ph4, initial cooperative recovery planning and scheduling; Ph6, refinement of cooperative recovery planning and scheduling; Ph8, request for counterpart carriers' *lightpath-SP-I*; and Ph9, final confirmation of these requests. Meanwhile, in phases Ph1, Ph3, Ph5, and Ph7, the PNE executes PNE-Side Matching Tasks (PSMT), which consist of: Ph1, PNE topology broadcast; Ph3, merging and broadcasting of price information; Ph5, initial matching, balancing, and scheduling of recovery tasks for carriers; and Ph7, refinement of the matching, balancing, and scheduling process.

In each phase, the corresponding optimization results are exchanged among stakeholders to enable the execution of planning subtasks in the next phase by other stakeholders (as illustrated on the left-hand side). Abstracted public information, such as the PNE topology, the pricing of connection services between PNE sites (e.g., in the form of lightpaths or IP-over-WDM connections), and the scheduling of PNE segment recovery, is broadcast to all stakeholders. In contrast, private information, such as requests for counterpart carriers' *lightpath-SP-I* and their subsequent confirmations, is exchanged only between the relevant carriers to preserve confidentiality.

Following the cooperative recovery planning process, individual carriers first provide *lightpath-SP-I* to each other. Then, they proceed to implement the assigned PNE segment recovery tasks and offer *lightpath-SP-II* to their counterparts as necessary. Details of the pricing mechanisms, along with related policy considerations, are discussed in Section 3.1. Readers are referred to [6] for further details on the planning subtasks of carriers and the PNE. The overall cooperative recovery planning process, including distributed and parallel subtasks and mechanisms for information sharing and exchange, has been modeled and implemented using our DLT-based MCP prototype [9], and is demonstrated in Section 4.

3. POLICY AND STRATEGY IN CARRIER COOPERATIVE RECOVERY PLANNING

3.1 Policy Consideration in Cooperative Recovery Planning

During the aforementioned cooperation, carriers may require the ability to define specific policies when offering network resources to counterpart carriers. In this section, building upon our previous work [6], we propose three potential policies. By appropriately setting the values of these policies, carriers can promote efficient recovery through cooperation while preventing the misuse of valuable resources. These policies can then be communicated to counterpart carriers through the aforementioned information exchange, allowing carriers to express their concerns regarding resource sharing.

(1) Price of *lightpath-SP-I* with survived resource

Each carrier a can define its price $p'^{a}_{i,j}$ of a *lightpath-SP-I* service between PNE sites <i, j>. This consists of a "regular price" and a "dummy price". The regular price stands for a price charged for the regular daily service. The dummy price (i.e., not a true price) as an auxiliary information abstracts the damage status of the carrier optical network. Namely, if the regular price and the dummy price of a *lightpath-SP-I* are the same, the *lightpath-SP-I* is then defined as available immediately over survived resources. Meanwhile, if a very higher dummy price is additionally declared, then the *lightpath-SP-I* is defined as unavailable (e.g., no survived resource is available and needs recovery first). This policy can be reflected in carriers pricing and notified to counterpart carriers (e.g., in phase Ph2). Upon receiving the price information, the counterpart carriers will optimize their recovery plans with CSPT by using the available *lightpath-SP-I* if it is beneficial to reduce the recovery time and cost. Note that, in our current work, it is expected that a "regular price" is charged for final services.

(2) Total number of wavelengths W employed in cooperative recovery

In our current work, we assume that carriers are willing to allocate a portion of their network resources for sharing in cooperative recovery. For instance, out of a total of T wavelengths in a carrier's optical network, a carrier may designate W wavelengths (where $W < T$) to be reserved specifically for resource sharing in cooperative recovery. Within these W wavelengths, support for *lightpath-SP-I* and *-II* can be provided to counterpart carriers. Note that, when a carrier (e.g., Carrier-A) asks the support of *lightpath-SP-II* of a counterpart carrier (e.g., Carrier-B), the capacity demand of Carrier-A must match the capacity supply of Carrier-B, i.e., the demand cannot be larger than the supply. Hence, this policy is a common concern of carriers and should be decided among carriers prior to the cooperative recovery planning.

(3) Maximum availability of *lightpath-SP-I* with survived resource between PNE sites

Extending the two aforementioned policies which have been modeled in CSPT [6], we introduce a new upper limit $U^{a}_{i,j}$ on the number of carrier-a's *lightpath-SP-I* connection services utilizing surviving resources between PNE sites <i, j>. This policy can be defined by carriers in conjunction with the aforementioned pricing policy and notified to counterpart carriers (e.g., during phase Ph2). Upon receiving this upper limit information, the counterpart carriers can incorporate $U^{a}_{i,j}$ as a new upper-bound constraint of the integer variable $\sigma^{a}_{i,j}$ (for solving the desired number of carrier-a's *lightpath-SP-I* between PNE sites <i, j>) in CSPT model [6] for restricting the demands on *lightpath-SP-I* between certain PNE sites. This policy can help provider carriers (i.e., those offering shared resources) prevent overuse or depletion of surviving resources by consumer carriers (i.e., those utilizing *lightpath-SP-I*) between specific PNE sites.

These three policy concerns are implemented and experimentally demonstrated in Section 4. The impact of these polices on cooperation is observed in Section 5.

3.2 Strategy in Cooperative Recovery Planning

Upon receiving policy information from provider carriers, namely, the aforementioned pricing and maximum availability of *lightpath-SP-I*, also as a consumer carrier each carrier performs recovery planning using CSPT, leveraging both *lightpath-SP-I* and *-II* of counterpart carriers to minimize recovery time and cost. Various strategies can be applied in recovery planning, such as prioritizing cost or recovery lead time. In our previous work [6], we introduced a Cost-First (CF) strategy in CSPT, which aims to minimize the combined cost of fiber link recovery and payments for acquiring *lightpath-SP-I*. In this study, we propose a novel Leadtime-First (LF) strategy that prioritizes minimizing the number of recovery tasks, thereby reducing the overall recovery lead time. For simplicity, only the objective functions of these two strategies are presented below, as both share the same set of constraints. Readers are referred to [6] for detailed descriptions of the CSPT model.

