

# Joint Placement and Allocation of VNF Nodes With Budget and Capacity Constraints

Gamal Sallam<sup>✉</sup>, Student Member, IEEE, and Bo Ji<sup>✉</sup>, Senior Member, IEEE, Member, ACM

**Abstract**—With the advent of Network Function Virtualization (NFV), network services that traditionally run on proprietary dedicated hardware can now be realized using Virtual Network Functions (VNFs) that are hosted on general-purpose commodity hardware. This new network paradigm offers a great flexibility to Internet service providers (ISPs) for efficiently operating their networks (collecting network statistics, enforcing management policies, etc.). However, introducing NFV requires an investment to deploy VNFs at certain network nodes (called VNF-nodes), which has to account for practical constraints such as the deployment budget and the VNF-node capacity. To that end, it is important to design a joint VNF-nodes placement and capacity allocation algorithm that can maximize the total amount of network flows that are fully processed by the VNF-nodes while respecting such practical constraints. In contrast to most prior work that often neglects either the budget constraint or the capacity constraint, we explicitly consider both of them. We prove that accounting for these constraints introduces several new challenges. Specifically, we prove that the studied problem is not only NP-hard but also non-submodular. To address these challenges, we introduce a novel relaxation method such that the objective function of the relaxed placement subproblem becomes submodular. Leveraging this useful submodular property, we propose two algorithms that achieve an approximation ratio of  $\frac{1}{2}(1 - 1/e)$  and  $\frac{1}{3}(1 - 1/e)$  for the original non-relaxed problem, respectively. Finally, we corroborate the effectiveness of the proposed algorithms through extensive evaluations using trace-driven simulations.

**Index Terms**—Network Function Virtualization (NFV), resource allocation, submodular optimization.

## I. INTRODUCTION

THE advent of Network Function Virtualization (NFV) has made it easier for Internet service providers (ISPs) to employ various types of functionalities in their networks. NFV requires the replacement of network functions that traditionally run on proprietary dedicated hardware with software modules, called Virtual Network Functions (VNFs), which run on general-purpose commodity hardware [1]. A wide variety of network functions (firewalls, intrusion detection systems, WAN optimizers, etc.) can be applied to flows passing through network nodes that host VNFs (called VNF-nodes). A flow must be fully processed at one or multiple VNF-nodes so that

Manuscript received November 25, 2019; revised August 3, 2020 and December 9, 2020; accepted January 31, 2021; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor J. Llorca. This work was supported in part by the NSF under Grant CNS-1651947. A preliminary version of this work has been presented at IEEE INFOCOM 2019. (Corresponding author: Bo Ji.)

Gamal Sallam is with the Department of Computer and Information Sciences, Temple University, Philadelphia, PA 19122 USA (e-mail: sallam9876@yahoo.com).

Bo Ji is with the Department of Computer Science, Virginia Tech, Blacksburg, VA 24061 USA (e-mail: boji@vt.edu).

Digital Object Identifier 10.1109/TNET.2021.3058378

the potential benefits introduced by NFV can be harnessed [2]. The new network paradigm enabled by NFV not only offers a great flexibility of introducing new network functions, but it also reduces capital and operational expenditure. Therefore, major ISPs have already started the process of transforming their technologies and operations to support NFV [3].

However, such moves often take place in multiple stages due to the budget limit; in each stage, only a subset of nodes can be selected for deploying/placing VNFs. Moreover, VNF instances typically have a limited capacity, which is shared for processing multiple passing flows. Therefore, given a deployment budget and capacity limit, it is of critical importance to choose a best subset of nodes to become VNF-nodes and to determine the optimal capacity allocation so as to maximize the amount of network traffic passing through them.

In contrast to most prior work that often neglects either the budget constraint (e.g., [4], [5]) or the capacity constraint (e.g., [2]), we explicitly consider both constraints and *formulate a joint problem of VNF-nodes placement and capacity allocation (VPCA)*. The VPCA problem has two main components: VNF-node placement and VNF-node capacity allocation, which are tightly coupled with each other. That is, deciding where to place the VNF-nodes depends on how the capacity of the VNF-nodes will be allocated; determining an optimal capacity allocation apparently depends on where the VNF-nodes are placed. The challenge posed by this problem is two-fold. First, *the placement and capacity allocation subproblems are both NP-hard*. Second, even if we assume that there is an oracle that can optimally solve the capacity allocation subproblem, *the placement subproblem is non-submodular* (a property that generally leads to efficient solutions for similar problems). This is in stark contrast to the previously studied problem without the capacity constraint [2], which has been shown to be submodular and can be approximately solved using efficient greedy algorithms.

To that end, we propose a new framework that integrates a *decomposition approach* with a *novel relaxation method*, enabling us to design efficient algorithms with constant approximation ratios for the studied VPCA problem. We summarize our key contributions as follows.

- First, we formulate the VPCA problem with budget and capacity constraints as an Integer Linear Program (ILP). Then, we provide an in-depth discussion about the new challenges introduced by the budget and capacity constraints. Specifically, we show that the placement and capacity allocation subproblems are both NP-hard. Further, we show that the objective function of the placement subproblem is not submodular.

- To address these challenges, we relax the requirement of fully processed flows and allow partially processed flows to be counted. This simple relaxation enables us to prove that the relaxed placement subproblem is submodular based on a *novel network flow reformulation* of the relaxed capacity allocation subproblem. Leveraging this useful submodular property, we design two efficient algorithms that achieve an approximation ratio of  $\frac{1}{2}(1 - 1/e)$  and  $\frac{1}{3}(1 - 1/e)$  for the original (non-relaxed) VPCA problem, respectively. To the best of our knowledge, *this is the first work that exploits this type of relaxation method to solve a non-submodular optimization problem with provable performance guarantees*.
- Finally, we evaluate the performance of the proposed algorithms using trace-driven simulations. The simulation results show that the proposed algorithms perform very closely to the optimal solution obtained from an ILP solver and better than another algorithm that iteratively selects the node with the highest volume of traffic traversing it [6].

The rest of the paper is organized as follows. First, we position our work compared to related work in Section II. Next, we describe the system model and problem formulation in Section III and discuss the challenges of the VPCA problem in Section IV. Then, we introduce the VPCA relaxation and reformulation in Section V and the proposed algorithms in Section VI. Then, we present the numerical results in Section VII. Finally, we conclude the paper and discuss future work in Section VIII.

## II. RELATED WORK

There has been a large body of work that studies the placement problem in different contexts such as NFV, SDN, and edge cloud computing. In NFV, a placement is usually considered at a scale of VNF instances, i.e., where and how many instances of each network function should be placed and allocated [4], [7]–[9]. Different objectives are considered in each of them. The problem of how to meet the demand from all of the flows with a minimum cost (e.g., in terms of the number of instantiated instances) is considered in [4], [10]. An extension of such work considers the setting where each flow must traverse a chain of network functions, instead of just one function, along a given route [11]. A similar problem is also considered in [7], [12] but for an online setting where flows arrive and leave in an online fashion. The work in [9] addresses the joint problem of VNF service chain placement and routing with the objective of minimizing total communication and computation resource cost. A dynamic version of this problem is considered in [8], where the goal is to ensure network stability while minimizing resource cost. Also, in [13], the authors consider the placement of a minimum number of nodes to achieve the original maximum flow under a given service function chaining constraint.

Note that the process of transitioning to NFV typically has two phases: the planning phase and the production phase. The planning phase is concerned about deciding where to introduce NFV to efficiently utilize the limited budget. Since this phase takes place before the actual deployment of the VNF-nodes, one can use historical traces to project flow demands across

the network. In the production phase, one can optimize flow admission and routing schemes decisions to efficiently utilize the available resources (e.g., [12]). In this work, we are mainly focused on the planning phase with budget and capacity constraints and assume that flow routes are fixed.

There are several studies that are highly relevant to our work. In [2], the authors consider the selection of a set of nodes to upgrade to SDN. By assuming that the SDN-nodes have an infinite capacity, they show that the problem is submodular. However, we show that with a capacity constraint (which is typically the case in practice), the problem becomes non-submodular. In addition, due to the capacity constraint, only a subset of flows traversing a VNF-node can be processed. Therefore, capacity allocation becomes a crucial component of the joint problem we consider. Similar to [14], which considers joint placement and scheduling in the edge clouds, we consider a new architecture with stateless network functions (see, e.g., [15]), which enables a fractional flow assignment over fixed routes. Similar to our problem, the problem considered in [14] is not submodular in general. While they can prove submodularity and provide an approximation algorithm for a special case, they develop a heuristic algorithm only for the general case. In contrast, we develop approximation algorithms for the general problem we consider. Specifically, we propose a new framework that enables us to address the challenge of non-submodularity and develop approximation algorithms for the general case. In [5], instead of considering a budget constraint, the authors aim to minimize the number of deployed middleboxes subject to the constraint that the length of the shortest path of any flow cannot exceed a certain threshold. They consider homogeneous flow demands (i.e., each flow requires one unit of processing capacity), which makes their capacity allocation subproblem solvable in polynomial time. In contrast, we consider heterogeneous flow demands, which renders the problem NP-hard. Moreover, while the objective function of their placement problem is submodular, ours is non-submodular. Different from the aforementioned studies, we consider both capacity and budget constraints that are of practical importance. Considering these practical constraints introduces new challenges discussed above.

In [5], the authors also extend their study to the case of heterogeneous flow demands. However, their proposed algorithm achieves a bicriteria approximation ratio only; specifically, the node capacities may be violated by a constant factor. The work of [16] considers the problem of joint service placement and request routing in Mobile Edge Computing networks. This work shows that the considered problem generalizes other well-studied problems, including that of [5] when flow demands are homogeneous. However, it is unclear whether the generalization applies to the case of heterogeneous flow demands and VNF-node deployment costs, which become relevant in the planning phase addressed in this work. Moreover, the proposed algorithm is based on randomized rounding and achieves a probabilistic bicriteria approximation ratio. Randomized rounding has also been employed to design bicriteria approximation algorithms for Virtual Network Embedding (see, e.g., [17], [18]). Being common in the studies on Virtual Network Embedding, a pre-determined VNF-node placement

is often assumed. In contrast to these studies, we consider the problem of joint VNF-node placement and capacity allocation and present algorithms that achieve constant approximation ratios and do not violate any constraint.

The concept of submodularity has been extensively studied in literature, starting with the seminal work in [19]. Submodular set functions exhibit the diminishing return property, which means that the value of adding an item to a set decreases as the size of the set increases. For problems with submodular objective function, several algorithms can be utilized to solve them efficiently [19], [20]. For non-submodular problems, several useful techniques, including weak submodularity [21], [22] and supermodularity [23], have been developed to address non-submodularity. As far as weak submodularity is concerned, a parameter  $\gamma \in (0, 1]$  is used to quantify how far the objective function is from being submodular. In such settings, approximation results have been established when cardinality constraint [21] or general matroid constraint [22] is considered; the guaranteed approximation ratios deteriorate gracefully as  $\gamma$  moves away from 1. Since weak submodularity is a relatively new concept, the approximation results remain unexplored when the constraint is of other form, such as knapsack. There is another similar concept called supermodularity [23]; the supermodular degree is proposed to measure the deviation from submodularity. The work in [23] introduces an algorithm that can be shown to be effective when the objective function has a small supermodular degree. Different from these techniques, we propose a novel framework that enables us to design efficient algorithms with constant approximation ratios for the non-submodular problem we consider; the achieved approximation ratios do not depend on problem parameters, such as  $\gamma$  in weak submodularity and the supermodular degree.

Recently, we have also extended our framework to more general settings with multiple network functions and multiple types of resources [24]. There are additional challenges in such general settings: it is unclear whether the relaxed placement subproblem is still submodular; the capacity allocation subproblem becomes a multi-dimensional generalization of the generalized assignment problem with assignment restrictions, which is much more challenging. Due to these new challenges, different algorithms and techniques are developed, and the derived approximation ratios are not constant in general.

### III. SYSTEM MODEL AND PROBLEM FORMULATION

We consider a network graph  $G = (\mathcal{V}, \mathcal{E})$ , where  $\mathcal{V}$  is the set of nodes, with  $V = |\mathcal{V}|$ , and  $\mathcal{E}$  is the set of edges connecting nodes in  $G$ . We have a set of flows  $\mathcal{F}$ , with  $F = |\mathcal{F}|$ . We use  $\lambda_f$  to denote the traffic rate of flow  $f \in \mathcal{F}$ . A node is called a VNF-node if it is able to support VNFs. Since ISPs have a limited budget to deploy VNFs in their networks, they can only choose a subset of nodes  $\mathcal{U} \subseteq \mathcal{V}$  to become VNF-nodes. We consider architectures with stateless network functions (e.g., [15]). A flow's state is stored in a data store; no matter where the flow is processed, the state can be accessed from the data store. Therefore, the traffic rate  $\lambda_f$  of each flow can be split and processed at multiple VNF-nodes. We use  $\lambda_f^v$  to denote the portion of flow  $f$  that is assigned to VNF-node  $v$  and use  $\boldsymbol{\lambda} \in \mathcal{R}^{F \times V}$  to denote the assignment matrix.

