

New Approaches for Virtual Private Network Design*

Friedrich Eisenbrand[†] Fabrizio Grandoni[‡] Gianpaolo Oriolo[§] Martin Skutella[¶]

Abstract

Virtual Private Network Design is the following NP-hard problem. We are given a communication network, represented as a weighted graph with thresholds on the nodes which represent the amount of flow that a node can send to and receive from the network. The task is to reserve capacities at minimum cost and to specify paths between every ordered pair of nodes such that all valid traffic-matrices can be routed along the corresponding paths. Recently, this network design problem has received considerable attention in the literature. It is motivated by the fact that the exact amount of flow which is exchanged between terminals is not known in advance and prediction is often illusive. The main contributions of this paper are as follows:

- Using Hu's 2-commodity flow theorem, we provide a new and considerably stronger lower bound on the cost of an optimum solution. With this lower bound we reanalyze a simple routing scheme which has been described in the literature many times, and provide an improved upper bound on its approximation ratio.
- We present a new randomized approximation algorithm. In contrast to earlier approaches from the literature, the resulting solution does not have tree structure. A combination of our new algorithm with the simple routing scheme yields an expected performance ratio 3.79 for virtual private network design. This is a considerable improvement of the previously best known 5.55 approximation result [Gupta, Kumar, Roughgarden (STOC'03)].
- Our VPND algorithm uses a Steiner tree approximation algorithm as a subroutine. It is known that an optimum Steiner tree can be computed in polynomial time if the number of terminals is logarithmic. Replacing the approximate Steiner tree computation with an exact one whenever the number of terminals is sufficiently small, we finally reduce the approximation ratio to 3.55. To the best of our knowledge, this is the first time that a non-trivial result from exact (exponential) algorithms leads to an improved polynomial-time approximation algorithm.

Key words. approximation algorithms, randomized algorithms, network design

AMS Subject Classifications. 68W20, 68W25

*Preliminary versions of these results appeared in the Proceedings of SODA'05 [7] and of ICALP'05 [8].

[†]Universität Dortmund, Fachbereich Informatik, 44221 Dortmund, Germany,

`friedrich.eisenbrand@cs.uni-dortmund.de`

[‡]Università di Roma "La Sapienza", Dipartimento di Informatica, Via Salaria 113, 00198 Roma, Italy, `grandoni@di.uniroma1.it`

[§]Università di Roma "Tor Vergata", Dipartimento di Ingegneria dell'Impresa, Via del Politecnico 1, 00165, Roma, Italy, `oriolo@disp.uniroma2.it`

[¶]Universität Dortmund, Fachbereich Mathematik, 44221 Dortmund, Germany, `martin.skutella@uni-dortmund.de`

1 Introduction

Consider a communication network which is represented by an undirected graph $G = (V, E)$ with edge weights $c : E \rightarrow \mathbb{R}_+$. Within this network there is a set of terminals $T \subseteq V$, which want to communicate with each other. However, the exact amount of traffic between pairs of terminals is not known in advance. Instead, each terminal $v \in T$ has an associated input and output threshold $b_{in}(v) \in \mathbb{Z}_{\geq 0}$ and $b_{out}(v) \in \mathbb{Z}_{\geq 0}$. A traffic matrix $D \in \mathbb{Q}_{\geq 0}^{T \times T}$ is *valid*, if it respects the lower and upper bounds on the incoming and outgoing traffic of the terminals, i.e., if the following holds for each terminal $i \in T$

$$\sum_{j \in T, j \neq i} D(i, j) \leq b_{out}(i) \quad \text{and} \quad \sum_{j \in T, j \neq i} D(j, i) \leq b_{in}(i).$$

The (*asymmetric*) *Virtual Private Network Design problem* (VPND) defined by G , c and T consists of finding capacities $u(e)$, $e \in E$, and paths P_{ij} for each ordered pair $(i, j) \in T \times T$ such that the following conditions hold.

- i) All valid traffic matrices can be routed without exceeding the installed capacities where all traffic from terminal i to terminal j is routed along path P_{ij} .
- ii) The total cost of the reservation $\sum_{e \in E} u(e)c(e)$ is minimal.

A reservation of capacities $u : E \rightarrow \mathbb{R}_+$ is a *tree reservation*, if the subgraph of G induced by the edges $e \in E$ with $u(e) > 0$ is a tree. A general reservation is sometimes referred to as a *graph reservation*.

The virtual private network design problem is NP-hard, by the following reduction from the Steiner tree problem [11]. Given an instance of the Steiner tree problem, pick a terminal which has to be connected with the other terminals in a Steiner tree. This terminal has thresholds $b_{in}(v) = 0$ and $b_{out}(v) = 1$. All other terminals u of the Steiner tree instance have $b_{in}(u) = 1$ and $b_{out}(u) = 0$. A minimum cost Steiner tree yields also an optimum reservation for this VPND-instance.