(1) Cost-First strategy [6]

$$\min\Bigg[-C_1\Big(\sum_{\alpha\in\Delta}\sum_{(s,d)\in R}\Gamma^{\alpha}_{s,d}A^{\alpha}_{s,d}\alpha^{\alpha}_{s,d}+\sum_{\alpha\in\Delta}\sum_{(i,j)\in\Pi}O^{\alpha}_{i,j}I^{\alpha}_{i,j}o^{\alpha}_{i,j}\Big)+C_2\sum_{i\in B}\mu_i+C_3\Big(\sum_{m,n\in V|T_{m,n}\neq inf}T_{m,n}\beta_{m,n}+\sum_{\alpha\in\Delta}\sum_{(i,j)\in\Psi}p'^{\alpha}_{i,j}\sigma^{\alpha}_{i,j}\Big)$$
$$+C_4\sum_{i,j\in V}\sum_{m,n\in V|U^{w}_{m,n}=0,or\,T_{m,n}\neq inf}\omega^{i,j}_{m,n}+a_{opt}\sum_{i,j\in V}\sum_{w\in W}v^{w}_{i,j}+a_{IP}\sum_{\alpha\in\Delta}\sum_{(s,d)\in R}\sum_{i,j\in V}\lambda^{s,d,\alpha}_{i,j}\Bigg] \quad (1)$$

As defined in objective function (1), this CF strategy solves the following problem: given the damage information in the carrier's multilayer network (e.g., with the optical and IP layers) and high-priority requests, maximize the number of satisfied high-priority requests (according to their priority) and minimize the combined cost of fiber link recovery and payments for acquiring *lightpath-SP-I* under the physical resource constraints of the carrier network. The six terms in objective function (1) (i.e., weighted with six coefficients C_1, C_2, C_3, C_4, a_{opt}, and a_{IP}, respectively) are detailed as follows: (i) maximize the satisfied customers' high-priority IP traffic demands ($\Sigma\Gamma_{sd}^{a} A_{sd}^{a} \alpha_{sd}^{a}$) and lightpath requests ($\Sigma O_{i,j}^{a} I_{i,j}^{a} o_{i,j}^{a}$), (ii) minimize the number of border nodes ($\Sigma\mu_i$) (to reduce management costs), (iii) minimize necessary (a) long-haul fiber links to restore ($\Sigma T_{m,n}\beta_{m,n}$) and (b) purchases of *lightpath-SP-I* ($\Sigma p'^{a}_{i,j} \sigma_{i,j}^{a}$) between the PNE sites from the counterpart carrier(s), (iv) minimize the hops of the multiple co-routed lightpaths ($\Sigma\omega_{m,n}^{i,j}$) (i.e., to offer the bandwidth higher than the capacity of one lightpath) in the optical network layer, (v) minimize the wavelength consumption in the optical network layer ($\Sigma v_{i,j}^{w}$), and (vi) minimize the total logical link bandwidth use in the upper IP layer ($\Sigma\lambda_{i,j}^{s,d,a}$). The coefficients C_1, C_2, C_3, C_4, a_{opt}, and a_{IP} separate the different terms into non-overlapping value ranges, e.g., $C_1 = 10^{12}$, $C_2 = 10^{10}$, $C_3 = 10^5$, $C_4 = 10^3$, $a_{opt} = 10$, and $a_{IP} = 1$.

In the third term, the cost of fiber link recovery and payments for acquiring *lightpath-SP-I* are combined and minimized, leading to a cost-first strategy.

(2) Leadtime-First strategy

$$\min\Bigg[-C_1\Big(\sum_{\alpha\in\Delta}\sum_{(s,d)\in R}\Gamma^{\alpha}_{s,d}A^{\alpha}_{s,d}\alpha^{\alpha}_{s,d}+\sum_{\alpha\in\Delta}\sum_{(i,j)\in\Pi}O^{\alpha}_{i,j}I^{\alpha}_{i,j}o^{\alpha}_{i,j}\Big)+C_2\sum_{i\in B}\mu_i+C_3\sum_{m,n\in V|T_{m,n}\neq inf}T_{m,n}\beta_{m,n}+C_5\sum_{\alpha\in\Delta}\sum_{(i,j)\in\Psi}p'^{\alpha}_{i,j}\sigma^{\alpha}_{i,j}$$
$$+C_4\sum_{i,j\in V}\sum_{m,n\in V|U^{w}_{m,n}=0,or\,T_{m,n}\neq inf}\omega^{i,j}_{m,n}+a_{opt}\sum_{i,j\in V}\sum_{w\in W}v^{w}_{i,j}+a_{IP}\sum_{\alpha\in\Delta}\sum_{(s,d)\in R}\sum_{i,j\in V}\lambda^{s,d,\alpha}_{i,j}\Bigg] \quad (2)$$

In addition to a CF strategy, some carriers may expect to reduce the recovery time as much as possible while temporarily ignoring the expense on acquiring *lightpath-SP-I*, namely, adopt an LF strategy. For the LF strategy, we define the objective function as shown in Eq. (2). Different from the CF strategy as shown in Eq. (1) we separately consider the fiber link recovery cost and expense on *lightpath-SP-I*. Namely, the third term minimizes the necessary long-haul fiber links to restore ($\Sigma T_{m,n}\ \beta_{m,n}$) with coefficient C_3, and the fourth term minimizes the purchases of *lightpath-SP-I* ($\Sigma p'_{i,j}{}^{a}\ \sigma_{i,j}{}^{a}$) between the PNE sites from the counterpart carrier(s) with an additional coefficient C_5. Moreover, we set the coefficients $C_3 = 10^8$ and $C_5 = 10^5$ for these two terms, respectively, to prioritize the third term first. Therefore, the LF strategy prioritizes minimizing the number of recovery tasks, finally reducing the overall recovery lead time which is prior to other concerns. Although the expense of *lightpath-SP-I* is of a lower priority compared to that in the CF strategy, through observation, we find that under different policy concerns LF outperforms CF significantly in terms of both recovery time and cost, which is discussed in Section 5.

4. DEMONSTRATION OF CARRIER COOPERATIVE RECOVERY

4.1 Experimental Scenario of Carrier Cooperative Recovery

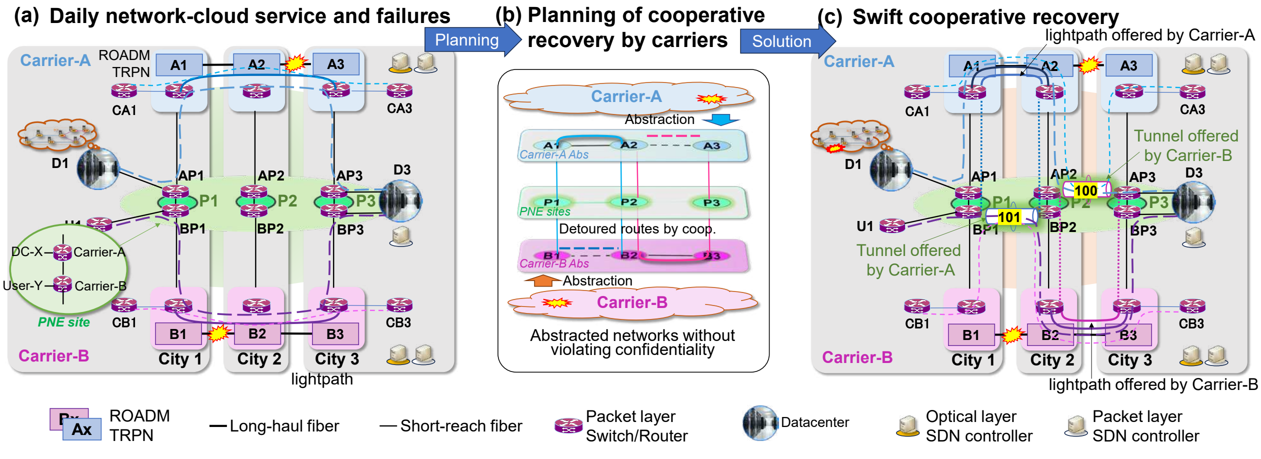

Figure 3: A Novel use cases of carrier cooperation/interconnection: (a) Daily network-cloud services; (b) Emergency planning of carrier cooperative recovery; and (c) Swift failure/disaster recovery of network-cloud service via carrier cooperation.

