

# Network Design via Core Detouring for Problems Without a Core

Fabrizio Grandoni<sup>1</sup> and Thomas Rothvoß<sup>2</sup>

<sup>1</sup> Dipartimento di Informatica, Sistemi e Produzione, Università di Roma Tor Vergata,  
Via del Politecnico 1, 00133 Roma, Italy, grandoni@disp.uniroma2.it

<sup>2</sup> Institute of Mathematics, EPFL, Lausanne, Switzerland,  
thomas.rothvoss@epfl.ch

**Abstract** Some of the currently best-known approximation algorithms for network design are based on random sampling. One of the key steps of such algorithms is connecting a set of source nodes to a random subset of them. In a recent work [Eisenbrand,Grandoni,Rothvoß,Schäfer-SODA'08], a new technique, *core-detouring*, is described to bound the mentioned connection cost. This is achieved by defining a sub-optimal connection scheme, where paths are detoured through a proper connected subgraph (core). The cost of the detoured paths is bounded against the cost of the core and of the distances from the sources to the core. The analysis then boils down to proving the *existence* of a convenient core.

For some problems, such as connected facility location and single-sink rent-or-buy, the choice of the core is obvious (i.e., the Steiner tree in the optimum solution). Other, more complex network design problems do not exhibit any such core. In this paper we show that core-detouring can be nonetheless successfully applied. The basic idea is constructing a convenient core by manipulating the optimal solution in a proper (not necessarily trivial) way. We illustrate that by presenting improved approximation algorithms for two well-studied problems: virtual private network design and single-sink buy-at-bulk.

**Key words:** Approximation algorithms, network design, virtual private network, buy-at-bulk, rent-or-buy, core detouring.

## 1 Introduction

In a seminal work, Gupta, Kumar, and Roughgarden [17] introduced a random-sampling-based framework to design and analyze approximation algorithms for network design. This way, they achieved improved approximation algorithms for three relevant network design problems: virtual private network design, single-sink rent-or-buy, and single-sink buy-at-bulk (see also [4,5,12,21]). Generalizations and adaptations of their approach were later successfully applied to several other problems, including multi-commodity rent-or buy [1,8,16], connected facility location [6], stochastic (online) Steiner tree [8,9,18], universal TSP [9,26] and many others.

One of the key ingredients in Gupta et al.’s approach [17] is connecting a set of source nodes to a randomly and independently sampled subset of them. The shortest-path distances from the source set to the sampled subset are then bounded against the cost of an optimum Steiner tree over the sampled nodes. In a recent work [6], Eisenbrand, Grandoni, Rothvoß, and Schäfer gave an improved analytical tool, *core detouring*, to bound the connection cost above. The crux of their method is designing a sub-optimal connection scheme, and bounding its cost. In their scheme connection paths are detoured through a proper connected subgraph (*core*). More formally, consider an undirected graph  $G$  with edge weights  $\{c_e\}_{e \in E}$ . We let  $\ell(v, u)$  be the shortest path distance between  $v$  and  $u$ , and  $\ell(v, U) := \min_{u \in U} \{\ell(v, u)\}$  for any  $U \subseteq V(G)$ . Let also  $c(E') := \sum_{e \in E'} c_e$  for any  $E' \subseteq E(G)$ . To lighten the notation, we sometimes use  $G'$  instead of  $E(G')$  or  $V(G')$ , where the meaning will be clear from the context.

**Theorem 1. (Core Detouring) [6]** *Given an undirected graph  $G = (V, E)$ , with edge weights<sup>3</sup>  $\{c_e\}_{e \in E}$ , clients  $C \subseteq V$ , a connected subgraph  $G'$ , a root  $z \in V(G')$  and  $p \in (0, 1]$ . Mark each client independently with probability  $p$ , and denote the marked clients by  $C'$ . Then  $E[\sum_{v \in C} \ell(v, C' \cup \{z\})] \leq (e^{-p|C|}|C| + \frac{0.8067}{p})c(G') + 2 \sum_{v \in C} \ell(v, G')$ .*

To have an intuition of the proof of the theorem, imagine  $G'$  as a cycle of length  $|C|$ , and think of the shortest paths  $\ell(v, G')$  as edges  $(v, f(v))$ , where  $f : C \rightarrow V(G')$  is a bijective function. This can be enforced by duplicating the edges of  $G'$ , computing an Euler tour of the new graph, and performing node duplications and edge contractions in a proper way. Now connect each client  $v \in C$  to the closest sampled client  $v' \in C'$  in terms of number of hops (disregarding edge weights). It is not hard to see that each edge  $(v, f(v))$  is used twice in expectation (once to approach  $G'$  and once to leave it). This accounts for the factor 2 in the upper bound. Moreover, each edge of  $G'$  is used by  $\frac{1}{2p}$  connection paths in expectation, which becomes  $\frac{1}{p}$  in the upper bound due to edge duplication. The refined factor  $(|C|e^{-p|C|} + 0.8067/p)$  is obtained by flow canceling, and a more involving analysis. We remark that in the argument above the connectivity of  $G'$  is crucial.

**Our Results and Techniques.** The Core Detouring Theorem is existential in flavor: it is sufficient to show the existence of a convenient core  $G'$ , of small cost and sufficiently close to the clients  $C$ . For some network design problems, a natural candidate is provided by the structure of the optimum solution. For example, the optimum solution for connected facility location and single-sink rent-or-buy contains a Steiner tree  $T$ . Applying the Core Detouring Theorem to  $T$  leads to improved approximation algorithms for those two problems [6].

In this paper we show that core-detouring can be successfully applied to other network design problems, where the optimum solution does not exhibit any convenient core. The basic idea here is constructing one such core by manipulating the optimal solution in a proper way. We illustrate that by presenting improved approximation algorithms for virtual private network design and

<sup>3</sup> Throughout this paper we use the terms weight and cost interchangeably.

single-sink buy-at-bulk. As we will see, the construction of a good core for the considered problems involves a few non-trivial ideas.

**Virtual Private Network Design (VPN).** VPN models scenarios where the traffic pattern is uncertain or rapidly changing, and henceforth the network must be able to support a family of traffic matrices. This family is implicitly expressed by upper bounding the amount of traffic which each node can send and receive.