We assume that the process of transitioning to NFV goes through two main phases: the planning phase and the production phase. The planning phase is concerned about deciding where to introduce NFV to efficiently utilize the limited budget. Since this phase takes place before the actual deployment of the VNF-nodes, we assume that we can utilize historical traces to project flow demands across the network. Then, in the production phase, we can employ online flow admission and routing schemes (e.g., [12]) to dynamically adjust flow routing to efficiently utilize the available resources. In this work, we are mainly focused on the planning phase. Therefore, we assume that the traffic of flow  $f$  will be sent along a predetermined path, which can be obtained from historical traces. We use  $\mathcal{V}_f$  to denote the set of nodes along this path. Alternatively, the nodes along the predetermined path of a flow can also be viewed as potential locations at which the flow will be processed, and routing between these nodes can be dynamically computed in an online fashion. We use  $\mathcal{F}_U$  to denote the set of all flows whose path has one or more nodes in a given set  $\mathcal{U}$ , i.e.,  $\mathcal{F}_U = \{f \in \mathcal{F} \mid \mathcal{V}_f \cap \mathcal{U} \neq \emptyset\}$ .

As we mentioned earlier, the benefits of processed traffic can be harnessed from fully processed flows, i.e., flows that have all of their traffic processed at VNF-nodes. Hence, when a flow traverses VNF-nodes and there is a sufficient capacity on these VNF-nodes to process all of its rate, i.e.,  $\sum_{v \in \mathcal{V}_f \cap \mathcal{U}} \lambda_f^v \geq \lambda_f$ , then the flow is counted as a processed flow.<sup>1</sup> Therefore, the total processed traffic can be expressed as follows:

$$J_1(\mathcal{U}, \boldsymbol{\lambda}) \triangleq \sum_{f \in \mathcal{F}} \lambda_f \mathbf{1}_{\{\sum_{v \in \mathcal{V}_f \cap \mathcal{U}} \lambda_f^v \geq \lambda_f\}}, \quad (1)$$

where  $\mathbf{1}_{\{\cdot\}}$  is the indicator function. Note that each VNF-node  $v$  has a limited processing capacity, denoted by  $c_v$ . Hence, the total traffic rate assigned to a node should satisfy the following capacity constraint:

$$\begin{cases} \sum_{f \in \mathcal{F}} \lambda_f^v \leq c_v, & \forall v \in \mathcal{U}, \\ \lambda_f^v = 0, & \forall f \in \mathcal{F} \text{ and } \forall v \notin \mathcal{U}. \end{cases} \quad (2)$$

We assume that the largest traffic rate of any flow is no larger than the smallest processing capacity of any node.<sup>2</sup> Also, we consider a limited budget, denoted by  $B$ , and require that the total cost of introducing VNF-nodes do not exceed  $B$ . We use  $b_v$  to denote the cost of making node  $v$  a VNF-node, which includes hardware and/or software installation cost and may also depend on the processing capacity. Hence, the total cost of VNF-nodes should satisfy the following budget constraint:

$$\sum_{v \in \mathcal{U}} b_v \leq B. \quad (3)$$

<sup>1</sup>For some flows, we can gain benefits by even processing a fraction of its traffic. In such cases, we have a mix of flows that need to be fully processed and that can be partially processed. The key challenges remain, and the proposed algorithms can be applicable with minimal modifications.

<sup>2</sup>While some studies (e.g., [4]) consider the placement of VNF instances and allow the flow rate to be larger than the capacity of a VNF instance, we consider the problem of placing VNF-nodes, each of which can host multiple VNF instances. Therefore, it is reasonable to assume that the capacity of such a VNF-node is larger than the rate of any flow.

The above budget constraint limits the number of nodes that can become VNF-nodes, and we may only have a subset of flows that traverse some VNF-nodes. Accounting for the above deployment budget and VNF capacity constraints, we consider a joint problem of VNF-nodes placement and capacity allocation (VPCA). The objective is to choose a best subset of nodes to become VNF-nodes and optimally allocate their capacities so as to maximize the total amount of fully processed traffic. We provide the mathematical formulation of the VPCA problem in the following:

$$\begin{aligned} & \underset{\mathcal{U} \subseteq \mathcal{V}, \lambda}{\text{maximize}} \quad J_1(\mathcal{U}, \lambda) \\ & \text{subject to (2), (3).} \end{aligned} \quad (P1)$$

#### IV. CHALLENGES OF VPCA

Here, we will identify the unique challenges of the VPCA problem formulated in (P1). We first decompose the VPCA problem into two subproblems: 1) placement: how to select a subset of nodes to become VNF-nodes and 2) capacity allocation: for a given set of VNF-nodes with fixed capacities, how to divide their capacity for processing a subset of flows. Then, we prove that both subproblems are NP-hard and that the placement subproblem is non-submodular. This is very different from similar problems neglecting the capacity constraint (2) [2], which have been shown to be submodular and can be approximately solved.

##### A. NP-Hardness

First, we present the formulations of the two subproblems. We start with the allocation subproblem because it will be used in the placement subproblem. For a given set of VNF-nodes  $\mathcal{U} \subseteq \mathcal{V}$ , let  $J_2^{\mathcal{U}}(\lambda)$  denote the total amount of fully processed traffic under flow assignment  $\lambda$ . Note that  $J_2^{\mathcal{U}}(\lambda)$  has the same expression as that of  $J_1(\mathcal{U}, \lambda)$  in Eq. (1). The superscript  $\mathcal{U}$  of  $J_2^{\mathcal{U}}(\lambda)$  is to indicate that it is associated with a given set of VNF-nodes  $\mathcal{U}$ . Then, the capacity allocation subproblem for a given set of VNF-nodes  $\mathcal{U}$  can be formulated as

$$\underset{\lambda: (2) \text{ is satisfied}}{\text{maximize}} \quad J_2^{\mathcal{U}}(\lambda). \quad (P2)$$

Let  $J_3(\mathcal{U}) \triangleq \max_{\lambda: (2) \text{ is satisfied}} J_2^{\mathcal{U}}(\lambda)$  denote the optimal value of problem (P2) for a given set of VNF-nodes  $\mathcal{U}$ . Then, the placement subproblem can be formulated as

$$\begin{aligned} & \underset{\mathcal{U} \subseteq \mathcal{V}}{\text{maximize}} \quad J_3(\mathcal{U}) \\ & \text{subject to (3).} \end{aligned} \quad (P3)$$

Note that in order to solve problem (P3), we need to solve problem (P2) to find the optimal  $\lambda$  for a given set of VNF-nodes  $\mathcal{U}$ . In the following theorem, we will show that both subproblems (P2) and (P3) are NP-hard.

*Theorem 1: The capacity allocation subproblem (P2) and the placement subproblem (P3) are both NP-hard.*

*Proof:* See Appendix A.  $\square$

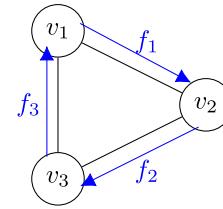


Fig. 1. An example to show non-submodularity of  $J_3(\mathcal{U})$ .

##### B. Non-Submodularity

Note that the objective function  $J_3(\mathcal{U})$  of the placement subproblem (P3) is a set function. At first glance, problem (P3) looks like a submodular maximization problem, which has been extensively studied in the literature and can be approximately solved using efficient algorithms [19], [20]. However, we will show that the objective function  $J_3(\mathcal{U})$  is generally non-submodular, which makes the placement subproblem (P3) and the overall problem (P1) much more challenging. We first give the definition of submodular functions.

*Definition 1: For a finite set of elements  $\mathcal{V}$ , a function  $H : 2^{\mathcal{V}} \rightarrow \mathbb{R}$  is submodular if for any subset  $\mathcal{V}_1 \subseteq \mathcal{V}_2 \subseteq \mathcal{V}$  and any element  $v \in \mathcal{V} \setminus \mathcal{V}_2$ , we have*

$$H(\mathcal{V}_1 \cup \{v\}) - H(\mathcal{V}_1) \geq H(\mathcal{V}_2 \cup \{v\}) - H(\mathcal{V}_2). \quad (4)$$

The above definition exhibits an important property of diminishing returns. In our problem, if the VNF-node capacity is infinite, i.e., there is no capacity constraint (2), then a flow  $f$  can always be fully processed as long as its path has at least one VNF-node, i.e.,  $\mathcal{V}_f \cap \mathcal{U} \neq \emptyset$ . In this case, the total processed traffic  $J_1(\mathcal{U}, \lambda)$  can be rewritten as

$$J'_1(\mathcal{U}) = \sum_{f \in \mathcal{F}} \lambda_f \mathbf{1}_{\{\mathcal{V}_f \cap \mathcal{U} \neq \emptyset\}}, \quad (5)$$

where the capacity allocation becomes irrelevant as it does not impact the value of function  $J'_1(\mathcal{U})$ . It has been shown in [2] that the function  $J'_1(\mathcal{U})$  is monotonically nondecreasing and submodular. In this special case, problem (P1) with objective function  $J'_1(\mathcal{U})$  can be approximately solved using efficient greedy algorithms.

However, using the example presented in Fig. 1, we show that the objective function  $J_3(\mathcal{U})$  is no longer submodular if the VNF-nodes have a limited capacity. Consider three flows: flow  $f_1$  with path  $v_1 \rightarrow v_2$ , flow  $f_2$  with path  $v_2 \rightarrow v_3$ , and flow  $f_3$  with path  $v_3 \rightarrow v_1$ . Assume that each VNF-node has a capacity of 3, and each flow has a traffic rate of 2. If node  $v_3$  is the only VNF-node, then it can only support one flow because its capacity is 3. Therefore, the marginal contribution of adding node  $v_3$  as a VNF-node to the empty set is  $J_3(\{v_3\}) - J_3(\emptyset) = 2 - 0 = 2$ . Now, assume that before making node  $v_3$  a VNF-node, node  $v_2$  is already a VNF-node, which can support one flow. By making node  $v_3$  a VNF-node, all three flows can be fully processed, and hence, the total processed traffic becomes 6, i.e., the marginal contribution of adding node  $v_3$  to the set  $\{v_2\}$  is  $J_3(\{v_2\} \cup \{v_3\}) - J_3(\{v_2\}) = 6 - 2 = 4 > J_3(\{v_3\}) - J_3(\emptyset) = 2$ . This violates the definition of submodular set functions in Eq. (4).

As we mentioned in Section II, there are several useful techniques that have been developed to handle non-submodular functions, such as weak submodularity [21], [22] and supermodularity [23]. When the notion of weak submodularity is considered, one uses a parameter<sup>3</sup>  $\gamma \in (0, 1]$  to measure the deviation of the objective function from being submodular. A larger value of  $\gamma$  is better. For the example presented in Fig. 1, we can show that the value of  $\gamma$  is 0.66. However, the value of  $\gamma$  for function  $J_3(\mathcal{U})$  could be small in general. Even if we can show that  $\gamma$  is relatively large, we are still faced with the following issues that hinder the application of weak submodularity to solving the VPCA problem. First, to our best knowledge, the case with knapsack constraint has not been studied yet in literature. Second, it is hard to analyze the value of  $\gamma$  in our case. This is due to the fact that the placement value function is the optimal value to the resource allocation subproblem, which is NP-hard (see Theorem 1). On the other hand, the supermodular degree characterizes the level of violation of submodularity for a set function. For problems with a bounded supermodular degree, the authors of [23] propose a greedy algorithm with performance guarantees for the case with a non-submodular objective function. However, the proposed greedy algorithm has two main limitations. First, its approximation ratio is a function of the supermodular degree, which, in our case, could be as large as the number of nodes in the network. Second, its complexity is exponential in the supermodular degree and could be prohibitively high when the supermodular degree is large.

Therefore, our problem  $(P1)$  is much more challenging than other similar problems studied in prior work, where the objective function is either submodular or weakly submodular with a large  $\gamma$ , or there is a bounded supermodular degree. To that end, in the next section we will address the aforementioned unique challenges by introducing a novel relaxation and a problem reformulation, which enable us to propose two algorithms with constant approximation ratios.

## V. RELAXATION AND REFORMULATION

In this section, we present a relaxation of the VPCA problem that allows partially processed flows to be counted. Further, we introduce a novel network flow reformulation of the relaxed capacity allocation subproblem. Both of these techniques will be utilized in designing two efficient approximation algorithms in the next section.

