



Multi-commodity Flow with In-Network Processing

Moses Charikar¹, Yonatan Naamad^{2(✉)}, Jenifer Rexford³, and X. Kelvin Zou⁴

¹ Stanford University, Stanford, CA, USA
moses@cs.stanford.edu

² Amazon.com, East Palo Alto, CA, USA
ynaamad@amazon.com

³ Princeton University, Princeton, NJ, USA
jrex@cs.princeton.edu

⁴ ByteDance, Seattle, WA, USA
xuanzou1991@gmail.com

Abstract. Modern networks run “middleboxes” that offer services ranging from network address translation and server load balancing to firewalls, encryption, and compression. In an industry trend known as Network Functions Virtualization (NFV), these middleboxes run as virtual machines on any commodity server, and the switches steer traffic through the relevant chain of services. Network administrators must decide how many middleboxes to run, where to place them, and how to direct traffic through them, based on the traffic load and the server and network capacity. Rather than placing *specific* kinds of middleboxes on each processing node, we argue that server virtualization allows each server node to host *all* middlebox functions, and simply vary the fraction of resources devoted to each one. This extra flexibility fundamentally changes the optimization problem the network administrators must solve to a new kind of multi-commodity flow problem, where the traffic flows consume bandwidth on the links as well as processing resources on the nodes. We show that allocating resources to maximize the processed flow can be optimized exactly via a linear programming formulation, and to arbitrary accuracy via an efficient combinatorial algorithm. Our experiments with real traffic and topologies show that a joint optimization of node and link resources leads to an efficient use of bandwidth and processing capacity. We also study a class of design problems that decide where to provide node capacity to best process and route a given set of demands, and demonstrate both approximation algorithms and hardness results for these problems.

Keywords: Multi-commodity flow · Middleboxes · Network Function Virtualization · Approximation algorithms · Hardness of approximation

Y. Naamad—This work was done while the author was at the Department of Computer Science, Princeton University.

1 Introduction

In addition to delivering data efficiently, modern networks often perform services on the traffic in flight to enhance security, privacy, or performance, or provide new features. Network administrators often install “middleboxes” such as firewalls, network address translators, server load balancers, Web caches, video transcoders, and devices that compress or encrypt the traffic. In fact, many networks have as many middleboxes as underlying routers or switches [29]. Often a single connection must traverse multiple middleboxes, and different connections may go through different sequences of middleboxes. For example, while Web traffic may go through a firewall followed by a server load balancer, video traffic may simply go through a transcoder. To keep up with the traffic demands, an organization may run multiple instances of the same middlebox. Deciding how many middleboxes to run, where to place them, and how to direct traffic through them, is a major challenge facing network administrators.

Until recently, each middlebox was a dedicated appliance, consisting of both software and hardware. Administrators typically installed these appliances at critical locations that naturally see most of the traffic, such as the gateway connecting a campus or company to the rest of the Internet. A network could easily have a long chain of these appliances at one location, forcing all connections to traverse every appliance—whether they need all of the services or not. In addition, placing middleboxes only at the gateway does not serve the organization’s many *internal* connections, unless the internal traffic is routed circuitously through the gateway.

Over the last few years, middleboxes have become increasingly *virtualized*, with the software service separate from the physical hardware—an industry trend called Network Functions Virtualization (NFV) [5, 22]. The network can “spin up” (or down) virtual machines on any physical server, as needed. This has led to a growing interest in good algorithms for optimizing the (i) *allocation* of middleboxes over a pool of server resources, (ii) *steering* of traffic through a suitable sequence of middleboxes based on a high-level policy, and (iii) *routing* of the traffic between the servers over efficient network paths [1, 16, 17, 19, 24].

Rather than solving these three optimization problems separately, we introduce—and solve—a joint optimization problem. Since server resources are fungible, we argue that each processing node could subdivide its resources *arbitrarily* across any of the middlebox functions, as needed. That is, the *allocation* problem is more naturally a question of what fraction of each node’s computational (or memory) resources to allocate to each middlebox function. Similarly, each connection can have its middlebox processing performed on any node, or set of nodes, that have sufficient resources. That is, the *steering* problem is more naturally a question of deciding which nodes should devote a share of their processing resources to a particular portion of the traffic. Hence, the joint optimization problem devolves to a new kind of *routing* problem, where we compute paths based on both the bandwidth and processing requirements of the traffic between each source-sink pair. That is, each flow from a source to a sink must

be allocated both (i) a certain amount of bandwidth on *every* link in its path and (ii) a *total* amount of computational across all of the nodes on its path.

In our *flow with in-network processing* problem, we have a flow demand with multiple sources and multiple sinks, and each flow requires a certain amount of processing. The required processing is proportional to the flow size and, without loss of generality, we assume one unit of flow requires one unit of processing. Each flow from a source to a sink is an aggregate flow of many connections, so the routing and processing for a flow are both divisible. In this model there are two types of constraints: edge capacity constraints and vertex capacity constraints, which represent link bandwidth and node processing capacity, respectively. A feasible flow pattern satisfies three conditions: (i) for each edge, the sum over all flows on that edge is bounded by that edge's capacity, (ii) for each node, the sum over all flows of in-network processing done at that node is bounded by the vertex capacity, and (iii) each flow must be allocated a total amount of node processing power equal to its size.

Although ignoring vertex capacity constraints reduces our class of problems to those of the standard multi-commodity flow variety, the introduction of these constraints yields a new class of problems that has not been studied before. This paper provides a systematic approach to this new class of network problems, applicable to both directed and undirected graphs.

In Sect. 2, we introduce the PROCESSED FLOW ROUTING class of problems, in which we discuss how to optimize processed flow routed in a fixed network. Next, we present two linear programming-based algorithms to find a maximum feasible multi-commodity flow with the additional processing constraints. We show that, like standard multi-commodity flow, the program can be written in two different equivalent ways: either with an exponentially-sized walk-based LP or with a polynomially-sized edge-based LP. The proof of equivalence of these two LPs requires a more careful argument than that for standard MCF. As an aside, we argue that this pair of LPs can also be adapted to optimize several other objective functions, such as those minimizing congestion. In Sect. 4, we briefly describe an experimental evaluation of this linear programming approach.

In Sect. 5, we design an efficient multiplicative weight update (MWU) algorithm that finds approximately optimal solutions to our walk-based linear program far more quickly than one could with the edge-based program paired with an off-the-shelf LP-solver.

In Sect. 6, we consider the MIDDLEBOX NODE PURCHASE class of problems, in which the goal is to optimally purchase processing power at various middleboxes. Prices for placing processing at the various nodes is given as part of the input, and may differ substantially from one location to the next. This class of problems comes has two natural variants:

1. MIN MIDDLEBOX NODE PURCHASE: given a set of flow demands, minimize cost while purchasing enough middlebox processing capacity so that all flow demands are simultaneously satisfiable (that is, jointly routable and steerable).

2. **BUDGETED MIDDLEBOX NODE PURCHASE**: given a set of flow demands and a budget of k dollars, spend at most $\$k$ on purchasing middlebox processing capacity while maximizing the fraction of the given demand that is simultaneously satisfiable.

Linear programs for both of these problems can be found in Sects. 6.2 and 6.4. For **MIN MIDDLEBOX NODE PURCHASE**, we show an $O(\log(n)/\delta^2)$ approximation for node costs and an associated multi-commodity flow that satisfies $(1 - \delta)$ fraction of the demands and satisfies all edge capacities, where n is the number of nodes. We show that in the directed case, the problem is hard to approximate better than a logarithmic factor, even if the demand requirements are relaxed. Additionally, we show that the undirected case is at least as hard to approximate as **VERTEX COVER**.

We also prove approximation and hardness results for **BUDGETED MIDDLEBOX NODE PURCHASE**. Although it's tempting to conjecture that the problem is an instance of **BUDGETED SUBMODULAR MAXIMIZATION**, one can construct instances on both directed and undirected graphs where the amount of routable processed flow is *not* submodular in the set of purchased nodes, so black-box submodular maximization techniques cannot be used here. We show an $\Omega(1/\log(n))$ approximation for both problems, as well as a constant factor approximation algorithm for undirected instances with a single source-sink pair. For the directed case, we show approximation hardness of $1 - 1/e$, and constant factor hardness for the undirected problem. Our results are summarized in the following Table 1

Table 1. Network design results

		Directed	Undirected
Budgeted	Approximation hardness	$\Omega(1/\log n)$.078 ^b
		$1 - 1/e$.999
Minimization	Approximation hardness	$O(\log n)$ ^a	$O(\log n)$ ^a
		$O(\log n)$	$2 - \epsilon$

^aAll demands are satisfied only up to an $(1 - \epsilon)$ fraction.

^bAssuming 1 source-sink pair. For multiple pairs, we adapt the $\Omega(1/\log n)$ -approximation digraph algorithm.

2 Flow Routing with In-Network Processing

2.1 Processed Flow Routing Problem

Network Function Virtualization (NFV) allows each node to function as a general-purpose server that can run any in-network processing task, such as transcoding, compression, and encryption. Such servers can reside anywhere in the network, from the leaf nodes (as in the case of traditional servers) to intermediary nodes (such as top-of-rack and spine switches).

Therefore, in our model, we treat all in-network processing as *homogeneous*, meaning that every node with a sufficient quantity of available computational




Fig. 1. The edge capacity is 10 for all edges and the node capacities are denoted in each node. Here, we can send maximum flow size 5, by routing it along the red arcs, have it processed at the nodes at the top, and then sent to T along the blue arcs. The capacity of the bottom middle edge forms the bottleneck here, as all flow must pass through it twice before reaching T . (Color figure online)

resources can be adapted to accomplish any processing task. In practice, this can be accomplished simply by spinning up a new virtual machine for that specific task as needed. We assume that all flows are both aggregate and sufficiently large that they can be treated as continuous quantities (and thus can be arbitrarily subdivided), and that the processing capacity of a given node can also be fractionally divided among a number of different flows.

Each flow is initially generated at a source fully unprocessed. By the time it reaches its destination, it needs to go through and get processed by one or more intermediate processing nodes with available computational resources. We assume that each unit of flow requires one unit of processing, meaning that for any given flow f , the total processing workload done on f by vertices along f 's flow path should equal the size f .

