

Slice-Aware Service Restoration with Recovery Trucks for Optical Metro-Access Networks

Sifat Ferdousi¹, Massimo Tornatore^{1,2}, Sugang Xu³, Yoshinari Awaji³, and Biswanath Mukherjee¹

¹University of California, Davis, USA ²Politecnico di Milano, Italy ³National Institute of Information and Communications Technology, Japan

Email: {sferdousi, mtornatore, bmukherjee}@ucdavis.edu, {xsg, yossy}@nict.go.jp

Abstract—Next-generation optical metro-access networks are expected to support end-to-end virtual network slices for critical 5G services. However, disasters affecting physical infrastructures upon which network slices are mapped can cause significant disruption in these services. Operators can deploy recovery units or trucks to restore services based on slice requirements. In this study, we investigate the problem of slice-aware service restoration in metro-access networks with specialized recovery trucks to restore services after a disaster failure. We model the problem based on classical vehicle-routing problem to find optimal routes for recovery trucks to failure sites to provide temporary backup service until the network components are repaired. Our proposed slice-aware service-restoration approach is formulated as a mixed integer linear program with the objective to minimize penalty of service disruption across different network slices. We compare our slice-aware approach with a slice-unaware approach and show that our proposed approach can achieve significant reduction in service-disruption penalty.

Index Terms—Metro-Access Networks; Network Slices; Vehicle Routing Problem; Service-Disruption Penalty.

I. INTRODUCTION

With growth of bandwidth-hungry cloud services and increase of data services over mobile networks, traffic in optical metro-access networks (MAN) is skyrocketing. As of 2017, 75% of total metro traffic is terminated within the metro network [1] (as video, data, and web content are increasingly sourced from within metro networks, due to adoption of ‘edge cloud’). Next-generation converged MANs are expected to support heterogeneous access for customers connected via both wireless and fixed network technologies [2] as envisioned for 5G services which can be even more bandwidth-hungry, computation-intensive, and latency-sensitive.

A major feature of network evolution toward 5G is network *slicing*, which allows a physical network infrastructure to be divided (sliced) into multiple logical networks to create independent application-centric networks [3]. A network slice may support one or many end-to-end services, and it may consist of components belonging to access, transport, core, and edge networks as capacity, computing, storage, and VNFs [3]. To address diverse service requirements of different applications, slices should have specialized features (slice types). For example, enhanced mobile broadband (eMBB) slice requires large bandwidth to support high-data-rate services. Reliability, latency, and security are critical for ultra-reliable and low-latency communications (uRLLC) slice to provide mission-critical services. Massive machine-type communications (mMTC) slice is concentrated in access network requiring massive and ubiquitous connectivity support [3].

MANs (comprising heterogeneous access technologies and an aggregation optical metro network) provide physical infrastructures to support: radio access networks (base stations with fronthaul/backhaul), fixed access networks (xDSL/cable/FTTx), and edge/metro cloud (edge DCs). 5G network slices, each with its respective resource and connectivity requirements, can be configured using resources from such physical components as shown in Fig. 1 [4]. In case of a disaster (shown in Fig. 1 as red shaded area in the physical infrastructure layer), multiple physical network components, each providing resources for different slices, may fail (marked with red crosses) and cause service disruptions across slices (uRLLC slice is disrupted in this case). Depending on the slice type and priority, each slice can have different sensitivity to disruption; e.g., slices with higher reliability requirements will have higher penalty for service downtime or disruption.

In this study, we assume that the affected network slices can be dynamically re-provisioned through software-defined networking (SDN)-inspired management of optical metro and access networks. However, even with dynamic reconfiguration capability of network slices, post-disaster service restoration may not be possible considering locality of services (e.g., coverage area of a base station) and limited network redundancy. Traditionally, core networks are equipped with high redundancy, while MANs, especially in lower hierarchical layers (those closer to users) usually do not have large enough user base to justify expensive redundancy for disaster resiliency. Hence, recovery approaches should evolve for MANs, especially with the emergence of slicing paradigm.