Figure 3(a) illustrates the experimental scenario of regular services which was created in our optical network testbed (in laboratory). With respect to the multi-layer optical transport networks, it was assumed that Carrier-A nodes A1 to A3 and Carrier-B nodes B1 to B3 were multilayer optical packet transport network nodes, including reconfigurable optical add/drop multiplexer (ROADM) (vendor-x), 100 Gbps transponders (vendor-y), and client-side packet switch/router (e.g., MPLS router). Client nodes (e.g., MPLS routers) CA1, CA3, and CB1, CB3, were connected with Carrier-A and -B networks, respectively. In each PNE site, e.g., P1, nodes AP1 and BP1 were the border nodes (e.g., MPLS routers, etc.) of Carrier-A and -B, respectively, which were directly interconnected as mentioned in Section 2.1. DCs D1, D2, and user node U1 were interconnected to Carrier-A and -B at MPLS layer via PNE sites P1 and P3, respectively. Carrier-A and -B established lightpaths between optical nodes <A1, A3> and <B1, B3>, respectively. And at the MPLS transport network layer, a 100 Gbps link was established between nodes <Ax, APx> and <Bx, BPx> for each lightpath. Communication services (e.g., MPLS label switched paths (LSPs)) were provided between carriers' client nodes <CA1, CA3>, and <CB1, CB3>, with Carrier-A and -B networks, respectively. Additionally, a DC-interconnection (DCI) link <D1, D3> and an user-DC connection <U1, D3> (e.g., LSPs) were established by Carrier-A and -B, respectively.

Through experiments, in addition to the use case of interconnection technology, e.g., for regular communication service in multi-carrier networks, we propose another novel use case from a resilience enhancement angle. For example, it was assumed that the fiber links <A2, A3> and <B1, B2> of Carrier-A and -B optical networks were broken in a disaster, and the recovery of fiber links and services were time consuming for individual carriers. The recovery time of each damaged fiber link in both carriers' networks was assumed "1" (in unit time, e.g., in hours, days or even weeks for a large-scale disaster [11]). The recovery cost was assumed "4" (unit cost, e.g., in local currency) for example, for each damaged fiber link, e.g., Carrier-A link <A2, A3> and Carrier-B link <B1, B2>, for illustration purpose. To quickly recover services, e.g., with the requests of <CA1, CA3>, <CB1, CB3>, <D1, D3> and <U1, D3>, Carrier-A and -B can conduct cooperative recovery. As shown in Fig. 3(b) a novel confidentiality-preserving carrier cooperative recovery planning was experimentally demonstrated to yield the solution for actual cooperative recovery, which is presented in Section 4.2. Based on this solution, as shown in Fig. 3(c) a swift and low-cost cooperative recovery in D-Plane was experimentally validated, which is presented in Section 4.3.

4.2 Carrier Cooperative Recovery Planning based on DLT without Violating Confidentiality

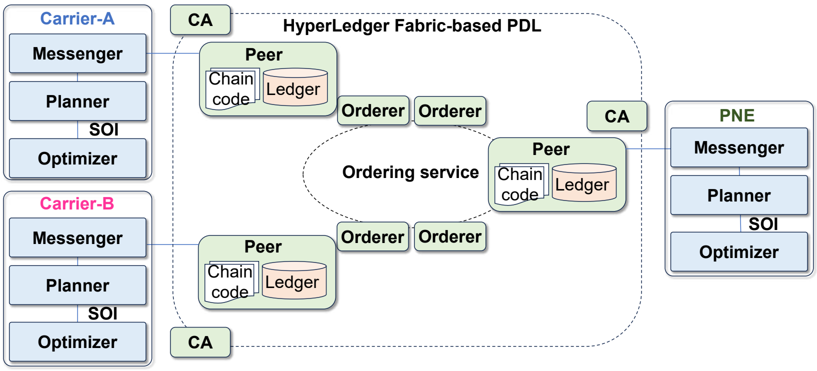

Figure 4: Experimental setup with a HyperLedger Fabric-based testbed.

Figure 4 shows the experimental setup which was built on a model-driven DLT-MCP prototype [9] for showcasing the cooperative recovery planning with a carrier cooperation scenario (described in Section 2), e.g., cooperation among Carrier-A, Carrier-B, and a third-party PNE. To support open and fair cooperation, public/private information sharing/exchange were enabled by a permissioned DLT subsystem, which was developed using a well-known DLT platform, HyperLedger Fabric (HLF) [12], including a set of CAs, Orderers (nodes for offering the ordering service and block generation, etc.), and Peers. Each stakeholder has three building blocks [9]: (1) Planner, a core building block that models and handles a sequence of actions (including information sharing/exchange and computational planning subtasks) in multiple phases for that stakeholder, as mentioned in Section 2 and illustrated in Fig. 2; (2) Optimizer, an optimization engine for executing a collection of recovery planning subtasks of the stakeholder, e.g., CSPT for carriers and PSMT for PNE [6]. The Planner calls the planning subtasks of the Optimizer via a Stakeholder Optimization Interface (SOI); (3) Messenger, a middleware element between the Planner of the stakeholder and the Peer of the DLT subsystem for public/private information sharing/exchange. Readers are referred to [9] for the details of the model-driven DLT-MCP. In the Optimizer, the recovery planning subtasks of Carrier-A and -B optical networks were modeled by CSPT with integer linear programming (ILP) taking account of the aforementioned policies and implementing a strategy of cooperation, e.g., CF and LF, as mentioned in Section 3. The recovery planning subtasks of PNE were modeled by PSMT with ILP for matching/balancing/scheduling of recovery tasks for carriers. Readers are referred to [6] for the details of the planning subtasks of carriers and PNE therein.

(1) Abstract view facilitating cooperation

Figure 3(b) illustrates a layered view of carrier optical networks in the affected area to facilitate carrier cooperative recovery planning. In the PNE layer, a 3-node and 2-link PNE reference topology was employed. In the underlying optical networks, Carrier-A and -B optical networks were abstracted to this PNE reference topology. In this abstract view, carriers' nodes A1 to A3 and B1 to B3 were interconnected via the PNE sites P1 to P3, respectively. It was assumed that in the cooperative recovery planning, each carrier used four wavelengths for cooperation, which was sufficient to accommodate the requests. A solution of carrier cooperative recovery planning is depicted in Fig. 3(b). Instead of waiting for the recovery of the fiber links <A2, A3> and <B1, B2> in carrier networks, Carrier-A and -B offered each other a *lightpath-SP-I* with the survived resources in the links <A1, A2> and <B2, B3>, respectively. By leveraging *lightpath-SP-I*, the traffic can be detoured via the survived resource in counterpart carrier networks. Therefore, all the services, e.g., <CA1, CA3>, <CB1, CB3>, <D1, D3> and <U1, D3> were able to recover quickly.

