Survivable Virtual Infrastructure Mapping With Dedicated Protection in Transport Software-Defined Networks [Invited]

Zilong Ye, Ankitkumar N. Patel, Philip N. Ji, and Chunming Qiao

***Abstract*—Efficiently mapping multiple virtual infrastructures (VIs) onto the same physical substrate with survivability is one of the fundamental challenges related to network virtualization in transport software-defined networks (T-SDNs). In this paper, we study the *survivable VI mapping* problem in T-SDNs with the objective of minimizing the VI request blocking probability. In particular, we address the subproblems of modulation selection and spectrum allocation in the process of provisioning optical channels to support virtual links, taking into consideration the optical layer constraints such as the transmission reach constraint and the spectral continuity constraint. We propose an auxiliary-graph-based algorithm, namely, parallel VI mapping (PAR), to offer dedicated protection against any single physical node or link failure. More specifically, the PAR algorithm can jointly optimize the assignments of mapping the primary and backup VIs by adopting the modified Suurballe algorithm to find the shortest pair of node-disjoint paths for each virtual link. Through extensive simulations, we demonstrate that the PAR algorithm can significantly reduce the VI request blocking probability and improve the traffic-carrying capacity of the networks, compared to the baseline *sequential VI mapping* approaches.**

***Index Terms*—Dedicated protection; Parallel VI mapping; Survivable virtual infrastructure mapping; Transport software-defined networks.**

I. INTRODUCTION

In transport software-defined networks (T-SDNs), the control plane and management intelligence are abstracted into a centralized controller while leaving behind a flexible and programmable optical transport data plane, such as bandwidth variable transponders and colorless–directionless–contentionless–gridless (CDCG) reconfigurable optical add/drop multiplexers (ROADMs) [1,2]. T-SDNs enable efficient and scalable network virtualization in optical transport by allowing multiple virtual infrastructures (VIs) to coexist on the same physical substrate and share physical node/link resources with guaranteed isolations [3–9]. From the physical substrate provider's standpoint, one of the fundamental challenges is how to make efficient use of the underlying resources through a procedure called VI mapping, which maps the virtual nodes and virtual links of a VI onto the physical substrate [10,11]. Furthermore, since natural disasters (e.g., earthquake) and intentional human attacks (e.g., fiber cuts or malicious attacks) may disrupt the services offered over such VIs, it is essential to provision survivability when mapping VIs over the physical substrate [12].

In this work, we investigate how to efficiently map the VIs with survivability against any single physical node or link failure over the T-SDNs, which is referred to as the *survivable VI mapping* problem. To provide fast recovery, we focus on studying the *protection* scheme (i.e., the backup VIs are precomputed and reserved in advance) rather than the *restoration* scheme. In particular, we take a first step to study the dedicated protection, and leave the shared protection as a future work. In dedicated protection, for each VI request, the primary and backup VIs need to be provisioned in a node-disjoint fashion, and the backup physical resources are not allowed to be shared. Dedicated protection is preferred because the service over the VIs can be immediately switched from its primary VI to the backup VI when a failure occurs, thus achieving a fast recovery.

Driven by the recent innovation of flexible grid optical transport [13], T-SDNs allow a dynamic and flexible allocation of the spacing and central frequency of the optical channels to gain a higher spectral efficiency. In such flexible optical transport, open challenges in provisioning survivable VI mapping are how to find primary and backup physical nodes for each virtual node, how to find node-disjoint primary and backup physical routes for each virtual link, and how to assign modulation formats and spectrum along these primary and backup routes while considering the spectral continuity constraint and the transmission reach constraint. The authors in [14–18] proposed several heuristic algorithms to solve the survivable VI mapping problem over fixed grid (WDM) optical transport; however, those solutions cannot be directly applied to solve the problem over flexible grid optical transport since

---

Manuscript received July 1, 2014; revised October 2, 2014; accepted October 9, 2014; published November 13, 2014 (Doc. ID 215088).

Z. Ye (e-mail: zilongye@buffalo.edu) and C. Qiao are with the Department of Computer Science and Engineering, University at Buffalo, The State University of New York, Buffalo, New York 14260, USA.

A. N. Patel and P. N. Ji are with NEC Laboratories of America, Princeton, New Jersey 08540, USA.

http://dx.doi.org/10.1364/JOCN.7.00A183

1943-0620/15/02A183-07$15.00/0

additional subproblems of modulation format selection and spectrum allocation are not addressed in these solutions.