**VIRTUAL PRIVATE NETWORK DESIGN (VPN).** Given an undirected graph  $G = (V, E)$ , with edge weights  $\{c_e\}_{e \in E}$ , a set  $S$  of senders and  $R$  of receivers, with out-traffic bounds  $\{b_s^+\}_{s \in S}$  and in-traffic bounds  $\{b_r^-\}_{r \in R}$ . A traffic matrix  $M = \{M_{sr}\}_{(s,r) \in S \times R}$  is feasible if  $\sum_{r \in R} M_{sr} \leq b_s^+$  and  $\sum_{s \in S} M_{sr} \leq b_r^-$  for all  $s \in S$  and  $r \in R$ . Find a minimum cost  $\sum_{e \in E} c_e x_e$  capacity reservation  $\{x_e\}_{e \in E}$  and paths  $\{P_{sr}\}_{(s,r) \in S \times R}$  such that, for every feasible traffic matrix  $M$ , it is possible to route a flow  $M_{sr}$  along path  $P_{sr}$  simultaneously for all  $(s, r) \in S \times R$ , without exceeding the capacity  $x_e$  reserved on each edge  $e$ .

Following the literature on the problem, w.l.o.g. and unless differently stated, we assume  $b_s^+ = b_r^- = 1$  for all  $s \in S, r \in R$ , and we let  $|S| \leq |R|$ . Possibly,  $S \cap R \neq \emptyset$ . We remark that a solution to VPN can be encoded by providing the paths  $P_{sr}$  only. In fact, for a given choice of the paths, the optimal choice for each  $x_e$  is the maximum cardinality of a matching in the bipartite graph  $G_e = (S \cup R, E_e)$ , where  $sr \in E_e$  iff  $e \in P_{sr}$ . The current best approximation ratio for VPN is 3.39 [2,5].

**Theorem 2.** *There is an expected 2.80-approximation algorithm for VPN.*

The proof of Theorem 2 is based on the following two main ideas. We first show that the input VPN instance  $\mathcal{I}$  is equivalent (in expectation) to a different VPN instance  $\mathcal{I}_s$  with the same set of receivers, and a unique random sender  $s$  with out-traffic bound  $|S|$ . Here we crucially exploit König's theorem: *in a bipartite graph the maximum cardinality of a matching equals the minimum cardinality of a vertex cover*. In particular, König's theorem implies that graph  $G_e$  has a vertex cover  $C_e$  of cardinality  $|C_e| = x_e$  for any given choice of the paths  $P_{sr}$ .

Consider the following problem:

**SINGLE-SINK RENT-OR-BUY (SROB).** Given an undirected graph  $G = (V, E)$ , with edge weights  $\{c_e\}_{e \in E}$ , a root  $z$ , a parameter  $M \in \mathbb{Q}^+$  and clients  $D \subseteq V$ . Find a set of paths  $\{P_{zv}\}_{v \in D}$  so as to minimize  $\sum_{e \in E} c_e \min\{M, |\{P_{zv} \mid e \in P_{zv}\}|\}$ .

We remark that an optimal solution to SROB consists of a Steiner tree containing the root and whose edges support at least  $M$  paths each, and a shortest path from each client to the Steiner tree. The second step of our proof is showing that, for any  $s$ ,  $\mathcal{I}_s$  is (deterministically) equivalent to an SROB instance  $\mathcal{I}'_s$  with clients  $D = R$ , root  $z = s$  and parameter  $M = |S|$ . We achieve the claimed approximation guarantee by applying the Core Detouring Theorem to the Steiner tree in  $OPT_{\text{SROB}}(\mathcal{I}'_s)$ <sup>4</sup>.

<sup>4</sup> For a problem  $\mathcal{P}$  and an instance  $\mathcal{I}$  of  $\mathcal{P}$ , we let  $OPT_{\mathcal{P}}(\mathcal{I})$  denote the optimal solution to  $\mathcal{P}$  on  $\mathcal{I}$ . We use  $OPT_{\mathcal{P}}(\mathcal{I})$  also to denote the cost of the optimal solution.

**Single-Sink Buy-at-Bulk (SSBB).** SSBB is another prototypical network design problem, which is used to model scenarios where the capacity is reserved on edges in a discrete fashion to support a given traffic matrix. (For a comparison, in the case of VPN the traffic is unknown but the capacity is installed in a continuous fashion). This is formalized by defining a set of cable types, each one characterized by a cost (per unit length) and a capacity. We are allowed to install  $n_{i,e} \geq 0$  copies of each cable type  $i$  on edge  $e$ .

**SINGLE-SINK BUY-AT-BULK (SSBB).** Given an undirected graph  $G = (V, E)$ , with edge weights  $\{c_e\}_{e \in E}$ , a set of source nodes  $D$  and a sink node  $r$ . Given a set of *cable types*  $1, 2, \dots, k$ , with capacities  $\mu_1 \leq \mu_2 \leq \dots \leq \mu_k$  and costs  $\sigma_1 \leq \sigma_2 \leq \dots \leq \sigma_k$ . Assume  $\delta_i := \frac{\sigma_i}{\mu_i}$  is a decreasing function of  $i$  (*economies of scale*). Find a cable installation  $\{n_{i,e}\}_{1 \leq i \leq k, e \in E}$ , with  $n_{i,e} \in \mathbb{N}$ , minimizing  $\sum_{i,e} c_e \sigma_i n_{i,e}$  and such that one unit of flow can be routed simultaneously from each source node to the sink without exceeding the capacity  $\sum_i \mu_i n_{i,e}$  on each edge  $e$ .

Depending on whether the flow originating at a given source can be routed along several paths or not, we distinguish between *splittable* SSBB (s-SSBB) and *unsplittable* SSBB (u-SSBB), respectively. The best-known approximation bounds for s-SSBB and u-SSBB are 23.93 [2,12] and 148.48 [2,21], respectively.