(2) Demonstration of information sharing without violating confidentiality

Figure 5(a) presents the nine-phase carrier cooperative recovery planning without violating confidentiality, as described in Section 2.2. Stakeholders performed their planning subtasks in designated phases and executed the public/private information sharing/exchange via the DLT-MCP. For example, abstract public information, such as the PNE reference topology (in Ph1), Price/Policy (in Ph2/3), PNE Segs recovery schedule and update (in Ph4/6), Matching/balancing and refinement of PNE Segs recovery (in Ph5/7) was shared among carriers and PNE. Note that, for public information sharing, the same information is received by all stakeholders via the *Public Chaincode* which is guaranteed by the DLT subsystem, achieving open and fair public information sharing [12]. Namely, no stakeholder is able to monopolize and manipulate the important public information in recovery (e.g., locations, topology, price/policy for applying the survived connection service of carriers) as all the public data in the distributed ledgers are traceable and auditable by all stakeholders. Meanwhile, private information, such as the initial/final request and confirmation of *lightpath-SP-I* (in Ph4/8/9) is exchanged only between the relevant carriers via the *Private Chaincode* of the underlying DLT subsystem to preserve confidentiality.

(3) Demonstration of policy

Figure 5(b) further demonstrates the policy of Carrier-A on resource sharing in phase Ph2 as an example. Carrier-A declared the same value “2” (as shown in the middle) for both the “regular price” and “dummy-price” (e.g., in unit price), indicating an immediately available *lightpath-SP-I* between <P1, P2>. Additionally, a policy of maximum availability of *lightpath-SP-I* was demonstrated with “max-lightpath-support-i”, showing an upper limit “4” on the number of this *lightpath-SP-I*. In contrast, as shown in Fig. 5(b) (bottom), Carrier-A declared a dummy price “54” higher than the price “4” to present the damaged services between <P1, P3>. Carrier-B declared similar policies which are not shown here for simplicity.

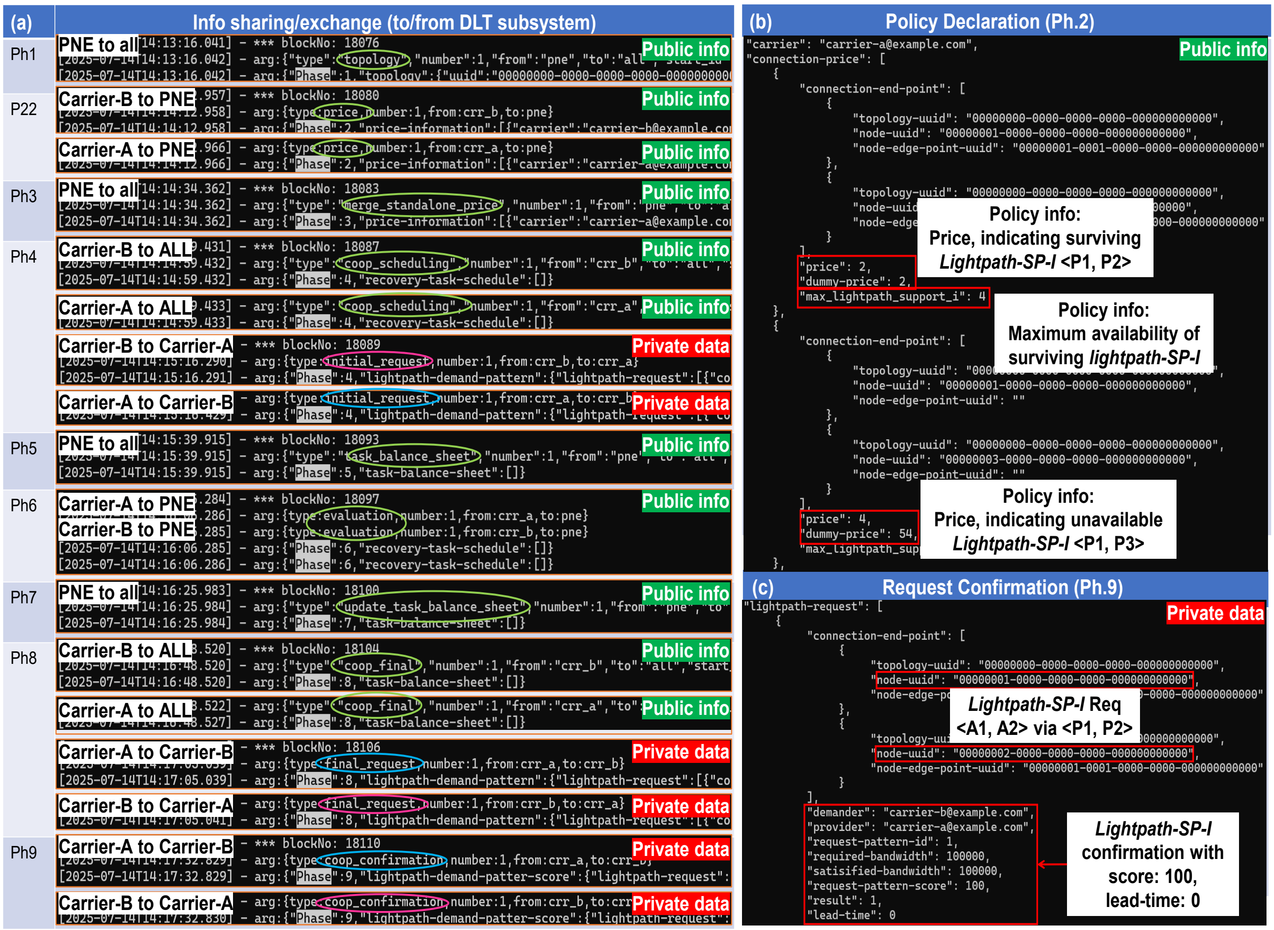

Figure 5: PNE-mediated carrier cooperative recovery planning: (a) Public/private information sharing/exchanging without violating confidentiality; (b) Policy declaration in phase Ph2; and (c) Final confirmation of the request for *lightpath-SP-I* in phase Ph9.

(4) Demonstration of the recovery planning (e.g., with CF strategy)

Upon receiving the PNE topology and price/policy information, Carrier-A and -B implemented the recovery planning (e.g., with a CF strategy in CSPT as an example) and progressively generated the final solution for swift recovery. For both Carrier-A and -B, the standalone recovery would require 1 unit time and 4 unit cost. On the country, applying a *lightpath-SP-I* of counterpart carriers, e.g., <P2, P3> of Carrier-B, and <P1, P2> of Carrier-A, respectively, could recover service immediately with a lower cost of 2, which was the optimal recovery plan. This is demonstrated in Fig. 5(c). For example, given a *lightpath-SP-I* request of <P1, P2> from Carrier-B in phase Ph8, Carrier-A confirmed the resource allocation within the reserved wavelengths W for cooperation (e.g., W = 4) in phase Ph9. Finally, Carrier-A notified Carrier-B with “satisfied-bandwidth” (e.g., 100,000 Mbps) and minimized “lead time” (i.e., the necessary recovery time for a desired service; where lead time “0” stands for an immediately available service).