### A. Relaxed VPCA Formulation

We first introduce the relaxed VPCA problem, which allows partially processed flows to be counted. In the relaxed VPCA problem, any fraction of flow  $f$  processed by VNF-nodes in  $\mathcal{V}_f \cap \mathcal{U}$  will be counted in the total processed traffic. That is, the relaxed  $J_1(\mathcal{U}, \lambda)$  can be expressed as follows:

$$R_1(\mathcal{U}, \lambda) \triangleq \sum_{f \in \mathcal{F}} \sum_{v \in \mathcal{V}_f \cap \mathcal{U}} \lambda_f^v. \quad (6)$$

<sup>3</sup>For a finite set of elements  $\mathcal{V}$ , a function  $H : 2^{\mathcal{V}} \rightarrow \mathbb{R}$  is  $\gamma$ -weakly submodular for some  $\gamma \in (0, 1]$  if for every two subsets  $\mathcal{V}_1, \mathcal{V}_2 \subseteq \mathcal{V}$ , we have  $\sum_{v \in \mathcal{V}_2} H(\{v\}|\mathcal{V}_1) \geq \gamma H(\mathcal{V}_2|\mathcal{V}_1)$  [22].

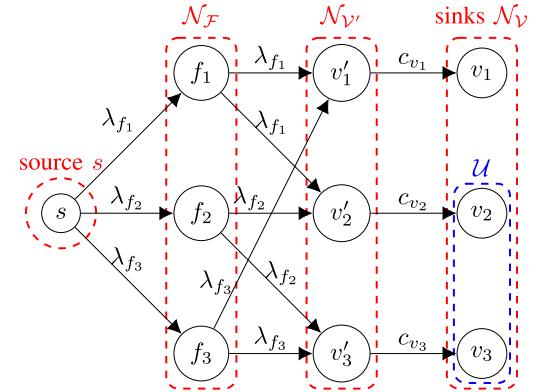


Fig. 2. An example of the constructed graph  $Z$  for the network in Fig. 1, where  $\mathcal{F} = \{f_1, f_2, f_3\}$ ,  $\mathcal{V} = \{v_1, v_2, v_3\}$ ,  $\mathcal{V}_{f_1} = \{v_1, v_2\}$ ,  $\mathcal{V}_{f_2} = \{v_2, v_3\}$ , and  $\mathcal{V}_{f_3} = \{v_1, v_3\}$ .

Apparently, the total processed traffic of flow  $f$  cannot exceed  $\lambda_f$ , i.e., the following constraint needs to be satisfied:

$$\sum_{v \in \mathcal{U}} \lambda_f^v \leq \lambda_f, \quad \forall f \in \mathcal{F}. \quad (7)$$

Then, the relaxed version of problem  $(P1)$  becomes

$$\begin{aligned} & \underset{\mathcal{U} \subseteq \mathcal{V}, \lambda}{\text{maximize}} \quad R_1(\mathcal{U}, \lambda) \\ & \text{subject to (2), (3), (7).} \end{aligned} \quad (Q1)$$

Next, we decompose problem  $(Q1)$ , in the same way as we did for problem  $(P1)$ , into placement and allocation subproblems. For a given set of VNF-nodes  $\mathcal{U} \subseteq \mathcal{V}$ , let  $\Lambda^{\mathcal{U}}$  be the set of all flow assignment matrices  $\lambda$  that satisfy the capacity constraint (2) and the flow rate constraint (7), and let  $R_2^{\mathcal{U}}(\lambda)$  be the total processed traffic, which has the same expression as that of  $R_1(\mathcal{U}, \lambda)$  but has  $\mathcal{U}$  in the superscript so as to indicate that this function is for a given set of VNF-nodes  $\mathcal{U}$ . Then, the capacity allocation subproblem for a given set of VNF-nodes  $\mathcal{U}$  can be formulated as

$$\underset{\lambda \in \Lambda^{\mathcal{U}}}{\text{maximize}} \quad R_2^{\mathcal{U}}(\lambda). \quad (Q2)$$

Now, let  $R_3(\mathcal{U}) \triangleq \max_{\lambda \in \Lambda^{\mathcal{U}}} R_2^{\mathcal{U}}(\lambda)$  denote the optimal value of problem  $(Q2)$  for a given set of VNF-nodes  $\mathcal{U}$ . Then, the placement subproblem can be formulated as

$$\begin{aligned} & \underset{\mathcal{U} \subseteq \mathcal{V}}{\text{maximize}} \quad R_3(\mathcal{U}) \\ & \text{subject to (3).} \end{aligned} \quad (Q3)$$

Note that although the relaxed placement subproblem  $(Q3)$  can still be shown to be NP-hard, we will prove that the objective function  $R_3(\mathcal{U})$  is monotonically nondecreasing and submodular. This useful submodular property allows us to approximately solve problem  $(Q3)$ . On the other hand, the relaxed capacity allocation subproblem  $(Q2)$  becomes an LP, which can be efficiently solved; alternatively, we can also solve  $(Q2)$  using a maximum flow algorithm (discussed at the end of Section VI-A).

### B. Network Flow Formulation

In this subsection, we introduce a novel network flow reformulation of problem  $(Q2)$ . The purpose of this reformulation

is two-fold: i) we will use it to prove that the objective function of the relaxed placement subproblem ( $Q3$ ) is submodular; ii) we will leverage it to develop a combinatorial algorithm for problem ( $Q2$ ) based on the efficient maxflow algorithms (e.g., [25]), which is also a key component of the approximation algorithms we will propose for the original VPCA problem.

For problem ( $Q2$ ), we reformulate a network flow problem by constructing a directed graph  $Z = (\mathcal{N}, \mathcal{L})$  as follows. The set of vertices  $\mathcal{N}$  consists of the following: an artificial source vertex  $s$ , set  $\mathcal{N}_{\mathcal{F}}$  consisting of flow-vertices  $f$  each corresponding to flow  $f \in \mathcal{F}$ , set  $\mathcal{N}_{\mathcal{V}}$  consisting of node-vertices  $v$  each corresponding to node  $v \in \mathcal{V}$ , and set  $\mathcal{N}_{\mathcal{V}'}$  consisting of node-vertices  $v'$  each corresponding to node  $v \in \mathcal{V}$ . Hence,  $\mathcal{N} = \{s\} \cup \mathcal{N}_{\mathcal{F}} \cup \mathcal{N}_{\mathcal{V}} \cup \mathcal{N}_{\mathcal{V}'}$ , where  $\mathcal{N}_{\mathcal{V}}$  consists of the sinks. Let  $(x, y)$  be an edge in  $\mathcal{L}$ , which is from  $x \in \mathcal{N}$  to  $y \in \mathcal{N}$ . The set of edges  $\mathcal{L}$  consists of the following: set  $\mathcal{L}_1$  consisting of edges  $(s, f)$  connecting the source vertex  $s$  to each flow-vertex  $f \in \mathcal{N}_{\mathcal{F}}$ , set  $\mathcal{L}_2$  consisting of edges  $(f, v')$  connecting each flow-vertex  $f \in \mathcal{N}_{\mathcal{F}}$  to each node-vertex  $v' \in \mathcal{N}_{\mathcal{V}'}$  corresponding to a node  $v \in \mathcal{V}$ , set  $\mathcal{L}_3$  consisting of edges  $(v', v)$  connecting each node-vertex  $v' \in \mathcal{N}_{\mathcal{V}'}$  to its corresponding node-vertex  $v \in \mathcal{N}_{\mathcal{V}}$ . We use  $c(x, y)$  to denote the capacity of edge  $(x, y)$ . Hence,  $\mathcal{L} = \mathcal{L}_1 \cup \mathcal{L}_2 \cup \mathcal{L}_3$ . An edge  $(s, f) \in \mathcal{L}_1$  has capacity  $\lambda_f$ ; an edge  $(f, v') \in \mathcal{L}_2$  has capacity  $\lambda_f$ ; an edge  $(v', v) \in \mathcal{L}_3$  has capacity  $c_v$ . Fig. 2 presents an example of the constructed graph  $Z$  for the network in Fig. 1.

Next, we describe flows over graph  $Z$ . Consider functions  $\varphi(x, y) : \mathcal{N} \times \mathcal{N} \rightarrow \mathbb{R}_+$ , where  $\mathbb{R}_+$  is the set of non-negative real numbers. We define  $\Phi(\mathcal{X}, \mathcal{Y}) \triangleq \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} \varphi(x, y)$  for  $\mathcal{X}, \mathcal{Y} \subseteq \mathcal{N}$ . An  $s$ - $\mathcal{V}$  flow is a function  $\varphi(x, y) : \mathcal{N} \times \mathcal{N} \rightarrow \mathbb{R}_+$  such that the following is satisfied:

- 1) Capacity constraints:  $\varphi(x, y) \leq c(x, y)$  for all pairs  $(x, y) \in \mathcal{N} \times \mathcal{N}$ . (Note that  $c(x, y) = 0$  if  $(x, y) \notin \mathcal{L}$ .)
- 2) Flow conservation: the net-flow at every non-source non-sink vertex  $x \in \mathcal{N} \setminus (\{s\} \cup \mathcal{N}_{\mathcal{V}})$  is zero, i.e.,  $\Phi(\mathcal{N}, \{x\}) - \Phi(\{x\}, \mathcal{N}) = 0$ .
- 3) Positive incoming flow: the net-flow at the source  $s$  is non-positive, i.e.,  $\Phi(\mathcal{N}, \{s\}) - \Phi(\{s\}, \mathcal{N}) \leq 0$ .
- 4) Positive outgoing flow: the net-flow at every sink  $t \in \mathcal{N}_{\mathcal{V}}$  is non-negative, i.e.,  $\Phi(\mathcal{N}, \{t\}) - \Phi(\{t\}, \mathcal{N}) \geq 0$ .

Let  $\overline{\mathcal{F}}$  be the set of all  $s$ - $\mathcal{V}$  flows over  $Z$ .

For a subset of sinks<sup>4</sup>  $\mathcal{U} \subseteq \mathcal{N}_{\mathcal{V}}$ , we define

$$F(\mathcal{U}) \triangleq \max_{\varphi \in \overline{\mathcal{F}}} (\Phi(\mathcal{N}, \mathcal{U}) - \Phi(\mathcal{U}, \mathcal{N})), \quad (8)$$

which is the maximum total net-flow at the sinks in  $\mathcal{U}$ . The maximum net-flow problem is to find an  $s$ - $\mathcal{V}$  flow (i.e., function  $\varphi$ ) that achieves the maximum in (8). In Lemma 1, we show the equivalence between the capacity allocation subproblem ( $Q2$ ) and the maximum net-flow problem (8).

*Lemma 1: The capacity allocation subproblem ( $Q2$ ) is equivalent to the maximum net-flow problem (8). Hence, for any given  $\mathcal{U} \subseteq \mathcal{V}$ , the optimal value of problem ( $Q2$ ) is equal to the maximum total net-flow at the sinks in  $\mathcal{U} \subseteq \mathcal{N}_{\mathcal{V}}$  of the*

<sup>4</sup>Note that each node  $v \in \mathcal{V}$  corresponds to a sink in  $\mathcal{N}_{\mathcal{V}}$ . Hence, by slightly abusing the notations, for any  $\mathcal{U} \subseteq \mathcal{V}$ , we also use  $\mathcal{U}$  to denote the corresponding subset of sinks in  $\mathcal{N}_{\mathcal{V}}$ .

---

**Algorithm 1** The RP-MCA and RP-GCA Algorithms

---

Input: set of nodes  $\mathcal{V}$ , set of flows  $\mathcal{F}$ , node capacities, node costs, flow rates, and budget  $B$ .  
Output: set of VNF-nodes  $\mathcal{U}$ , capacity allocation  $\lambda$ .

- 1: **Relaxed Problem:** relax function  $J_1(\mathcal{U}, \lambda)$  to become  $R_1(\mathcal{U}, \lambda)$ ;
- 2: **Placement Subproblem:** solve problem ( $Q3$ ) using the Submodular Greedy algorithm or the Enumeration-based Greedy algorithm, described in Section VI-A, to obtain  $\mathcal{U}$ .
- 3: **Capacity Allocation:** use either the MCA algorithm (Algorithm 2) or the GCA algorithm (Algorithm 3) to obtain capacity allocation  $\lambda$ .

---

associated graph  $Z$ , i.e.,

$$R_3(\mathcal{U}) = F(\mathcal{U}). \quad (9)$$

*Proof:* See Appendix B. □

## VI. PROPOSED ALGORITHMS

In this section, we design two efficient algorithms that can achieve constant approximation ratios for the VPCA problem ( $P1$ ). The main idea is to utilize the relaxation introduced in the previous section, which allows partially processed flows to be counted. By doing so, we can show that the relaxed placement subproblem is submodular based on the network flow reformulation of the relaxed capacity allocation subproblem. In this case, the relaxed placement subproblem can be approximately solved using efficient greedy algorithms. Moreover, the relaxed allocation subproblem becomes a Linear Program (LP), which can also be solved efficiently in polynomial time. However, the solution to the relaxed problem is for the case where any fraction of the processed flows is counted. In order to obtain a solution for the original VPCA problem ( $P1$ ), where only the fully processed flows are counted, we propose two approximation algorithms by modifying the solution to the relaxed capacity allocation subproblem: the first one is based on a maximum flow algorithm, and the second one is based on a greedy algorithm.