This problem can be modeled mathematically as follows. We are given an (un)directed graph $G = (V, E)$ along with edge capacities $B : E \rightarrow \mathbb{R}^+$, vertex capacities $C : V \rightarrow [0, \infty)$, and a collection of flows of varying commodities $D = \{(s_1, t_1, k_1), (s_2, t_2, k_2), \dots\} \subseteq V \times V \times \mathbb{R}^+$. While the edge capacities are used in the same way as in a standard multi-commodity flow problem, we also require that each unit of flow undergo a total of one unit of processing at intermediate vertices. In particular, while edge capacities limit the *total* amount of flow that may pass through an edge, vertex capacities only bottleneck the amount of processing that may be done at a given vertex, regardless of the total amount of flow that uses the vertex as an intermediate node. The goal is then either to route as much flow as possible, or to satisfy all flow demand subject to a congestion-minimization objective function. For concreteness, this paper focuses on maximizing the total amount of flow we can send between the source-destination pairs while satisfying edge and node capacity constraints. In practice we can also extend our results to other objective functions such as minimizing the weighted sum of congestion at edges and nodes. For ease of exposition, we focus our attention almost entirely on directed instances. The results for undirected graphs follow analogously (Table 2).

Table 2. Variables in the optimization solutions

Variable	Description
V	Set of nodes in a graph
E	Set of edges in a graph
$B(e)$	Edge capacity for edge e
$C(v)$	Node capacity for node v
D	The set of flow demands
$\delta^+(v)$	The edges leaving vertex v
$\delta^-(v)$	The edges entering vertex v
P	The set of 2-walks from sources to destinations
$p_{i,\pi}^v$	2-walk-based; the amount of flow i from s_i to t_i exactly using 2-walk π and processed at v
$f_i(e)$	Edge-based; the amount of flow i that traverses e on its way from s_i to t_i
$w_i(e)$	Edge-based; the amount of unprocessed flow i that traverses e on its way from s_i to t_i
$p_i(v)$	Edge-based; the amount of processing done at node v for the i th flow

2.2 A 2-Walk-Based Solution

We now describe a *2-walk-based* formulation of the problem. A *2-walk* from s to t is a route between s and t that visits each vertex (and thus each edge) at most *twice*.

The approach we take is analogous to *path-based* solutions for the traditional multicommodity flow (MCF) problem, with the key difference that, unlike paths, our 2-walks may visit vertices and edges more than once. Additionally, a 2-walk may traverse the same undirected edge in both directions.

To express the *2-walk-based* linear program, we introduce one variable $p_{i,\pi}^v$ for each {2-Walk}–vertex–demand triplet, representing the total amount of flow from s_i, t_i exactly utilizing walk π and processed at v . Note here the set P of 2-walks is an enumeration of all possible 2-walks in the graph, which can be *exponential* in size. The LP is then formulated as follows:

$$\begin{aligned}
& \text{MAXIMIZE} && \sum_{i=1}^{|D|} \sum_{\pi \in P} \sum_{v \in \pi} p_{i,\pi}^v \\
& \text{SUBJECT TO} && \\
& p_{i,\pi} = \sum_{v \in \pi} p_{i,\pi}^v && \forall i \in [|D|], \forall \pi \in P \\
& \sum_{i=1}^{|D|} \sum_{\substack{\pi \in P \\ \ni e}} p_{i,\pi} \leq B(e) && \forall e \in E \\
& \sum_{i=1}^{|D|} \sum_{\pi \in P} p_{i,\pi}^v \leq C(v) && \forall v \in V
\end{aligned}$$

$$\begin{aligned} \sum_{\pi \in P} \sum_{v \in \pi} p_{i,\pi}^v &\leq k_i && \forall i \in [|D|] \\ p_{i,\pi}^v &\geq 0 && \forall i \in [|D|], \forall \pi \in P, \forall v \in V \end{aligned}$$

While the first constraint enforces that all flows are fully processed, the second and third constraints ensure that no edge or vertex is over-saturated.

3 An Edge-Based Polynomially-Sized LP

Although the 2-walk-based solution exactly solves our MCF with in-network processing problem, the LP may be exponentially sized and thus even writing it down (let alone solving it) leaves us with an exponential worst-case running time. In Sect. 3.1, we present a polynomially-sized (and thus polytime-solvable) edge-based linear program for this problem. We then follow this up by a proof of correctness in Sect. 3.2.

3.1 The Edge-Based Solution

A standard technique for solving the traditional MCF problem relies on constructing a polynomially-sized *edge-based* LP whose set of feasible solutions equals that of an exponentially-sized path-based LP. Analogously, we establish a polynomial-sized edge-based LP corresponding to the *2-walk-based* LP introduced previously.

$$\text{MAXIMIZE} \quad \sum_{i=1}^{|D|} \sum_{e \in \delta^+(s_i)} f_i(e) \quad (2a)$$

$$\text{SUBJECT TO} \quad (2b)$$

$$\sum_{e \in \delta^-(v)} f_i(e) = \sum_{e \in \delta^+(v)} f_i(e) \quad \forall i \in [|D|], \forall v \in V \setminus \{s_i, t_i\} \quad (2c)$$

$$p_i(v) = \sum_{e \in \delta^-(v)} w_i(e) - \sum_{e \in \delta^+(v)} w_i(e) \quad \forall i \in [|D|], \forall v \in V \setminus \{s_i\} \quad (2d)$$

$$\sum_{i=1}^{|D|} f_i(e) \leq B(e) \quad \forall e \in E \quad (2e)$$

$$\sum_{i=1}^{|D|} p_i(v) \leq C(v) \quad \forall v \in V \quad (2f)$$

$$\sum_{e \in \delta^+(s_i)} f_i(e) \leq k_i \quad \forall i \in [D] \quad (2g)$$

$$w_i(e) \leq f_i(e) \quad \forall i \in [D], \forall e \in E \quad (2h)$$

$$w_i(e) = f_i(e) \quad \forall i \in [D], \forall e \in \delta^+(s_i) \quad (2i)$$

$$w_i(e) = 0 \quad \forall i \in [D], \forall e \in \delta^-(t_i) \quad (2j)$$

$$w_i(e), p_i(v) \geq 0 \quad \forall i \in [D], \forall e \in E \quad (2k)$$

The LP constraints can be interpreted as follows. Constraint (2c) is a flow conservation constraint: at any non-terminal node of flow i , the amount of flow i that enters the node equals the amount that leaves it. Constraint (2d) is a *processing conservation* constraint, ensuring that the total amount of flow (processed or unprocessed) going through a node remains unchanged, although the quantity of each might change if the node processes any of the flow. Constraints (2e) and (2f) ensure that we don't exceed edge and node capacities. Constraint (2g) ensures that we don't route more flow than is requested between any demand pair. Constraint (2h) ensures that the amount of work yet to be done on a flow does not exceed the size of the flow itself, while (2i) and (2j) ensure that all flows leave the sources unprocessed and arrive to the destinations fully processed.

3.2 Proof of Equivalence to the 2-Walk-Based LP

While the construction of the edge-based LP is not particularly difficult, it is not obvious that the edge-based solution actually solves the problem in question. We need to prove the *correctness* of the edge-based LP. A priori, solutions to the edge-based LP here may not be decomposable to a valid routing pattern at all. We provide an efficient algorithm converting feasible solutions to the edge-based LP into corresponding solutions to the 2-walk-based program, proving both that the edge-based LP is correct and that the actual flow paths can be recovered in polynomial time as well. We summarize this result in the following theorem.

Theorem 1. *The edge-based formulation provides a polynomial-sized linear program solving the Maximum Processed Flow problem. Further, the full routing pattern can be extracted from the LP solution by decomposing it into its composing s_i, t_i 2-walks in $O(|V| \cdot |E| \cdot |D| \cdot \log |V|)$ time.*

Notably, as the reduction maps the set of feasible solutions to the edge-based LP to equivalent feasible solutions of the 2-walk-based LP, the same technique can also be used to show the equivalence of the two corresponding programs when the objective function is changed to optimize some other linear quantity, such as the amount of congestion.

We now describe the algorithm in more detail and prove its correctness. The first part of the proof involves showing how to construct a solution to the flow-based LP when there is exactly one s_i, t_i pair. Extracting the corresponding flow paths and iterating this procedure for each demand pair eventually extracts all s_i, t_i flows, giving us a solution to the multicommodity problem.

The flow extraction argument proceeds in two steps. First, we simplify the solution by removing extraneous loops that do not affect the optimal solution. Next, we show that the existence of any residual flow in the graph (i.e., the existence of some strictly positive $f_i(e)$) implies that there exists at least one valid 2-walk we can efficiently extract while maintaining feasibility of all constraints for the updated residual graph. As we show, a linear number of extractions suffices to remove all flow from the solution. We provide a complete algorithm in Algorithm 1.

Algorithm 1. 2-Walk Decomposition**Data:** $G'(V, E)$, $w(e)$, $f(e)$ for $\forall e \in E$ and $p(v)$ for $\forall v \in V$ **Result:** $f(\pi)$, $p(\pi, v)$ with $v \in \pi$ **Algorithm** TwoWalkConstruction(s, t, v)//Construct 2-walk from $s \rightarrow v$ and $v \rightarrow t$ From v , run a backward traversal, each time picking an incoming edge e maximizing $\rho(e) = w(e)/f(e)$ From v , run a forward traversal, each time picking an outgoing edge minimizing $\rho(e) = w(e)/f(e)$.**return** π **Algorithm** FlowPlacement(s, t)**while** there exists a v with $p(v) > 0$ **do** $\pi \leftarrow$ TwoWalkConstruction(s, t, v) $f' \leftarrow \min_{e \in \pi, e \text{ precedes } v} f^1(e)$ $f'' \leftarrow \min_{e \in \pi, e \text{ succeeds } v} f^2(e)$ $p_\pi^v \leftarrow \min\{f', f'', p(v)\}$ **for** $u \in \pi$ and $u \neq v$ **do** $p_\pi^u = 0$ **end** $C(v) \leftarrow C(v) - p_\pi^v$ $p(v) \leftarrow p(v) - p_\pi^v$ **for** $e \in \pi$ **do** $f(e) \leftarrow f(e) - p_\pi^v$ $B(e) \leftarrow B(e) - p_\pi^v$ **end****end**

Removing Extraneous Loops. Suppose we are given a nonempty solution to the *edge-based* LP for an instance with graph $G(V, E)$. We focus on some (arbitrarily chosen) commodity i with positive flow in this solution, and drop subscripts to let $f(e)$, $w(e)$, and $p(v)$ denote $f_i(e)$, $w_i(e)$, and $p_i(v)$, respectively. To assist with our exposition, we restrict our attention to the subgraph G' which excludes all edges for which $f(e) = 0$. For each edge e in this subgraph, we also associate two new variables, $f^1(e)$ and $f^2(e)$ denoting the amount of unprocessed and processed flow passing through this edge, respectively. Thus, by definition, $f^1(e) = w(e)$ and $f^2(e) = f(e) - w(e)$.