In the experiments, PC servers (Xeon Gold 5115 2.4-GHz 20-core CPU, 128 GB memory) were employed for stakeholders. The running time of this nine-phase carrier cooperative recovery planning, including the time to solve the planning subtasks in Optimizers of Carrier-A, -B and PNE (with IBM CPLEX), was around 4.25 minutes, which was acceptable. Total generated blocks in the DLT subsystem were 34 blocks, each with 10 MB limit.

Experiments of other cases with different damage situations of Carrier-A and -B networks (including the heavy damage where all the fiber links were damaged and needed recovery task sharing and balancing among carriers) were conducted. These case studies successfully revealed the benefit of cooperation in terms of reduction of recovery time and cost. The results are omitted herein for simplicity.

4.3 Demonstration of Swift D-Plane Recovery with *Lightpath-SP-I*

In this section, we perform the experimental validation of carrier cooperative recovery of D-Plane, which is based on the solution yielded by carrier cooperative recovery planning with DLT-MCP as demonstrated in Section 4.2. In our current work, in the D-Plane of MPLS layer, the MPLS D-Plane functionality was emulated by using an open source software Open vSwitch (OVS) running on PC servers. In the C/M-Plane, OpenDaylight (ODL) [13] running on PC servers was used to emulate the software defined networking (SDN) controllers of Carrier-A and -B to establish LSPs. Environments and experiments which are close to the carrier-grade MPLS networks are envisioned as future work.

Figure 6 demonstrates the emergency post-disaster carrier cooperative recovery with the scenario as illustrated in Fig. 3(c). Namely, Carrier-A and -B offered each other a *lightpath-SP-I*, between <P1, P2> and <P2, P3> at the packet layer, respectively to achieve swift and low-cost recovery. During the cooperative recovery, Carrier-A and -B collaborated to configure lightpaths in the optical layer and establish the appropriate LSPs in the packet layer, which is demonstrated below.

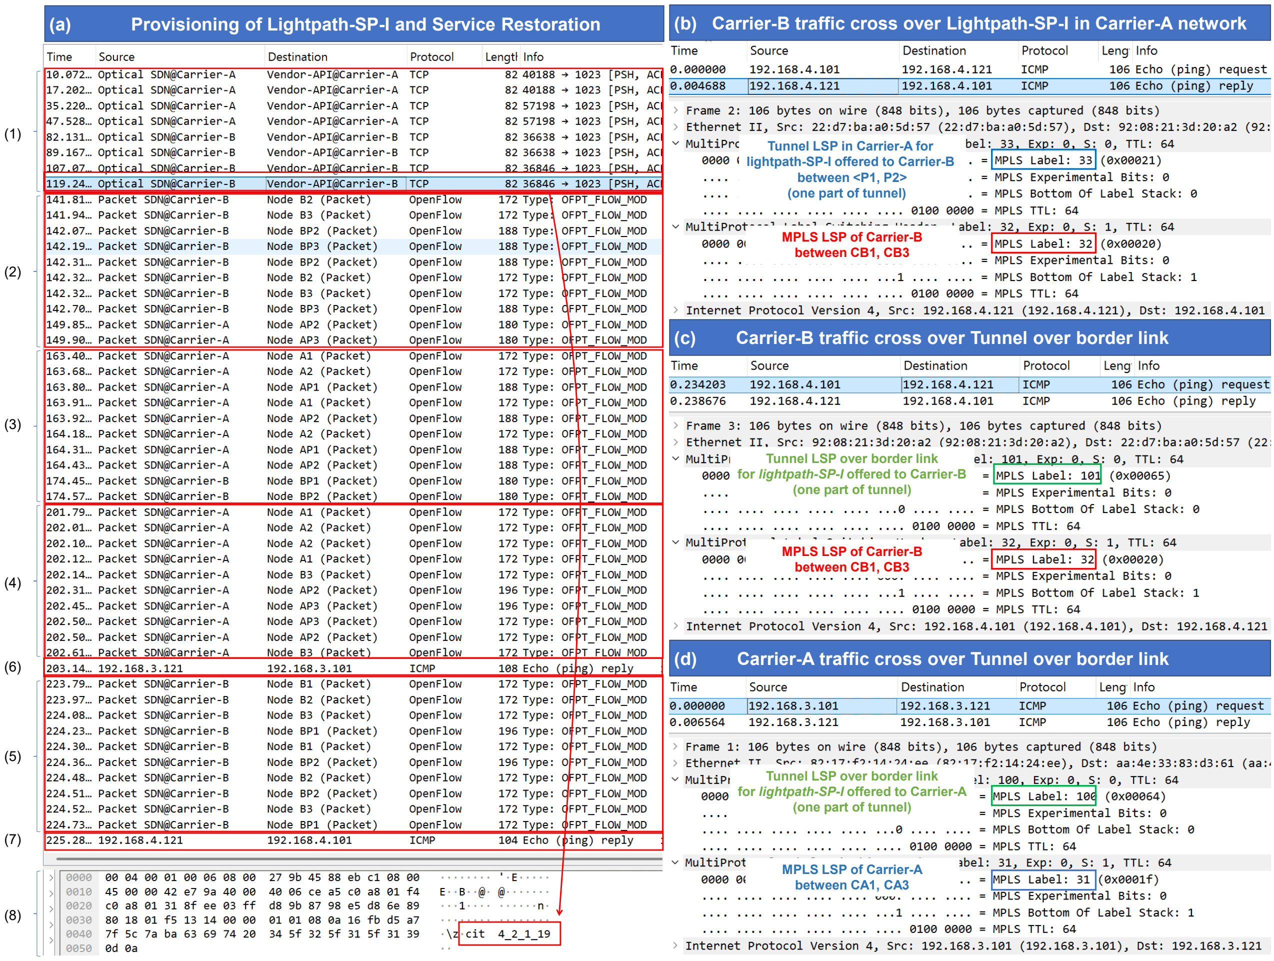

Figure 6: Captured messages demonstrating the provisioning of *lightpath-SP-I* and LSPs in the multi-layer networks of Carrier-A and -B: (a) Sequence of recovery exploiting *lightpath-SP-I*; (b) Captured packet of Carrier-B detoured traffic in Carrier-A network at the interface between <AP2, A2>; (c) Captured packet of Carrier-B transit traffic at the border link <AP2, BP2>; and (d) Captured packet of Carrier-A transit traffic at the border link <AP2, BP2>.