In this paper, we study the survivable VI mapping problem over T-SDNs with the objective of minimizing the VI request blocking probability. In particular, we address the subproblems of modulation selection and spectrum allocation in the process of mapping virtual links. We propose an auxiliary-graph-based algorithm, called the parallel VI mapping (PAR) algorithm, which adopts the modified Suurballe's algorithm [19] to find a parallel mapping of the primary and backup VIs, and offers dedicated protection against any single physical node or link failure. In PAR, the assignments of mapping the primary and backup VIs are jointly optimized, thus achieving a cost-efficient way of provisioning survivability in the process of VI mapping. Simulation results show that the PAR algorithm can significantly reduce the VI request blocking probability and enhance the traffic-carrying capacity of the networks, compared to the traditional sequential VI mapping (SEQ) approaches. It is worth noting that the proposed PAR algorithm can be applied in the SDN centralized controller to automatically and intelligently manage the resource allocation in the physical substrate from a global view.

The organization of the paper is as follows. We first discuss the related works in Section II, and then introduce the problem models and formulations in Section III. In Section IV, we present the proposed auxiliary-graph based heuristic algorithm. In Section V, we compare the performance of the proposed algorithms. Section VI concludes the paper.

II. Related Works

Recently, the VI mapping problem has been extensively studied in both general IP networks [10,11] and optical networks [3–9]. In [10,11], the authors proposed mixed integer linear programming models and efficient heuristic algorithms to solve the VI mapping problem with the objective of maximizing the revenue over time. The authors in [3–9] studied the VI mapping problem in optical networks, taking into consideration the optical layer constraints such as the transmission reach constraint and the spectral continuity/conflict constraints. However, none of these works considered the survivability issues in provisioning VIs. Recent works [14–18] started to investigate how to provision VI mapping with survivability against physical node or link failures. In particular, the authors in [14] proposed an efficient shared protection scheme that relies on constructing an enhanced topology with survivability in the virtual layer, and the work in [17] allowed using migration to recover from a single physical node failure. However, these works have not addressed the subproblems of modulation format selection and spectrum allocation in the process of mapping virtual links. The most closely related work [18] proposed an efficient algorithm to provision survivable VI mapping in flexible grid optical networks. However, the solution offered resilience against only a single link failure, while the possible node failure was not considered.

In this paper, we study the survivable VI mapping problem in flexible grid based T-SDNs, in order to offer dedicated protection against any single physical node or link failure. We address the subproblems of modulation selection and spectrum allocation in the process of virtual link mapping, while considering the transmission reach constraint and the spectral continuity/conflict constraints.

III. Problem Models and Formulation

In this section, we first introduce the problem model that includes the physical substrate model, the VI request model, and the survivable VI mapping process. We then formulate the survivable VI mapping problem mathematically.

A. Physical Substrate Model

We are given a physical substrate $G_p(N_p, L_p)$, as shown in Fig. 1, where N_p is the set of physical nodes and L_p is the set of physical links. For a given physical node $n \in N_p$, it is associated with node resources of $C_n^k = \{c_n^i | i = [1, k]\}$, where k is the number of types of node resources (e.g., $k = 3$, including computing/switching/storage resources) and c_n^i is the capacity of node resource of type i at physical node n. The variable rate transponders can offer a set of modulation formats M and each modulation format $m \in M$ has a spectral efficiency of e_m b/s/Hz and can transmit over up to r_m km due to the physical layer impairment. For a given physical link $l \in L_p$, it is associated with a spectrum capacity of S_l GHz, and the spacing of a spectrum slot is H GHz.

B. VI Request Model

We are also given a set of VI requests $D_v(N_v, L_v, Z_x^k, B_y)$, as shown in Fig. 1, where N_v is the set of virtual nodes and L_v is the set of virtual links. For each virtual node $x \in N_v$, it is associated with a node resource demand of $Z_x^k = \{z_x^i | i = [1, k]\}$, where z_x^i represents the requested amount of node resources of type i from virtual node x. For each virtual link $y \in L_v$, it is associated with an amount of bandwidth demand, denoted by B_y.

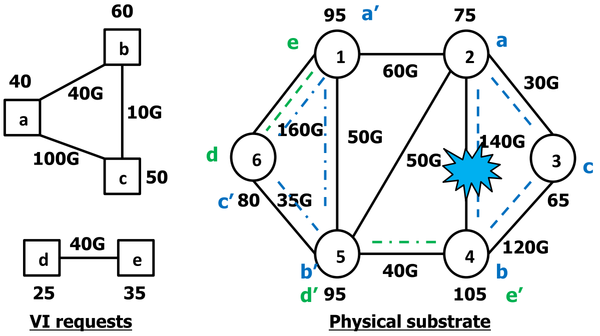

Fig. 1. Survivable VI mapping.