In the immediate aftermath of a disaster, the utmost priority for network operators is to recover the network as soon as possible to minimize service downtime. In practice, to speed up post-disaster recovery, it may be possible (with proper equipment) to provide temporary backup service while repair work is going on. Hence, network operators are building disaster recovery systems [5]–[10] that restore connectivity, storage, and computing facilities using specialized recovery trucks as “temporary relief”. This can mitigate service disruption for end-users during repair process of failed components. Examples of such trucks include: portable cell sites or cell on trucks (which can connect to backhaul network feed or via satellite links, if backhaul is also failed), communication trucks providing satellite-based Ethernet/Wi-Fi services, mobile DC trucks, portable power generators, etc. Disaggregation-based portable recovery functional units called first-aid unit (FAU) [10] can also be developed for quick optical network recovery.




Fig. 1: 5G network slices running on a common underlying metro-access network.

In this work, we study post-disaster service restoration in MANs with recovery trucks considering slice requirements. Our aim is to schedule deployment of specialized recovery trucks at failure locations to minimize service disruption in slices based on respective priority (hence, *slice-aware*). While trucks may not guarantee pre-disaster QoS level, due to limited capacity, they can provide reduced level of services (i.e., degraded services) for the duration of recovery process. We model the problem of slice-aware service restoration with recovery trucks based on classical *vehicle routing problem*.

Vehicle routing problem is the generalization of the well-known traveling salesman problem, and it determines the optimal route used by a group of vehicles when serving a group of users [11], [12]. Since the problem involves finding routes for different recovery trucks destined for different types of nodes, our problem introduces an added dimension of vehicle heterogeneity. We formulate the problem as an optimization model to find the optimal routes for the recovery trucks to restore services in a disrupted MAN with the objective to minimize the penalty of service disruption across network slices. We show how taking into account slice requirements is fundamental to minimize the impact of service disruption.

II. SLICE-AWARE SERVICE RESTORATION

Scheduling post-disaster recovery is a complex task involving several dimensions such as physical constraints, limited resource availability, service requirements, etc. 5G services can impose additional dimensions of service requirements and priorities of different network slices.

A. Modeling

1) *Network Modeling*: In this study, we consider an optical MAN with heterogeneous network components supporting different network slices. In the example of Fig. 2(a), we consider 10 network slices, S1, S2, ..., S10 (shown as different shaded areas), with different resource requirements. The slices are supported by various components in the network, such as

metro nodes, base stations, edge DCs, etc. Slice constructions are based on different application requirements, e.g., slice S4 is configured based on computing and storage resources of DC1 and DC2 (to provide cloud services), radio access network resources of BS1 and BS2 (to provide cellular services), fixed access network resources of FA1 (to provide high-speed network connection), and optical backhaul resources (switching and bandwidth) of MN1, AN1-AN4, and associated links. Due to a disaster, failure of a network component (marked with red crosses) may affect multiple virtual network slices, e.g., failures of BS2 and DC2 disrupt services of slice S4 (these failures also disrupt services of other slices).

One research challenge is how to efficiently route and deploy recovery trucks in post-disaster MANs to minimize service-disruption *penalty* across network slices. For truck-routing decision, we can consider vehicular road maps (e.g., a road network) of a metro-access area, identifying routes that can be used for accessing the physical components (an example is shown in Fig. 2(b)). Each node in the MAN (Fig. 2(a)) corresponds to a node in the road network. Thus, each failed component (a failed node/link) in the MAN is a *failure site* (physical location of actual failure) or a failed node in the road network. Each failed link (MN1-MN2, MN5-AN5, and AN6-AN9) in the MAN is a failed node (L1, L2, and L3) in the road network (shown as circles in Fig. 2(b)) in addition to other failed nodes. Information collection on failure status of nodes/links (e.g., by employing surviving wired/wireless resources outside MAN [10]) is beyond the scope of this study. Links in road network are roads used for routing the recovery trucks. We consider roads that are accessible in a post-disaster scenario; monitoring drones or surviving distributed sensor networks may provide updates on road conditions after disaster (which is beyond the scope of this work). Also, there is a central depot (denoted by green square) for a fleet of trucks.