(1) Lightpath provisioning: Carrier-A and -B established a lightpath with surviving resource between <A1, A2> and <B2, B3>, respectively, connecting the corresponding client-side nodes at the MPLS layer. These lightpaths were configured with Carrier-A and -B optical-layer SDN controllers in our optical network testbed. The captured messages for the lightpath provisioning are shown in Fig. 6(a), block (1). A vendor-specific ROADM configuration command was executed and shown in block (8) for example.

(2) Tunnel creation (by Carrier-A and -B): At the MPLS layer, with the ODL SDN controllers in Carrier-A and -B networks, two LSPs were established as tunnels through nodes AP1, A1, A2, AP2, and BP2, B2, B3, BP3, respectively. Then, Carrier-A and -B configured their border nodes AP1, AP2, and BP1, BP2, respectively, to create two tunnels with MPLS Labels 100 and 101, as shown in Fig. 3(c). This process is demonstrated in block (2) and (3).

(3) Tunnel utilization (by Carrier-A and -B): At the MPLS layer, with these two tunnels, Carrier-A established LSPs to restore the communication service between <CA1, CA3> and <D1, D3>. Carrier-B established LSPs for the service between <CB1, CB3> and <U1, D3>. For simplicity, only the provisioning of LSPs between <CA1, CA3> and <CB1, CB3> are presented in Fig. 6(a), blocks (4) and (5). The LSP provisioning between <D1, D3> and <U1, D3> are not shown. The restored reachability between <CA1, CA3> and <CB1, CB3> was verified with ICMP PING, as shown in block (6) and (7), respectively.

Figures 6(b)-(d) further demonstrate the established tunnel through cooperation among Carrier-A and -B by capturing packets at the interface of the border node AP2 of Carrier-A. Figure 6(b) shows the tunnel LSP (with MPLS Label = 33) in the Carrier-A network using the *lightpath-SP-I*, which is highlighted in blue color. The nested LSP (with MPLS Label = 32) established by Carrier-B was embedded in the tunnel showing the detoured traffic of Carrier-B between <CB1, CB3>, which is highlighted in red color. Figure 6(c) shows the tunnel LSP (with MPLS Label =101, highlighted in green color) and the aforementioned nested LSP (with MPLS Label = 32) between the border nodes of Carrier-A and -B networks. Similarly, Fig. 6(d) shows the tunnel LSP (with MPLS Label = 100, highlighted in green color) and a nested LSP (with MPLS Label = 31 for the detoured traffic of Carrier-A between <CA1, CA3>) over the border link <AP2, BP2>.

Through the experiments, we proved that no confidential information (e.g., route, topology, etc.) of carriers is shared during the network control for D-Plane recovery. The cooperative recovery of D-Plane from the lightpath provisioning to the final reachability confirmation took around 3.5 min. Including the time of cooperative recovery planning (around 4.25 minutes) shown in Section 4.2, the communication service restoration was achieved in minutes, which was shorter than that of non-cooperative recovery (e.g., requiring 1 unit time, which cloud be in hours, days or even weeks).

Note that, to first verify the concept of carrier cooperative recovery, for simplicity, it was assumed that the corresponding emergency recovery of the C/M-Plane of stakeholders has been conducted, e.g., with wireless communication technology [14]. In D-Plane, the routes and labels of LSPs in the MPLS layer of carrier networks and border nodes (e.g., in PNE sites) were manually configured for multi-entity interconnection. Orchestration of these actions of individual carriers is desirable for automated configuration, which is an open problem for future research.

5. ILLUSTRATIVE NUMERICAL RESULTS OF POLICY/STRATEGY IN CARRIER COOPERATIVE RECOVERY

Numerical evaluations were conducted to observe the impact of three policy concerns (i.e., pricing, total number of wavelengths W employed in cooperative recovery, and maximum availability of *lightpath-SP-I*) on two strategies in cooperative recovery planning (i.e., the CF and LF strategies which are described in Section 3). Cooperation was conducted between Carrier-A and Carrier-B, and a PNE, as presented in Fig. 1. In the affected area, Carrier-A network consisted of 12 nodes, with one abstracted outside node 0; 2 border node candidates 1 and 2; the other 9 inside nodes (e.g., one node per city); and 17 bidirectional fiber links. Similarly, Carrier-B network consisted of 9 nodes and 14 fiber links. A subset of the Japan photonic network topology [10] was employed as the PNE reference topology with 6 PNE sites and 7 PNE Segs. We studied three *damage situations* in carrier networks. (i) Heavy damage: 66% fiber links were damaged in each carrier's network; (ii) Mixed damage: the damaged fiber links were 66% and 33% in Carrier-A (heavily damaged) and Carrier-B (lightly damaged) networks, respectively. (iii) Light damage: 33% fiber links were damaged in each carrier network. The damaged fiber links were randomly selected such that they had a strong correlation [6], i.e., if a fiber link fails in Carrier-A network, the co-located fiber link(s) (if exists) in Carrier-B fails simultaneously with a high probability, e.g., 0.8. For each damaged fiber link, the recovery cost was randomly selected in [1, 7]. 30 instances were randomly generated for each *damage situation*.

To simulate the emergency recovery of the highest priority traffic, we generated 13 high priority IP-over-WDM connection requests (carrying, on average, 100 Gbps among nodes) for all the carriers. For *lightpath-SP-I* and *-II*, the capacity of the lightpath was set to 100 Gbps. The coefficients in objective (1) for the CF strategy were set as $C_1 = 10^{12}$, $C_2 = 10^{10}$, $C_3 = 10^5$, $C_4 = 10^3$, $a_{\text{opt}} = 10$, and $a_{\text{IP}} = 1$ to separate the different terms into non-overlapping value ranges. The coefficients in objective (2) for the LF strategy were set as $C_3 = 10^8$ and $C_5 = 10^5$. The optimization instances (for CSPT, CSPT-Scheduling, and PSMT) were solved by IBM CPLEX, on a PC (Xeon Gold 6142 2.6-GHz CPU, 196 GB memory).

Three recovery schemes were executed for each instance. (i) *Standalone* (the benchmark): no cooperation was performed; (ii) *Survived coop*: only support of *lightpath-SP-I* with survived resource was performed in cooperation (i.e., in phase Ph4). (iii) *Advanced coop*: in addition to *Survived coop*, PNE Segs recovery task sharing/balancing/refinement were further performed (i.e., all the nine-phase cooperative recovery planning) (see Section 2 for details). In *Survived coop* and *Advanced coop*, the responses of CF and LF strategies under different policy concerns were observed, respectively. The average results are reported below.