C. Survivable VI Mapping Problem

To offer dedicated protection against any single physical node or link failure, for each virtual node, we need to find distinct primary and backup physical nodes that have at least the requested node resources (e.g., physical node 1 and 2 for virtual node a, as shown in Fig. 1). In addition, for each virtual link, we need to determine the node-disjoint primary and backup physical routes [e.g., physical route (2,4) and (1,5) for virtual link (a, b), as shown in Fig. 1], select modulation formats based on the distance of these paths, and conduct spectrum allocation while considering the transmission reach constraint and the spectral continuity/conflict constraints. A VI request is considered blocked if any of its virtual nodes or virtual links cannot be mapped. The objective is to find a mapping of the primary and backup VIs over the physical substrate such that the VI request blocking probability is minimized.

D. Notations

The following parameters are given in the model.

- δ_n^k: the node resource capacity of type k at physical node $n \in N_p$.
- $\gamma_{i,j}$: the spectrum capacity of physical link $(i,j) \in L_p$.
- $\lim(m)$: the reachability of using modulation format m.
- H: the spacing of a spectrum slot.
- Φ: the total number of VI requests.
- $\sigma_{r,v}^k$: the node resource demand of type k from virtual node v in VI request r.
- $\lambda_{r,s,d}$: the bandwidth demand of virtual link (s, d) in VI request r.

To formulate the problem mathematically, we define the following decision variables, which specify the assignments of virtual node mapping, virtual link routing, modulation selection, and spectrum allocation.

- $\theta_{r,v}^n$: the primary mapping assignment of a given virtual node. $\theta_{r,v}^n = 1$, if virtual node v from VI request r is mapped to physical node n as a primary location; otherwise, $\theta_{r,v}^n = 0$.
- $\omega_{r,v}^n$: the backup mapping assignment of a given virtual node. $\omega_{r,v}^n = 1$, if virtual node v from VI request r is mapped to physical node n as a backup location; otherwise, $\omega_{r,v}^n = 0$.
- $\xi_{r,s,d}^{i,j}$: the primary route for a given virtual link. $\xi_{r,s,d}^{i,j} = 1$, if virtual link (s, d) from VI request r traverses physical link (i,j) in its primary route; otherwise, $\xi_{r,s,d}^{i,j} = 0$.
- $\zeta_{r,s,d}^{i,j}$: the backup route for a given virtual link. $\zeta_{r,s,d}^{i,j} = 1$, if virtual link (s, d) from VI request r traverses physical link (i,j) in its backup route; otherwise, $\zeta_{r,s,d}^{i,j} = 0$.
- $\mu_{r,s,d}$: the modulation format selected for the primary route for a given virtual link (s, d) in VI request r.
- $\eta_{r,s,d}$: the modulation format selected for the backup route for a given virtual link (s, d) in VI request r.
- $\alpha_{r,s,d}^w$: the spectrum slot used by the primary route of a given virtual link. $\alpha_{r,s,d}^w = 1$, if spectrum slot w is used by the primary route of virtual link (s, d) from VI request r; otherwise, $\alpha_{r,s,d}^w = 0$.
- $\beta_{r,s,d}^w$: the spectrum slot used by the backup physical route of a given virtual link; $\beta_{r,s,d}^w = 1$, if the spectrum slot w is used by the backup route of virtual link (s, d) from VI request r; otherwise, $\beta_{r,s,d}^w = 0$.
- ϕ_r^b: the VI request that is blocked. $\phi_r^b = 1$, if request r is blocked; otherwise, $\phi_r^b = 0$.

E. Constraints

In this subsection, we first list all the constraints and then explain them afterward.

1) Virtual Node Mapping:

- **One-to-one mapping constraint:**

$$0 \le \sum_n \theta_{r,v}^n \le 1, \quad \forall\, r, v, \tag{1}$$

$$0 \le \sum_n \omega_{r,v}^n \le 1, \quad \forall\, r, v. \tag{2}$$

- **Node-disjoint constraint:**

$$0 \le \sum_{v \in N_r} (\theta_{r,v}^n + \omega_{r,v}^n) \le 1, \quad \forall\, r, n. \tag{3}$$

- **Capacity constraint:**

$$\sum_{r,v} \sigma_{r,v}^k \cdot (\theta_{r,v}^n + \omega_{r,v}^n) \le \delta_n^k, \quad \forall\, n, k. \tag{4}$$

Equations (1) and (2) specify that each virtual node can be mapped only to at most one primary physical node and one backup physical node, i.e., virtual node splitting is not allowed. Equation (3) ensures that i) virtual nodes from the same VI request cannot be mapped to the same physical node, and ii) primary physical nodes and backup physical nodes from the same VI requests cannot be mapped to the same physical node. The capacity constraint for each type of node resource at any physical node is described in Eq. (4).