2) *Recovery Truck Modeling*: We assume that recovery trucks provide both repair (i.e., repair resources and repair

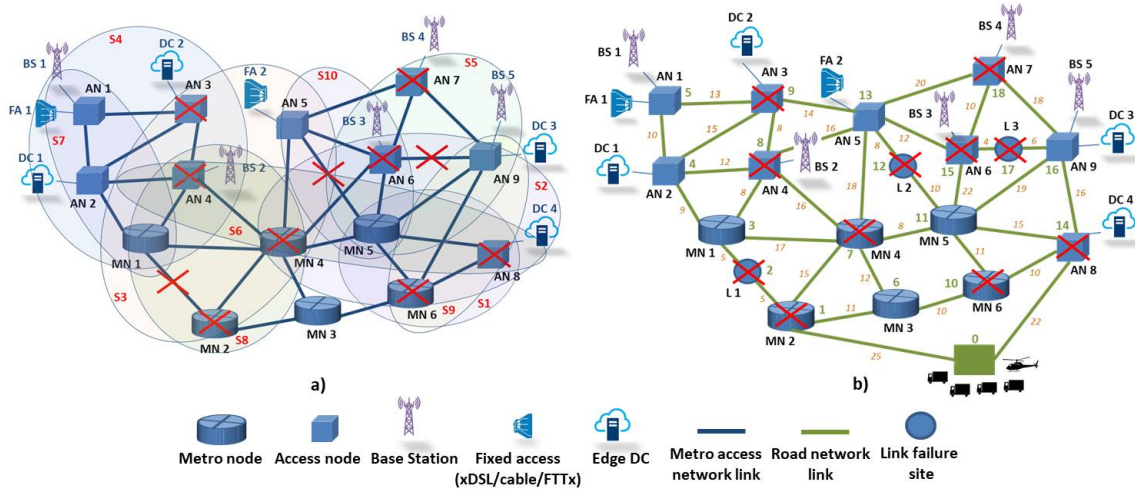


Fig. 2: Sample optical metro-access network. a) Different network slices mapped on physical network infrastructure. b) Physical road network (with distances) considered for truck routing including failure sites and the central depot for recovery trucks.

crew are available to execute the repair work) and temporary backup services of limited capacity. The goal is to route recovery trucks to locations of failed nodes to repair the components *and* provide backup services *while* the components are being repaired. After a disaster event, recovery trucks leave the depot destined for designated failure site(s) and then move to new destination(s) after the failed nodes are served by the trucks. Once a truck reaches a failure site, repair work begins, and the truck (i.e., the service unit) is deployed to provide “temporary relief” for the duration of repair work. In this study, we consider recovery of failed nodes (e.g., using FAUs for optical node recovery), links (e.g., quickly creating emergency detours for affected fibers with multi-vendor network interconnection approach [10] aided by FAUs), base stations, and edge DCs.

3) *Routing and Deployment Strategy*: Since number of trucks can be limited, recovery trucks should be routed such that high-priority slices are restored first (“slice-aware truck routing”). The routing and deployment strategy takes into account service priority of each slice based on its service requirements, criticality, etc. (e.g., uRLLC slices can have higher priority compared to other slices). In Fig. 2(a), suppose slice S4 has highest priority (i.e., highest service-disruption penalty); hence, failed nodes supporting S4 (BS2 and DC2) have precedence for service restoration. In particular, a failed network component of type r can be served only by a recovery truck of corresponding type r (e.g., portable cell sites for base stations). If we have only one mobile DC truck to restore DC functionalities in the network, it should be deployed at AN3 even though AN8 is closer to the depot and could restore DC4 functionalities earlier. If two mobile DC trucks are available, then trucks for DC2 and DC4 can be routed in parallel, and functionalities in the closest DC will be restored first. The problem is to decide how to distribute available trucks of each type to corresponding multiple failure sites.