As in solutions to the edge-based linear program for the standard multicommodity flow problem, solutions to our edge-based LP may introduce closed loops (that is, directed cycles along which a positive amount of flow is routed). In traditional MCF, such loops are easily shown to be *non-essential*, and can be easily removed from a feasible solution without affecting its correctness. As illustrated in Fig. 1, such loops may actually be critical in solutions to our variant, and handling such cases takes additional care. Thus, instead of arguing that cycles can be removed (so that the flows form a set of paths), we show how to ensure that no vertex may be visited more than twice (and thus the flows form a set of

2-walks). In particular, we show how to cancel out all cycles along which each edge contains f^1 flow, as well as all cycles along which each edge carries f^2 flow.

Lemma 1. *Any closed loop for which every edge contains f^1 (resp. f^2) flow can be removed without affecting the total (s, t) flow.*

Proof. This argument proceeds similarly to the flow cancellation arguments in the traditional MCF setting. For any loop l containing a positive amount of f_1 flow, reducing both $f(e)$ and $w(e)$ on the constituent edges by $\min_{e' \in l} f^1(e')$ ensures that all constraints in the LP remain satisfied. For loops containing a positive amount of f_2 flow, similarly reducing just $f(e)$ suffices.

Extracting 2-Walks. Suppose extraneous loops have been removed using the process described in Lemma 1. Define $\rho_e = \frac{w(e)}{f(e)} = \frac{f^1(e)}{f^1(e)+f^2(e)}$. By Lemma 1, every cycle with a positive $f(e)$ on each edge contains at least one edge with $\rho = 1$ and another with $\rho = 0$. We now repeat the following until all flow is removed from the graph. Select a vertex v that is allocated processing (i.e., $p(v) > 0$), and run a backwards traversal from v , at each step selecting the incoming edge with the largest fraction of unprocessed flow (i.e., maximizing $\rho(e)$) until we reach s . Similarly, run a forward traversal from v to t along edges minimizing ρ . This route will be our “flow-2-walk”. The amount of flow routable along this flow-2-walk is the minimum of three quantities: (1) the smallest amount of unprocessed flow sent on each edge of the $s \rightsquigarrow v$ path, (2) the smallest amount of processed flow sent along each edge of the $v \rightsquigarrow t$ path, and (3) the amount of processing still to be done at v (i.e., $p(v)$). We then extract this flow-2-walk from the solution by decreasing each LP variable accordingly. Complete pseudo-code for this algorithm is given in Algorithm 1.

Lemma 2 (2-Walk Extraction). *Algorithm 1 can always generate a 2-walk with non-zero flow from source to sink if there exists any v where $p(v) > 0$. Further, the number of iterations needed of Algorithm 1 is bounded by $O(|E|)$, each of which can be made to take $O(|V| \log |V|)$ time. Thus, the total running time is $O(|E| \cdot |V| \log |V|)$.*

Proof. The removal of extraneous cycles guarantees that no 2-walk can visit the same vertex more than twice. Now suppose that a vertex v has $p(v) > 0$. By constraint (2d), the f^1 flow on some incoming edge and the f^2 flow on some outgoing edge must both be positive. By a combination of constraints (2c), (2d), and (2j), the reverse traversal from v to s must succeed: it cannot get “stuck” at a vertex u with no in-edge with positive f^1 flow. Similarly, the forward traversal from v to t must find a path with positive f^2 on each edge. Subtracting the minimum of all of the reverse path’s observed f^1 values, all of the forward path’s observed f^2 values, and $p(v)$ from each of those variables ensures that all variables remain nonnegative. Further, as this operation is monotone and it decreases one of the variables to 0, repeating this must remove all flow from the graph in at most $|V| + 2|E| = O(|E|)$ iterations. By initially constructing


a priority queue for each vertex on the f^1 , f^2 , and ρ values of its neighboring edges and updating them accordingly, the forward and backward traversals can be found in $|V| \log |V|$ time, each.

We can generalize the above approach to the multicommodity problem by treating each of the commodities independently. Namely, sequentially applying the above algorithm to remove flow 2-walks for each of the $|D|$ demand pair gives us a solution to the multicommodity problem without violating any of the LP constraints. Thus, we get the $O(|V| \cdot |E| \cdot |D| \cdot \log |V|)$ running time promised in the statement of Theorem 1.

4 Evaluations

We ran several experiments to address (i) how well the LP fares against “naive” algorithms, and (ii) the in-practice running time for an edge-based LP solution.

Throughput Improvement. To determine how well the LP fares against simple approaches, we compare it to a “naive” algorithm that first routes flow without vertex capacities in mind, and then processes as much flow as possible on the flow paths it initially routed. This is a variant of the *path-selection* approach used in [15]. While there are simple examples where the naive algorithm performs extremely poorly in theory, we seek to study the performance in practice.



(a) Amount of flow that the two algorithms could process given various node processing capacities.

(b) The ratio of the demanded flow processable by the naive algorithm to that processable by the LP, plotted at various processing capacities.

Fig. 2. Experimental results showing how much of the demand both the naive and the (optimal) LP-based algorithm could successfully route and process given the Abilene traffic matrices.

We ran both algorithms on 150 randomly sampled traffic matrices provided by the TOTEM project [31] for the Abilene network in 2004. As these datasets

don't include vertex processing capacities, we compared the two algorithms on a wide range of values, with processing capacities assigned according to one of two distributions: either they all have the same capacity (the *all* case) or exactly half of them have the prescribed capacity and the other half have zero (the *half* case). The results are diagrammed in Fig. 2.

Experimental analysis shows that while the LP and the “naive” algorithm fare similarly when the network is low on processing capacity and thus node-throttled, or, in the *all* case, high on node capacity and thus bottlenecked by the link capacities and the demand itself, the LP has a distinct advantage in between the two extremes when either resource could become the bottleneck when the flows are not routed efficiently. Additionally, the experiments show that the naive algorithm suffers when processing is not uniformly distributed among the nodes even in the high-capacity case, as many of the initial flow paths might go entirely through nodes without any processing capacity and thus fail to get processed. Our experiments show that using the exact algorithm gives an improvement of up to 30% over the naive approach if processing power is available at all nodes, and up to 80% if the processing power is only placed at half of the nodes.

Runtime Analysis. Although the edge-based LP provides a polynomial running-time guarantee, it may still be too slow in practice on large graphs. To study the empirical performance of linear programming, we also run the LP solver over a number of topologies acquired from SNDLib [23]. For each of the topologies in Table 3, enough processing capacity was evenly distributed among a random sample of half of all nodes so that the total processing capacity equals half of the total demand. Although the implementation, hardware, and choice of solver were not optimized for running time, the table below indicates that the time to solve the LP grows quickly with the input size.

Table 3. Time to solve the edge based LP for various topologies. All values are averaged over 15 runs of CoinLP [13] on a 3.3 GHz Intel i5 2500k processor.

Network	$ V $	$ E $	Time (sec)
abilene	12	15	1.91
dfn-bwin	10	45	3.08
atlanta	15	22	5.28
dfn-gwin	11	47	13.91
geant	22	36	23.69
france	25	45	44.38
india35	35	57	105.89

The cost of solving this LP even on small topologies justifies the use of the faster multiplicative weight algorithm instead. The MWU algorithm has a

running time of roughly $\tilde{O}(|D|*|E|^2/\epsilon^2)$, which on sparse graphs is roughly equal to just the number of variables in the edge-based LP (as opposed to the time needed to actually solve it). While the algorithm is only approximately optimal, choosing an appropriate value of ϵ (say, $\epsilon = 0.1$) can grant a better running time while still significantly outperforming the naive algorithm.

5 Multiplicative Weights Based Approximation Algorithm

We first briefly overview the MWU method in Sect. 5.1. Next, we describe how to apply the MWU method to our model including processing vertices. The proof of correctness is given in Sect. 5.3;

5.1 Multiplicative Weight Update for Traditional MCF

In the traditional multiplicative weights algorithm for multicommodity flow, there an “expert” is assigned to each edge, each of which is initially assigned a sufficiently small weight. The algorithm then iteratively finds s_i, t_i walks minimizing the sum of weighted utilization of their edges and adds together scaled down versions of these paths to eventually construct a solution. When a path is chosen, all experts corresponding to edges along the path have their weight increased by a multiplicative factor, making it less likely that we repeat our selection of the edges. This process is repeated until some expert’s weight surpasses the value 1, corresponding to a fully utilized edge. When this happens, all paths are scaled down by the weight of the largest expert to ensure that no capacities are exceeded. One then shows that the final result is within a $(1 - \epsilon)$ factor of the maximum multicommodity flow.

5.2 Formulation and Analysis

Although we derive the same $(1 - \epsilon)$ approximation factor for our problem, the analysis of our multiplicative weights algorithm is quite different from that of traditional multicommodity flow. Intuitively, this is because vertex capacities are inherently very different from edge capacities: while a flow 2-walk reduces the remaining capacity on *all* edges it traverses, it only reduces the capacity for *one* of its vertices. Thus, we set up a different update condition, as well as a different method for picking the best flow 2-walks for each round.

Setup. For each edge e , we have a constraint $\sum_{\pi} p_{\pi} \leq B(e)$, where p_{π} is the amount of flow sent on 2-walk π . For each vertex, the corresponding constraint is $\sum_{\pi} p_{\pi}^v \leq C(v)$, where p_{π}^v is the amount of flow on 2-walk π that is processed at v . For each of these two sets of constraints, we associate one expert, (which we call \hat{e} and \hat{v}), whose weights are denoted by $w_{\hat{e}}$ and $q_{\hat{v}}$, respectively.

Consider a feasible solution to the 2-walk-based LP. The feasible solution consists of variables of the form p_{π} and p_{π}^v . In this section, we abuse notation

and let the variable p denote a feasible solution to the LP, at which point p_π and p_π^v become bound variables for each π and v (that is, p can be thought of as a dictionary containing the aforementioned set of variables). Further, define $A(p)$ as the objective function value of p , i.e. $A(p) = \sum_{v \in V} \sum_{\pi \in P} p_\pi^v$.