**Theorem 3.** *There is an expected 20.41-approximation algorithm for s-SSBB.*

The best-known approximation algorithms for s-SSBB [12,17,21] are based on random sampling. These algorithms consist of a sequence of aggregation rounds. In their analysis, in order to upper bound the cost of cables installed at round  $t$ , it is convenient to consider the cost paid by the optimum solution to install cables of type larger than  $s$ , for a proper  $s$ . That subset of cables induces a graph  $G_s$  which is in general disconnected. This rules out a direct application of the Core Detouring Theorem. For this reason, we developed the following generalized version of that theorem, which applies also to the case  $G'$  is disconnected. For a given subgraph  $G'$ , we let  $\ell_{G'}(v, w)$  be the distance from  $v$  to  $w$  in the graph resulting from the contraction (of the connected components) of  $G'$ .

**Theorem 4. (Multi-Core Detouring)** *Given an undirected graph  $G = (V, E)$ , with edge weights  $\{c_e\}_{e \in E}$ , clients  $C \subseteq V$ , a subgraph  $G'$ , a root  $z$  and  $p \in (0, 1]$ . Mark each client independently with probability  $p$ , and denote the marked clients by  $C'$ . Then  $E[\sum_{v \in C} \ell(v, C' \cup \{z\})] \leq (e^{-p|C|}|C| + \frac{0.8067}{p})c(G') + 2 \sum_{v \in C} \ell_{G'}(v, z)$ .*

The theorem above is achieved by embedding the shortest paths in the contracted graph into the original graph, and considering the graph  $G''$  induced by edges crossed by a large enough number of paths. This graph is connected and contains the root. The Core Detouring Theorem is then applied to  $G''$ . With respect to our applications, we can think of Theorem 4 as a way of extracting from the optimum solution (providing  $G'$ ) a convenient core.

We describe a simple polynomial-time procedure, inspired by an existential proof by Karger and Minkoff [22], to transform any solution  $S$  to s-SSBB into

a tree solution<sup>5</sup>  $U$  on the same input instance, while increasing the cost of the solution at most by a factor 2.

**Theorem 5.** *For any given solution  $S$  to an  $s$ -SSBB instance, there is a polynomial-time procedure to construct a tree solution  $U$  for the same instance of cost at most twice the original cost.*

Being  $U$  feasible for the corresponding  $u$ -SSBB instance, we obtain the following corollary.

**Corollary 1.** *Given a  $\rho$ -approximation algorithm for  $s$ -SSBB, there is a  $2\rho$ -approximation algorithm for  $u$ -SSBB. In particular, there is a  $2 \cdot 20.41 = 40.82$  approximation algorithm for  $u$ -SSBB.*

The results concerning VPN and SSBB are described in Sections 2 and 3, respectively.

**Related Work.** VPN was independently defined by Fingerhut et al. [7], and by Duffield et al. [3] and since then, studied by various authors in several variations. The version that we refer to is also called *asymmetric* VPN. This problem is NP-hard even when we restrict to tree solutions [15]. Constant approximation algorithms are presented in [5,15,17,28]. It is known that the optimum solution is not always a tree. Curiously, the algorithms in [15,17] construct a tree solution, while the current best algorithm in [5] does not. We will use a variant of the latter algorithm to achieve our improved bound.

A 2-approximation is known [5] for the *balanced* case  $|S| = |R|$ , which improves on the 3-approximation in [20]. In [20] it is proved that an optimal tree solution for the balanced case can be computed in polynomial time. Very recently [24], it has been shown that the optimal solution is not a tree even in the balanced case, and that the problem remains NP-hard in that special case as well. In the same paper, a  $(2 + \varepsilon)$ -approximation for the *almost balanced* case  $|R|/|S| = O(1)$  was stated, for an arbitrary but fixed  $\varepsilon > 0$ .

In *symmetric* VPN the traffic is undirected, and each terminal  $v$  has a unique threshold  $b_v$  which upper bounds the amount of traffic which  $v$  is responsible for. In [15] a 2-approximation is given for symmetric VPN. In the same paper the authors show that an optimal tree solution can be computed in polynomial time. The so-called *VPN conjecture* states that symmetric VPN always has an optimal tree solution, and hence can be solved in polynomial time. In a breakthrough paper [11], this conjecture was recently proved to be true (see also [13,19] for former proofs of the conjecture on ring networks, which introduce part of the ideas used in [11]).

SSBB has been extensively studied in the literature. It is NP-hard, e.g., by reduction from the Steiner tree problem. Meyerson, Munagala, and Plotkin [23] gave an  $O(\log n)$  approximation for  $s$ -SSBB. Garg, Khandekar, Konjevod, Ravi, Salman, and Sinha [10] described an  $O(k)$  approximation, where  $k$  is the number of cable types. The first constant approximation is due to Guha, Meyerson, and Munagala [14]: the approximation ratio of their algorithm is roughly

<sup>5</sup> A tree solution is a solution where edges with non-zero capacity induce a tree.

**Figure 1** Algorithm `vpn` for VPN.

- 
- (U1) Choose a receiver  $r^* \in R$  uniformly at random.
  - (U2) Mark each receiver uniformly at random with probability  $\frac{\alpha}{|S|}$ . Term  $R'$  as the marked receivers.
  - (U3) For each  $s \in S$ , compute a  $\rho_{st}$ -approximative Steiner tree  $T_s$  spanning  $\{s, r^*\} \cup R'$  and install cumulatively 1 unit of capacity on  $T_s$ .
  - (U4) Install 1 unit of capacity cumulatively on the shortest path from each receiver  $r$  to the closest node in  $R' \cup \{r^*\}$ .
- 

2000. This approximation was reduced to 216 by Talwar [27]. Gupta, Kumar, and Roughgarden [17] described an improved 76.8 approximation algorithm, based on random sampling. Refining their approach, the approximation was later reduced to 65.49 by Jothi and Raghavachari [21], and eventually to 24.92 by Grandoni and Italiano [12].

The unsplittable case u-SSBB is less studied, though probably more interesting from an application point of view. The algorithm by Talwar is a 216-approximation for u-SSBB as well. Unfortunately, this is not the case for the following improved random-sampling algorithms (i.e., those algorithms do not guarantee that the flow is unsplittable). Jothi and Raghavachari [21] show how to transform the 76.8 approximation algorithm for s-SSBB by Gupta et al. [17] into a  $2 \cdot 76.8 = 153.6$  approximation algorithm for u-SSBB. Their approach is algorithm-specific: it would not work with an (even slightly) different algorithm. (In particular, it cannot be applied to our s-SSBB algorithm nor to the s-SSBB algorithms in [12,21]). In contrast, the reduction from Corollary 1 can use any s-SSBB algorithm as a black box.