Number of trucks required for full recovery of a failed node i of type r , w_i^r , depends on repair capacity of each truck k of type r . Since total number of trucks of type r , K^r , is limited, all w_i^r trucks may not be deployed at node i at the same time; based on “slice-aware” routing, the trucks may arrive at

node i at different times. Hence, total repair requirement of node i may not be satisfied at once. Every recovery truck k of type r also has limited service capability (e.g., antenna on a portable cell site has limited coverage) and hence, depending on the capacity of the failed components, trucks may provide full/partial services. We assume that a truck k of type r can restore functionalities at a failed node i as a ratio, c_i^r , of the node’s full capacity (e.g., a *cell on truck* may provide about 30% of a base station’s services). Repair time of a failed node is fixed based on repair efficiency of trucks of type r and hence, duration of the temporary service, q^r , is also fixed and is equal to repair time. After a truck leaves a failure site, failed node i can be fully or partially repaired depending on the repair capacity of truck k . In case of partial repair, additional trucks will arrive for further repair of node i , during which node i will provide partial services on its own until fully restored. Hence, these trucks serve as *temporary replacement* for the failed components in the network.

Duration in which a failed node i is fully non-operational is the travel time taken by the *first* recovery truck to arrive at node i and the deployment time of its service unit. Essentially, *effective downtime* of node i depends on duration and amount of service unavailability. When a recovery truck repairs a portion of node i , node i can become partially operational and only the non-operational portion of node i contributes to *effective downtime*. For example, node i requires w_i^r number of trucks for full recovery. If a truck k arrives at node i at time $A_i^{r,k}$ and repairs node i for q^r units of time, after the repair is complete at time $(A_i^{r,k} + q^r)$, $\frac{1}{w_i^r}$ part of node i becomes operational. In addition, truck k provides c_i^r amount of service (temporarily) at node i during repair time of q^r . In this way, node i can gradually become fully operational upon receiving services from w_i^r trucks. Hence, effective downtime of node i , Z_i^r (explained in details later), indicates the fraction of node i that is non-operational over time except $(c_i^r \cdot q^r)$, during which the recovery truck provides temporary service at node i .

A slice s is *fully* restored when functionalities of all failed nodes (indicating bandwidth, computing, and storage re-

sources) required to support s , are restored by recovery trucks. Note that slice s can be partially restored when supporting nodes are restored (at least partially). Functionalities of failed nodes are restored only after the trucks are deployed at the sites. Hence, for a slice s , total time required by trucks to be deployed at various components supporting s determines how soon slice s can be restored. Penalty of service disruption of slice s is the effective downtime of supporting nodes weighted by priority of slice s . Each node i can provide resources to a slice s based on its functional weight, $\beta_i^{s,r}$, for slice s . If node i supports multiple slices, then its resources are distributed as a ratio to the slices based on $\beta_i^{s,r}$. Suppose, in Fig. 2(a), for DC2 supporting slices S3 and S4, DC resources can be dedicated to S3 and S4 in 7:3 ratio. Hence, for a slice s , service-disruption penalty, P^s , is determined by cumulative effective downtime of each supporting node i , Z_i^r , weighted by node i 's functional weight for s , $\beta_i^{s,r}$, and priority of s , α^s . Our objective is to minimize service-disruption penalty across all slices.

B. Problem Formulation

The slice-aware service-restoration problem can be summarized as follows: given a physical road network topology, a set of failed nodes in the road network (corresponding to failed components in the MAN), a set of slices in the MAN (with corresponding physical network mapping and service requirements), and a set of recovery trucks, decide routing and deployment strategy for the trucks such that service-disruption penalty of all slices is minimized. We formulate the optimization problem as a mixed integer linear program (MILP), by adapting formulation for vehicle routing problem [11], [12], with the objective to minimize cumulative penalty of service disruption, P^s , across all network slices as follows:

• Given:

- $G(V, E)$: Physical road network topology for vehicle routing with set of physical nodes (corresponding to nodes and failed links in the MAN), V , and set of physical links (roads), E .
- τ : Number of node (and recovery truck) types.
- $\{0\}$: Central depot for recovery trucks.
- V^r : Set of physical nodes of type $r = 1, 2, \dots, \tau$. $V = \{0\} \cup_{r=1}^{\tau} V^r$.
- S : Set of virtual slices mapped on network G .
- V^s : Set of physical nodes $V^s \subseteq V$ supporting network slice $s \in S$.
- \bar{V} : Set of failed physical nodes, $\bar{V} \subseteq V$.
- \bar{V}^r : Set of failed physical nodes of type $r = 1, 2, \dots, \tau$. $\bar{V}^r = V^r \cap \bar{V}$, $r = 1, 2, \dots, \tau$.
- $\bar{V}^{s,r}$: Set of failed physical nodes of type $r = 1, 2, \dots, \tau$ supporting slice $s \in S$. $\bar{V}^{s,r} = \bar{V}^r \cap V^s$.
- F : Fleet of heterogeneous recovery trucks.
- K^r : Total number of recovery trucks of type $r = 1, 2, \dots, \tau$. $|F| = \sum_{r=1}^{\tau} K^r$.
- c_i^r : Amount of temporary service provided by trucks of type $r = 1, 2, \dots, \tau$ at failed node $i \in \bar{V}^r$ as a ratio of the node's full capacity.
- w_i^r : Units of recovery trucks of type $r = 1, 2, \dots, \tau$ required for full repair at failed node $i \in \bar{V}^r$.

- q^r : Duration of service provided by recovery trucks of type $r = 1, 2, \dots, \tau$.
- $t_{i,j}$: Travel time of recovery trucks, which is proportional to distance between nodes $i \in V$ and $j \in V$.
- $\beta_i^{s,r}$: Functional weight of node $i \in V^{s,r}$ in slice $s \in S$.
- α^s : Priority of slice $s \in S$.

• Variables:

- $X_{i,j}^{r,k} \in \{0, 1\}$: 1 if node $i \in V$ is served after node $j \in V$ by recovery truck $k = 1, 2, \dots, K^r$ of type $r = 1, 2, \dots, \tau$; 0 otherwise.
- $Y_i^{r,k} \in \{0, 1\}$: 1 if recovery truck $k = 1, 2, \dots, K^r$ is deployed at failed node $i \in \bar{V}^r$; 0 otherwise.
- $A_i^{r,k} \geq 0$: Arrival time of recovery truck $k = 1, 2, \dots, K^r$ of type $r = 1, 2, \dots, \tau$ at node $i \in V$.
- $Z_i^r \geq 0$: Effective service downtime of node $i \in \bar{V}^r$.
- $P^s \geq 0$: Penalty of service disruption in slice $s \in S$.

• Objective:

$$\min \sum_{s \in S} P^s \quad (1)$$

• Constraints:

$$\sum_{j \in V} X_{0,j}^{r,k} \leq 1, \quad r = 1, 2, \dots, \tau, k = 1, 2, \dots, K^r \quad (2)$$

$$\sum_{i \in V} X_{i,l}^{r,k} - \sum_{j \in V} X_{l,j}^{r,k} = 0, \quad (3)$$

$$\forall l \in V, r = 1, 2, \dots, \tau, k = 1, 2, \dots, K^r$$

$$\sum_{k=1}^{K^r} Y_i^{r,k} = w_i^r, \quad \forall i \in \bar{V}^r, r = 1, 2, \dots, \tau \quad (4)$$

$$Y_i^{r,k} \leq \sum_{j \in V} X_{j,i}^{r,k}, \quad (5)$$

$$\forall i \in \bar{V}^r, r = 1, 2, \dots, \tau, k = 1, 2, \dots, K^r$$

$$A_j^{r,k} \geq (A_i^{r,k} + t_{i,j}) \cdot X_{i,j}^{r,k} + q^r \cdot Y_i^{r,k}, \quad (6)$$

$$\forall i \in V, \forall j \in V, r = 1, 2, \dots, \tau, k = 1, 2, \dots, K^r$$

$$Z_i^r = \sum_{k=1}^{K^r} \left(\left(\frac{A_i^{r,k} + q^r}{w_i^r} \right) - (c_i^r \cdot q^r) \right) \cdot Y_i^{r,k}, \quad (7)$$

$$\forall i \in \bar{V}^r, r = 1, 2, \dots, \tau$$

$$P^s = \sum_{r=1}^{\tau} \sum_{i \in \bar{V}^{s,r}} (\alpha^s \cdot \beta_i^{s,r} \cdot Z_i^r), \quad \forall s \in S \quad (8)$$