For an expert \hat{e} and feasible solution p , define the gain $M(\hat{e}, p)$ by $M(\hat{e}, p) = \frac{1}{B(\hat{e})} \sum_{\pi \ni \hat{e}} p_\pi$. This can be thought of as the fraction of \hat{e} 's capacity actually utilized by the feasible solution. For each expert \hat{v} , we define the gain $M(\hat{v}, p)$ by $M(\hat{v}, p) = \frac{1}{C(\hat{v})} \sum_{\pi \ni \hat{v}} p_\pi^v$, which corresponds to the fractional utilization of \hat{v} 's processing capacity.

Let \mathcal{D} be the probability distribution over experts in which the probability of choosing a given expert is proportional to its weight. For a fixed p , the expected gain of a random variable sampled from \mathcal{D} is

$$M(\mathcal{D}, p) = \frac{\sum_e w_{\hat{e}} M(\hat{e}, p) + \sum_v q_{\hat{v}} M(\hat{v}, p)}{\sum_e w_{\hat{e}} + \sum_v q_{\hat{v}}}$$

We first make two observations:

Observation 1: For any feasible solution p , $0 \leq M(\mathcal{D}, p) \leq 1$. This is because $M(\hat{e}, p) \leq 1$ and $M(\hat{v}, p) \leq 1$ for all e and v .

Observation 2: For any feasible solution p and weights w, q , if $\pi^* = \operatorname{argmin}_\pi (\sum_{e \in \pi} w_{\hat{e}}/B(e) + \min_{\hat{v} \in \pi} q_{\hat{v}}/C(v))$, then

$$M(\mathcal{D}, p) \geq \frac{A(p) (\sum_{\hat{e} \in \pi^*} w_{\hat{e}}/B(e) + \min_{\hat{v} \in \pi^*} q_{\hat{v}}/C(v))}{\sum_e w_{\hat{e}} + \sum_v q_{\hat{v}}}$$

This is due to the fact that:

$$\begin{aligned} M(\mathcal{D}, p) &= \frac{\sum_e w_{\hat{e}} M(\hat{e}, p) + \sum_v q_{\hat{v}} M(\hat{v}, p)}{\sum_e w_{\hat{e}} + \sum_v q_{\hat{v}}} \\ &= \frac{\sum_\pi (p_\pi (\sum_e w_{\hat{e}}/B(e)) + \sum_{v \in \pi} p_\pi^v q_{\hat{v}}/C(v))}{\sum_e w_{\hat{e}} + \sum_v q_{\hat{v}}} \\ &\geq \frac{\sum_\pi (p_\pi (\sum_e w_{\hat{e}}/B(e) + \min_{\hat{v} \in \pi} q_{\hat{v}}/C(v)))}{\sum_e w_{\hat{e}} + \sum_v q_{\hat{v}}} \\ &\geq \frac{\sum_\pi p_\pi \min_\pi (\sum_{e \in \pi} w_{\hat{e}}/B(e) + \min_{\hat{v} \in \pi} q_{\hat{v}}/C(v))}{\sum_e w_{\hat{e}} + \sum_v q_{\hat{v}}} \\ &\geq \frac{A(p) (\sum_{\hat{e} \in \pi^*} w_{\hat{e}}/B(e) + \min_{\hat{v} \in \pi^*} q_{\hat{v}}/C(v))}{\sum_e w_{\hat{e}} + \sum_v q_{\hat{v}}} \end{aligned}$$

Where π^* is the path minimizing the argmin in the statement of the observation. Thus, in each round, we aim to find the π^* minimizing this value. Conditioned on us being able to do so, the rest of the MWU algorithm proceeds as follows:

1. We initialize all expert weights $\{w_{\hat{e}}\}$ and $\{q_{\hat{v}}\}$ to $1/\delta$, where $\delta = (1 + \epsilon)((1 + \epsilon) \cdot |E|)^{-1/\epsilon}$. This choice of δ will be justified in the analysis of Sect. 5.3.

2. At each step t , given weights w_e^t and q_v^t on the experts, we pick the flow-2-walk p^t minimizing the quantity $\sum_{e \in \pi} \frac{w_e}{B(e)} + \min_{v \in \pi} \frac{q_v}{C(v)}$. An efficient algorithm for finding such a 2-walk is given in Sect. 5.2.
3. Given the 2-walk p^t chosen in the previous step, we treat this as a feasible solution to the instance, giving expert \hat{j} a gain of $M(\hat{j}, p^t)$. Consequently, the weight $w_{\hat{e}}$ or $q_{\hat{v}}$ of each expert j is increased by a multiplicative factor of $M(\hat{j}, p^t)$.
4. The algorithm stops when one of the weights $w_{\hat{e}}$ or $q_{\hat{v}}$ is larger than 1. Once the algorithm terminates, we scale down the flow p^t computed at each round by a factor of $\log_{1+\epsilon} \frac{1+\epsilon}{\delta} = 1 - \frac{\ln \delta}{\ln 1+\epsilon}$, and return the set of all flow-2-walks p^t .

Computing the Optimal Path. To compute the 2-walk π^t with minimum cost, we use a dynamic programming algorithm reminiscent of Dijkstra's shortest path algorithm. Given a graph $G(V, E)$, with weights $w(e)$ on edges, weights $n(v)$ on nodes, and some source-sink pair s, t , we are interested in computing the following quantity

$$\text{opt}(s, t) := \underset{\pi=(s, \dots, t), v \in \pi}{\text{argmin}} \quad \text{cost}(\pi, v) \quad (3)$$

where $\text{cost}(\pi, v)$ is defined as

$$\text{cost}(\pi, v) := \left(\sum_{e \in \pi} w(e) + n(v) \right)$$

We compute $\text{opt}(s, t)$ in two stages. First, for every v , we upper bound the value of $\text{opt}(s, v)$ by $n(v)$ plus the shortest distance from s to v . Afterwards, we use dynamic programming to iteratively decrease these upper bounds. Full details are given in Algorithm 2.

Algorithm 2. Optimal 2-Walk Algorithm

Require: Graph $G = (V, E)$ with edge weights $w(e)$, node weights $n(v)$, and a designated source s .

return $r(v) = \text{opt}(s, v)$ for every $v \in V$.

Use Dijkstra's algorithm to compute the shortest path $d(v)$ between s and v .

Initialize $r(v) \leftarrow d(v) + n(v)$ for all $v \in V$. $S \leftarrow \{s\}$.

while $S \neq V$ **do**

Let $u^* = \underset{v \in V \setminus S}{\text{argmin}} r(v)$. Add u^* to S .

For all neighbors z of u^* that are not already in S , let $r(z) \leftarrow \min\{r(u^*) + w(u, z), r(z)\}$

end while

Update. Suppose now that the 2-walk π with smallest cost has been computed. One of two things may bottleneck the amount of processed flow that can be sent along π : either the edge capacity of some edge e , or the processing capacity of some vertex v . We consider the two cases separately

If the bottleneck is edge-based, i.e. $\sum_{v \in \pi^t} C(v) \geq \min_{e \in \pi^t} B(e)$, then let $e^t = \operatorname{argmin}_{e \in \pi^t} B(e)$, and let the chosen flow 2-walk p^t be the one satisfying

$$p_{\pi}^{t,v} = \begin{cases} \frac{C(v)}{\sum_{v \in \pi^t} C(v)} \cdot B_{e^t} & \text{if } \pi = \pi^t, v \in \pi^t \\ 0 & \text{otherwise} \end{cases}$$

On the other hand, if $\sum_{v \in \pi^t} C(v) < \min_{e \in \pi^t} B(e)$, select p^t to satisfy

$$p_{\pi}^{t,v} = \begin{cases} C(v) & \text{if } \pi = \pi^t, v \in \pi^t \\ 0 & \text{otherwise} \end{cases}$$

5.3 Proof of the $(1 - \epsilon)$ Approximation

Let T be the number of rounds taken until we hit the stopping criterion, and let $\bar{p} = \sum_{t=1}^T p^t$ be the total amount of flow selected after T rounds. By the guarantee of the multiplicative update method (Theorem 2.5 in [2]), we have that for any e and any v

$$\begin{aligned} \sum_{t=1}^T M(\mathcal{D}^t, p^t) &\geq \frac{\ln(1 + \epsilon)}{\epsilon} M(\hat{e}, \bar{p}) - \frac{\ln m}{\epsilon} \\ \sum_{t=1}^T M(\mathcal{D}^t, p^t) &\geq \frac{\ln(1 + \epsilon)}{\epsilon} M(\hat{v}, \bar{p}) - \frac{\ln m}{\epsilon} \end{aligned}$$

Since at time T , $w_{\hat{e}}^T = w_{\hat{e}}^0(1 + \epsilon)^{M(\hat{e}, \bar{p})}$, and $q_{\hat{v}}^T = w_{\hat{v}}^0(1 + \epsilon)^{M(\hat{v}, \bar{p})}$, and the stopping rule ensures that at there exists e or v such that $w_{\hat{e}}^T \geq 1$ or $q_{\hat{v}}^T \geq 1$, we have that either there exists an e such that $M(\hat{e}, \bar{p}) \geq \frac{\ln 1/\delta}{\ln(1+\epsilon)}$ or there exists v such that $M(\hat{v}, \bar{p}) \geq \frac{\ln 1/\delta}{\ln(1+\epsilon)}$. Therefore, by the guarantee of the MWU method, we have that

$$\sum_{t=1}^T M(\mathcal{D}^t, p^t) \geq \frac{\ln 1/\delta}{\epsilon} - \frac{\ln m}{\epsilon}$$

We now attempt to bound the left-hand-side of the preceding inequality. Note that

$$\begin{aligned} M(\mathcal{D}^t, p^t) &= \frac{\sum_e w_e^t M(\hat{e}^t, p^t) + \sum_v q_v^t M(\hat{v}^t, p^t)}{\sum_e w_e^t + \sum_v q_v^t} \\ &= \frac{A(p^t) \cdot (\sum_{e \in \pi^t} w_e^t / B(e) + \min_{v \in \pi^t} q_v^t / C(v))}{\sum_e w_e^t + \sum_v q_v^t} \end{aligned}$$

By the definition of π^t and Observation 2, we have

$$\begin{aligned} M(\mathcal{D}^t, p^t) &= \frac{A(p^t)(\sum_{e \in \pi^t} w_e^t/B(e) + \min_{v \in \pi^t} q_v^t/C(v))}{\sum_e w_e^t + \sum_v q_v^t} \\ &\leq A(p^t)/A(p^{opt}) \end{aligned}$$

Combining these inequalities, we get that

$$A(\bar{p})/A(p^{opt}) \geq \sum_{t=1}^T M(\mathcal{D}^t, p^t) \geq \frac{\ln 1/(|E| \cdot \delta)}{\epsilon}$$

Fixing any edge e , its expert's initial weight is $1/\delta$ and its expert's final weight is at most $1+\epsilon$. Thus, \bar{p} passes at most $B(e) \log_{1+\epsilon}((1+\epsilon)/\delta)$ flow through it. Similarly, for each v , at most $C(v) \log_{1+\epsilon}((1+\epsilon)/\delta)$ units of processing are assigned to it. In other words, scaling down all p^t flows by $\log_{1+\epsilon}(1+\epsilon)/\delta$ will result in a feasible flow. Letting $p' = \bar{p}/\log_{1+\epsilon} \frac{1+\epsilon}{\delta}$, we get

$$\begin{aligned} A(p')/A(p^{opt}) &\geq A(\bar{p})/\left(A(p^{opt}) \log_{1+\epsilon} \frac{1+\epsilon}{\delta}\right) \\ &\geq \frac{\ln(1/(|E| \cdot \delta))}{\epsilon} / \log_{1+\epsilon} \frac{1+\epsilon}{\delta} \end{aligned}$$

Taking $\delta = (1+\epsilon)((1+\epsilon)m)^{-1/\epsilon}$, we have that

$$\frac{A(p')}{A(p^{opt})} \geq (1-\epsilon),$$

giving the promised approximation factor.