(1) Impact of price and total number of wavelengths on the CF strategy

Figure 7 plots the recovery lead time with the CF strategy under different price values of services (e.g., the *lightpath-SP-I*). Figures 7(a)-(c) show that, for both *Survived coop* and *Advanced coop* when $W = 4$ and the price was higher, the benefit of cooperation was reduced in all three *damage situations.* This is because, in CF minimizing the expenses for purchasing *lightpath-SP-I* is of a high priority, a higher price suppresses the utilization of a *lightpath-SP-I*. Figures 7(d)-(f) show the same trend in the cases of more wavelengths ($W = 8$) in cooperation. When the wavelength was comparatively large, more traffic could be accommodated in a smaller number of fiber links (e.g., with more wavelengths). A lower requirement for fiber links resulted in less recovery tasks. Hence, the recovery time of standalone recovery was reduced, which was helpful to achieve efficient recovery. However, interestingly, in this situation it was difficult for cooperation (with CF) to employ *lightpath-SP-I*. This is because, to further reduce a fiber recovery task, more *lightpath-SP-I* should be purchased. Consequently, the expense of *lightpath-SP-I* increased, which was avoided by CF. This finding reveals that the CF strategy is sensitive to both price and W, hence, we need a more efficient strategy to achieve swift recovery.

(2) Impact of policy on the LF strategy

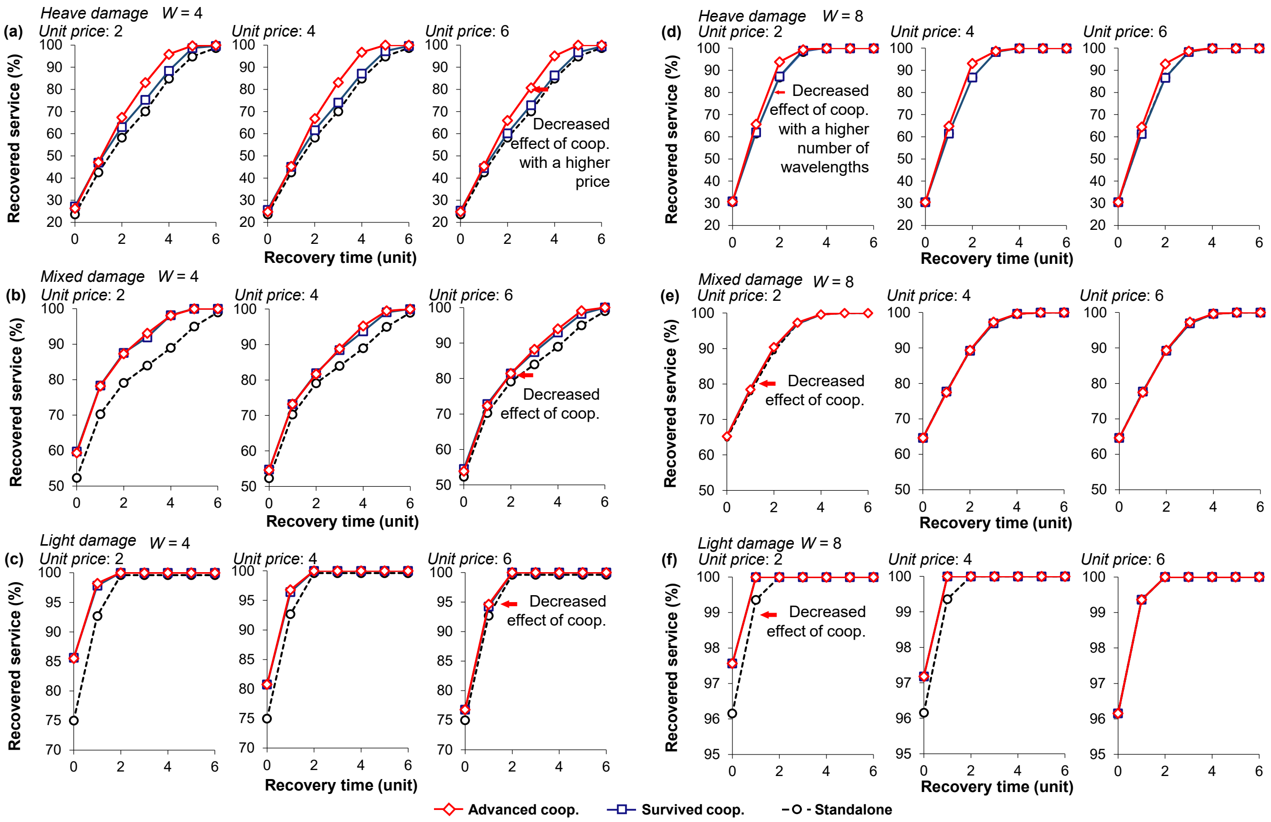

Figure 7: Impact of policy concerns on the CF strategy: (a) Heavy damage, $W = 4$; (b) Mixed damage, $W = 4$; (c) Light damage, $W = 4$; (d) Heavy damage, $W = 8$; (e) Mixed damage, $W = 8$; and (f) Light damage, $W = 8$.

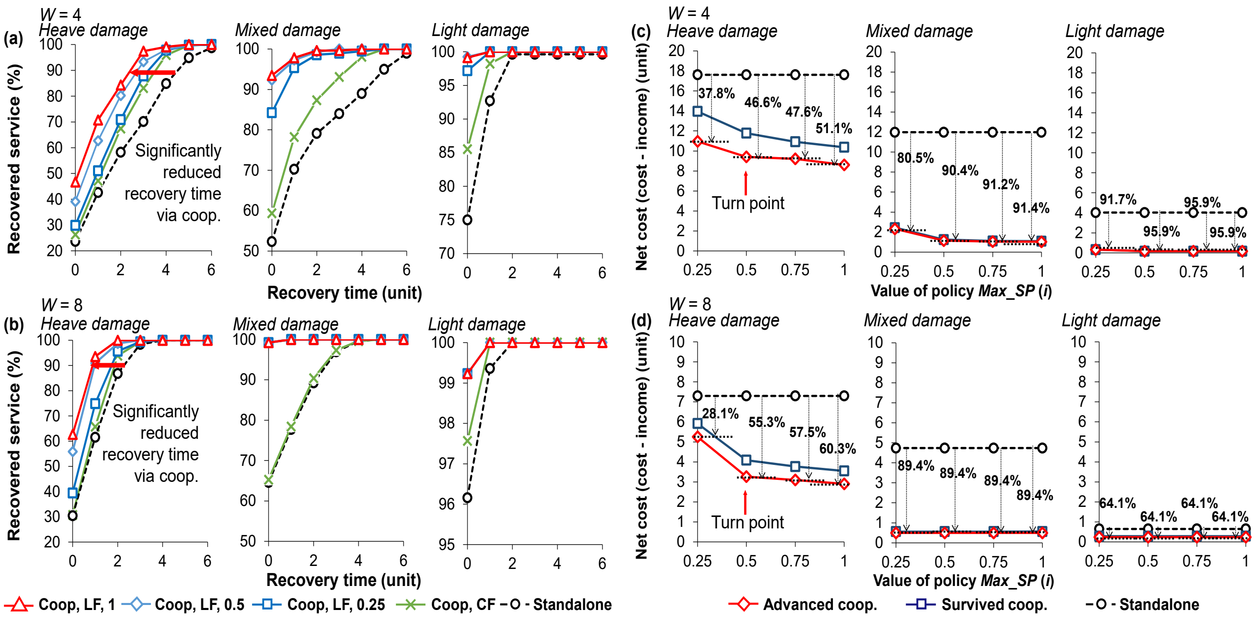

Figure 8: Impact of policies on the LF strategy: (a) Recovery time comparison of LF and CF, $W = 4$; (b) Recovery time comparison of LF and CF, $W = 8$; (c) Recovery cost of LF under the policy of maximum availability of *lightpath-SP-I*, $W = 4$; and (d) Recovery cost of LF, $W = 8$.