The objective function in Eqn. (1) minimizes cumulative service-disruption penalty. P^s represents penalty of service disruption in slice s , and Eqn. (1) represents overall service disruption associated with all slices $s \in S$. Eqn. (2) states that each recovery truck leaves the depot and arrives at a node, if required. Eqn. (3) states that each recovery truck leaves a node and goes to the next node. In this way, flow conservation is satisfied for a truck at each node. Eqns. (2) and (3) enforce that, since repair truck leaves each node, upon completing repairs, it returns back to depot. Eqn. (4) ensures that total demand (required units of recovery trucks) at each node is fulfilled. Eqn. (5) indicates that a recovery truck can be deployed at a node if that truck is routed through that node. Eqn. (6) represents arrival time of a recovery truck at a failed node j . If the truck arrives at node j from node i , its arrival

time at node j is bounded by travel time from node i to j and service time at node i . It sets a minimum time for beginning of service at a node in a determined route. Effective service downtime of a failed node is represented in Eqn. (7) based on repair completion time, $(A_i^{r,k} + q^r)$, by each of the w_i^r trucks and service time, $(c_i^r \cdot q^r)$, provided by each truck. Eqn. (8) determines penalty of service disruption in a slice s , P^s , as total penalty incurred by each failed node i of type r in s , during its effective downtime, Z_i^r , based on its functional weight in s , $\beta_i^{s,r}$. It is then multiplied by priority of slice s , α^s . Note that Eqns. (6) and (7) are non-linear and can be linearized as follows (constant M is a large number) [11]:

$$A_j^{r,k} \geq (A_i^{r,k} + q^r \cdot Y_i^{r,k} + t_{i,j}) - M(1 - X_{i,j}^{r,k}), \quad (9)$$

$$\forall i \in V, \forall j \in V, r = 1, 2, \dots, \tau, k = 1, 2, \dots, K^r$$

$$A_j^{r,k} \leq M \sum_{i \in V} X_{i,j}^{r,k}, \quad (10)$$

$$\forall j \in V, r = 1, 2, \dots, \tau, k = 1, 2, \dots, K^r$$

$$Z_i^r = \sum_{k=1}^{K^r} \left[\left(\frac{A_i^{r,k} + q^r}{w_i^r} \right) - \left(c_i^r \cdot q^r \cdot Y_i^{r,k} \right) - M \cdot \left(1 - Y_i^{r,k} \right) \right], \quad \forall i \in \bar{V}^r, r = 1, 2, \dots, \tau \quad (11)$$

To evaluate our slice-aware service-restoration approach, we consider slice-unaware service restoration with the objective to minimize total travel time (proportional to distance) by recovery trucks as in classical vehicle routing problem. Hence, routing and deployment decision does not consider the impact of service-disruption penalty of slices. The objective function is shown below and other constraints remain same.

$$\min \sum_{r=1}^{\tau} \sum_{k=1}^{K^r} \sum_{i \in V} \sum_{j \in V} X_{i,j}^{r,k} \cdot t_{i,j} \quad (12)$$

III. ILLUSTRATIVE NUMERICAL EXAMPLES

To evaluate the performance of our proposed slice-aware service-restoration approach, we simulated the sample network with 10 slices as in Fig. 2(a). The corresponding road network for truck routing is shown with distances in Fig. 2(b). Node 0 represents the central depot for recovery trucks. We consider four types of failed components (i.e., $\tau = 4$): network nodes ($r = 1$), network links ($r = 2$), base stations ($r = 3$), and edge DCs ($r = 4$), and four corresponding types of recovery trucks. We assume service provider knows the following information: total number of available trucks of type r , K^r ; number of trucks required at each failed node i of type r , w_i^r ; amount of temporary service provided at each failed node i by trucks of type r , c_i^r ; and duration of service provided by trucks of type r , q^r . The simulation parameters are given in Table I.