2) Virtual Link Mapping:

- **Flow conservation constraint:**

$$\sum_{i \in N_p} \xi_{r,s,d}^{i,n} - \sum_{j \in N_p} \xi_{r,s,d}^{n,j} = \theta_{r,d}^n - \theta_{r,s}^n, \quad \forall\, r, s, d, n, \tag{5}$$

$$\sum_{i \in N_p} \zeta_{r,s,d}^{i,n} - \sum_{j \in N_p} \zeta_{r,s,d}^{n,j} = \omega_{r,d}^n - \omega_{r,s}^n, \quad \forall\, r, s, d, n. \tag{6}$$

- **Node-disjoint constraint:**

$$\sum_{i \in N_p} (\xi_{r,s,d}^{i,n} + \xi_{r,s,d}^{n,i}) \cdot \sum_{i \in N_p} (\zeta_{r,s,d}^{i,n} + \zeta_{r,s,d}^{n,i}) = 0, \quad \forall\, r, s, d, n. \tag{7}$$

- **Bandwidth requirement constraint:**

$$\sum_k \alpha^w_{r,s,d} = \lceil \frac{\lambda_{r,s,d}}{H \cdot \mu_{r,s,d}} \rceil, \quad \forall\, r,s,d, \tag{8}$$

$$\sum_k \beta^w_{r,s,d} = \lceil \frac{\lambda_{r,s,d}}{H \cdot \eta_{r,s,d}} \rceil, \quad \forall\, r,s,d. \tag{9}$$

- **Transmission reach constraint:**

$$\sum_{i,j} \xi^{i,j}_{r,s,d} \le \lim(\mu_{r,s,d}), \quad \forall\, r,s,d, \tag{10}$$

$$\sum_{i,j} \zeta^{i,j}_{r,s,d} \le \lim(\eta_{r,s,d}), \quad \forall\, r,s,d. \tag{11}$$

- **Capacity constraint:**

$$\sum_{r,s,d,w} (\alpha^w_{r,s,d} + \beta^w_{r,s,d}) \le \gamma_{i,j}, \quad \forall\, i,j. \tag{12}$$

- **Spectrum conflict constraint:**

$$0 \le \sum_{r,s,d} (\xi^{i,j}_{r,s,d} \cdot \alpha^w_{r,s,d} + \zeta^{i,j}_{r,s,d} \cdot \beta^w_{r,s,d}) \le 1, \quad \forall\, i,j,w. \tag{13}$$

- **Spectrum consecutive constraint:**

$$0 \le \frac{1}{2} \cdot \sum_w |\alpha^{w+1}_{r,s,d} - \alpha^w_{r,s,d}| \le 1, \quad \forall\, r,s,d, \tag{14}$$

$$0 \le \frac{1}{2} \cdot \sum_w |\beta^{w+1}_{r,s,d} - \beta^w_{r,s,d}| \le 1, \quad \forall\, r,s,d. \tag{15}$$

Equations (5) and (6) ensure that a primary (backup) route can be found to connect the two primary (backup) physical nodes that host the two end virtual nodes of a given virtual link. Equation (7) ensures that, for a given VI, its primary VI is node-disjoint with its backup VI. Equations (8) and (9) specify that enough spectrum slots are reserved along the primary and backup routes for each virtual link. Equations (10) and (11) ensure that the selected modulation formats for the primary and backup routes satisfy the transmission reach constraint. Equation (12) represents that the spectrum capacity cannot be exceeded in each physical link. Equation (13) specifies that a given spectrum slot can be used only by at most one physical (primary or backup) route. Equations (14) and (15) specify that the provisioned optical channel has to use a set of consecutive spectrum slots. The spectral continuity constraint is inherently guaranteed in the definition of $\alpha^w_{r,s,d}$ and $\beta^w_{r,s,d}$.

F. Objective

The objective is to minimize the VI request blocking probability, which is shown as follows:

$$\min \frac{\sum_r \phi^b_r}{\Phi}. \tag{16}$$

IV. AUXILIARY-GRAPH-BASED HEURISTIC ALGORITHM

In this section, we first propose an auxiliary-graph-based heuristic algorithm, called the *parallel VI mapping* (PAR) algorithm, to efficiently solve the survivable VI mapping problem. For the comparison purpose, we then present a baseline algorithm, called the *sequential VI mapping* (SEQ) algorithm. The detailed workflow of the proposed algorithms are described as follows.

A. Parallel VI Mapping Algorithm

The proposed PAR algorithm is an auxiliary-graph-based algorithm that adopts the modified Suurballe's algorithm to find the shortest pair of node-disjoint paths as the primary and backup routes for each virtual link. Compared to the traditional way of sequentially mapping the primary and the backup VIs, the PAR algorithm can jointly optimize the assignments of primary and backup VI mapping.