Note that in each iteration, we either increase the weight of one w_e by a factor of $(1+\epsilon)$, or increase all of the q_v 's on a path π^t by a factor of $(1+\epsilon)$. Since each w_e and each q_v can only be increased by such a factor at most $\frac{\ln 1/(|E| \cdot \delta)}{\epsilon}$ times before its weight exceeds 1, the total running time T is bounded by $(|V| + |E|) \frac{\ln 1/(|E| \cdot \delta)}{\epsilon} \cdot T_{sp} = O(|E| \log |V|/\epsilon^2 \cdot T_{sp})$, where T_{sp} is the time it takes Algorithm 2 to compute the optimal path for each the $|D|$ flows. As a single flow takes time $O(|E| + |V| \log V)$ using Fibonacci heaps, we can compute the 2-walk for each of the flows in time $O(|D| \cdot (|E| + |V| \log |V|))$. Thus, the total running time is $O(|D| \cdot |E| \cdot (|E| + |V| \log |V|) \cdot \log^2 |V|/\epsilon^2)$.

6 Middlebox Node Purchase Optimization

We now discuss the network design problems mentioned in the introduction. Although such problems can be modeled in multiple ways, we limit our discussion to the case where each vertex v has a potential processing capacity C , which can only be utilized if v is ‘‘purchased’’. Flow processed elsewhere can be routed through v regardless of whether or not v is purchased. We give results for both directed and undirected variants of two versions of the network design problem:

1. The *minimization* version (MIN MIDDLEBOX NODE PURCHASE), where the goal is to pick the smallest set of vertices such that all flow is routable.
2. The *maximization* version (BUDGETED MIDDLEBOX NODE PURCHASE), in which we try to maximize the amount of routable flow while subject to a budget constraint of k .

Formally, the input to MIN MIDDLEBOX NODE PURCHASE is a (di)graph $G = (V, E)$ with nonnegative costs q_v on its vertices, a potential processing capacity $C : V \rightarrow [0, \infty)$, and a collection of (s_i, t_i) pairs with demands R_i . The goal is to select a set $T \subseteq V$ of vertices such that all demands are satisfied. BUDGETED MIDDLEBOX NODE PURCHASE is given the same collection of inputs along with a budget integer k , and the goal is to route as much of the demand as possible.

All four problems (maximization or minimization, directed or undirected), are NP-hard.

6.1 Approximation Hardness for Directed Min Middlebox Node Purchase

We now prove that directed MIN MIDDLEBOX NODE PURCHASE is NP-hard to approximate to a factor better than $(1 - \epsilon) \ln n$ by showing an approximation-preserving reduction from SET COVER, a problem already known to have the aforementioned $(1 - \epsilon) \ln n$ hardness [8].

Given a SET COVER instance with set system $\mathcal{S} = \{S_1, S_2, \dots\}$ and universe of elements \mathcal{U} , we create one vertex v_S for each $S \in \mathcal{S}$ and one vertex w_u for each $u \in \mathcal{U}$. Further, we create one source vertex s and one sink vertex t , where t demands $|\mathcal{U}|$ units of processed flow from s . We add one capacity- n arc from s to each v_S , and one capacity-1 arc from each w_u to t . We then add a capacity-1 arc from each v_S to w_u whenever $S \ni u$. Finally, we give each v_S vertex n units of processing capacity at a cost of 1 each.

In order for t to get $|\mathcal{U}|$ units of flow, each w_u must get at least one unit of processed flow itself. Thus, at least one of its incoming v_S neighbors must be able to process flow. Therefore, this instance of directed MIN MIDDLEBOX NODE PURCHASE can be seen as the problem of purchasing as few of the v_S vertices so that each w_u vertex has one (or more) incoming v_S vertex. This provides a direct one-to-one mapping between solutions to our constructed instance and the initial SET COVER instance, and the values of the solutions are conserved by the mapping. Therefore, we have an approximation-preserving reduction between the two problems, and directed MIN MIDDLEBOX NODE PURCHASE acquires the known $(1 - \epsilon) \ln n$ inapproximability of SET COVER, summarized in the following result:

Theorem 2. *For every $\epsilon > 0$, it is NP-hard to approximate directed MIN MIDDLEBOX NODE PURCHASE to within a factor of $(1 - \epsilon) \ln n$.*

Note that this construction provides the same hardness even when all demands are only to be satisfied up to a $(1 - \delta)$ fraction, showing the asymptotic tightness of the approximation factor in Theorem 3.




Fig. 3. Approximation-preserving reduction from SET COVER and MAX k -COVERAGE to directed MIN MIDDLEBOX NODE PURCHASE and directed BUDGETED MIDDLEBOX NODE PURCHASE. Solid edges have infinite capacity, dashed edges have capacity 1. v_i vertices have infinite processing potential, at a cost of 1 each.

6.2 Bicriterion Approximation Algorithm for Directed and Undirected Min Middlebox Node Purchase

We first describe an algorithm for directed MIN MIDDLEBOX NODE PURCHASE that satisfies all flow requirements up to a factor of $1 - \delta$ fraction with expected cost bounded by $O(\log n/\delta^2)$ times the optimum.

We begin our approximation algorithm for directed MIN MIDDLEBOX NODE PURCHASE by modifying the 2-walk-based LP formulation with additional variables x_v corresponding to whether or not processing capacity at vertex v has been purchased. We further give a polynomial sized edge-based LP formulation with flow variables $f_i^{1,v}(e)$ and $f_i^{2,v}(e)$ for each commodity i , each vertex $v \in V$ and each edge $e \in E$. The variables $f_i^{1,v}(e)$ correspond to the (processed) commodity i flow that has been processed by vertex v : these variables describe a flow from v to t_i . The variables $f_i^{2,v}(e)$ correspond to the (unprocessed) commodity i flow that will be processed by vertex v : these variables describe a flow from s_i to v . See Fig. 4 for the full linear program.

Given an optimal solution to this LP, we pick vertices to install processing capacity on by randomized rounding: pick vertex v with probability x_v . if x_v is picked, then all flows processed by v are rounded up in the following way: $\hat{F}_i^{j,v}(e) = f_i^{j,v}(e)/x_v$ for all $i \in [|D|]$, $j \in \{1, 2\}$, $e \in E$. If v is not picked, then all flows processed by v are set to zero, i.e. $\hat{F}_i^{j,v}(e) = 0$.

By design, $E[\hat{F}_i^{j,v}(e)] = f_i^{j,v}(e)$. In the solution produced by the rounding algorithm, the total flow through edge e is $\sum_{v \in V} \sum_{i=1}^{|D|} ((\hat{F}_i^{1,v}(e) + \hat{F}_i^{2,v}(e)))$. This is a random variable whose expectation is at most $B(e)$, and is the sum of independent random variables, one for each vertex v . The constraints of the LP ensure that if v is selected, then the total processing done by vertex v is at most $C(v)$. Further, the total contribution of vertex v to the flow on edge e does not exceed the capacity $B(e)$, i.e. $\sum_{i=1}^{|D|} (\hat{F}_i^{1,v}(e) + \hat{F}_i^{2,v}(e)) \leq B(e)$. Also,

2-Walk-based formulation:

$$\begin{aligned}
& \text{MINIMIZE } \sum_{v \in V} q_v x_v \\
& \text{SUBJECT TO} \\
& x_v \leq 1 \quad \forall v \in V \\
& p_{i,\pi} = \sum_{v \in \pi} p_{i,\pi}^v \quad \forall i \in [|D|], \pi \in P \\
& \sum_{\pi \in P} p_{i,\pi} \geq R_i \quad \forall i \in [|D|] \\
& \sum_{i=1}^{|D|} \sum_{\pi \in P} p_{i,\pi} \leq B(e) \quad \forall e \in E \\
& \sum_{i=1}^{|D|} \sum_{\pi \in P} p_{i,\pi}^v \leq C(v) x_v \quad \forall v \in V \\
& \sum_{i=1}^{|D|} \sum_{\pi \in P} p_{i,\pi}^v \leq B(e) x_v \quad \forall e \in E, v \in V \\
& \sum_{\pi \in P} p_{i,\pi}^v \leq R_i x_v \quad \forall i \in [|D|], v \in V, \\
& p_{i,\pi}^v \geq 0 \quad \forall i \in [|D|], \pi \in P, v \in \pi \\
& x_v \geq 0 \quad \forall v \in V
\end{aligned}$$