The virtual private network design problem was independently defined by Fingerhut et al. [10], and by Duffield et al. [6] and since then, studied by various authors in several variations, which we next discuss. In the following list, the last one (*AsymG*) is the one which we refer to as VPND.

(*SymT*) Symmetric thresholds, tree reservation: In this variant, each terminal $i \in T$ has only one threshold $b(i)$, which is an upper bound on the cumulative amount of traffic that terminal i can send or receive.

The task is to find an optimal tree reservation, which supports all valid traffic matrices. Gupta et al. [11] show that (*SymT*) is polynomially solvable.

(*SymG*) Symmetric thresholds, graph reservation: This variant is defined in the same way as (*SymT*), except that the capacity reservation can be arbitrary and not necessarily a tree. Gupta et al. [11] present a 2-approximation for (*SymG*). It is not known whether *SymG* is NP-hard.

(*BalT*) Balanced thresholds, tree reservation: The thresholds are *balanced*, which means that $\sum_{v \in T} b_{in}(v) = \sum_{v \in T} b_{out}(v)$. The reservation has to be a tree. Italiano et al. [15] show that this variant can be solved in polynomial time.

(*BalG*) Balanced thresholds, graph reservation: The same as (*BalT*), except that an arbitrary graph reservation is allowed.

(*AsymT*) Asymmetric thresholds, tree reservation: This problem is NP-hard [11]. Constant approximation algorithms are presented in [11, 12]. Interestingly, while the algorithm in [11] is deterministic, the algorithm in [12] is randomized and seems difficult to de-randomize.

(*AsymG*) Asymmetric thresholds, graph reservation: This is the VPND problem defined above. We have seen that this problem is NP-hard. The randomized 5.55 approximation result presented in [12] in fact compares the computed tree solution to an optimal graph reservation.

Simplifying assumptions and a lower bound

Following [12], we make some simplifying assumptions without loss of generality. By duplicating nodes, we can assume that each terminal is either a *sender* s , with $b_{out}(s) = 1$ and $b_{in}(s) = 0$, or a *receiver* r , with $b_{out}(r) = 0$ and $b_{in}(r) = 1$. This simplifying assumption is feasible as long as we make sure that the selected paths in our solution between copies of a terminal v and copies of a terminal u are all equal. The algorithms presented in this paper can easily be adapted such as to run in polynomial time even when the thresholds are not polynomially bounded and to satisfy the consistence property described above. Let \mathcal{S} and \mathcal{R} be the set of senders and the set of receivers, respectively. By $S = |\mathcal{S}|$ and $R = |\mathcal{R}|$ we denote the corresponding cardinalities. By symmetry reasons, we always assume $R \geq S$.

We can now interpret VPND as follows. Let $B = (\mathcal{S} \cup \mathcal{R}, E^B)$ be the complete bipartite graph with nodes partitioned into senders and receivers. We have to reserve capacities on the edges of G and we have to specify a set of paths \mathcal{P} in graph G containing one path P_{sr} for each edge $sr \in E^B$ such that each bipartite matching of B can be routed along these paths. In other words, for each edge $e \in E$, the reservation $u(e)$ has to satisfy the following condition:

$$|\{P_{rs} \in \mathcal{P} \mid e \in P_{rs} \text{ and } rs \in M\}| \leq u(e) \quad \text{for each matching } M \text{ in } B. \quad (1)$$

Notice that for a fixed set of paths \mathcal{P} , an optimal reservation of capacity is the component-wise minimal u satisfying (1). (In particular, given \mathcal{P} , the integral capacity $u(e)$ of edge e can be obtained by a maximum bipartite matching computation.) Thus, a solution to VPND can be encoded by only specifying a set of paths \mathcal{P} in G .

The cost of a bipartite matching between senders and receivers in the metric closure of G is obviously a lower bound on OPT , the value of an optimum solution to the VPND -instance. We denote the shortest path distance between nodes u and v of G by $\ell(u, v)$. Thus, if edges (r, s) in B are assigned weights $\ell(r, s)$, then the cost of any matching in B is a lower bound on OPT . This lower bound is used in the analysis of all previous constant factor approximation algorithms for VPND .

Lemma 1 ([12]). *Let $B = (\mathcal{S} + \mathcal{R}, E^B)$ be the complete bipartite graph on the senders and receivers with edge weights $\ell : E^B \rightarrow \mathbb{R}_+$ given by the shortest path distances in the graph G . Then, the weight of any matching in B is a lower bound on OPT .*

Contribution of this paper

The design of good approximation results usually requires two main ingredients: Cleverly constructed algorithms and thoroughly chosen lower bounds on the optimum such that the quality of the computed solutions can be assessed in terms of the lower bounds. We considerably advance the state of the art of approximating VPND by making contributions to both ingredients.

In Section 2 we present a new lower bound which strengthens the one stated in Lemma 1. We prove that the weight of *any matching* (not necessarily bipartite) on the union of the senders and at most S receivers is at most OPT . The edge-weights in the matching are again the shortest path distances in G . This new lower bound relies on an interesting interrelation between a special case of VPND and 2-commodity flows. Its proof is based on an application of Hu's 2-commodity flow theorem [13].