We use RP-MCA and RP-GCA to denote the algorithms we develop by combining the Relaxed Placement with the Maximum-flow-based Capacity Allocation and the Greedy Capacity Allocation, respectively. We show that the RP-MCA and RP-GCA algorithms achieve an approximation ratio of  $\frac{1}{2}(1 - 1/e)$  and  $\frac{1}{3}(1 - 1/e)$ , respectively. We describe the algorithms in a unified framework presented in Algorithm 1. The difference is in the capacity allocation subproblem (line 3), where RP-MCA algorithm uses a Max-flow-based Capacity Allocation (MCA) algorithm presented in Algorithm 2, while RP-GCA algorithm uses a Greedy Capacity Allocation (GCA) algorithm presented in Algorithm 3.

### A. Proposed Placement Algorithms

In this subsection, we first prove in Lemma 2 that the objective function  $R_3(\mathcal{U})$  of the relaxed placement subproblem ( $Q3$ ) is monotonically nondecreasing and submodular. Then,

using the property of submodularity, we propose two greedy algorithms for solving the placement subproblem.

*Lemma 2: The function  $R_3(\mathcal{U})$  is monotonically nondecreasing and submodular.*

*Proof:* The function  $R_3(\mathcal{U})$  is monotonically nondecreasing because adding an additional VNF-node does not reduce the amount of flows that can be fully processed.

Next, we prove that the function  $R_3(\mathcal{U})$  is submodular. The proof follows from the network flow reformulation introduced in Section V-B. Applying Lemma 1 and the max-flow min-cut theorem (see, e.g., [26, pp. 348–349]) immediately gives

$$R_3(\mathcal{U}) = \min_{\mathcal{X} \subseteq \mathcal{N}: s \in \mathcal{X}, \mathcal{U} \subseteq \mathcal{N} \setminus \mathcal{X}} \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{N} \setminus \mathcal{X}} c(x, y). \quad (10)$$

One then obtains that  $R_3(\mathcal{U})$  is submodular (as the partial minimization of a cut function) [26, p. 230].  $\square$

Because of this useful submodular property, problem (Q3) can be approximately solved using efficient greedy algorithms. Next, we consider two cases of problem (Q3): uniform VNF-node costs (Case I, a special case) and heterogeneous VNF-node costs (Case II, a general case).

In Case I, the VNF-nodes have uniform costs, i.e.,  $b_v = b$  for all  $v \in \mathcal{V}$ . Then, the budget constraint (3) can be expressed as a cardinality constraint, i.e.,  $|\mathcal{U}| \leq k$ , where  $k = \lfloor B/b \rfloor$ . In this case, we can use a simple *Submodular Greedy* (SG) algorithm to approximately solve problem (Q3). In the SG algorithm, we start with an empty solution of VNF-nodes  $\mathcal{U}$ ; in each iteration, we add a node that has the maximum marginal contribution to  $\mathcal{U}$ , i.e., a node that leads to the largest increase in the value of the objective function. If multiple nodes have the same marginal contribution, we break ties by selecting a node uniformly at random. We repeat the above procedure until  $k$  VNF-nodes have been selected. This solution has been shown to achieve an approximation ratio of  $(1 - 1/e)$  [19]. However, this algorithm does not guarantee to have the same approximation ratio for the case of heterogeneous VNF-node costs [20].

In Case II, the VNF-nodes have heterogeneous costs, i.e., the costs of VNF-nodes are different. For this case, an *Enumeration-based Greedy* (EG) algorithm has been proposed in [20], which can be shown to achieve the same approximation ratio of  $(1 - 1/e)$ , but with a higher running time complexity compared to the SG algorithm. The EG algorithm has two phases. In Phase I, it samples all node subsets of cardinality one or two that satisfy the budget constraint, picks the one with the largest value of the objective function  $R_3$ , and stores this temporary solution in  $\mathcal{U}_1$ . In Phase II, the algorithm samples all node subsets of cardinality three and augments each of these subsets with nodes that maximize the relative marginal contribution  $(R_3(\mathcal{V}' \cup \{u\}) - R_3(\mathcal{V}'))/b_u$ , in a greedy manner. The budget constraint must also be satisfied throughout this procedure. Then, it selects the augmented subset with the largest value of the objective function  $R_3$  and stores it in  $\mathcal{U}_2$ . The final solution will be the better one between  $\mathcal{U}_1$  and  $\mathcal{U}_2$ , i.e., the one that achieves a larger value of the objective function  $R_3$ .

Note that although the value of function  $R_3(\mathcal{U})$  can be obtained using an LP solver, we can alternatively compute

it using the network flow formulation presented in Section V-B as follows. For the constructed graph  $Z$ , we connect all the sink vertices corresponding to nodes  $\mathcal{U}$  to an artificial sink vertex  $d$ . Then, the value of  $R_3(\mathcal{U})$  is the maximum flow from vertex  $s$  to vertex  $d$  in graph  $Z$ , which can be computed using several efficient algorithms (see, e.g., [25]). In Lemma 3, we restate the results of [19], [20] about the approximation ratio of the SG and EG algorithms.

*Lemma 3: Both the SG and EG algorithms achieve an approximation ratio of  $(1 - 1/e)$ .*

*Proof:* The proofs can be found in [19] and [20] for the SG algorithm and the EG algorithm, respectively.  $\square$

### B. Proposed Capacity Allocation Algorithms

While the solution of problem (Q3) allows partially processed flows to be counted, only fully processed flows will be counted in the original problem (P1). To that end, we propose two algorithms to modify the capacity allocation of VNF-nodes  $\mathcal{U}$  so as to ensure fully processed flows and provide certain performance guarantees. The first algorithm is based on the network flow formulation, and the second one is based on a simple greedy approach. We develop these algorithms by modifying two algorithms for the multiple knapsack problem with assignment restrictions (MKAR) [27]. However, we want to point out that there is a key difference between our studied VPCA problem and the MKAR problem: in the VPCA problem, a flow can be split and assigned to more than one VNF-node, while in the MKAR problem, an item (corresponding to a flow in our problem) cannot be split and must be assigned to at most one knapsack (corresponding to VNF-node in our problem). Because of this key difference, an optimal solution for the VPCA problem generally has a larger value compared to that of the MKAR problem. Therefore, the algorithms developed for the MKAR problem need to be modified so as to yield a better performance.

First, we introduce some additional notations for the algorithms that will be described below. We use  $\mathcal{U}_f$  to denote the nodes on the path of flow  $f$  that are included in  $\mathcal{U}$ , i.e.,  $\mathcal{U}_f = \mathcal{V}_f \cap \mathcal{U}$ . Let  $c'_v$  denote the remaining capacity of VNF-node  $v$ , and let  $c'_{\mathcal{U}_i}$  denote the total remaining capacity of the set of VNF-nodes in  $\mathcal{U}_i$ , i.e.,  $c'_{\mathcal{U}_i} = \sum_{v \in \mathcal{U}_i} c'_v$ . In what follows, we will introduce the MCA algorithm and the GCA algorithm.

1) *Maximum-Flow-Based Capacity Allocation (MCA):* We first present the MCA algorithm (Algorithm 2), a capacity allocation algorithm based on the network flow formulation. The MCA algorithm has two phases. In Phase I, MCA makes allocation decisions by rounding a fractional flow assignment obtained by solving problem (Q2); in Phase II, the remaining VNF-node capacities are allocated in a greedy manner.

Phase I: Let  $\lambda_{\mathcal{U}}$  be a flow assignment obtained from an optimal basic solution<sup>5</sup> of problem (Q2), which can be obtained by solving a maximum flow problem as discussed earlier. We use  $y_f^v \triangleq \lambda_f^v / \lambda_f$  to denote the fraction of flow  $f$  assigned to VNF-node  $v$  in the obtained solution  $\lambda_{\mathcal{U}}$ . The

<sup>5</sup>A basic feasible solution is a solution that cannot be expressed as a convex combination of two feasible solutions.

**Algorithm 2** The MCA Algorithm

---

Input: set of VNF-nodes  $\mathcal{U}$ , set of flows  $\mathcal{F}_{\mathcal{U}}$ , flow rates, and VNF-node capacities.  
 Output: Capacity allocation  $\lambda$ .

**Phase I:**

- 1: Obtain a basic optimal solution  $\lambda_{\mathcal{U}}$ ;
- 2:  $y_f^v \triangleq \lambda_f^v / \lambda_f$ , for all  $\lambda_f^v$  in  $\lambda_{\mathcal{U}}$ ;
- 3: Assign each flow  $f$  with  $y_f^v = 1$  to VNF-node  $v$ ;
- 4: Construct  $G'$  for the unassigned flows with positive  $y_f^v$ ;
- 5: **while**  $G'$  is not empty **do**
- 6:   **while** there is a singleton VNF-node in  $G'$  **do**
- 7:     Perform the rounding in Step 1;
- 8:   **end while**
- 9:   Perform the rounding in Step 2;
- 10: **end while**

**Phase II:**

  - 11: **for** each flow  $f$  in  $\mathcal{F}_{\mathcal{U}}$  that is not assigned yet **do**
  - 12:   **if**  $c'(\mathcal{U}_f) \geq \lambda_f$  **then**
  - 13:     Assign flow  $f$  to a subset of VNF-nodes in  $\mathcal{U}_f$ ;
  - 14:   **end if**
  - 15: **end for**

---

algorithm begins with a temporary assignment of every flow  $f$  with  $y_f^v = 1$  to the corresponding VNF-node  $v$ . For the remaining flows, we do the following. Let  $G' = (\mathcal{F}', \mathcal{V}', \mathcal{E}')$  be a bipartite graph constructed as follows. For each  $\lambda_f^v \in \lambda_{\mathcal{U}}$ , if  $0 < y_f^v < 1$ , we add a flow vertex  $f$  to the set  $\mathcal{F}'$ , a VNF-node vertex  $v$  to the set  $\mathcal{V}'$ , and an edge, with weight  $y_f^v$ , connecting flow vertex  $f$  to VNF-node vertex  $v$ , to the set  $\mathcal{E}'$ . Note that graph  $G'$  cannot have a cycle because  $\lambda_{\mathcal{U}}$  is a basic feasible solution [27, Lemma 5]. After constructing graph  $G'$ , we repeatedly apply the following two steps to graph  $G'$  until it becomes empty. As a result, the modified flow assignment  $y_f^v$  will become either zero or one.

Step 1: For each VNF-node  $v \in \mathcal{V}'$  that has only one incident flow  $f$  (called a singleton VNF-node), we modify its capacity allocation as follows. Let  $r_v$  denote the total amount of flow rates assigned to VNF-node  $v$  and let  $r'_v$  be the portion of  $r_v$  contributed by fully assigned flows. Note that  $r_v = r'_v + \lambda_f^v$ . If  $r'_v \geq \lambda_f^v$ , then we set  $y_f^v$  to zero. Now, VNF-node  $v$  has no incident edges to it, so we remove it from  $G'$ . In this case, the value of solution  $\lambda_{\mathcal{U}}$  will be reduced by  $\lambda_f^v$ , which is no greater than  $\frac{1}{2}r_v$ . If  $r'_v < \lambda_f^v$ , then we unassign the flows temporarily assigned to VNF-node  $v$  and assign flow  $f$  to VNF-node  $v$  instead, i.e., set  $y_f^v$  to one, and cancel the other fractions of flow  $f$  assigned to other VNF-nodes. This is feasible because the rate of any flow is assumed to be no larger than the minimum VNF-node capacity. Then, we remove VNF-node  $v$ , flow  $f$ , and the associated edges from  $G'$ . In this case, the value of solution  $\lambda_{\mathcal{U}}$  will be reduced by at most  $r'_v$ , which is no greater than  $\frac{1}{2}r_v$ . We repeat Step 1 until no singleton VNF-node exists. Then, we go to Step 2.