SROB [15,22,25] is the special case of SSBB where there are only two cable types, one of very small capacity and cost per unit capacity  $\delta_1 = 1$ , and the other of cost  $\sigma_2 = M \geq 1$  and very large capacity. The current best approximation ratio for SROB is 2.92 [6].

The recent result in [2], trivially implies improved approximation factors 3.39, 2.80, 23.93, and 148.48 for VPN, SROB, s-SSBB, and u-SSBB, respectively.

## 2 Virtual Private Network Design

In this section we present our improved 2.80-approximation algorithm for VPN, hence proving Theorem 2. Having in mind that for any fixed  $\delta > 0$ , there is a  $(2 + \delta \frac{|R|}{|S|})$ -approximation algorithm for VPN [24] (recall that  $|R| \geq |S|$ ), we may assume that  $|S| \leq \varepsilon |R|$  for an arbitrarily small  $\varepsilon > 0$ .

Algorithm `vpn`, which is a slight adaptation of the VPN algorithm in [5], is described in Figure 1. The quantity  $\alpha$  is a positive constant to be fixed later. The best-known approximation factor for the Steiner tree problem is denoted by  $\rho_{st}$ . Currently  $\rho_{st} < 1.39$  [2].

For a given VPN instance  $\mathcal{I}$  and any sender  $s \in S$ , we let  $\mathcal{I}_s$  be the VPN instance with the same receiver set as  $\mathcal{I}$ , and with sender set  $\{s\}$ , where the

out-traffic bound for  $s$  is  $|S|$ . (Recall that  $b_s^+ = 1$  for the original problem). The following lemma crucially exploits König's Theorem.

**Lemma 1.**  $\sum_{s \in S} OPT_{\text{VPN}}(\mathcal{I}_s) \leq |S| \cdot OPT_{\text{VPN}}(\mathcal{I})$ .

*Proof.* We show that, for a random sender  $s^*$ ,  $E[OPT_{\text{VPN}}(\mathcal{I}_{s^*})] \leq OPT_{\text{VPN}}(\mathcal{I})$ . Let  $\{P_{sr}\}_{(s,r) \in S \times R}$  be the optimal paths for  $\mathcal{I}$  and let  $\{x_e\}_{e \in E}$  be the induced capacities. Consider the solution to  $\mathcal{I}_{s^*}$  induced by paths  $\{P_{s^*r}\}_{r \in R}$  and let  $\{x'_e\}_{e \in E}$  be the corresponding capacity reservation. Consider the bipartite graph  $G_e = (S \cup R, E_e)$ , with  $sr \in E_e$  iff  $e \in P_{sr}$ . Let  $C_e \subseteq S \cup R$  be a vertex cover for  $G_e$  of size  $x_e$  (which exists by König's theorem). Clearly  $x'_e \leq \min\{|N_e(s^*)|, |S|\}$ , whereby  $N_e(s^*)$  are the nodes adjacent to  $s^*$  in  $G_e$ . Let us show that  $E[x'_e] \leq x_e$ . The event  $\{s^* \in S \cap C_e\}$  happens with probability  $\frac{|S \cap C_e|}{|S|}$ . In this case we can upper bound  $x'_e$  with  $|S|$ . In the complementary case we can upper bound  $x'_e$  with  $|N_e(s^*)| \leq |R \cap C_e|$ , where we exploit the fact that  $s^*$  can be only adjacent to nodes of  $R \cap C_e$  (otherwise  $C_e$  would not be a vertex cover). Altogether

$$E[x'_e] \leq \frac{|S \cap C_e|}{|S|} \cdot |S| + \left(1 - \frac{|S \cap C_e|}{|S|}\right) \cdot |R \cap C_e| \leq |S \cap C_e| + |R \cap C_e| = |C_e| = x_e.$$

The claim follows since

$$E[OPT_{\text{VPN}}(\mathcal{I}_{s^*})] \leq E\left[\sum_{e \in E} c_e x'_e\right] \leq \sum_{e \in E} c_e x_e = OPT_{\text{VPN}}(\mathcal{I}). \quad \square$$

Let  $\mathcal{I}'_s$  be the SROB instance with clients  $D = R$ , root  $z = s$  and parameter  $M = |S|$ .

**Lemma 2.**  $OPT_{\text{SROB}}(\mathcal{I}'_s) = OPT_{\text{VPN}}(\mathcal{I}_s)$ .

*Proof.* Consider any set of paths  $\{P_{sr}\}_{r \in R}$ . It is sufficient to show that, for any given edge  $e$ , the corresponding cost associated to edge  $e$  is the same in the SROB and VPN case. The cost paid in the SROB case is by definition  $c(e) \cdot \min\{M, |\{P_{sr} : e \in P_{sr}\}|\}$ . With respect to VPN, consider the set of receivers  $R_e := \{r \in R \mid e \in P_{sr}\}$  using edge  $e$ . In the worst case a traffic matrix routes  $k' := \min\{|S|, |R_e|\}$  units of flow along  $e$ . Hence, also in this case the cost associated to  $e$  is  $c(e) \cdot \min\{|S|, |R_e|\} = c(e) \cdot \min\{M, |\{P_{sr} : e \in P_{sr}\}|\}$ .  $\square$

Let  $C_s$  be the Steiner tree in  $OPT_{\text{SROB}}(\mathcal{I}'_s)$ , and  $U_{s,r}$  the shortest path from  $r \in R$  to  $C_s$ . Define  $\Sigma_S := \sum_{s \in S} c(C_s)$  and  $\Sigma_C := \frac{1}{M} \sum_{s \in S} \sum_{r \in R} c(U_{s,r})$ . First we upper bound the cost of the  $M = |S|$  Steiner trees computed by  $\text{vpn}$ .