Figure 8 plots the performance of a novel strategy LF which is insensitive to price. Figures 8(a)-(b) show the recovery time comparison of LF and CF. For simplicity, only the final results of *Advanced coop* (coop, for short) are presented. Herein, for the observation of policy "maximum availability of *lightpath-SP-I*", the value of maximum availability of *lightpath-SP-I* is defined as a ratio of the allowed number of *lightpath-SP-I* to *W* (denoted as "value of policy *Max_SP* (*i*)"), where a value "1" stands for that all *W* number of supports *lightpath-SP-I* are applicable. We can find that LF outperformed CF significantly, especially when the maximum availability of *lightpath-SP-I* policy was relaxed. For different *damage situations*, by applying more *lightpath-SP-I*, LF (e.g., LF, 1) could recover more service with a shorter time. Especially, in light damage almost 99% of services could be restored immediately through cooperation. When we restrict the policy of maximum availability of *lightpath-SP-I*, e.g., by setting the value from 1 to 0.25, the performance of LF degraded gradually. Figures 8(c)-(d) show the net cost of recovery (total cost – income for offering *lightpath-SP-I* and *-II*). We can clearly see a significant cost down brought by LF. For example, when the value of the maximum availability of *lightpath-SP-I* was set e.g., 0.75 and 0.5, more than 50%, 80%, and 60% cost down could be achieved by LF in heavy, mixed and light damage situations, respectively. Note that, the purpose of introducing this policy is to prevent overuse or depletion of surviving resources between certain PNE sites. As shown in Figures 8(c)-(d), a turn point of this policy can be found at 0.5, revealing that with an appropriately configured policy, the overuse or depletion of surviving resources between certain PNE sites can be avoided while not scarifying the performance of cooperation.

Through experimental validation, we successfully addressed the questions raised in Section 1 and demonstrated that our proposed approach meets the stated requirements. (1) As presented in Section 4.2, we showcased, for the first time, a confidentiality-preserving carrier cooperative recovery planning mechanism using DLT-MCP. Only abstract public information, such as the PNE reference topology and the pricing/policy of connection services, was shared among stakeholders. Furthermore, with the extended experimental validation of both the C/M-Plane and D-Plane in Section 4.3, we confirmed that no confidential stakeholder information was disclosed during D-Plane recovery. (2) In Sections 4.2 and 4.3, we demonstrated that service restoration could be achieved within minutes through cooperative recovery based on an existing direct interconnection model, significantly faster than non-cooperative recovery, which could take hours, days, or even weeks. (3) In Section 3, we extended our previous work by investigating policy-based mechanisms to offer carriers more flexible control over resource sharing. Additionally, we proposed a novel Leadtime-First strategy for cooperative recovery planning. Numerical evaluations in Section 5 show that, by appropriately applying the LF strategy and specifying relevant policies, carriers can significantly reduce both recovery time and cost through cooperation. This represents another efficient and practical use case of state-of-the-art interconnection technology.

6. CONCLUSION

To promote network-cloud service and innovation, and to enrich massive enterprise users, interconnection and cooperation among carriers and data center providers become inevitable and crucial to form a large-scale common network-cloud ecosystem. Interconnection technology is a key in this evolution and attracting more attention. In this work, beyond the large-scale regular service we propose another use case of the next-generation interconnection technology: swift and low-cost recovery through carrier cooperation in case of large-scale resource crunch. Extending our previous research on modeling and platform design, with an exiting direct interconnection model in the industry, we showcased a confidentiality-preserving carrier cooperative recovery planning mechanism using DLT-MCP. Through experiments in D-Plane recovery, we successfully demonstrated a proof-of-concept of carrier cooperative recovery. Moreover, to offer carriers more options in resource control during cooperation, we investigated three policies and proposed a novel Leadtime-First strategy in carrier recovery planning to significantly improve the performance of carrier cooperation, enabling swift and low-cost recovery.

ACKNOWLEDGMENT

This work is supported in part by National Science Foundation: US-Japan JUNO3 project (2210384). We sincerely thank Mr. Weiping Ren and Mr. Taiga Suzuki for the help in experiments.

REFERENCES

1. IOWN Global Forum, “Open APN architecture PoC reference,” 2022, https://iowngf.org/technology/.
2. Telecom Infra Project, https://telecominfraproject.com/.
3. Optical Internetworking Forum, “400ZR, Management/CMIS,” 2022, https://www.oiforum.com/.
4. OpenROADM, “OpenROADM MSA specification ver 5.1,” 2022, http://openroadm.org/.
5. H. Nishizawa et al., “Fast WDM provisioning with minimal probing: the first field experiments for DC exchanges,” IEEE/OSA J. Opt. Commun. Netw., Vol. 16, No. 2, 233–242, 2024.
6. S. Xu, S. Sahoo, S. Ferdousi, N. Yoshikane, M. Shiraiwa, Y. Hirota, M. Tornatore, T. Tsuritani, Y. Awaji, and B. Mukherjee, “Scheme of carrier cooperation with coordinated scheduling for faster and lower-cost failure/disaster recovery,” IEEE/OSA J. Opt. Commun. Netw., Vol. 16, No.5, B45–B59 2024.
7. S. Sahoo, S. Xu, S. Ferdousi, Y. Hirota, M. Tornatore, Y. Awaji, and B. Mukherjee, “Service restoration in multi-entity network-cloud ecosystems: how to cooperate?,” IEEE Commun. Mag., Vol. 63, 129–135, 2025.
8. S. Sahoo, S. Xu, S. Ferdousi, Y. Hirota, M. Tornatore, Y. Awaji, and B. Mukherjee, “Post-disaster cloud-service restoration through datacenter-carrier cooperation,” IEEE/OSA J. Opt. Commun. Netw., Vol. 17, No. 8, 700–712, 2025.
9. S. Xu, S. Sahoo, S. Ferdousi, M. Shiraiwa, Y. Hirota, M. Tornatore, Y. Awaji, and B. Mukherjee, “Multi-entity cooperation platform facilitating network-cloud recovery,” Vol. 17, No. 9, D53–D72, 2025.
10. JPNM, https://www.ieice.org/cs/pn/eng/jpnm_en.html/.
11. “Technical Report on Telecommunications and Disaster Mitigation,” ITU-T FG-DR&NRR, Version 1.0 (06/2013).
12. HyperLedger Fabric, https://www.hyperledger.org/.
13. OpenDaylight, https://www.opendaylight.org/.
14. M. Shiraiwa, et al., “Experimental demonstration of disaggregated emergency optical system for quick disaster recovery,” IEEE/OSA J. Lightwave Technol., Vol. 36, No. 15, 3083–3096, 2018.