Edge-based formulation:

$$\begin{aligned}
& \text{MINIMIZE } \sum_{v \in V} q_v x_v \\
& \text{SUBJECT TO} \\
& x_v \leq 1 \quad \forall v \in V \\
& \sum_{e \in \delta^-(u)} f_i^{j,v}(e) = \sum_{e \in \delta^+(u)} f_i^{j,v}(e) \\
& \quad \forall i \in [|D|], j \in \{1,2\}, v \in V, \forall u \in V \setminus \{s_i, t_i, v\} \\
& \sum_{e \in \delta^-(v)} f_i^{2,v}(e) = \sum_{e \in \delta^+(v)} f_i^{1,v}(e) \quad \forall i \in [|D|], v \in V \\
& \sum_{v \in V} \sum_{e \in \delta^+(s_i)} f_i^{2,v}(e) \geq R_i \quad \forall i \in [|D|] \\
& \sum_{i=1}^{|D|} \sum_{v \in V} (f_i^{1,v}(e) + f_i^{2,v}(e)) \leq B(e) \quad \forall e \in E \\
& \sum_{i=1}^{|D|} \sum_{e \in \delta^-(v)} f_i^{2,v}(e) \leq C(v) x_v \quad \forall v \in V \\
& \sum_{i=1}^{|D|} (f_i^{1,v}(e) + f_i^{2,v}(e)) \leq B(e) x_v \quad \forall e \in E, v \in V \\
& \sum_{e \in \delta^+(s_i)} f_i^{2,v}(e) \leq R_i x_v \quad \forall i \in [|D|], v \in V \\
& f_i^{2,v}(e) = 0 \quad \forall i \in [|D|], v \in V, e \in \delta^-(s_i) \\
& f_i^{1,v}(e) = 0 \quad \forall i \in [|D|], v \in V, e \in \delta^+(t_i) \\
& p_i^{1,v}(e), p_i^{2,v}(e), x_v \geq 0 \quad \forall i \in [|D|], v \in V, e \in E
\end{aligned}$$

Fig. 4. Linear programs from Sect. 3 adapted to middlebox placement problems.

the total contribution of vertex v to the commodity i flow is at most R_i , i.e.

$$\sum_{e \in \delta^+(s_i)} \hat{F}_i^{2,v}(e) \leq R_i.$$

We repeat this randomized rounding process $t = O(\log(n)/\epsilon^2)$ times. Let $g^k(e)$ denote the total flow along edge e , and h_i^k denote the total amount of commodity i flow in the solution produced by the k th round of the randomized rounding process. The following lemma follows easily by Chernoff-Hoeffding bounds:

AQ3

Lemma 3

$$Pr \left[\sum_{k=1}^t g^k(e) \geq (1 + \epsilon)t \cdot B(e) \right] \leq e^{-t\epsilon^2/3} \quad \forall e \in E \quad (4)$$

$$Pr \left[\sum_{k=1}^t h_i^k \leq (1 - \epsilon)t \cdot R_i \right] \leq e^{-t\epsilon^2/2} \quad \forall i \in [|D|] \quad (5)$$

We set $t = O(\log(n)/\epsilon^2)$ so that the above probabilities are at most $1/n^3$ for each edge $e \in E$ and each commodity i . With high probability, none of the associated events occurs. The final solution is constructed as follows: A vertex is purchased if it is selected in any of the t rounds of randomized rounding. Thus the expected cost of the solution is at most $t = O(\log(n)/\epsilon^2)$ times the LP optimum. We consider the superposition of all flows produced by the t solutions and scale down the sum by $t(1 + \epsilon)$. This ensures that the capacity constraints are satisfied. Note that the vertex processing constraints are also satisfied by

the scaled solution. The total amount of commodity i flow is at least $\frac{1-\epsilon}{1+\epsilon}R_i \geq (1-2\epsilon)R_i$. Hence we get the following result:

Theorem 3. *For directed MIN MIDDLEBOX NODE PURCHASE, there is a polynomial time randomized algorithm that satisfies all flow requirements up to factor $1 - \delta$ and produces a solution that respects all capacities, with expected cost bounded by $O(\log(n)/\delta^2)$ times the optimal cost.*

We can modify the LP to simulate the inclusion of an undirected edge with capacity $B(e)$ by adding the constraints for two arcs between its endpoints with capacity $B(e)$ each, as well as an additional constraint requiring that the sum of flows over these two arcs is bounded by $B(e)$. The analysis done above carries through line-by-line, giving the following result.

Theorem 4. *For undirected MIN MIDDLEBOX NODE PURCHASE, there is a polynomial time randomized algorithm that satisfies all flow requirements up to factor $1 - \delta$ and produces a solution that respects all capacities, with expected cost bounded by $O(\log(n)/\delta^2)$ times the optimal cost.*

6.3 Approximation Hardness for Undirected Min Middlebox Node Purchase

We now show an approximation preserving reduction from MIN VERTEX COVER to undirected MIN MIDDLEBOX NODE PURCHASE, proving that the latter problem is **UGC**-hard to approximate within a factor of $2 - \epsilon$ for any $\epsilon > 0$ [18], and **NP**-hard to approximate within a factor of 1.36 [7].

The construction is simple. Given a VERTEX COVER instance with graph $G = (V, E)$, we create an identical graph with each vertex v demanding one unit of processed flow from each of its neighbors, and each edge's capacity is 2. Further, each vertex has n units of processing potential, at a cost of 1. Because the total demand equals the sum of all edge capacities, each unit of flow sent must use exactly one unit of edge capacity, i.e. all flow paths have length exactly one. Thus, the set of solutions exactly corresponds to vertex covers, with one unit of flow going each way across each edge, from source to sink and either to or from its point of processing. The unit costs ensure that the objective value equals the number of vertices picked, and thus that the optimal solution to this undirected MIN MIDDLEBOX NODE PURCHASE instance equals that of the original MIN VERTEX COVER. The conclusion, summarized below, follows.

Theorem 5. *Approximating undirected MIN MIDDLEBOX NODE PURCHASE is at least as hard as approximating MIN VERTEX COVER. In particular, it is **NP**-hard to approximate within a factor of 1.36 and **UGC**-hard to approximate within a factor of $2 - \epsilon$, for any $\epsilon > 0$.*

6.4 Approximation Algorithm for Directed Budgeted Middlebox Node Purchase

The algorithm here proceeds similarly to that in Sect. 6.2. The LPs we use are the natural maximization variant of those used for the minimization problem, with the added restriction that we only use a 1/2 fraction of the budget. It is easy to see that this additional restriction does not reduce the objective value of the optimal LP solution by more than an 1/2-fraction. We also assume (without loss of generality) that no vertex has cost greater than the budget. The LPs are formulated as follows:

2-Walk-based formulation:

$$\begin{aligned}
& \text{MAXIMIZE } \sum_{i=1}^{|D|} \sum_{\pi \in P} p_{i,\pi} \\
& \text{SUBJECT TO} \\
& \sum_{v \in V} c_v x_v \leq k/2 \\
& x_v \leq 1 \quad \forall v \in V \\
& p_{i,\pi} = \sum_{v \in \pi} p_{i,\pi}^v \quad \forall i \in [|D|], \pi \in P \\
& \sum_{\pi \in P} p_{i,\pi} \geq R_i \quad \forall i \in [|D|] \\
& \sum_{i=1}^{|D|} \sum_{\pi \in P} p_{i,\pi} \leq B(e) \quad \forall e \in E \\
& \sum_{i=1}^{|D|} \sum_{\pi \in P} p_{i,\pi}^v \leq C(v) x_v \quad \forall v \in V \\
& \sum_{i=1}^{|D|} \sum_{\pi \in P} p_{i,\pi}^v \leq B(e) x_v \quad \forall e \in E, v \in V \\
& \sum_{\pi \in P} p_{i,\pi}^v \leq R_i x_v \quad \forall i \in [|D|], v \in V \\
& p_{i,\pi}^v \geq 0 \quad \forall i \in [|D|], \pi \in P, v \in \pi \\
& 0 \leq x_v \leq 1 \quad \forall v \in V
\end{aligned}$$

Edge-based formulation:

$$\begin{aligned}
& \text{MAXIMIZE } \sum_{v \in V} \sum_{i=1}^{|D|} \sum_{e \in \delta^-(v)} f_i^{2,v}(e) \\
& \text{SUBJECT TO} \\
& \sum_{v \in V} c_v x_v \leq k/2 \\
& \sum_{e \in \delta^-(u)} f_i^{j,v}(e) = \sum_{e \in \delta^+(u)} f_i^{j,v}(e) \\
& \quad \quad \quad \forall i \in [|D|], j \in \{1,2\}, v \in V, \forall u \in V \setminus \{s_i, t_i, v\} \\
& \sum_{e \in \delta^-(v)} f_i^{2,v}(e) = \sum_{e \in \delta^+(v)} f_i^{1,v}(e) \quad \forall i \in [|D|], v \in V, \\
& \sum_{v \in V} \sum_{e \in \delta^+(s_i)} f_i^{2,v}(e) \geq R_i \quad \forall i \in [|D|] \\
& \sum_{i=1}^{|D|} \sum_{v \in V} (f_i^{1,v}(e) + f_i^{2,v}(e)) \leq B(e) \quad \forall e \in E \\
& \sum_{i=1}^{|D|} \sum_{e \in \delta^-(v)} f_i^{2,v}(e) \leq C(v) x_v \quad \forall v \in V \\
& \sum_{i=1}^{|D|} (f_i^{1,v}(e) + f_i^{2,v}(e)) \leq B(e) x_v \quad \forall e \in E, v \in V \\
& \sum_{e \in \delta^+(s_i)} f_i^{2,v}(e) \leq R_i x_v \quad \forall i \in [|D|], v \in V \\
& f_i^{2,v}(e) = 0 \quad \forall i \in [|D|], v \in V, e \in \delta^-(s_i) \\
& f_i^{1,v}(e) = 0 \quad \forall i \in [|D|], v \in V, e \in \delta^+(t_i) \\
& p_i^{1,v}(e), p_i^{2,v}(e), x_v \geq 0 \quad \forall i \in [|D|], v \in V, e \in E \\
& 0 \leq x_v \leq 1 \quad \forall v \in V
\end{aligned}$$

If purchasing a single vertex allows us to route a $1/(2 \ln n)$ fraction of the objective value of the above LP, we purchase only this vertex. Otherwise, we can remove the potential for processing at each vertex v with $c_v \geq k/\ln n$ and resolve the LP to get a solution with objective value at least half as large as before. Thus, from now on we can assume that no c_v exceeds $k/\ln n$ and therefore that the optimal LP solution puts support on at least a $1/\ln n$ fraction of the x_v s (at a cost of 2 in our approximation factor). We will call the objective value of this modified linear program $\text{OPT}_{\text{LP}'}$.

Again, we pick the vertices on which to install processing capacity on by randomized rounding: each vertex v is picked with probability x_v . If x_v is picked, then all flows processed by v are rounded so that $\hat{F}_i^{j,v}(e) = f_i^{j,v}(e)/(4x_v \ln n)$ for all $i \in [|D|], j \in \{1,2\}, e \in E$. If v is not picked, then all flows processed by v are set to zero, i.e. $\hat{F}_i^{j,v}(e) = 0$.