Step 2: In this step, we will perturb the fractional values of some edges in  $G'$  to make one of them either zero or one. The perturbation is designed such that the capacity and assignment constraints are not violated and the total assigned

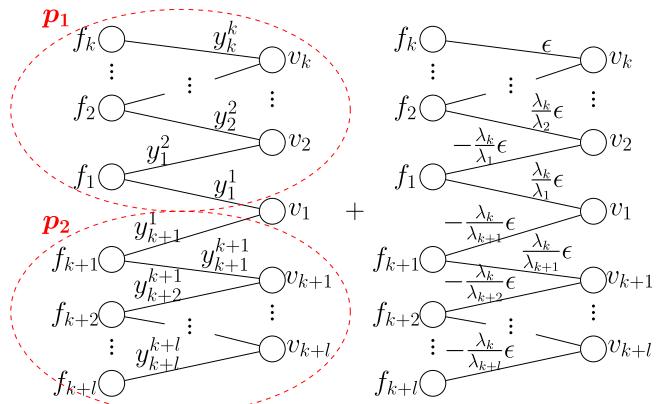


Fig. 3. An example of the edges perturbation.

traffic remains the same. We describe the perturbation procedure in the following. Consider a VNF-node  $v_1 \in \mathcal{V}'$  that has a degree of at least two. Let  $(v_1, f_1)$  and  $(v_1, f_{k+1})$  denote two of the incident edges to VNF-node  $v_1$ . Let  $\mathbf{p}_1$  and  $\mathbf{p}_2$  denote the longest paths starting from VNF-node  $v_1$  through edges  $(v_1, f_1)$  and  $(v_1, f_{k+1})$ , respectively; such paths exist because  $G'$  is a forest. Here, we use  $y_i^j$  to denote the fractional value of flow  $i$  assigned to VNF-node  $j$  and use  $\lambda_j$  to denote the rate of flow  $j$ . Let  $\mathbf{y}_1 = (y_1^1, y_1^2, \dots, y_1^k)$  denote the fractional flow assignment on the edges of path  $\mathbf{p}_1$ , and let  $f_1, \dots, f_k$  be the flow nodes of path  $\mathbf{p}_1$ . Similarly, let  $\mathbf{y}_2 = (y_{k+1}^1, y_{k+1}^2, \dots, y_{k+1}^{k+l-1})$  denote the fractional flow assignment on the edges of path  $\mathbf{p}_2$ , and let  $f_{k+1}, \dots, f_{k+l}$  denote the flow nodes of path  $\mathbf{p}_2$ . We perturb  $\mathbf{y}_1$  by adding to it  $\mathbf{y}'_1 = (\frac{\lambda_k}{\lambda_1} \epsilon, -\frac{\lambda_k}{\lambda_1} \epsilon, \frac{\lambda_k}{\lambda_2} \epsilon, \dots, -\frac{\lambda_k}{\lambda_{k-1}} \epsilon, \epsilon)$ , and we perturb  $\mathbf{y}_2$  by adding to it  $\mathbf{y}'_2 = (-\frac{\lambda_k}{\lambda_{k+1}} \epsilon, \frac{\lambda_k}{\lambda_{k+1}} \epsilon, -\frac{\lambda_k}{\lambda_{k+2}} \epsilon, \dots, \frac{\lambda_k}{\lambda_{k+l-1}} \epsilon, -\frac{\lambda_k}{\lambda_{k+l}} \epsilon)$ . We increase  $\epsilon$  until one fractional value  $y_f^v$  becomes zero or one, and if one, i.e.,  $y_f^v = 1$ , then we assign flow  $f$  to the corresponding VNF-node  $v$ . An example to illustrate this step is shown in Fig. 3. In this new solution, at least one edge is removed from  $G'$ . We repeat the perturbation procedure until at least one VNF-node becomes a singleton, and then we go back to Step 1. If  $G'$  becomes empty, we start Phase II.

Phase II: We leverage the property that the traffic of a flow can be split and processed at multiple VNF-nodes. That is, after Phase I, we pick an unassigned flow  $f$  and check if the total remaining capacity of VNF-nodes  $\mathcal{U}_f$  is no smaller than  $\lambda_f$ . If so, we split flow  $f$  so that the remaining capacities of some VNF-nodes in  $\mathcal{U}_f$  can be used to fully process flow  $f$  and assign flow  $f$  to a subset of these VNF-nodes. We repeat this procedure until no more flow can be assigned.

We use  $OPT(Q2, \mathcal{U})$  to denote the total traffic assigned to a given set of VNF-nodes  $\mathcal{U}$  by an optimal solution to problem  $(Q2)$ . Also, we use  $\pi_{MCA}^{\mathcal{U}}$  to denote the total traffic assigned to VNF-nodes  $\mathcal{U}$  by the MCA algorithm. The approximation ratio of the MCA algorithm is stated in the following Lemma.

**Lemma 4:** *The MCA algorithm has an approximation ratio of  $1/2$ , i.e.,  $\pi_{MCA}^{\mathcal{U}} \geq \frac{1}{2}OPT(Q2, \mathcal{U})$ .*

*Proof:* See Appendix C.  $\square$

2) *Greedy Capacity Allocation (GCA):* While the MCA algorithm achieves an approximation ratio of  $1/2$ , it has a relatively high complexity of  $O(F^2V^2)$  (refer to Table I for

**Algorithm 3** The GCA Algorithm

---

Input: set of VNF-nodes  $\mathcal{U}$ , set of flows  $\mathcal{F}_{\mathcal{U}}$ , flow rates, and VNF-node capacities.  
 Output: Capacity allocation  $\lambda$ .

- 1: Sort flows  $\mathcal{F}_{\mathcal{U}}$  in a nonincreasing order of their flow rates;
- 2: **Phase I:**
- 3: **for** each flow  $f$  in the sorted set  $\mathcal{F}_{\mathcal{U}}$  **do**
- 4:   **if** there is a VNF-node  $v$  in  $\mathcal{U}_f$  such that  $c'_v \geq \lambda_f$  **then**
- 5:     Set  $\lambda_f^v = \lambda_f$ ;
- 6:     Set  $c'_v = c'_v - \lambda_f$ ;
- 7:   **end if**
- 8: **end for**
- 9: **Phase II:**
- 10: **for** each flow  $f$  in the sorted set  $\mathcal{F}_{\mathcal{U}}$  that is not assigned yet **do**
- 11:   **if**  $c'(\mathcal{U}_f) \geq \lambda_f$  **then**
- 12:     Assign flow  $f$  to a subset of VNF-nodes in  $\mathcal{U}_f$ ;
- 13:   **end if**
- 14: **end for**

---

the complexity analysis). This high complexity may render the MCA algorithm unsuitable for certain scenarios in practice. To that end, we propose the GCA algorithm, a simple greedy capacity allocation algorithm that has a much lower complexity of  $O(FV)$ . A lower complexity of the GCA algorithm is achieved at the cost of a slightly worse approximation ratio of  $1/3$  (Lemma 5). However, the approximation ratio of the GCA algorithm can be improved to  $2/5$  (cf. Lemma 6 below) if an additional mild assumption (Assumption 1) holds. The GCA algorithm has two phases. In Phase I, we sort flows of  $\mathcal{F}_{\mathcal{U}}$  in a nonincreasing order of their flow rates. Then, we iteratively go through the sorted list and assign each flow to any VNF-node in  $\mathcal{U}_f$  if it has a sufficient capacity. In Phase II, the remaining capacities of the VNF-nodes can be allocated in a similar way to Phase II of the MCA algorithm by leveraging the property that a flow can be processed at multiple VNF-nodes. However, here the remaining flows need to be considered according to the order in the sorted list  $\mathcal{F}_{\mathcal{U}}$ . The GCA algorithm is presented in Algorithm 3.

In Lemma 5, we state the result about the approximation ratio of the GCA algorithm. We use  $\pi_{\text{GCA}}^{\mathcal{U}}$  to denote the total traffic assigned to VNF-nodes  $\mathcal{U}$  by the GCA algorithm.

*Lemma 5: The GCA algorithm has an approximation ratio of  $1/3$ , i.e.,  $\pi_{\text{GCA}}^{\mathcal{U}} \geq \frac{1}{3} \text{OPT}(\mathcal{Q}_2, \mathcal{U})$ .*

*Proof:* See Appendix D.  $\square$

Further, we show in Lemma 6 that the approximation ratio of the GCA algorithm can be improved to  $2/5$  when an additional mild assumption (Assumption 1) holds.

*Assumption 1: Assume that all the VNF-nodes in  $\mathcal{U}$  have the same capacity and that every flow  $f$  in  $\mathcal{F}_{\mathcal{U}}$  traverses at least two VNF-nodes in  $\mathcal{U}$ , i.e.,  $|\mathcal{V}_f \cap \mathcal{U}| \geq 2$ .*

*Lemma 6: Suppose that Assumption 1 holds. Then, the GCA algorithm has an improved approximation ratio of  $2/5$ , i.e.,  $\pi_{\text{GCA}}^{\mathcal{U}} \geq \frac{2}{5} \text{OPT}(\mathcal{Q}_2, \mathcal{U})$ .*

*Proof:* See Appendix E.  $\square$

**C. Main Results**

We state our main results in Theorems 2 and 3.

*Theorem 2: The RP-MCA algorithm has an approximation ratio of  $\frac{1}{2}(1 - 1/e)$  for problem (P1).*

*Proof:* The RP-MCA algorithm has two main components: 1) VNF-nodes placement and 2) capacity allocation. For the relaxed placement subproblem ( $\mathcal{Q}_3$ ), we use  $\pi_G^{\mathcal{U}}$  to denote the value of the optimal relaxed allocation for the set of VNF-nodes  $\mathcal{U}$  selected by the SG algorithm or the EG algorithm. Also, we use  $\text{OPT}(P)$  to denote the optimal value of any problem ( $P$ ). We have the following result:

$$\begin{aligned} \pi_G^{\mathcal{U}} &\stackrel{(a)}{\geq} (1 - 1/e)\text{OPT}(8) \\ &\stackrel{(b)}{=} (1 - 1/e)\text{OPT}(8) \\ &\stackrel{(c)}{\geq} (1 - 1/e)\text{OPT}(\mathcal{P}_1), \end{aligned} \quad (11)$$

where (a) is due to Lemma 3, (b) holds because an optimal capacity allocation is assumed for the objective function of problem ( $\mathcal{Q}_3$ ), and (c) holds because problem ( $\mathcal{P}_1$ ) is a relaxed version of problem ( $P_1$ ).

The second component of the RP-MCA algorithm is the capacity allocation using the MCA algorithm for the set of VNF-nodes  $\mathcal{U}$  selected by the SG or EG algorithm. We have the following result:

$$\begin{aligned} \pi_{\text{MCA}}^{\mathcal{U}} &\stackrel{(a)}{\geq} \frac{1}{2} \text{OPT}(\mathcal{Q}_2, \mathcal{U}) \\ &\stackrel{(b)}{=} \frac{1}{2} \pi_G^{\mathcal{U}} \\ &\stackrel{(c)}{\geq} \frac{1}{2} (1 - 1/e) \text{OPT}(\mathcal{P}_1), \end{aligned} \quad (12)$$

where (a) comes from the approximation ratio of the MCA algorithm in Lemma 4, (b) holds because when  $\pi_G^{\mathcal{U}}$  is obtained for problem ( $\mathcal{Q}_3$ ) using the greedy algorithms, an optimal capacity allocation (i.e., an optimal solution to problem ( $\mathcal{Q}_2$ ) associated with the considered  $\mathcal{U}$ ) is assumed for the objective function, and (c) holds from Eq. (11). Therefore, the result of Theorem 2 follows.  $\square$

*Theorem 3: The RP-GCA algorithm has an approximation ratio of  $\frac{1}{3}(1 - 1/e)$  for problem (P1).*

*Proof:* The proof follows the same argument as in the proof of Theorem 2. Since the GCA algorithm achieves an approximation ratio of  $1/3$  for the capacity allocation (Lemma 5), the proof proceeds exactly the same except that we need to replace  $1/2$  with  $1/3$  in Eq. (12).  $\square$

Table I summarizes the complexity of our proposed algorithms. In literature on submodular optimization, the complexity of algorithms for submodular functions is often measured through the number of function evaluations. The function evaluation itself is usually assumed to be conducted by an oracle, and thus its complexity is not taken into account [28]. We followed this approach here. Note that we can utilize other alternative algorithms to the EG algorithm to improve the running time substantially but with a slightly worse approximation ratio [20], [28]. We provide more discussions about the complexity analysis in Appendix F.

TABLE I

APPROXIMATION RATIOS AND TIME COMPLEXITIES OF THE PROPOSED ALGORITHMS. \*THESE ARE THE APPROXIMATION RESULTS FOR THE GCA ALGORITHM WHEN ASSUMPTION 1 HOLDS.  $\dagger$ THIS IS THE NUMBER OF FUNCTION EVALUATIONS USED IN THE SUBMODULAR OPTIMIZATION

Setting	Algorithm	Approximation	Complexity
Homogeneous VNF costs	RP-MCA	$\frac{1}{2}(1 - 1/e)$	$O(kV)^\dagger + O(F^2V^2)$
	RP-GCA	$\frac{1}{3}(1 - 1/e)$ $\frac{2}{5}(1 - 1/e)^*$	$O(kV)^\dagger + O(FV)$
Heterogeneous VNF costs	RP-MCA	$\frac{1}{2}(1 - 1/e)$	$O(V^5)^\dagger + O(F^2V^2)$
	RP-GCA	$\frac{1}{3}(1 - 1/e)$ $\frac{2}{5}(1 - 1/e)^*$	$O(V^5)^\dagger + O(FV)$

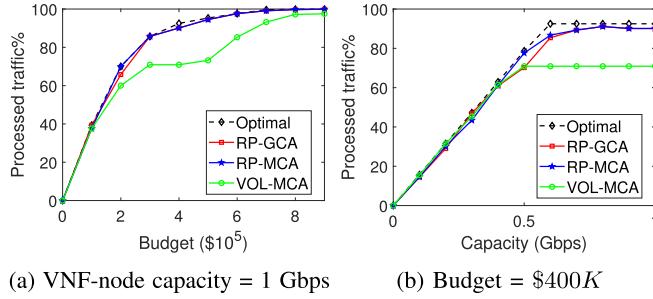


Fig. 4. Evaluation on the Abilene dataset.