**Lemma 3.**  $E[\sum_{s \in S} c(T_s)] \leq \rho_{st} \cdot \Sigma_S + \rho_{st}(\alpha + \varepsilon) \cdot \Sigma_C$ .

*Proof.* Recall that  $M/|R| = |S|/|R| \leq \varepsilon$ . For each  $s$  take the core  $C_s$  and attach the path  $U_{s,r}$  for all  $r \in R$ . Each  $U_{s,r}$ ,  $r \in R$ , is used with probability at most  $\frac{\alpha}{M} + \frac{1}{|R|}$ , thus there exists a Steiner tree over  $\{s, r^*\} \cup R$  of expected cost  $c(C_s) + \left(\frac{\alpha}{M} + \frac{1}{|R|}\right) \sum_{r \in R} c(U_{s,r})$ . Multiplying this quantity by the Steiner tree approximation factor, and summing over all  $s$  we obtain

$$\sum_{s \in S} E[c(T_s)] \leq \rho_{st} \left( \sum_{s \in S} c(C_s) + \left(\frac{\alpha}{M} + \frac{1}{|R|}\right) \sum_{s \in S} \sum_{r \in R} c(U_{s,r}) \right) \leq \rho_{st} \Sigma_S + \rho_{st}(\alpha + \varepsilon) \Sigma_C. \quad \square$$

We next bound the cost of connecting each receiver to the closest node in  $R' \cup \{r^*\}$  via the Core Detouring Theorem.

**Lemma 4.**  $E[\sum_{r \in R} \ell(r, R' \cup \{r^*\})] \leq \frac{0.81}{\alpha} \Sigma_S + 2\Sigma_C$ .

*Proof.* Let  $s \in S$ . Applying the Core Detouring Theorem with  $C = R$ ,  $G' = C_s$ ,  $z = r^*$  and  $p = \alpha/M$ ,  $E[\sum_{r \in R} \ell(r, R' \cup \{r^*\})] \leq (e^{-\frac{\alpha}{M}|R|}|R| + \frac{0.8067}{\alpha/M})c(C_s) + 2 \sum_{r \in R} c(U_{s,r}) \leq \frac{0.81M}{\alpha}c(C_s) + 2 \sum_{r \in R} c(U_{s,r})$ , where we use the assumption  $|R| \gg |S| = M$ . Averaging this bound over all  $s$ , we obtain

$$E\left[\sum_{r \in R} \ell(r, R' \cup \{r^*\})\right] \leq \frac{0.81M}{\alpha} \sum_{s \in S} \frac{1}{M} c(C_s) + 2 \frac{1}{M} \sum_{s \in S} \sum_{r \in R} c(U_{s,r}) = \frac{0.81}{\alpha} \Sigma_S + 2\Sigma_C. \square$$

**Theorem 6.** *For a suitable choice of  $\alpha$  and  $|R|/|S|$  large enough, Algorithm `vpn` gives an expected 2.80 approximation for VPN.*

*Proof.* From Lemmas 1 and 2,  $\Sigma_S + \Sigma_C = 1/M \cdot \sum_{s \in S} (M c(C_s) + \sum_{r \in R} c(U_{s,r})) = 1/M \cdot \sum_{s \in S} OPT_{\text{SROB}}(\mathcal{I}_s) = 1/M \cdot \sum_{s \in S} OPT_{\text{VPN}}(\mathcal{I}_s) \leq OPT_{\text{VPN}}$ . By Lemmas 3 and 4, the expected cost of the solution computed by the algorithm is

$$(\rho_{st} \cdot \Sigma_S + \rho_{st}(\alpha + \varepsilon)\Sigma_C) + \left(\frac{0.81}{\alpha} \Sigma_S + 2\Sigma_C\right)^{\alpha=0.5748} \leq 2.80(\Sigma_C + \Sigma_S) \leq 2.80 \cdot OPT_{\text{VPN}}. \square$$

Theorem 2 then follows.

### 3 Single-Sink Buy-at-Bulk

In this section we present our improved algorithms for SSBB. We start with the proof of the Multi-Core Detouring Theorem. Then we present our results for the unsplittable and splittable case.

#### Multi-Core Detouring.

*Proof (of Theorem 4).* Let  $\phi_p(C, G') := 2 \sum_{v \in C} \ell_{G'}(v, z) + \gamma c(G')$  with  $\gamma := (|C| \cdot e^{-p|C|} + \frac{0.8067}{p})$ . We will find a connected subgraph  $G''$  of  $G$  with  $z \in V(G'')$ , having  $\phi_p(C, G'') \leq \phi_p(C, G')$ . The claim then follows by applying the Core Detouring Theorem to  $G''$ . Let  $P_{vz}$  be the path, attaining the length  $\ell_{G'}(v, z)$ , i.e. it is a shortest  $v$ - $z$  path in  $G$ , where edges in  $G'$  account with cost 0. Since these paths are shortest paths, we may assume that  $\bigcup_{v \in C} P_{vz}$  induces a tree  $T$ , rooted at  $z$ . For  $e \in T$ , let  $m_e := |\{v \in C \mid e \in P_{vz}\}|$  be the number of  $v$ - $z$  paths that contain  $e$ . Let  $G''$  be the graph, induced by the edges  $e \in T$  with  $m_e \geq \gamma/2$  ( $G'' := \{z\}$  if no such edge exists). Moving from a leaf of  $T$  to the root  $z$ , the quantity  $m_e$  can only increase, hence the subgraph  $G''$  is connected and  $z \in V(G'')$ . To upperbound  $\phi_p(C, G'')$ , we still use  $P_{vz}$  as  $v$ - $z$  path, even if  $\ell_{G''}(v, z)$  is attained by a different path. Consider any edge  $e \in T$ . If  $e \in G'$ , then  $e$  contributes cost  $\gamma c_e$  to  $\phi_p(C, G')$ , otherwise it contributes  $2m_e c_e$ . Note that  $\gamma c_e \leq 2m_e c_e$  iff  $m_e \geq \gamma/2$ . By the definition of  $G''$ , the contribution of  $e$  to  $\phi_p(C, G'')$  is  $\min\{2m_e c_e, \gamma c_e\}$ , which is never larger than the contribution to



---

**Figure 2** Algorithm *sssbb*.

**(S1)** Select a subset of cable types  $i(1), \dots, i(k')$  in increasing order of capacity, where  $i(1) = 1$  and  $i(k') = k$ .