By design, $E[\hat{F}_i^{j,v}(e)] = f_i^{j,v}(e)/(4 \ln n)$ and thus the total amount of flow processed, P , satisfies $E[P] = E \left[\sum_{v \in V} \sum_{i=1}^{|D|} \sum_{e \in \delta^-(v)} \hat{F}_i^{2,v}(e) \right] = \text{OPT}_{\text{LP}'}/(4 \ln n)$. In the solution produced by the rounding algorithm, the total flow through edge e is $\sum_{v \in V} \sum_{i=1}^{|D|} ((\hat{F}_i^{1,v}(e) + \hat{F}_i^{2,v}(e)))$. This sum of random variables is $\hat{B}(e) = B(e)/(4 \ln n)$ in expectation. Letting $g(e)$ denote the flow along edge e , standard bounds give

Lemma 4

$$\Pr \left[g(e) \geq (4 \lg n) \cdot \hat{B}(e) \right] \leq e^{-4 \ln n} = n^{-4} \quad \forall e \in E \quad (8)$$

$$\Pr \left[P \leq (1/4) \cdot (1/(4 \lg n)) \cdot \text{OPT}_{\text{LP}'} \right] \leq e^{-4 \ln n} = n^{-4} \quad \forall e \in E \quad (9)$$

so by the union bound, with probability higher than $1 - 1/n$ every edge is assigned $\leq B(e)$ total flow and the amount of flow processed and routed is within a $1/16 \ln n$ factor of $\text{OPT}_{\text{LP}'}$.

Finally, by Markov's inequality, the original budget constraint is satisfied with probability at least $1/2$. Combining this with Lemma 4, the algorithm fails with probability at most $1/2 + 1/n$. Repeating the algorithm $O(\log n)$ times and taking the best feasible solution therefore provides an $\Omega(1/\log n)$ approximation with probability at least $1 - 1/\text{poly}(n)$. This can be summarized in the following result:

Theorem 6. *For directed BUDGETED MIDDLEBOX NODE PURCHASE, there is a polynomial-time randomized algorithm producing an $\Omega(1/\log(n))$ approximation.*

We can also apply this algorithm to undirected instances by adding additional constraints the as we did in Sect. 6.2, with the analysis carrying through as before. Thus, we attain the following:

Theorem 7. *For undirected BUDGETED MIDDLEBOX NODE PURCHASE, there is a polynomial-time randomized algorithm producing an $\Omega(1/\log(n))$ approximation.*

6.5 Approximation Algorithm for Undirected Budgeted Middlebox Node Purchase

We now show that the undirected BUDGETED MIDDLEBOX NODE PURCHASE admits a constant-factor approximation algorithm when restricted to a single source s . Let $\text{OPT}(G, k)$ denote the value of the optimal solution to an instance with graph G and budget k . Our algorithm works by splitting the problem into both a *processing step* and a *routing step*. The algorithm begins by reserving a

1/2 fraction of each edge for use in the processing step and the remaining 1/2 fraction for use in the routing step. Calling the reserved-capacity graphs G_{proc} and G_{route} , respectively, the algorithm proceeds as follows:

Processing Step. A well known fact in capacitated network design is that the maximum amount of flow routable (sans processing) from a set $S \subseteq V$ of source vertices to a single sink forms a monotone, submodular function in S [4]. Although this problem is usually defined in the context of sources that can produce an arbitrary amount of flow (should the network support it), we can bottleneck each source s_i into producing at most some c_i units of flow by replacing it with a pair of vertices connected by a capacity c_i edge, without changing the submodularity of the routable flow function, $f_G(S)$. For the purpose of this lemma, redefining s as our “sink” and the set P of processing nodes as our source set S , we immediately attain that the function $f_G(P)$ is submodular, where $P \subset V$ is the set of nodes purchased for processing.

Let H be a copy of G_{proc} with all edge capacities halved. Because f_H is a submodular function, the problem of using our budget to purchase a set $P \subseteq V$ of processing nodes so to maximize $f_H(P)$ is simply an instance of a monotone, submodular maximization subject to knapsack constraints. Such problems are known to admit simple $(1 - 1/e)$ -approximation algorithms [30]. Let $P(H, k)$ be the optimal solution to this *processable flow problem* on H with budget k and $\text{ALG}_1(H, K)$ denote the value of the solution found by our algorithm. Because $P(H, k)$ is an upper bound on $\text{OPT}(H, k)$ (indeed, the former is simply an instance of the former without the need to account for post-processing routing), the $(1 - 1/e)$ approximation we get has value at least equal to $(1 - 1/e)$ times the value of $\text{OPT}(H, k)$. In particular

$$\begin{aligned} \text{ALG}_1(H, k) &\geq (1 - 1/e)P(H, k) \\ &\geq (1 - 1/e)\text{OPT}(H, k) \\ &\geq (1 - 1/e)(1/2)\text{OPT}(G_{\text{proc}}, k) \\ &\geq (1 - 1/e)(1/2)(1/2)\text{OPT}(G, k) \\ &= (1 - 1/e)/4 \cdot \text{OPT}(G, k) \end{aligned}$$

Further, because our solution only uses at most half of the capacity of any edge in G_{proc} , we can use the remaining, unused half of the capacities to route all flow we managed to process back to s .

Routing Step. All flow residing in s after the end of the processing step is already processed, all of it can be routed directly to the sinks using the 1/2 fraction of edge capacities we reserved for G_{route} . Because multiplying all edge capacities by 1/2 reduces the amount of routable flow by the same (multiplicative) amount, we can route at least $(1/2) \min(\text{ALG}_1(H, k), \text{MAXFLOW}_G(s, t))$ units of the processed flow from s to t . As $\text{MAXFLOW}_G(s, t)$ is a (trivial) upper bound on $\text{OPT}(G, k)$, this means we can route at least $(1/2)(1 - 1/e)/4\text{OPT}(G, k)$ units of the processed flow from s to the sinks, giving a $(1 - 1/e)/8 > .078$ approximation algorithm.

Thus, we get the following theorem:

Theorem 8. *For undirected BUDGETED MIDDLEBOX NODE PURCHASE with a single source, there is a deterministic polynomial time algorithm that produces a solution that can route at least $(1 - 1/e)/8 \approx .078$ times the optimal objective solution.*

6.6 Approximation Hardness for Directed Budgeted Middlebox Node Purchase

We now prove that directed BUDGETED MIDDLEBOX NODE PURCHASE is NP-hard to approximate to a factor of $1 - 1/e + \epsilon$. To show this, we reduce from MAX K-COVER, which is known to have the same hardness result [12].

Given a MAX K-COVER instance with set system \mathcal{S} and universe of elements \mathcal{U} , we create one vertex v_S for each $S \in \mathcal{S}$ and one vertex w_u for each $u \in \mathcal{U}$. Further, we create one source vertex s and one sink vertex t , where t demands $|\mathcal{U}|$ units of processed flow from s . We add one capacity- n arc from s to each v_S , and one capacity-1 arc from each w_u to t . We then add a capacity-1 arc from each v_S to w_u whenever $S \ni u$. Finally, we give each v_S vertex n units of processing capacity at a cost of 1 each. The budget for the instance is k – the same as the budget for the MAX-K-COVER instance. A diagram of the reduction is given in Fig. 3.

When flow is routed maximally, each w_u contributes 1 unit of flow to the total $s - t$ flow if and only if it has a neighbor v_S that was chosen to be active. Otherwise, this vertex does not help contribute towards the $s - t$ flow. Thus, this instance of directed BUDGETED MIDDLEBOX NODE PURCHASE can be seen as the problem of buying k different v_S vertices so to maximize the number of distinct w_u vertices to which they are adjacent. Thus, there is a direct one-to-one mapping between solutions to our constructed instance and the initial MAX K-COVER instance, and the values of the solutions are conserved by the mapping. Therefore, we have an approximation-preserving reduction between the two problems, and directed BUDGETED MIDDLEBOX NODE PURCHASE acquires the known $(1 - 1/e + \epsilon)$ inapproximability of MAX K-COVER. The result can be summarized as follows

Theorem 9. *For every $\epsilon > 0$, it is NP-hard to approximate directed BUDGETED MIDDLEBOX NODE PURCHASE to within a factor of $1 - 1/e + \epsilon$.*

6.7 Approximation Hardness for Undirected Budgeted Middlebox Node Purchase

We show that for some fixed $\epsilon_0 > 0$, the undirected version of BUDGETED MIDDLEBOX NODE PURCHASE is NP-hard to approximate within a factor of $1 - \epsilon$, implying that the problem does not admit a PTAS unless $\mathbf{P} = \mathbf{NP}$. We make no attempt to maximize the value ϵ_0 .

We show this hardness by reducing from MAX BISECTION on degree-3 graphs, shown to be hard to approximate within a factor of .997 in [3]¹. Let $G = (V, E)$ be the input to the degree-3 MAX BISECTION instance. For each $v_i \in V$, create two vertices, u_i and w_i , joined by an edge with capacity 3. We also add a capacity-1 edge between u_i and u_j whenever v_i and v_j are adjacent in G . Each w_i vertex demands 3 units of flow from every u_j (including when $i = j$). Further, every u_i vertex can be given $3|V|$ units of processing capacity (or, equivalently, ∞ units) at a cost of 1, and the instance's budget is set to $|V|/2$.

The intuition behind the construction is as follows. With a budget of $|V|/2$, we can purchase exactly half of the u_i vertices (and all budget is used up without loss of generality); our bisection will be between the purchased u_i s and the unpurchased ones. Let b be the number of edges in any such bisection. Each w_i adjacent to a purchased u_i can have 3 units of its demand satisfied by flow originating from and processed by u_i , and the only edge connecting w_i to the rest of the graph ensures w_i can never receive more than 3 units of flow regardless. Thus, such w_i s are maximally satisfied, and contribute $3|V|/2$ units to our objective value. The remaining w_i s must have their processed flow routed to them via edge via the b capacity-1 edges in the bisection (and, indeed, every edge in the bisection will carry 1 unit of flow when routed optimally, as witnessed by the solution where each unprocessed u_i receives flow on each cut-edge and routes it directly to w_i), so the total amount of demand satisfied by the w_i adjacent to unpurchased vertices is exactly b , so the objective value of a solution with b edges in the bisection is exactly $3|V|/2 + b$.