## VII. NUMERICAL RESULTS

In order to evaluate the performance of the proposed algorithms, we consider real-world network topologies and traffic statistics. We compare the proposed algorithms with the following baselines: 1) optimal solution: we can solve problem (*P*1) optimally using Gurobi [29], an ILP solver, for the presented instances. Recall that the VPCA problem is NP-hard in general (Theorem 1). Although we are able to obtain the optimal solution for the problem instances we consider here, it may take a prohibitively large amount of time to obtain the optimal solution for some other problem instances. 2) VOL-MCA [6]: this scheme selects the nodes with the highest traffic volume that traverses them. For the selected nodes, we allocate their capacity using the proposed MCA algorithm. We evaluate the performance of each algorithm based on the percentage of the processed traffic, which is defined as the ratio between the total volume of the traffic fully processed by the VNF-nodes and the total traffic volume. We run the simulations on a PC with Intel Core i7-7700 processor and 32GB memory.

### A. Evaluation Datasets and Simulation Parameters

1) *Abilene Dataset*: We consider the Abilene dataset [30] collected from an educational backbone network in North America. The network consists of 12 nodes and 144 flows. Each flow rate was recorded every five minutes for 6 months. The OSPF weights were recorded, which allows us to compute the shortest path of each flow based on these weights. In our experiments, we set the flow rate to the recorded value of the first day at 8:00 pm.

2) *SNDlib Datasets*: We also consider two datasets from SNDlib [31]: Cost266 and Ta2. Cost266 has 37 nodes and 1332 flows; Ta2 has 65 nodes and 1869 flows. For Cost266, the routing cost of each link is available, which can be used

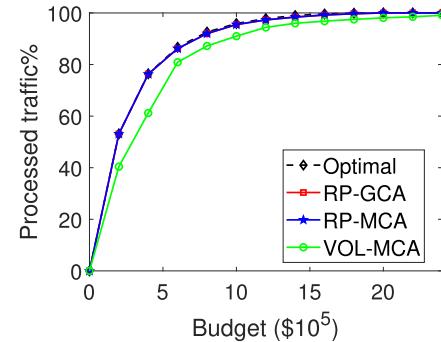


Fig. 5. Evaluation on the Cost266 datasets.

to compute the shortest path of each flow. For Ta2, we use the hop-count-based shortest path. We set the capacity of each VNF-node to 1 Gbps as the default value.

### B. Evaluation Results

We start with the Abilene dataset, where we set the cost of a VNF-node to \$100K and vary the processing capacity between 100 Mbps and 1 Gbps. Also, we vary the total budget between \$100K and \$900K, with an increment of \$100K. Fig. 4(a) shows the percentage of processed traffic for the considered algorithms. We can see that both the RP-MCA and RP-GCA algorithms perform almost the same as the optimal solution and have up to 20% improvement over the VOL-MCA algorithm. Note that as the budget increases, the total processed traffic increases under all the considered algorithms. However, while the proposed algorithms need a budget of \$500K to process around 95% of the flows, the VOL-MCA algorithm actually requires around \$800K to process the same amount. While the total amount of processed traffic can be improved by deploying more VNF-nodes, another option is to consider provisioning more capacities at each VNF-node. In other words, we can deploy fewer but more powerful VNF-nodes, which can improve the percentage of the total processed traffic as well. We can make this observation in Fig. 4(b), where we vary the capacity of each VNF-node between 100 Mbps and 1 Gbps. The total processed traffic reaches around 92% and saturates at that point. This happens because the path of some flows may not contain any of the VNF-nodes placed by the algorithms we consider.

Furthermore, we evaluate the proposed algorithms on datasets with a larger number of nodes and flows. We also set the cost of each VNF-node to \$100K and the processing capacity to 1 Gbps. We start with the Cost266 dataset, which consists of 37 nodes and 1332 flows. Since we have a large number of flows, we select flows uniformly at random. Specifically, we select 1000 flows and repeat each experiment 10 times to gain more statistical significance. We also vary the budget between \$100K and \$2.4M with an increment of \$200K. Fig. 5 shows that the proposed algorithms still exhibit superior performance compared to the VOL-MCA algorithm and match the performance of the optimal solution. An interesting observation is that for this specific instance, we only need to deploy six VNF-nodes to process around 90% of the traffic. This suggests that we may gradually transition to the NFV paradigm with a lower total cost and consolidate

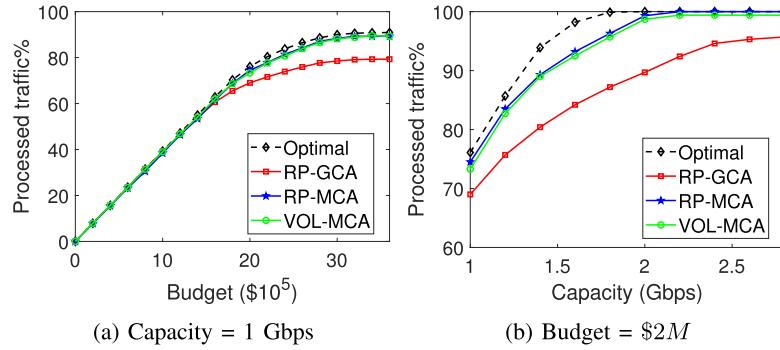


Fig. 6. Evaluation on the Ta2 datasets.

the NFV functionalities at a small portion of nodes. The later point of having a small number of VNF-nodes is advantageous in terms of reducing the management burden, but that could increase the risk of having a single point of failure. In our future work, we will investigate how to adjust the proposed algorithms to ensure resilience against hardware failures at the VNF-nodes. Note that VNF-node capacities are not a bottleneck because flow rates of the Cost266 dataset are relatively small. Therefore, increasing the capacity of VNF nodes shows no impact on the total processed traffic.

Finally, we consider a denser topology Ta2, which consists of 65 nodes and 1869 flows. Similarly, we select 1500 flows uniformly at random and repeat each experiment 10 times. We present the results in Fig. 6(a). The simulation results show that both RP-MCA and VOL-MCA have a similar performance, which is also close to the optimal. An interesting observation is that all the considered algorithms tend to select the same set of nodes, which results in a similar performance of the RP-MCA and VOL-MCA as both of them use the same capacity allocation algorithm. In addition, the GCA algorithm performs worse than the MCA algorithm for this particular instance, which results in a worst overall performance under the RP-GCA algorithm. The maximum total processed traffic reaches around 92% and saturates at that point. The reason is that while all the nodes on the paths of the flows have been selected, they do not have sufficient processing capacity to process all the flows. In this case, selecting more other nodes (due to an increased budget) does not help. This motivates us to study the impact of increasing the VNF-node capacity in Fig. 6(b). We consider a fixed budget of \$2M and vary the capacity of each node between 1 Gbps and 2.8 Gbps with an increment of 200 Mbps. We can see that as the capacity increases, the total processed traffic increases and can reach around 100% when the capacity is more than 2 Gbps.

## VIII. CONCLUSION AND FUTURE WORK

In this paper, we studied the problem of deploying VNF-nodes and allocating their capacities. We showed how to overcome the non-submodularity of the problem by introducing a novel relaxation method. By utilizing a decomposition of the problem and a novel network flow reformulation, we were able to prove submodularity of the relaxed placement subproblem and develop efficient algorithms with constant approximation ratios for the original problem. Through extensive evaluations based on trace-driven simulations, we showed that the pro-

posed algorithms have a performance close to the optimal solution and better than a heuristic algorithm.

Our work also raises several interesting questions that are worth investigating as future work. First, we would like to consider the problem of joint placement, routing, and capacity allocation and investigate the impact of dynamic routing on the objective function. Despite the expected additional challenges due to routing, especially with integral resources and service function chaining as shown in [9], we plan to further investigate whether our proposed framework can be extended to more general settings with joint placement and routing. In this regard, recent work in [32] provides multi-criteria approximation algorithms for minimum cost joint placement and integral routing of service function chains under storage, computation, and communication capacity constraints. Understanding if our framework can be used to provide efficient approximations in similar settings with budget constraints is of interest for future work. Second, while we have assumed that each VNF-node has a fixed capacity in this paper, it is worth studying the optimization of the VNF-node capacity. A straightforward extension is to consider the setting where each VNF-node hosts multiple servers, each of which has a different cost and a different processing capacity. Our proposed framework can be modified to accommodate the extension: we consider multiple replications of each node, called virtual nodes, each of which corresponds to a server with its own cost and processing capacity. The path of each flow now needs to be updated to include the corresponding virtual nodes as well. We further elaborate on the proposed extension and present some additional simulation results in our online technical report [33].

## APPENDIX A PROOF OF THEOREM 1

*Proof:* We start by proving that the allocation subproblem ( $P2$ ) is NP-hard. The proof is by a reduction from a special case of the single knapsack (SK) problem, where for each item the profit and the weight are identical. In the SK problem, we have a knapsack  $k$  and a set of items  $\mathcal{I}$ . The knapsack has a capacity  $W$ , and each item  $i \in \mathcal{I}$  has a weight of  $w_i$ , which is the same as the profit. The objective is to find a subset of items  $\mathcal{I}' \subseteq \mathcal{I}$  that has the maximum total profit and can be packed in the knapsack without exceeding its capacity. Given an arbitrary instance  $\mathcal{A} = (k, \mathcal{I})$  of the SK problem, we construct an instance  $\mathcal{D} = (\mathcal{V}, \mathcal{F})$  of the allocation problem

(P2). The set  $\mathcal{V}$  has only one node  $v_1$  with a capacity that is equal to the capacity of the knapsack  $k$ . Each flow  $f \in \mathcal{F}$  corresponds to an item in  $i \in \mathcal{I}$ . A flow  $f$  has a traffic rate  $\lambda_f$  that is equal to the corresponding item weight  $w_i$ . Node  $v_1$  is the only VNF-node, and all the flows traverse node  $v_1$ . If we can solve the instance  $\mathcal{D}$  of problem (P2), the subset of flows assigned to node  $v_1$ , which has the maximum total traffic rate, can be mapped to the corresponding items and solve the instance  $\mathcal{A}$  of SK problem. Similarly, a solution for instance  $\mathcal{A}$  of the SK problem can be mapped to a solution for instance  $\mathcal{D}$  by simply mapping the selected items  $\mathcal{I}'$  to the corresponding flows that solve the instance  $\mathcal{D}$  of problem (P2).

Next, we prove the NP-hardness of the placement subproblem (P3). The proof is by a reduction from the budgeted maximum coverage (BMC) problem. In the BMC problem, we have a set of points  $\mathcal{M}$  and a set of candidate locations  $\mathcal{S}$ . Each point  $m \in \mathcal{M}$  has a weight of  $w_m$ . Each location  $s \in \mathcal{S}$  has a cost of  $b_s$  and covers a subset of points  $\mathcal{M}_s \subseteq \mathcal{M}$ . The objective is to select a subset of locations  $\mathcal{S}' \subseteq \mathcal{S}$  such that the total weight of the points covered by at least one location in  $\mathcal{S}'$  is maximized while the total cost of the selected locations does not exceed a given budget  $B$ . Given an arbitrary instance  $\mathcal{A} = (\mathcal{M}, \mathcal{S}, B)$  of BMC, we will construct an instance  $\mathcal{D} = (\mathcal{F}, \mathcal{V}, B)$  of problem (P3) as follows. Each flow  $f \in \mathcal{F}$  corresponds to a point  $m \in \mathcal{M}$ ; the rate of a flow is equal to the weight of the corresponding point. Each node  $v \in \mathcal{V}$  corresponds to a location  $s \in \mathcal{S}$ ; the cost of a node is equal to that of the corresponding location. The path of a flow consists of the nodes corresponding to the locations that cover the point corresponding to this flow. The deployment budget of the instance  $\mathcal{D}$  is equal to the budget of the instance  $\mathcal{A}$ . All the nodes have an infinite capacity. We will show that a solution for the instance  $\mathcal{D}$  exists if and only if a solution for the instance  $\mathcal{A}$  exists. If we can solve the instance  $\mathcal{A}$  of BMC, the subset of locations  $\mathcal{S}' \subseteq \mathcal{S}$  that solves  $\mathcal{A}$  of BMC can be mapped to the corresponding nodes in  $\mathcal{V}$  to become VNF-nodes and solves the instance  $\mathcal{D}$  of problem (P2). Similarly, if we solve the instance  $\mathcal{D}$ , then the obtained set of VNF-nodes can be mapped to the corresponding subset of locations that solve the instance  $\mathcal{A}$  of BMC.  $\square$

## APPENDIX B PROOF OF LEMMA 1

*Proof:* Recall that  $R_3(\mathcal{U}) = \max_{\lambda \in \Lambda^{\mathcal{U}}} R_2^{\mathcal{U}}(\lambda)$ , where  $\Lambda^{\mathcal{U}}$  is the set of assignments satisfying the capacity constraint (2) and the flow rate constraint (7). It suffices to show the following:

- (A) for any assignment  $\lambda \in \Lambda^{\mathcal{U}}$ , one can construct an  $s$ - $\mathcal{V}$  flow  $\varphi \in \overline{\mathcal{F}}$  such that  $R_2^{\mathcal{U}}(\lambda) = \Phi(\mathcal{N}, \mathcal{U}) - \Phi(\mathcal{U}, \mathcal{N})$ ;
- (B) for any  $s$ - $\mathcal{V}$  flow  $\varphi \in \overline{\mathcal{F}}$ , one can construct an assignment  $\lambda \in \Lambda^{\mathcal{U}}$  such that  $R_2^{\mathcal{U}}(\lambda) = \Phi(\mathcal{N}, \mathcal{U}) - \Phi(\mathcal{U}, \mathcal{N})$ .