**(S2)** For  $t = 0, 1, \dots, k'$ :

**(Collection)** Let  $D_t$  be the set of nodes with positive demand. Each node in  $D_t$  is marked with probability  $p_t = \alpha \sigma_{i(t)} / \sigma_{i(t+1)}$  (probability 1 if  $t = 0$ ). Let  $D'_t$  be the set of marked nodes. Each node sends its demand to the closest node in  $D'_t \cup \{r\}$  along a shortest path, using cables of type  $i(t)$  (type 1 for  $t = 0$ ). Let  $d'_t(w)$  be the new demand collected at each  $w \in D'_t \cup \{r\}$ .

**(Aggregation)** If  $t < k'$ , compute a  $\rho_{st}$ -approximate Steiner tree  $T_t$  on  $D'_t \cup \{r\}$ . Apply the aggregation algorithm to  $T_t$  with  $U = \mu_{i(t+1)}$  and  $x(w) = d'_t(w) \pmod{\mu_{i(t+1)}}$  for each terminal node  $w$ . The corresponding flow is supported by installing cables of type  $i(t+1)$  (at most one for each edge of  $T_t$ ). Let  $d''_t(w)$  be the new demand aggregated at each node  $w$ .

**(Redistribution)** If  $t < k'$ , for each node  $w \in D'_t \cup \{r\}$ , consider the subset of nodes  $D_t(w) \subseteq D_t$  that sent their demand to  $w$  during the collection step (including  $w$  itself, if  $w \neq r$ ). Uniformly select a random subset  $\tilde{D}_t(w)$  of  $D_t(w)$  of cardinality  $d''_t(w) / \mu_{i(t+1)}$ . Send  $\mu_{i(t+1)}$  units of flow back from  $w$  to each node in  $\tilde{D}_t(w)$  along shortest paths, installing cables of type  $i(t+1)$ .

---

$\phi_p(C, G')$ . The claim follows by applying the Core Detouring Theorem to the core  $G''$ :

$$\begin{aligned} E \left[ \sum_{v \in C} \ell(v, C' \cup \{z\}) \right] &\leq \phi_p(C, G'') \leq \sum_{e \in T \setminus G''} 2m_e c_e + \sum_{e \in G''} \gamma c_e \\ &= \sum_{e \in T} \min\{2m_e c_e, \gamma c_e\} \leq \sum_{e \in T \setminus G'} m_e c_e + \sum_{e \in G'} \gamma c_e = \phi_p(C, G'). \square \end{aligned}$$

**From Splittable to Unsplittable Flows.** We first state the following simple lemma, which is implicitly given in [27].

**Lemma 5.** [27] *Let  $c^*(x)$  be the minimum-cost of a cable installation supporting a capacity reservation  $x = \{x_e\}_{e \in E}$ . Then there is a polynomial-time computable concave function  $g(\cdot)$  such that  $c^*(x) \leq g(x) \leq 2c^*(x)$ .*

The choice of  $g(\cdot)$  here is  $g(x) := \sum_{e \in E} c(e) f(x_e)$  with  $f(z) := \min_{i=1, \dots, k} \{\sigma_i + \delta_i \cdot z\}$  for  $z > 0$  and  $f(0) := 0$ .

For an arbitrary concave cost function  $g(\cdot)$ , Karger and Minkoff [22] showed that there is always an optimum solution inducing a tree. It is not hard (just slightly technical) to turn their existential proof into a polynomial-time procedure to transform any given solution into a tree solution without increasing its cost with respect to  $g(\cdot)$ . The proof of Theorem 5, which is omitted here for lack of space, is obtained by combining that procedure with the concave function provided by Lemma 5.

**The Splittable Case.** Our algorithm *sssbb* for s-SSBB, which is described in Figure 2, is a slight variant of the algorithms in [12]. By adding dummy clients

in the sink, we can assume that  $|D|$  is a multiple of all the capacities  $\mu_i$ . The quantity  $\alpha$  is a proper constant to be fixed later. The *aggregation algorithm* [17] is a randomized procedure to aggregate a given set of demands  $x(v) \in [0, U]$  on the nodes  $v \in V(T)$  of a given tree  $T$ , under the assumption that the sum of the demands is a multiple of  $U > 0$ . This is obtained by moving demands over  $T$  such that: (1) The amount of flow along each edge of  $T$  is at most  $U$ , (2) The final demand  $x'(v)$  at each node is either 0 or  $U$ , and (3) The expected demand at each node is preserved, that is:  $\Pr[x'(v) = U] = x(v)/U$ .

The algorithm initially selects a subset of  $k'$  cable types  $i(1), i(2), \dots, i(k')$ . For notational convenience, we assume  $\sigma_{i(k'+1)} = \infty$  and  $i(0) = 0$ . Then there is a sequence of  $k' + 1$  rounds. In each round the demand of the clients (which is initially 1 for each client) is aggregated in a smaller and smaller subset of clients. At the beginning of round  $t \geq 1$  the demand at each client is in  $\{0, \mu_{i(t)}\}$ . Each round  $t$  consists of three steps. Initially the demand is collected at a random subset of aggregation points (Collection Step). Then a Steiner tree is computed on the aggregation points, and the demand is aggregated along such tree via the aggregation algorithm in multiples of  $\mu_{i(t+1)}$  (Aggregation Step). This is possible since the sum of the  $d'_t(w)$ 's, and hence of the  $x(w)$ 's, is a multiple of  $\mu_{i(t+1)}$ . Eventually, the aggregated demand is redistributed back to the source nodes (Redistribution Step). Only cables of type  $i(t)$  and  $i(t+1)$  are used in round  $t$ . At the end of the round the demand at each client is in  $\{0, \mu_{i(t+1)}\}$ .