In PAR, for a given virtual node, two candidate physical nodes are found that can provide sufficient node resources to host the given virtual node. For a given virtual link, a pair of node-disjoint candidate physical routes is found between these selected physical nodes by applying the modified Suurballe's algorithm on an auxiliary graph as shown in Fig. 2(a). The auxiliary graph is constructed by adding a pair of auxiliary source/destination nodes and a set of auxiliary links that connect the auxiliary source/destination with the candidate physical nodes of the virtual nodes that belong to the virtual link.

The above steps ensure that, for a given virtual node (link), its primary physical node (route) is node-disjoint

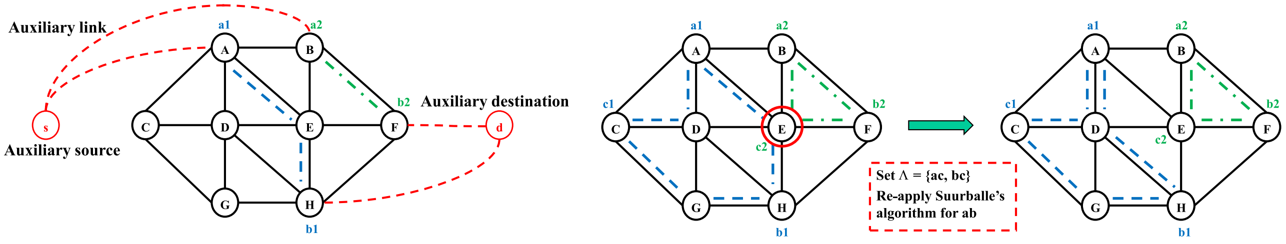

(a) (b)

Fig. 2. PAR algorithm. (a) Auxiliary graph and Suurballe's algorithm. (b) Parallel mapping of the primary and backup VIs.

with its backup physical node (route). Since the PAR algorithm aims to offer dedicated protection, we still need to make sure that, considering the whole virtual topology of a VI, none of the primary physical nodes (routes) is overlapping with any backup physical node (route). To do so, after finding a pair of node-disjoint physical routes for each virtual link, a set Λ is found that includes the maximum number of virtual links whose node-disjoint routes are partitioned into two separate sets, and the routes in one of the sets are also node-disjoint with the routes of the other set. As shown in Fig. 2(b), virtual links ac and bc are in set Λ, while virtual link ab is not included because a_1b_1 has overlapped with a_2c_2 and b_2c_2 at physical node E. After set Λ is found, the functionality (primary or backup) of the candidate physical nodes/routes for the virtual links in set Λ can be determined according to the criteria that the shortest route among the node-disjoint paths is considered the primary virtual link, and the other route is the backup virtual link. As mentioned in Step 5 [also as shown in Fig. 2(b)], for the virtual links that are not included in set Λ, we will reapply the modified Suurballe's algorithm to find the pair of node-disjoint paths. In addition to the assignments of primary and backup routing of virtual links, the subproblems of modulation selection and spectrum allocation are addressed in Step 6 and Step 7, respectively, while the transmission reach constraint and the spectral continuity constraint will be considered. The detailed steps of the PAR algorithm are shown in Algorithm 1.

Algorithm 1 Parallel VI Mapping Algorithm

Step 1: For a given virtual node, find two physical nodes that have not been assigned with any of the virtual nodes from the same VI, and have the first and second highest available resources.
Step 2: For a given virtual link, construct an auxiliary graph as follows. Establish an auxiliary source and an auxiliary destination, and establish auxiliary links between the auxiliary source/destination and the candidate physical nodes that the virtual link's two end nodes are mapped to.
Step 3: Use the modified Suurballe's algorithm to find a pair of node-disjoint paths on the auxiliary graph.
Step 4: Find set Λ and determine the functionality for the physical nodes/links of the virtual links in set Λ.
Step 5: For virtual links that are not included in set Λ, generate an auxiliary graph and apply the modified Suurballe's algorithm as follows. First, remove all the physical nodes and physical links along the backup routes on which existing virtual links are mapped, and find the shortest route between the auxiliary source and destination, which is considered a primary route. Next, restore the removed physical nodes and physical links of the backup routes and remove the physical nodes and physical links of the already mapped primary routes, and find the second shortest route as the backup route.
Step 6: According to the path distance, select a modulation format that meets the transmission reach constraint.
Step 7: According to each virtual link's bandwidth demand and the spectral efficiency of the selected modulation format, determine the number of spectrum slots needed, and allocate the spectrum at the lowest available spectrum slots, considering the spectral continuity/conflict constraint.
Step 8: If any of the above steps fails, block the VI demand and check the next one.