Let b_{OPT} denote the number of edges cut by the optimal bisection. It is a well-known fact that $b_{\text{OPT}} \geq |E|/2 = 3|V|/4$. By the theorem of [3] it is **NP**-hard to distinguish instances with $3|V|/2 + b_{\text{OPT}}$ units of satisfiable demand from those with only $3|V|/2 + (1 - .003)b_{\text{OPT}}$, giving an inapproximability ratio of

$$\begin{aligned} \frac{3|V|/2 + (1 - .003)b_{\text{OPT}}}{3|V|/2 + b_{\text{OPT}}} &= 1 - \frac{.003b_{\text{OPT}}}{3|V|/2 + b_{\text{OPT}}} \\ &= 1 - \frac{.003}{3|V|/(2b_{\text{OPT}}) + 1} \\ &\leq 1 - \frac{.003}{3|V|/(2 \cdot 3|V|/4) + 1} \\ &= 1 - \frac{.003}{2 + 1} \\ &= .999 \end{aligned}$$

This calculation is summarized in the following result:

Theorem 10. *It is **NP**-hard to approximate undirected BUDGETED MIDDLE-BOX NODE PURCHASE to within a factor better than .999.*

¹ To be precise, this paper shows the aforementioned hardness for MAX CUT. A simple approximation preserving reduction from MAX CUT to MAX BISECTION can be derived by looking at maximum cuts of the graph formed by 2 disjoint copies of the MAX CUT instance graph.

7 Related Work

Network Function Optimization. In software-defined networking, SIMPLE [24] and FlowTags [11] take advantage of switches with fine-grained rule support. Both approaches focus on how to use the constrained TCAM size, a hardware limitation to support fine-grained policy. Neither approach attempts to solve the joint optimization of the capacity constraints for both servers and switches. Slick [1] offers a high-level control program that specifies custom processing on precise subset of flows. It also assumes the server processing power is heterogeneous, and uses heuristic approaches for the underlying placement, routing, and steering.

Network Function Consolidation. CoMB [28] and Click [21] both consolidate network functions into applications or a VM images, and consider server machines that can each run multiple instances of different network functions. Both focus on improving the performance on single nodes, and treat network functions homogeneously. Neither covers a network-wide optimization.

Network Function Migration and Reroute. OpenNF [14] and Split-Merge [25] leverage the SDN controller to manage the network function's state migration and the network function's flow migration. Both focus on reallocating resources and rerouting flows when either a node or a link is over-utilized. While their solution focuses on fixing congestion when it occurs, ours focuses on figuring out how to avoid congestion in the first place.

Network Function Online Request Model. Recently, Even, Medina, and Patt-Shamir [9] studied an online request admission problem in the same multi-commodity flow with processing setting that we study. In their work, requests arrive online and specify a processing pipeline for flow between a source and sink; intermediate nodes in the pipeline may be any subset of nodes in the underlying graph. The goal is to accept as many such flow requests as possible while ensuring that accepted requests are assigned flow paths that satisfy capacity constraints. In this setting, the authors show an $O(k \log(kn))$ -competitive online algorithm for instances with length- k pipelines.

Routing and Middlebox optimization. A couple of recent papers consider approximation algorithms for path computation and service placement [10] and Service Chain and Virtual Network Embeddings [26, 27]. Both papers use randomized rounding of a linear programming relaxation of the problem. Both of these works differ from our paper in that packets between demand pairs are not splittable, and thus must be sent along *paths* rather than *flows*. Other recent papers provide approximation algorithms for variants of MIN MIDDLEBOX NODE PURCHASE with no hard edge constraints [6, 20]. In [20], the authors independently derive the same SET COVER-based hardness construction for their problem variant.

References

1. Anwer, B., Benson, T., Feamster, N., Levin, D.: Programming Slick network functions. In: Proceedings of Symposium on SDN Research, June 2015
2. Arora, S., Hazan, E., Kale, S.: The multiplicative weights update method: a meta-algorithm and applications. *Theor. Comput.* **8**(1), 121–164 (2012)
3. Berman, P., Karpinski, M.: On some tighter inapproximability results (extended abstract). In: Wiedermann, J., van Emde Boas, P., Nielsen, M. (eds.) ICALP 1999. LNCS, vol. 1644, pp. 200–209. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-48523-6_17
4. Chakrabarty, D., Krishnaswamy, R., Li, S., Narayanan, S.: Capacitated network design on undirected graphs. In: Raghavendra, P., Raskhodnikova, S., Jansen, K., Rolim, J.D.P. (eds.) APPROX/RANDOM -2013. LNCS, vol. 8096, pp. 71–80. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40328-6_6
5. Chiosi, M., et al.: Network functions virtualisation: introductory white paper. In: SDN and OpenFlow World Congress, October 2012
6. Cohen, R., Lewin-Eytan, L., Naor, J.S., Raz, D.: Near optimal placement of virtual network functions. In: IEEE Conference on Computer Communications (INFOCOM), pp. 1346–1354. IEEE (2015)
7. Dinur, I., Safra, S.: On the hardness of approximating minimum vertex cover. *Ann. Math.* **162**, 439–485 (2005)
8. Dinur, I., Steurer, D.: Analytical approach to parallel repetition. In: Proceedings of the Annual ACM Symposium on Theory of Computing, pp. 624–633. ACM, New York (2014). <https://doi.org/10.1145/2591796.2591884>. <http://doi.acm.org/10.1145/2591796.2591884>
9. Even, G., Medina, M., Patt-Shamir, B.: Competitive path computation and function placement in SDNs. arXiv preprint [arXiv:1602.06169](https://arxiv.org/abs/1602.06169) (2016)
10. Even, G., Rost, M., Schmid, S.: An approximation algorithm for path computation and function placement in SDNs. In: Suomela, J. (ed.) SIROCCO 2016. LNCS, vol. 9988, pp. 374–390. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-48314-6_24
11. Fayazbakhsh, S.K., Chiang, L., Sekar, V., Yu, M., Mogul, J.C.: Enforcing network-wide policies in the presence of dynamic middlebox actions using flowtags. In: 11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2014), pp. 543–546. USENIX Association, Seattle, April 2014. <https://www.usenix.org/conference/nsdi14/technical-sessions/presentation/fayazbakhsh>
12. Feige, U.: A threshold of $\ln n$ for approximating set cover. *J. ACM (JACM)* **45**(4), 634–652 (1998)
13. Forrest, J.: Clp: Coin-or linear program solver. In: DIMACS Workshop on COIN-OR, pp. 17–20, July 2006
14. Gember-Jacobson, A., et al.: OpenNF: enabling innovation in network function control. In: Proceedings of the ACM Conference on SIGCOMM, pp. 163–174. ACM (2014). <https://doi.org/10.1145/2619239.2626313>. <http://doi.acm.org/10.1145/2619239.2626313>
15. Heorhiadi, V., Reiter, M.K., Sekar, V.: Accelerating the development of software-defined network optimization applications using SOL. arXiv preprint [arXiv:1504.07704](https://arxiv.org/abs/1504.07704) (2015)
16. Heorhiadi, V., Reiter, M.K., Sekar, V.: Simplifying software-defined network optimization using sol. In: 13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2016), pp. 223–237. USENIX Association,

- Santa Clara, March 2016. <https://www.usenix.org/conference/nsdi16/technical-sessions/presentation/heorhiadi>
17. Jin, Y., Wen, Y., Westphal, C.: Towards joint resource allocation and routing to optimize video distribution over future internet. In: IFIP Networking Conference (IFIP Networking) 2015, 1–9 May 2015. <https://doi.org/10.1109/IFIPNetworking.2015.7145311>
 18. Khot, S., Regev, O.: Vertex cover might be hard to approximate to within $2-\epsilon$. *J. Comput. Syst. Sci.* **74**(3), 335–349 (2008)
 19. Li, X., Qian, C.: A survey of network function placement. In: 13th IEEE Annual Consumer Communications Networking Conference (CCNC), pp. 948–953, January 2016. <https://doi.org/10.1109/CCNC.2016.7444915>
 20. Lukovszki, T., Rost, M., Schmid, S.: Approximate and incremental network function placement. *J. Parallel Distrib. Comput.* **120**, 159–169 (2018)
 21. Martins, J., et al.: Clickos and the art of network function virtualization. In: 11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2014), pp. 459–473. USENIX Association, April 2014. <https://www.usenix.org/conference/nsdi14/technical-sessions/presentation/martins>
 22. OPNFV: OPNFV: an open platform to accelerate NFV, Linux Foundation. <https://www.opnfv.org/>
 23. Orłowski, S., Wessälly, R., Pióro, M., Tomaszewski, A.: Sndlib 1.0—survivable network design library. *Networks* **55**(3), 276–286 (2010)
 24. Qazi, Z.A., Tu, C.C., Chiang, L., Miao, R., Sekar, V., Yu, M.: SIMPLE-fying middlebox policy enforcement using SDN. In: Proceedings of ACM SIGCOMM, pp. 27–38. ACM (2013). <https://doi.org/10.1145/2486001.2486022>. <http://doi.acm.org/10.1145/2486001.2486022>
 25. Rajagopalan, S., Williams, D., Jamjoom, H., Warfield, A.: Split/merge: system support for elastic execution in virtual middleboxes. In: Presented as Part of the 10th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2013), pp. 227–240. USENIX, Lombard (2013). <https://www.usenix.org/conference/nsdi13/technical-sessions/presentation/rajagopalan>
 26. Rost, M., Schmid, S.: Charting the complexity landscape of virtual network embeddings. In: IFIP Networking, May 2018. <http://eprints.cs.univie.ac.at/5580/>
 27. Rost, M., Schmid, S.: Virtual network embedding approximations: leveraging randomized rounding. In: IFIP Networking, May 2018. <http://eprints.cs.univie.ac.at/5579/>
 28. Sekar, V., Egi, N., Ratnasamy, S., Reiter, M.K., Shi, G.: Design and implementation of a consolidated middlebox architecture. In: Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation, NSDI 2012, p. 24. USENIX Association (2012). <http://dl.acm.org/citation.cfm?id=2228298.2228331>
 29. Sherry, J., Hasan, S., Scott, C., Krishnamurthy, A., Ratnasamy, S., Sekar, V.: Making middleboxes someone else’s problem: network processing as a cloud service. In: Proceedings of the ACM SIGCOMM 2012 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication, SIGCOMM 2012, pp. 13–24. ACM (2012). <https://doi.org/10.1145/2342356.2342359>. <http://doi.acm.org/10.1145/2342356.2342359>
 30. Sviridenko, M.: A note on maximizing a submodular set function subject to a knapsack constraint. *Oper. Res. Lett.* **32**(1), 41–43 (2004)
 31. Uhlig, S., Quoitin, B., Leprope, J., Balon, S.: Providing public intradomain traffic matrices to the research community. *ACM SIGCOMM Comput. Commun. Rev.* **36**(1), 83–86 (2006)