We assume repairs for link failures to be relatively simpler (e.g., a fiber cut) compared to repairs of base stations, DCs, etc. Hence, values for w_i^r are generally 1, meaning that, for links repairs, one recovery truck is sufficient for full recovery. c_i^r is assigned a value $0 < n < 1$ as a ratio of full capacity of node i , depending on the serving capability of trucks of type r . q^r ranges from 15 to 50 units of time depending on node type and repair efficiency of trucks of type r . We assume negligible

truck deployment time. Travel times, $t_{i,j}$, are proportional to road distances, as shown in Fig. 2(b). Functional weight of node i for slice s , $\beta_i^{s,r}$, is assigned a value $0 < n < 1$ depending on number of slices supported by node i . Priority of each slice s , α^s , is assigned a value $0 < n < 1$ using a uniform distribution. *All values of evaluation metrics reported in our results are normalized to the lowest value.*

In Table II, we compare the cumulative service-disruption penalty between different approaches studied in this work. We show that, compared to slice-unaware approach, our slice-aware approach provides about 29% savings in penalty. We also study the impact of providing temporary relief through recovery trucks during the repair time for the slice-aware approach. We show that providing temporary relief along with repair can significantly reduce (about 38%) service-disruption penalty compared to providing repair *without* temporary relief (i.e., $c_i^r = 0$) since some services (at degraded level) can be restored early. Hence, we see that even the slice-unaware approach (which allows some services to be available during the repair time through temporary relief) performs somewhat better than the slice-aware approach *without* temporary relief, even though it is not optimized for penalty reduction.




Fig. 3: Comparison of service-disruption penalty per network slice between slice-aware and slice-unaware with temporary relief, and slice-aware without temporary relief approaches.

We compare penalty of service disruption of each network slice for slice-aware, slice-unaware, and slice-aware *without* temporary relief approaches, in Fig. 3. Our slice-aware approach achieves significant reduction (up to 60%) in service-disruption penalty compared to slice-unaware approach. Compared to other slices, cumulative penalty of slice S2 is higher since it has more failed network nodes (and also higher priority). Penalty reduction greater than 40% is achieved in slices S3, S4, and S5. Although some failure sites are relatively farther away from the depot, recovery trucks were deployed earlier at those sites for service restoration due to higher priority of the slices. Hence, services were restored earlier for these slices in the slice-aware approach, compared to slice-unaware approach, which minimizes only distance and does not consider priority. We also find that, with *both* repair and temporary relief provided by recovery trucks during repair work, slice-aware approach achieves significant reduction (up to 46%) in service-disruption penalty compared to providing *only* repair (as in general network recovery approaches). Note that, in case of slices S3, S4, and S5, slice-aware approach

TABLE I: Simulation parameters for service-restoration approaches in sample network. Number of node (and recovery truck) types, $\tau = 4$.

Component	r	K^r	w_i^r	q^r	c_i^r
Network node	1	4	$w_1^1 = 3, w_7^1 = 2, w_{10}^1 = 2$	40	$c_1^1 = 0.4, c_7^1 = 0.4, c_{10}^1 = 0.35$
Network link	2	3	$w_5^2 = 1, w_{12}^2 = 1, w_{17}^2 = 1$	15	$c_5^2 = 0.8, c_{12}^2 = 0.7, c_{17}^2 = 0.9$
Base station	3	2	$w_8^3 = 2, w_{15}^3 = 1, w_{18}^3 = 3$	25	$c_8^3 = 0.3, c_{15}^3 = 0.3, c_{18}^3 = 0.25$
Edge DC	4	2	$w_9^4 = 1, w_{14}^4 = 2$	50	$c_9^4 = 0.2, c_{14}^4 = 0.2$

TABLE II: Comparison of cumulative service-disruption penalty between slice-aware, slice-unaware, and slice-aware without temporary relief approaches.

	Slice-aware	Slice-unaware	Slice-aware w/o temp. relief
Cumulative penalty	29.04	41.09	47.88

without temporary relief performs better than slice-unaware approach even if it does not provide temporary relief during repair because slice-unaware approach suffers greatly for not considering high priority of the slices (as described earlier).