B. Sequential VI Mapping Algorithm

For the purpose of comparison, we implement a SEQ algorithm, which maps the primary and backup VIs in sequential order. First, an existing VI mapping algorithm is applied to map the primary VI. Second, we remove the physical nodes and physical links along the primary route and reapply the VI mapping algorithm again to map the backup VI. In these procedures, the node-load-balance-first algorithm and the link-load-balance-first algorithm from [7] are used, which are denoted by SEQ-N and SEQ-L, respectively.

V. Performance Evaluation

In this section, we will evaluate the performance of the proposed PAR and SEQ algorithms in the 24-node U.S. mesh network. We will focus on the VI request blocking probability as the main performance metrics. Table I shows the default values of the parameters used in the simulation, and Table II shows the spectral efficiency and transmission reach of each type of modulation format used in the simulation.

In the simulation, each physical node consists of three types of node resources (e.g., computing, switching, and storage resources) with capacity of 1500 units per type of resource. Each physical link has a spectrum capacity of 4 THz with a spacing of 12.5 GHz per spectrum slot.

TABLE I
Default Setting of the Simulation Parameters

Parameter	Value
No. of physical nodes	24
No. of node resources in a physical node	3
Capacity of each type of node resources	1500
No. of modulations offered	3
No. of physical links	43
Spectrum capacity of each physical link	4 THz
Spacing of each spectrum slot	12.5 GHz
No. of VI demands	200
No. of virtual nodes in a VI demand	2–5
Max. node demand of a virtual node	30
Bandwidth demand of a virtual link	1000 Gb/s

TABLE II
Efficiency and Reachability of Each Modulation

Modulation	Spectral Efficiency (b/s/Hz)	Reachability (km)
PM-BPSK	1.6	8000
PM-QPSK	3.2	3000
PM-16QAM	6.4	1000