Note that Part (A) implies  $\max_{\lambda \in \Lambda^{\mathcal{U}}} R_2^{\mathcal{U}}(\lambda) \leq \max_{\varphi \in \overline{\mathcal{F}}} (\Phi(\mathcal{N}, \mathcal{U}) - \Phi(\mathcal{U}, \mathcal{N}))$  and Part (B) implies  $\max_{\lambda \in \Lambda^{\mathcal{U}}} R_2^{\mathcal{U}}(\lambda) \geq \max_{\varphi \in \overline{\mathcal{F}}} (\Phi(\mathcal{N}, \mathcal{U}) - \Phi(\mathcal{U}, \mathcal{N}))$ , which lead to Eq. (9).

We first show Part (A). For any assignment  $\lambda \in \Lambda^{\mathcal{U}}$ , we construct a function  $\varphi \in \overline{\mathcal{F}}$  in the following manner:

- (i) set  $\varphi(s, f) = \sum_{v \in \mathcal{V}_f} \lambda_f^v$  for each edge  $(s, f) \in \mathcal{L}_1$ ;
- (ii) set  $\varphi(f, v') = \lambda_f^v$  for each edge  $(f, v') \in \mathcal{L}_2$ ;
- (iii) set  $\varphi(v', v) = \sum_{f \in \mathcal{F}} \lambda_f^v$  for each edge  $(v', v) \in \mathcal{L}_3$ .

Note that  $\lambda_f^v = 0$  for all  $v \notin \mathcal{U}$ . It is easy to verify that constraints (2) and (7) imply that the constructed function  $\varphi$  is an  $s$ - $\mathcal{V}$  flow. Further, the following is also satisfied:

$$\begin{aligned}
 R_2^{\mathcal{U}}(\lambda) &\stackrel{(a)}{=} \sum_{f \in \mathcal{F}} \sum_{v \in \mathcal{V}_f \cap \mathcal{U}} \lambda_f^v \\
 &\stackrel{(b)}{=} \sum_{f \in \mathcal{F}} \sum_{v \in \mathcal{V}_f} \lambda_f^v \\
 &\stackrel{(c)}{=} \sum_{f \in \mathcal{F}} \varphi(s, f) \\
 &\stackrel{(d)}{=} \sum_{f \in \mathcal{N}_{\mathcal{F}}} \varphi(s, f) \\
 &\stackrel{(e)}{=} \Phi(\{s\}, \mathcal{N}) - \Phi(\mathcal{N}, \{s\}) \\
 &\stackrel{(f)}{=} \Phi(\mathcal{N}, \mathcal{N}_{\mathcal{V}}) - \Phi(\mathcal{N}_{\mathcal{V}}, \mathcal{N}) \\
 &\stackrel{(g)}{=} \Phi(\mathcal{N}, \mathcal{U}) - \Phi(\mathcal{U}, \mathcal{N}),
 \end{aligned} \tag{13}$$

where (a) is from the definition of  $R_2^{\mathcal{U}}(\lambda)$ , (b) is from  $\lambda_f^v = 0$  for all  $v \notin \mathcal{U}$ , (c) is from (i), (d) is from the one-to-one mapping between  $\mathcal{F}$  and  $\mathcal{N}_{\mathcal{F}}$ , (e) is from the definition of  $\Phi(\{s\}, \mathcal{N})$  and  $\Phi(\mathcal{N}, \{s\}) = 0$ , (f) holds because the net-flow at the source  $s$  plus the net-flow at the sinks is equal to zero, and (g) holds because no flow goes to the sinks in  $\mathcal{N}_{\mathcal{V}} \setminus \mathcal{U}$ .

We now show Part (B). For any  $s$ - $\mathcal{V}$  flow  $\varphi \in \overline{\mathcal{F}}$ , we first obtain another  $s$ - $\mathcal{V}$  flow  $\varphi' \in \overline{\mathcal{F}}$  by deleting all the flow going to the sinks in  $\mathcal{N}_{\mathcal{V}} \setminus \mathcal{U}$ . Note that this procedure does not change the net-flow at the sinks in  $\mathcal{U}$ , i.e.,  $\Phi'(\mathcal{N}, \mathcal{U}) - \Phi'(\mathcal{U}, \mathcal{N}) = \Phi(\mathcal{N}, \mathcal{U}) - \Phi(\mathcal{U}, \mathcal{N})$ , where  $\Phi'$  corresponds to  $\varphi'$ . Then, we construct an assignment  $\lambda \in \Lambda^{\mathcal{U}}$  by simply setting  $\lambda_f^v = \varphi'(f, v')$  for every  $f \in \mathcal{F}$  and every  $v \in \mathcal{U}$ . It is easy to verify that the definition of the  $s$ - $\mathcal{V}$  flow implies that constraints (2) and (7) are satisfied for assignment  $\lambda$ . Finally, following the same steps in Eq. (13), we can show  $R_2^{\mathcal{U}}(\lambda) = \Phi'(\mathcal{N}, \mathcal{U}) - \Phi'(\mathcal{U}, \mathcal{N})$ , and thus,  $R_2^{\mathcal{U}}(\lambda) = \Phi(\mathcal{N}, \mathcal{U}) - \Phi(\mathcal{U}, \mathcal{N})$ .

Combining Parts (A) and (B) completes the proof.  $\square$

## APPENDIX C PROOF OF LEMMA 4

*Proof:* In Phase I, the algorithm repeatedly alternates between Step 1 and Step 2. Step 1 will be repeated  $m$  times, where  $m \leq |\mathcal{V}|$ . In each repetition of Step 1, the capacity allocation of a given VNF-node  $v$  is modified, and the value of  $OPT(Q2, \mathcal{U})$  is reduced by at most  $\frac{1}{2}r_v$ . In Step 2, the perturbation does not change the value of the total assigned traffic. Therefore, the following is satisfied after Phase I:

$$\begin{aligned}
 \pi_{\text{MCA}}^{\mathcal{U}} &\geq OPT(Q2, \mathcal{U}) - \frac{1}{2} \sum_{i=1}^m r_{v_i} \\
 &\geq OPT(Q2, \mathcal{U}) - \frac{1}{2} OPT(Q2, \mathcal{U}) \\
 &= \frac{1}{2} OPT(Q2, \mathcal{U}).
 \end{aligned} \tag{14}$$

That is, the total traffic assigned to VNF-nodes  $\mathcal{U}$  after Phase I is  $\pi_{\text{MCA}}^{\mathcal{U}} \geq \frac{1}{2}OPT(Q2, \mathcal{U})$ . In Phase II, the total assigned traffic will either increase or remain the same in the worst case. Therefore, the result of the Lemma follows.  $\square$

#### APPENDIX D PROOF OF LEMMA 5

*Proof:* Let  $\mathcal{F}' \subseteq \mathcal{F}_{\mathcal{U}}$  denote the set of unassigned flows after the end of algorithm 3 and  $\mathcal{U}' = \cup_{f \in \mathcal{F}' \setminus \mathcal{U}} \mathcal{U}_f$  be the set of candidate VNF-nodes for the unassigned flows  $\mathcal{F}'$ . We remind the reader that we use  $c_u$  (resp.  $c_{\mathcal{U}_i}$ ) to denote the capacity of VNF-node  $u$  (resp. VNF-nodes  $\mathcal{U}_i$ ). Similarly, we use  $r_u$  (resp.  $r_{\mathcal{U}_i}$ ) to denote the total traffic assigned to VNF-node  $u$  (resp. VNF-nodes  $\mathcal{U}_i$ ). In Lemma 7, we start by showing that the total assigned traffic to any VNF-node in  $\mathcal{U}'$  is at least half of its total capacity, i.e.,  $r_u \geq \frac{1}{2}c_u$  for any VNF-node  $u$  in  $\mathcal{U}'$ , and use that to prove the  $1/3$  approximation ratio of Lemma 5.

**Lemma 7:** *After Phase I of the GCA algorithm, it holds that for any VNF-node  $u$  in  $\mathcal{U}'$ ,  $r_u \geq \frac{1}{2}c_u$ .*

*Proof:* We will prove this by contradiction. Let say, for the sake of contradiction, that there is an unassigned flow  $f$  for which there is a VNF-node  $u$  in  $\mathcal{U}_f$  such that  $r_u < \frac{1}{2}c_u$ . This means the flows assigned to VNF-node  $u$  have traffic rate less than  $\frac{1}{2}c_u$ . Furthermore, since flow  $f$  was not assigned to VNF-node  $u$ , its traffic rate has to be greater than half the capacity of node  $u$ , i.e.,  $\lambda_f > \frac{1}{2}c_u$ . However, that contradicts our algorithm where flows with the highest traffic rate are considered first, and thus flow  $f$  would be assigned to VNF-node  $u$  instead of some of the already assigned flows.  $\square$

Next, the proof of Lemma 5 proceeds as follows. The maximum traffic that can be assigned by any algorithm to VNF-nodes  $\mathcal{U}$  has the following upper bound:

$$\begin{aligned} OPT(Q2, \mathcal{U}) &\stackrel{(a)}{\leq} \pi_{\text{GCA}}^{\mathcal{U}} + c_{\mathcal{U}'} \\ &\stackrel{(b)}{\leq} \pi_{\text{GCA}}^{\mathcal{U}} + 2r_{\mathcal{U}'} \\ &\stackrel{(c)}{\leq} \pi_{\text{GCA}}^{\mathcal{U}} + 2\pi_{\text{GCA}}^{\mathcal{U}} \\ &= 3\pi_{\text{GCA}}^{\mathcal{U}}, \end{aligned} \quad (15)$$

where (a) holds because the maximum traffic that can be assigned by an optimal solution is at most the sum of the traffic of the assigned flows, which is  $\pi_{\text{GCA}}^{\mathcal{U}}$ , and the maximum possible traffic that can be assigned for the unassigned flows, which is  $c_{\mathcal{U}'}$ ; (b) holds from Lemma 7 because the total traffic assigned to VNF-nodes in  $\mathcal{U}'$  by Algorithm 3 is at least half of their total capacity, i.e.,  $r_{\mathcal{U}'} \geq \frac{1}{2}c_{\mathcal{U}'}$ ; (c) holds because  $r_{\mathcal{U}'}$  is upper bounded by  $\pi_{\text{GCA}}^{\mathcal{U}}$ . This completes the proof.  $\square$

#### APPENDIX E PROOF OF LEMMA 6

*Proof:* We first present Lemma 8, but before that we repeat Assumption 1 here to ease the proof navigation.