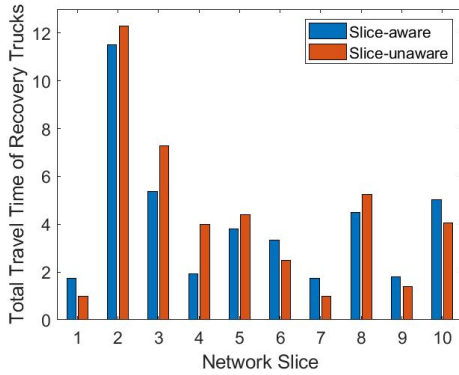


Fig. 4: Comparison of total travel time of recovery trucks between slice-aware and slice-unaware approaches.

Next, we compare total travel time taken by recovery trucks to restore each network slice between slice-aware and slice-unaware approaches in Fig. 4. Slice-unaware approach only minimizes travel time (proportional to distances) of recovery trucks to each failed node without considering slices. Hence, failed nodes closer to the depot are recovered before other nodes. Thus, different slices are affected differently as nodes supporting a slice can be at varying distances. We find that, for about 50% of slices, our approach required additional travel time of trucks to reduce penalty compared to slice-unaware approach. For the other 50% of slices, for which service-disruption penalty was greatly reduced, travel time is lower. Overall, slice-aware approach required about 40% less travel time compared to slice-unaware approach.

We also compare total service-restoration time of each network slice between slice-aware and slice-unaware approaches in Fig. 5. Total service-restoration time represents how early services in each slice were restored, and is derived from arrival times of recovery trucks, indicating the beginning of providing services. We show that our approach provides significant savings (up to 46%) in service-restoration time, compared to slice-unaware approach.

IV. CONCLUSION

We studied post-disaster service restoration in optical MANs supporting 5G slices using recovery trucks. Disasters affecting physical infrastructures can disrupt virtual network slices




Fig. 5: Comparison of total service-restoration time per network slice between slice-aware and slice-unaware approaches.

mapped on the physical network. We proposed a slice-aware service-restoration approach, using recovery trucks to provide temporary relief during repair work, with the objective to minimize service-disruption penalty across different slices. We modeled the problem as a vehicle routing problem to find optimal routes for recovery trucks to failure sites. We showed that, compared to a slice-unaware approach, which only minimizes travel time of recovery trucks, our approach can achieve significant reduction in service-disruption penalty and savings in service-restoration time. We also showed that providing temporary relief is crucial for fast service restoration.

ACKNOWLEDGEMENT

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

REFERENCES

- [1] "Strategic White Paper: Metro network traffic growth: An architecture impact study," tech. rep., Nokia Bell Labs (Alcatel-Lucent), 2013.
- [2] M. Ruffini, et al., "Access and metro network convergence for flexible end-to-end network design [invited]," IEEE/OSA JOCN, vol. 9, no. 6, pp. 524-535, June 2017.
- [3] K. Makhijani, et al., "Network slicing use cases: Network customization and differentiated services," IETF Internet-Draft, June 2017.
- [4] Jose Ordonez-Lucena, et al., "Network slicing for 5G with SDN/NFV: Concepts, architectures, and challenges," IEEE Comm. Mag., vol. 55, no. 5, pp. 80-87, May 2017.
- [5] T. Sakano, et al., "Bringing movable and deployable networks to disaster areas: development and field test of MDRU," IEEE Network, vol. 30, no. 1, pp. 86-91, Jan. 2016.
- [6] "Requirements on the improvement of network resilience and recovery with movable and deployable ICT resource units," tech. rep., ITU-T Focus Group, 2014.
- [7] "Issue Brief: Disaster response," tech. rep., AT&T, 2017.
- [8] K. T. Morrison, "Rapidly recovering from the catastrophic loss of a major telecommunications office," IEEE Comm. Mag., vol. 49, no. 1, pp. 28-35, Jan. 2011.
- [9] "Best Practices: Network continuity overview," tech. rep., AT&T, 2005.
- [10] M. Shiraiwa, et al., "Experimental demonstration of disaggregated emergency optical system for quick disaster recovery," IEEE/OSA JLT, vol. 36, no. 15, pp. 3083-3096, Aug. 2018.
- [11] J.F. Cordeau, et al., "Chapter 6 vehicle routing," Handbooks in Operations Research and Management Science, vol. 14, pp. 367-428, Dec. 2007.
- [12] T. Caric, et al., Vehicle Routing Problem, InTech, 2008.