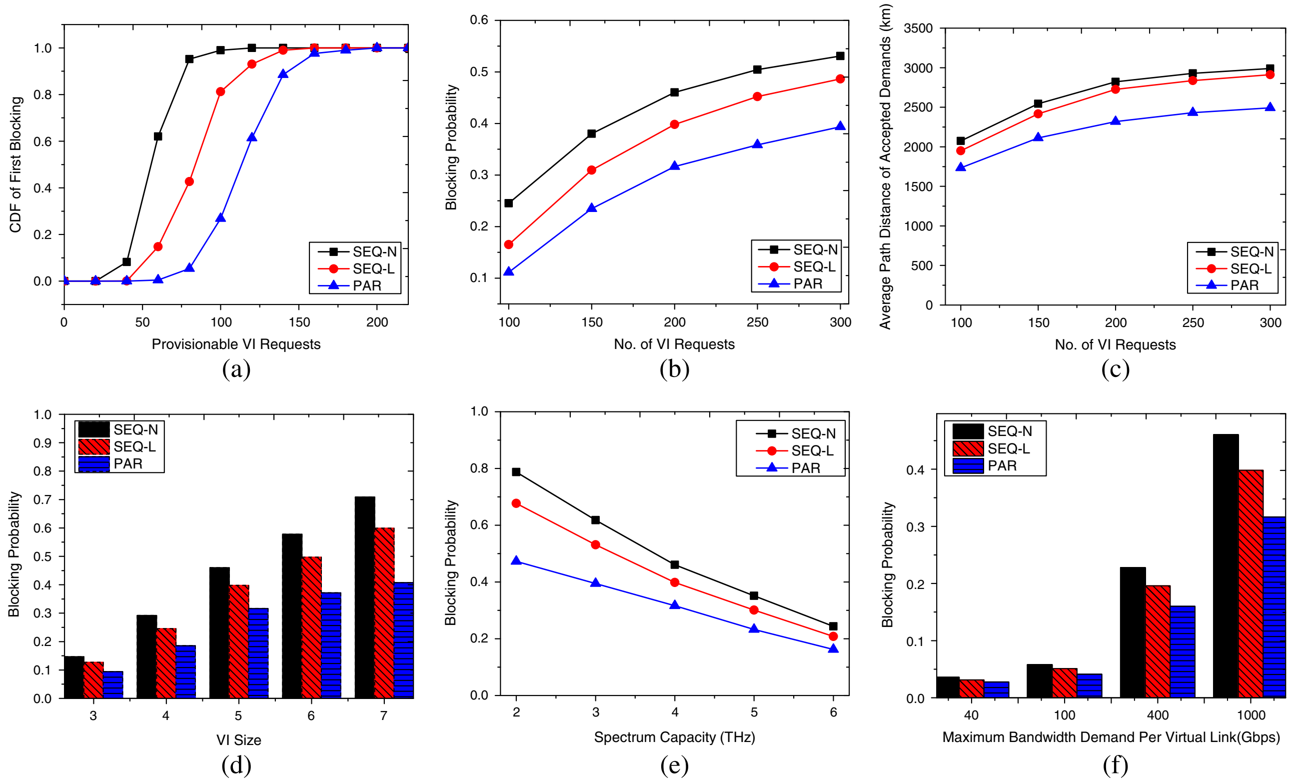

Fig. 3. Simulation results. (a) CDF of first blocking. (b) Blocking probability versus number of VI requests. (c) Average path distance of accepted requests versus number of VI requests. (d) Blocking probability versus VI size. (e) Blocking probability versus spectrum capacity of each physical link. (f) Blocking probability versus maximum bandwidth demand of each virtual link.

The variable rate transponder can offer a set of modulation formats including PM-BPSK, PM-QPSK, and PM-16QAM, with spectral efficiencies of {1.6, 3.2, 6.4} b/s/Hz and transmission reaches of {8000, 3000, 1000} km, respectively. The VI requests arrive according to the incremental traffic pattern in which a virtual topology consists of two to five interconnected virtual nodes. Each virtual node demands three types of node resources, each of which is less than 30 units with uniform distribution. Each virtual link has a bandwidth demand among {10, 40, 100, 400, 1000} Gb/s with uniform distribution. For a given parameter setting, we randomly generate 1000 test cases, and the statistics in this section are the average results.

Figure 3(a) shows the cumulative distribution function (CDF) of the first blocking over the number of provisionable requests (i.e., the number of requests that can be accepted before blocking occurs for the first time). We can see that SEQ-L performs better than SEQ-N because link load balance is given first priority in SEQ-L. We can also observe that PAR performs the best. Compared to SEQ-N and SEQ-L, the average performance (when 50% of first blocking happens) improvement is 113% and 33%, respectively. The reason is that PAR can jointly optimize the assignments of mapping the primary and backup VIs so that physical node and link resources can be used in a more efficient way to accept more future VI requests.

Figure 3(b) shows the blocking probability (the number of blocking VI requests divided by the total number of VI requests) with different numbers of VI requests. We can see that the blocking probability of PAR is the lowest, and PAR can provision 25% and 13% more demands than SEQ-N and SEQ-L, respectively. The reason behind Fig. 3(b) is that PAR adopts a parallel mapping of primary and backup VIs, and uses the modified Suurballe's algorithm to find the shortest pair of node-disjoint paths for each virtual link, which reduces the average channel distance and, thus, potentially increases the opportunity to accept future demands. To better understand that, we analyze the average path distance of the accepted requests with different numbers of VI requests in Fig. 3(c). We can see that the average path distance of the accepted requests of PAR is smaller than that of SEQ-N and SEQ-L.

Next, we study how the VI size will affect the VI request blocking probability. As shown in Fig. 3(d), as the VI size increases, the blocking probability for all the proposed heuristic algorithms increases. However, the blocking performance improvement of PAR is further enhanced, compared to SEQ-N and SEQ-L. We can infer that PAR is more efficient at accommodating VI requests with more complex virtual topology, compared to SEQ-N and SEQ-L.

We further evaluate how the spectrum capacity of a physical link will affect the VI request blocking probability in Fig. 3(e). We can see that as the spectrum capacity increases, the blocking probability decreases, as we expected. The relative performance difference between PAR and SEQ-N and SEQ-L is obvious, which demonstrates that

PAR outperforms SEQ-N and SEQ-L. This is because the PAR algorithm adopts a parallel mapping of primary and backup VIs.

Figure 3(f) shows the blocking probability with different maximum bandwidth demands (per virtual link). We can observe that as the maximum bandwidth demand increases, the blocking probability increases significantly. We can also observe that the relative performance difference between PAR and SEQ-N and SEQ-L increases. The increases of maximum bandwidth demand will emphasize the advantage of using PAR because the channel distance provisioned by PAR is smaller, thus leading to a lower amount of bandwidth product, compared to SEQ-N and SEQ-L.