**Assumption 1:** *Assume that all the VNF-nodes in  $\mathcal{U}$  have the same capacity and that every flow  $f$  in  $\mathcal{F}_{\mathcal{U}}$  traverses at least two VNF-nodes in  $\mathcal{U}$ , i.e.,  $|\mathcal{V}_f \cap \mathcal{U}| \geq 2$ .*

**Lemma 8:** *If Assumption 1 holds, then for any unassigned flow  $f \in \mathcal{F}'$ , it holds that for any pair of VNF-nodes  $(u, v)$  in  $\mathcal{U}_f$ ,  $r_{\{u,v\}} \geq \frac{2}{3}c_{\{u,v\}}$ .*

*Proof:* Since the capacity of all VNF-nodes is the same by Assumption 1, we will use the symbol  $c$  to denote the capacity of any VNF-node. We prove this lemma by contradiction. Assume that there is an unassigned flow  $f$  for which there is a pair of VNF-nodes  $(u, v)$  in  $\mathcal{U}_f$  such that  $r_{\{u,v\}} < \frac{2}{3}c_{\{u,v\}}$ . In this case, the rate of flow  $f$  has to be greater than  $\frac{2}{3}c$ ; otherwise, it would fit on the combined remaining capacities of VNF-nodes  $u$  and  $v$  and would be assigned in Phase II of the algorithm. However, this also means that the flows assigned to VNF-nodes  $u$  and  $v$  have a rate that is less than  $\frac{2}{3}c$ , which contradicts our algorithm where flows with larger traffic rate will be considered first, and if possible get assigned.  $\square$  The rest of the proof of Lemma 6 follows the same argument as in the proof of Lemma 5 with the difference that we have  $r_{\mathcal{U}'} \geq \frac{2}{3}c_{\mathcal{U}'}$  by Lemma 8. The other parts of the proof are the same. That is,

$$\begin{aligned} OPT(Q2, \mathcal{U}) &\leq \pi_{\text{GCA}}^{\mathcal{U}} + c_{\mathcal{U}'} \\ &\leq \pi_{\text{GCA}}^{\mathcal{U}} + \frac{3}{2}r_{\mathcal{U}'} \\ &\leq \pi_{\text{GCA}}^{\mathcal{U}} + \frac{3}{2}\pi_{\text{GCA}}^{\mathcal{U}} \\ &\leq \frac{5}{2}\pi_{\text{GCA}}^{\mathcal{U}}. \end{aligned} \quad (16)$$

This completes the proof.  $\square$

#### APPENDIX F COMPLEXITY ANALYSIS

In this section, we analyze the complexity of the algorithms presented in Table I. Each of the proposed algorithms has two sequential components: placement and capacity allocation. We analyze the complexity of each component in the following.

**Complexity of the Placement Algorithms.** We have two placement algorithms: the SG algorithm and the EG algorithm. For the SG algorithm, in each iteration, we select a new node, which requires  $O(V)$  functions evaluations. Since we can select at most  $k$  nodes (due to the budget limit), we need  $O(kV)$  function evaluations in total. For the EG algorithm, we have two phases: in Phase I, we need  $O(V^2)$  function evaluations to evaluate all subsets of size one or two; in Phase II, we need  $O(V^5)$  function evaluations to evaluate all subsets of size three and augment each subset in a greedy manner. Hence, the overall complexity of the EG algorithm is  $O(V^5)$ .

**Complexity of the Capacity Allocation Algorithms.** We have two capacity allocation algorithms: the MCA algorithm and the GCA algorithm. For the MCA algorithm, we have two phases. In Phase I, we start by solving a maximum flow problem for a graph with  $O(F + V)$  vertices, which has a complexity of  $O(F^3)$  if solved using the Push-relabel algorithm [25]. Then, the algorithm proceeds by repeatedly alternating between implementing Step 1 and Step 2. Step 1 is executed for at most  $V$  times; each execution of Step 1 has a complexity of  $O(F)$ . Therefore, the overall complexity of Step 1 is  $O(FV)$ . For Step 2, we remove at least one edge at a time using the edge perturbation. Since the total number of edges is at most  $O(FV)$ , each execution of Step 2 has a complexity of  $O(FV)$ . Hence, the complexity of repeating Step 2 for at most  $O(FV)$  edges is  $O(F^2V^2)$ . The overall complexity

of Phase I is  $O(F^2V^2)$  as Step 2 dominates. In Phase II, for each of the (at most  $F$ ) unassigned flows, the algorithm tries to assign it to a subset of nodes in  $\mathcal{V}$ , so Phase II has a complexity of  $O(FV)$ . The overall complexity of the MCA algorithm is  $O(F^2V^2)$  as Phase I dominates. For the GCA algorithm, in Phase I, the sorting operation has a complexity of  $O(F \log F)$ , and assigning each flow to only one VNF-node has a complexity of  $O(FV)$ . Phase II is the same as that of the MCA algorithm, which has a complexity of  $O(FV)$ . Therefore, the complexity of the GCA algorithm is  $O(FV)$ .

## REFERENCES

- [1] M. Chiosi *et al.*, "Network functions virtualisation: An introduction, benefits, enablers, challenges and call for action," in *Proc. SDN OpenFlow World Congress*, vol. 48, 2012, pp. 1–16.
- [2] K. Poularakis, G. Iosifidis, G. Smaragdakis, and L. Tassiulas, "One step at a time: Optimizing SDN upgrades in ISP networks," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, May 2017, pp. 1–9.
- [3] Amdocs, "Bringing NFV to life-technological and operational challenges in implementing NFV," White Paper 1, 2016.
- [4] Y. Sang, B. Ji, G. R. Gupta, X. Du, and L. Ye, "Provably efficient algorithms for joint placement and allocation of virtual network functions," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, May 2017, pp. 1–9.
- [5] T. Lukovszki, M. Rost, and S. Schmid, "Approximate and incremental network function placement," *J. Parallel Distrib. Comput.*, vol. 120, pp. 159–169, Oct. 2018.
- [6] D. K. Hong, Y. Ma, S. Banerjee, and Z. M. Mao, "Incremental deployment of SDN in hybrid enterprise and ISP networks," in *Proc. Symp. SDN Res.*, Mar. 2016, pp. 1–7.
- [7] M. Shi, X. Lin, S. Fahmy, and D.-H. Shin, "Competitive online convex optimization with switching costs and ramp constraints," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, Apr. 2018, pp. 1835–1843.
- [8] H. Feng, J. Llorca, A. M. Tulino, and A. F. Molisch, "Optimal dynamic cloud network control," *IEEE/ACM Trans. Netw.*, vol. 26, no. 5, pp. 2118–2131, Oct. 2018.
- [9] H. Feng, J. Llorca, A. M. Tulino, D. Raz, and A. F. Molisch, "Approximation algorithms for the NFV service distribution problem," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, May 2017, pp. 1–9.
- [10] Y. Chen, J. Wu, and B. Ji, "Virtual network function deployment in tree-structured networks," in *Proc. IEEE 26th Int. Conf. Netw. Protocols (ICNP)*, Sep. 2018, pp. 132–142.
- [11] A. Tomassilli, F. Giroire, N. Huin, and S. Perennes, "Provably efficient algorithms for placement of service function chains with ordering constraints," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, Apr. 2018, pp. 774–782.
- [12] T. Lukovszki and S. Schmid, "Online admission control and embedding of service chains," in *Structural Information and Communication Complexity*. Cham, Switzerland: Springer, 2015, pp. 104–118.
- [13] G. Sallam, G. R. Gupta, B. Li, and B. Ji, "Shortest path and maximum flow problems under service function chaining constraints," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, Apr. 2018, pp. 2132–2140.
- [14] T. He, H. Khamfroush, S. Wang, T. La Porta, and S. Stein, "It's hard to share: Joint service placement and request scheduling in edge clouds with sharable and non-sharable resources," in *Proc. IEEE 38th Int. Conf. Distrib. Comput. Syst. (ICDCS)*, Jul. 2018, pp. 365–375.
- [15] M. Kablan, A. Alsdais, E. Keller, and F. Le, "Stateless network functions: Breaking the tight coupling of state and processing," in *Proc. 14th USENIX Symp. Netw. Syst. Design Implement. (NSDI)*, 2017, pp. 97–112.
- [16] K. Poularakis, J. Llorca, A. M. Tulino, I. Taylor, and L. Tassiulas, "Service placement and request routing in MEC networks with storage, computation, and communication constraints," *IEEE/ACM Trans. Netw.*, vol. 28, no. 3, pp. 1047–1060, Jun. 2020.
- [17] M. Rost and S. Schmid, "Virtual network embedding approximations: Leveraging randomized rounding," *IEEE/ACM Trans. Netw.*, vol. 27, no. 5, pp. 2071–2084, Oct. 2019.
- [18] B. Nemeth, Y.-A. Pignolet, M. Rost, S. Schmid, and B. Vass, "Cost-efficient embedding of virtual networks with and without routing flexibility," in *Proc. IFIP Netw. Conf. (Netw.)*, Jun. 2020, pp. 476–484.
- [19] G. L. Nemhauser and L. A. Wolsey, "Maximizing submodular set functions: Formulations and analysis of algorithms," in *North-Holland Mathematics Studies*. Amsterdam, The Netherlands: Elsevier, 1981, vol. 59, pp. 279–301.
- [20] S. Khuller, A. Moss, and J. Naor, "The budgeted maximum coverage problem," *Inf. Process. Lett.*, vol. 70, no. 1, pp. 39–45, Apr. 1999.
- [21] A. Das and D. Kempe, "Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection," 2011, *arXiv:1102.3975*. [Online]. Available: <http://arxiv.org/abs/1102.3975>
- [22] L. Chen, M. Feldman, and A. Karbasi, "Weakly submodular maximization beyond cardinality constraints: Does randomization help greedy?" 2017, *arXiv:1707.04347*. [Online]. Available: <http://arxiv.org/abs/1707.04347>
- [23] M. Feldman and R. Izsak, "Constrained monotone function maximization and the supermodular degree," 2014, *arXiv:1407.6328*. [Online]. Available: <http://arxiv.org/abs/1407.6328>
- [24] G. Sallam, Z. Zheng, and B. Ji, "Placement and allocation of virtual network functions: Multi-dimensional case," in *Proc. IEEE 27th Int. Conf. Netw. Protocols (ICNP)*, Oct. 2019, pp. 1–11.
- [25] A. V. Goldberg and R. E. Tarjan, "Efficient maximum flow algorithms," *Commun. ACM*, vol. 57, no. 8, pp. 82–89, Aug. 2014.
- [26] F. Bach, "Learning with submodular functions: A convex optimization perspective," *Found. Trends Mach. Learn.*, vol. 6, nos. 2–3, pp. 145–373, 2013.
- [27] M. Dawande *et al.*, "Approximation algorithms for the multiple knapsack problem with assignment restrictions," *J. Combinat. Optim.*, vol. 4, no. 2, pp. 171–186, 2000.
- [28] W. Li, "Nearly linear time deterministic algorithms for submodular maximization under knapsack constraint and beyond," 2018, *arXiv:1804.08178*. [Online]. Available: <http://arxiv.org/abs/1804.08178>
- [29] *Gurobi*. Accessed: Feb. 14, 2021. [Online]. Available: <https://www.gurobi.com/>
- [30] *Abilene Dataset*. Accessed: Feb. 14, 2021. [Online]. Available: <https://www.cs.utexas.edu/~yzhang/research/AbileneTM/>
- [31] S. Orłowski, R. Wessäly, M. Pióro, and A. Tomaszewski, "SNDlib 1.0—Survivable network design library," *Netw. Int. J.*, vol. 55, no. 3, pp. 276–286, 2010.
- [32] K. Poularakis, J. Llorca, A. M. Tulino, and L. Tassiulas, "Approximation algorithms for data-intensive service chain embedding," in *Proc. 21st Int. Symp. Theory, Algorithmic Found., Protocol Design Mobile Netw. Mobile Comput.*, Oct. 2020, pp. 131–140.
- [33] G. Sallam and B. Ji, "Joint placement and allocation of VNF nodes with budget and capacity constraints," 2019, *arXiv:1901.03931*. [Online]. Available: <http://arxiv.org/abs/1901.03931>



**Gamal Sallam** (Student Member, IEEE) received the B.Sc. degree in information technology and systems from Cairo University in 2011, the M.Sc. degree in computer networks from the King Fahd University of Petroleum and Minerals in 2016, and the Ph.D. degree in computer science from Temple University in 2020. His current research interest includes resource allocation in network function virtualization. He received the Outstanding Research Assistant Award from the Department of Computer and Information Sciences and from the College of

Science and Technology, Temple University, in 2018 and 2019, respectively.



**Bo Ji** (Senior Member, IEEE) received the B.E. and M.E. degrees in information science and electronic engineering from Zhejiang University, Hangzhou, China, in 2004 and 2006, respectively, and the Ph.D. degree in electrical and computer engineering from The Ohio State University, Columbus, OH, USA, in 2012. He is currently an Associate Professor with the Department of Computer Science, Virginia Tech, Blacksburg, VA, USA. Prior to joining Virginia Tech, he was an Associate/Assistant Professor with the Department of Computer and Information Sciences, Temple University, from July 2014 to July 2020. He was also a Senior Member of the Technical Staff with AT&T Labs, San Ramon, CA, USA, from January 2013 to June 2014. His research interests include modeling, analysis, control, and optimization of computer and network systems, such as wired and wireless networks, large-scale IoT systems, high performance computing systems and data centers, and cyber-physical systems. He is a member of the ACM. He received the National Science Foundation (NSF) CAREER Award in 2017 and the NSF CISE Research Initiation Initiative (CRII) Award in 2017. He was also a recipient of the IEEE INFOCOM 2019 Best Paper Award. He currently serves on the Editorial Board for the IEEE/ACM TRANSACTIONS ON NETWORKING, IEEE TRANSACTIONS ON NETWORK SCIENCE AND ENGINEERING, IEEE INTERNET OF THINGS JOURNAL, and IEEE OPEN JOURNAL OF THE COMMUNICATIONS SOCIETY.