It remains to specify how the cable types  $i(1), \dots, i(k')$  are chosen. Differently from prior work on the topic, we use a randomized cable selection rule. Let  $i(1) = 1$ . Given  $i(t)$ ,  $1 < i(t) < k$ ,  $i(t+1)$  is chosen as follows. Let  $i'(t) > i(t)$  and  $i''(t) > i(t)$  be the smallest indexes such that  $\frac{\delta_{i'(t)}}{\delta_{i(t)}} \leq \frac{1}{\beta}$  and  $\frac{\sigma_{i''(t)}}{\sigma_{i(t)}} \geq \beta$ , respectively. Here  $\beta > 1$  is a constant to be fixed later. If no proper  $i'(t)$  (resp.  $i''(t)$ ) exists, we let  $i'(t) = k$  (resp.  $i''(t) = k$ ). If  $i'(t) \geq i''(t)$ ,  $i(t+1) = i'(t)$ . Otherwise,  $i(t+1) = i''(t) - 1$  with probability  $\frac{\sigma_{i''(t)} - \beta \sigma_{i(t)}}{\sigma_{i''(t)} - \sigma_{i''(t)-1}}$ , and  $i(t+1) = i''(t)$  otherwise. Note that, as required  $i(1) = 1$  and  $i(k') = k$ . The proof of the following lemma will appear in the final version of the paper.

**Lemma 6.** *For any  $t \in \{1, 2, \dots, k' - 2\}$  and  $h \in \{0, 1, \dots, k' - t - 1\}$ , and for any  $s \in \{1, 2, \dots, k\}$ ,  $i(t') < s \leq i(t' + 1)$ : (a)  $\delta_{i(t+h)} \leq \frac{1}{\beta^h} \delta_{i(t)}$ ; (b)  $E[\sigma_{i(t+h)}] \geq \beta^h E[\sigma_{i(t)}]$ ; (c)  $E[\min\{\frac{\sigma_{i(t'+1)}}{\sigma_s}, \frac{\delta_{i(t')}}{\delta_s}\}] \leq \beta$ .*

Let  $A_t$  be the cost of the  $t$ -th round,  $t \in \{0, 1, \dots, k'\}$ . Let moreover  $A_t^c$ ,  $A_t^a$ , and  $A_t^r$  denote the collection, aggregation, and redistribution costs of the  $t$ -th round,  $t \in \{1, \dots, k' - 1\}$  respectively. By  $OPT(s)$  we denote the cost paid by the optimum solution for cables of type  $s$ . The proof of the following lemma is implicitly given in [12].

**Lemma 7.** [12] *For  $t' \in \{1, \dots, k'\}$  and  $t \in \{1, \dots, k' - 1\}$ : (1)  $\Pr[d \in D_{t'} | v \in D_0] = \frac{1}{\mu_{i(t'')}}$ ; (2)  $A_0 \leq \rho_{st} \sum_s \frac{\sigma_{i(1)}}{\sigma_s} OPT(s)$ ; (3)  $E[A_{k'}] \leq \sum_s \frac{\delta_{i(k')}}{\delta_s} OPT(s)$ ; (4)  $E[A_t^a] \leq E\left[\sum_s \min\left\{\rho_{st} \alpha \frac{\delta_{i(t)}}{\delta_s}, \rho_{st} \frac{\sigma_{i(t+1)}}{\sigma_s}\right\} OPT(s)\right]$ ; (5)  $E[A_t^r] \leq E\left[\frac{\delta_{i(t+1)}}{\delta_{i(t)}} A_t^c\right]$ .*

Hence it remains to bound  $E[A_t^c]$ . Following [12,17], it is not hard to show that  $E[A_t^c] \leq \frac{2}{\alpha} E[A_t^a]$ . We next present an improved bound based on the Multi-Core Detouring Theorem. By adding dummy demands at the root, we can assume that  $|D_t| = |D|/\mu_{i(t)} \gg p_t = \alpha\sigma_{i(t)}/\sigma_{i(t+1)}$  for all  $t$ , and consequently  $|D_t|e^{-p_t|D_t|} + 0.8067/p_t \leq 0.8068/p_t$ .

**Lemma 8.** For  $t = 1, \dots, k' - 1$ ,  $E[A_t^c] \leq E[\sum_s \min\{2\frac{\delta_{i(t)}}{\delta_s}, \frac{0.8068}{\alpha} \frac{\sigma_{i(t+1)}}{\sigma_s}\} OPT(s)]$ .

*Proof.* Let  $j \in \{1, 2, \dots, k\}$  be an integer value to be fixed later. We denote by  $G_j$  the graph induced by the edges where  $OPT$  installs at least one cable of type  $s > j$ . Note that this graph might be disconnected. By the Multi-Core Detouring Theorem applied to  $C = D_t$ ,  $z = r$ ,  $p = p_t$  and  $G' = G_j$ ,  $E[A_t^c] := E[\sigma_{i(t)} \sum_{d \in D_t} \ell(d, D_t' \cup \{r\})] \leq E[\sigma_{i(t)} (2 \sum_{d \in D_t} \ell_{G_j}(d, r) + \frac{0.8068}{p_t} c(G_j))]$ .

By definition,  $E[\frac{0.8068}{p_t} \sigma_{i(t)} c(G_j)] = E[\frac{0.8068 \sigma_{i(t+1)}}{\alpha} c(G_j)] \leq E[\frac{0.8068}{\alpha} \sum_{s > j} \frac{\sigma_{i(t+1)}}{\sigma_s} OPT(s)]$ . By Lemma 7.1,  $\Pr[d \in D_t | d \in D] = \frac{1}{\mu_{i(t)}}$ . Then  $E[2\sigma_{i(t)} \sum_{d \in D_t} \ell_{G_j}(d, r)] = E[2\frac{\sigma_{i(t)}}{\mu_{i(t)}} \sum_{d \in D} \ell_{G_j}(d, r)] = E[2\delta_{i(t)} \sum_{d \in D} \ell_{G_j}(d, r)]$ . Let  $L_{t,j}$  be the cost of routing the flow as in  $OPT$ , but paying zero on the edges of  $G_j$  and  $\delta_{i(t)}$  per unit of flow on the remaining edges. Then trivially  $\delta_{i(t)} \sum_{d \in D} \ell_{G_j}(d, r) \leq L_{t,j}$ . In turn,  $OPT$  pays at least  $\delta_s$  per unit flow on each cable of type  $s \leq j$ , which implies  $L_{t,j} \leq \sum_{s \leq j} \frac{\delta_{i(t)}}{\delta_s} OPT(s)$ . We can conclude that  $E[2\sigma_{i(t)} \sum_{d \in D_t} \ell_{G_j}(d, r)] \leq E[2 \sum_{s \leq j} \frac{\delta_{i(t)}}{\delta_s} OPT(s)]$ . Altogether  $E[A_t^c] \leq E[2 \sum_{s \leq j} \frac{\delta_{i(t)}}{\delta_s} OPT(s) + \frac{0.8068}{\alpha} \sum_{s > j} \frac{\sigma_{i(t+1)}}{\sigma_s} OPT(s)]$ . Since, deterministically,  $\delta_{i(t)}/\delta_s$  is decreasing in  $t$  while  $\sigma_{i(t+1)}/\sigma_s$  is increasing in  $t$ , the claim follows by choosing  $j$  properly.  $\square$