VI. CONCLUSION

Survivability is an essential and practical concern in the process of VI mapping. In this work, we have studied the survivable VI mapping problem in T-SDNs with the objective of minimizing the VI request blocking probability. We have addressed the subproblems of modulation selection and spectrum allocation in the process of virtual link mapping. In particular, we have proposed an auxiliary-graph-based algorithm, called the PAR algorithm, to offer dedicated protection against any single physical node or link failure. The PAR algorithm can jointly optimize the assignment of mapping the primary and backup VIs by adopting the modified Suurballe's algorithm to find the shortest pair of node-disjoint paths for each virtual link. Through extensive simulations, we have demonstrated that the PAR algorithm can significantly improve the blocking performance of VI requests, thus enhancing the traffic-carrying capacity of networks. The proposed PAR algorithm can be implemented as an optimization tool run in the SDN centralized controller to efficiently manage the resource allocation in the optical transport data plane.

ACKNOWLEDGMENT

This work is supported in part by DTRA project HDTRA1-09-0032.

REFERENCES

[1] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "OpenFlow: Enabling innovation in campus networks," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, pp. 69–74, 2008.

[2] P. N. Ji, "Software defined optical network," in *Proc. ICOCN*, 2012, pp. 31–34.

[3] A. Pages, J. Perello, S. Spadoro, J. A. Garcia-Espin, J. F. Riera, and S. Figuerola, "Optimal allocation of virtual optical networks for the future Internet," in *Proc. ONDM*, 2012.

[4] S. Peng, R. Nejabati, and D. Simeonidou, "Impairment-aware optical network virtualization in single-line-rate and mixed-line-rate WDM networks," *J. Opt. Commun. Netw.*, vol. 5, no. 4, pp. 283–293, 2013.

[5] C. Gao and J. Jue, "Virtual optical network embedding considering mixed transparent and translucent virtual links," in *Proc. OFC*, 2013, paper JW2A.63.

[6] A. N. Patel, P. N. Ji, Y. Huang, and T. Wang, "Distance-adaptive virtual infrastructure embedding in software-defined optical networks," in *Proc. OECC*, 2013, paper WQ3-2.

[7] Z. Ye, A. Patel, P. Ji, C. Qiao, and T. Wang, "Virtual infrastructure embedding over software-defined flex-grid optical networks," in *Proc. IEEE GLOBECOM*, 2013, pp. 1204–1209.

[8] L. Gong and Z. Zhu, "Virtual optical network embedding (VONE) over elastic optical networks," *J. Lightwave Technol.*, vol. 32, no. 3, pp. 450–460, 2014.

[9] X. Wang, Q. Zhang, I. Kim, P. Palacharla, and M. Sekiya, "Flexible virtual network provisioning over distance-adaptive flex-grid optical networks," in *Proc. OFC*, 2014, paper W3H.3.

[10] M. Yu, Y. Yi, J. Rexford, and M. Chiang, "Rethinking virtual network embedding: Substrate support for path splitting and migration," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, pp. 17–29, 2008.

[11] N. M. Mosharaf, M. R. Rahman, and R. Boutaba, "Virtual network embedding with coordinated node and link mapping," in *Proc. INFOCOM*, 2009, pp. 783–791.

[12] Z. Ye, A. N. Patel, P. N. Ji, and C. Qiao, "Survivable virtual infrastructure mapping over transport software-defined networks (T-SDN)," in *Proc. OFC*, 2014, paper M3H.1.

[13] M. Jinno, H. Takara, B. Kozicki, Y. Tsukishima, Y. Sone, and S. Matsuoka, "Spectrum-efficient and scalable elastic optical path network: Architecture, benefits, and enabling technologies," *IEEE Commun. Mag.*, vol. 47, no. 11, pp. 66–73, 2009.

[14] H. Yu, V. Anand, C. Qiao, and G. Sun, "Enhancing virtual infrastructure to survive facility node failures," in *Proc. OFC*, 2011, paper JWA4.

[15] B. Guo, C. Qiao, J. Wang, H. Yu, Y. Zuo, J. Li, Z. Chen, and Y. He, "Survivable virtual network design and embedding to survive a facility node failure," *J. Lightwave Technol.*, vol. 32, no. 3, pp. 483–493, 2014.

[16] H. Yu, V. Anand, C. Qiao, and G. Sun, "Cost efficient design of survivable virtual infrastructure to recover from facility node failures," in *Proc. GLOBECOM*, 2011.

[17] H. Yu, V. Anand, C. Qiao, and D. Hao, "Migration based protection for virtual infrastructure survivability for link failure," in *Proc. OFC*, 2011, paper OTuR2.

[18] W. Xie, J. P. Jue, Q. Zhang, X. Wang, Q. She, P. Palacharla, and M. Sekiya, "Survivable virtual optical network mapping in flexible-grid optical networks," in *Proc. ICNC*, 2014, pp. 221–225.

[19] J. Suurballe and R. E. Tarjan, "A quick method for finding shortest pairs of disjoint paths," *Networks*, vol. 14, pp. 325–336, 1984.