The proof of Theorem 3, omitted here due to space constraints, follows easily by combining Lemmas 6, 7, and 8, and imposing  $\beta = 2.80$  and  $\alpha = 0.531$ .

## References

1. L. Becchetti, J. Könemann, S. Leonardi, and M. Pál. Sharing the cost more efficiently: improved approximation for multicommodity rent-or-buy. *ACM Transactions on Algorithms*, 3(2), 2007.
2. J. Byrka, F. Grandoni, T. Rothvoß, and L. Sanità. An Improved LP-based Approximation for Steiner tree. In *STOC*, 2010. Best Paper Award. To appear.
3. N. G. Duffield, P. Goyal, A. Greenberg, P. Mishra, K. K. Ramakrishnan, and J. E. van der Merwe. A flexible model for resource management in virtual private networks. In *SIGCOMM*, 95–108, 1999.
4. F. Eisenbrand and F. Grandoni. An improved approximation algorithm for virtual private network design. In *SODA*, 928–932, 2005.
5. F. Eisenbrand, F. Grandoni, G. Oriolo, and M. Skutella. New approaches for virtual private network design. *SIAM Journal on Computing*, 37(3):706–721, 2007.
6. F. Eisenbrand, F. Grandoni, T. Rothvoß, and G. Schäfer. Approximating connected facility location problems via random facility sampling and core detouring. In *SODA*, 1174–1183, 2008.

7. J. A. Fingerhut, S. Suri, and J. S. Turner. Designing least-cost nonblocking broadband networks. *Journal of Algorithms*, 24(2):287–309, 1997.
8. L. Fleischer, J. Könemann, S. Leonardi, and G. Schäfer. Simple cost sharing schemes for multicommodity rent-or-buy and stochastic steiner tree. In *STOC*, 663–670, 2006.
9. N. Garg, A. Gupta, S. Leonardi, and P. Sankowski. Stochastic analyses for online combinatorial optimization problems. In *SODA*, 942–951, 2008.
10. N. Garg, R. Khandekar, G. Konjevod, R. Ravi, F. Salman, and A. Sinha. On the integrality gap of a natural formulation of the single-sink buy-at-bulk network design problem. In *IPCO*, 170–184, 2001.
11. N. Goyal, N. Olver, and F. B. Shepherd. The VPN conjecture is true. In *STOC*, 443–450, 2008.
12. F. Grandoni and G. F. Italiano. Improved approximation for single-sink buy-at-bulk. In *ISAAC*, 111–120, 2006.
13. F. Grandoni, V. Kaibel, G. Oriolo, M. Skutella. A short proof of the VPN Tree Routing Conjecture on ring networks. *Operations Research Letters*, 36(3): 361–365, 2008.
14. S. Guha, A. Meyerson, and K. Munagala. A constant factor approximation for the single sink edge installation problem. *SIAM Journal on Computing*, 38(6): 2426–2442, 2009.
15. A. Gupta, J. Kleinberg, A. Kumar, R. Rastogi, and B. Yener. Provisioning a virtual private network: a network design problem for multicommodity flow. In *STOC*, 389–398, 2001.
16. A. Gupta, A. Kumar, M. Pal, and T. Roughgarden. Approximation via cost-sharing: simpler and better approximation algorithms for network design. *Journal of the ACM*, 54(3): 11, 2007.
17. A. Gupta, A. Kumar, and T. Roughgarden. Simpler and better approximation algorithms for network design. In *STOC*, 365–372, 2003.
18. A. Gupta and M. Pál. Stochastic Steiner trees without a root. In *ICALP*, 1051–1063, 2005.
19. Cor A. J. Hurkens, J. C. M. Keijsper, and L. Stougie. Virtual Private Network Design: A Proof of the Tree Routing Conjecture on Ring Networks. *SIAM Journal on Discrete Mathematics*, 21(2): 482–503, 2007.
20. G. F. Italiano, S. Leonardi, and G. Oriolo. Design of trees in the hose model: The balanced case. *Operations Research Letters*, 34(6): 601–606, 2006.
21. R. Jothi and B. Raghavachari. Improved approximation algorithms for the single-sink buy-at-bulk network design problems. In *SWAT*, 336–348, 2004.
22. D. R. Karger and M. Minkoff. Building steiner trees with incomplete global knowledge. In *FOCS*, 613–623, 2000.
23. A. Meyerson, K. Munagala, and S. Plotkin. Cost-distance: two metric network design. In *FOCS*, 624–630, 2000.
24. T. Rothvoß and L. Sanità. On the complexity of the asymmetric VPN problem. In *APPROX*, 326–338, 2009.
25. C. Swamy and A. Kumar. Primal–dual algorithms for connected facility location problems. In *APPROX*, 256–269, 2002.
26. D. B. Shmoys and K. Talwar. A Constant Approximation Algorithm for the a priori Traveling Salesman Problem. In *IPCO*, 331–343, 2008.
27. K. Talwar. The single-sink buy-at-bulk LP has constant integrality gap. In *IPCO*, 475–486, 2002.
28. A. van Zuylen. Deterministic Sampling Algorithms for Network Design. In *ESA*, 830–841, 2008.