In Section 3 we employ the new lower bound in order to show that the following simple algorithm achieves performance ratio $1 + R/S$: Find a vertex $v \in V$ which minimizes the sum of the shortest path distances between v and the union of senders and receivers; cumulatively install a capacity of one on each such shortest path. One interesting consequence of this result is that (*BalG*), VPND with balanced thresholds and graph reservation, has a 2-approximation. Thus our result improves upon the 3-approximation of Italiano et al. [15] for this problem, and generalizes the 2-approximation for (*SymG*) by Gupta et al. [11].

In Section 4 we present a new randomized algorithm for VPND. The algorithm chooses a random subset of receivers and connects each sender via its own Steiner tree to this subset. The remaining receivers are then connected to the randomly chosen subset of receivers by shortest paths. Due to the Steiner trees for each individual sender, the resulting solution has in general no tree structure. In contrast to our new approach, the previous algorithm by Gupta et al. [12] constructs only one random “high-bandwidth core” which is a Steiner tree with high capacity. All previous approximation algorithms for VPND produce tree solutions.

In Section 4 we show also that our new algorithm in combination with the simple algorithm from above yields a 3.79 randomized approximation algorithm. The previously best known algorithm [12] achieves performance ratio 5.55.

Our VPND algorithm uses a Steiner tree approximation algorithm as a subroutine. Dreyfus and Wagner [5] showed how to compute optimum Steiner trees in polynomial time when the number of terminals is logarithmic. In Section 5, by replacing the approximate Steiner tree computation with an exact one whenever the number of terminals is sufficiently small, we eventually obtain a 3.55 approximation algorithm. To the best of our knowledge, this is the first time a result from exact (exponential) algorithms leads to an improved polynomial-time approximation algorithm, which might be of independent interest.

Related work

We have seen that the Steiner tree problem is a special case of VPND. The current best approximation ratio for the Steiner tree problem is $\rho_{st} < 1.55$ [18]. A related problem is *buy at bulk* network design (see, e.g. [1, 2]). Here, there is a fixed demand $d_{i,j}$ between each pair of nodes in the graph, specifying the amount of flow which has to be sent from i to j . The costs of the capacities however is a concave function on the amount purchased, which reflects “economies of scale”. Gupta et al. [12] consider the single source buy-at-bulk network design problem and present a simple constant factor approximation algorithm.

Another important issue in this context is to cope with arc failures in the network [3, 4]. Italiano, et al. [16] consider the problem of restoring the network, when at most one arc in a tree-solution to VPND might fail and provide a constant factor approximation algorithm.

It is conjectured that (*SymG*) can be solved in polynomial time. It is in fact conjectured that an optimal solution to (*SymT*) is also an optimal solution to (*SymG*). Hurkens et al. [14] show that this conjecture holds for rings. The authors also describe an integer programming formulation for VPND, which proves that a fractional version can be solved in polynomial time. This fractional version allows to specify several paths for each sender-receiver pair and requires the fraction for each of these paths, which describes how the commodity has to be split.

2 A new lower bound via Hu’s 2-commodity flow theorem

This section is devoted to proving a new lower bound on the cost of an optimal solution to VPND. Generalizing Lemma 1, we prove that the cost of an arbitrary (not necessarily bipartite) matching between

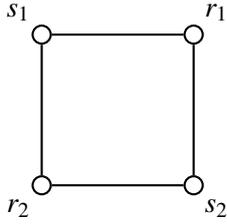


Figure 1: An undirected graph with unit capacities on the edges. While there exist integral 2-commodity flows satisfying unit demands for terminal pairs $\{s_1, r_1\}, \{s_2, r_2\}$ and for terminal pairs $\{s_1, r_2\}, \{s_2, r_1\}$, there is only a half-integral 2-commodity flow for terminal pairs $\{s_1, s_2\}, \{r_1, r_2\}$.

terminals in $\mathcal{S} \cup \mathcal{R}'$ is at most OPT , for any subset of receivers $\mathcal{R}' \subseteq \mathcal{R}$ of cardinality $|\mathcal{R}'| = S$. The proof of this result is based on Hu’s classical 2-commodity flow theorem [13].

Theorem 1 (Hu’s 2-commodity flow theorem). *Let $G = (V, E)$ be a graph and let s_1, r_1, s_2, r_2 be pairs of vertices of G ; let $u : E \rightarrow \mathbb{R}_+$ be a capacity function on the edges and let $d_1, d_2 \in \mathbb{R}_+$. There exists a (fractional) 2-commodity flow of value d_1, d_2 if and only if the cut condition is satisfied, i.e., if and only if $u(\delta(U)) \geq d_1 \chi_1(U) + d_2 \chi_2(U)$ for each $U \subseteq V$. Here $\delta(U)$ denotes the cut induced by U and, for $i \in \{1, 2\}$, $\chi_i(U) = 1$ if the cut $\delta(U)$ separates s_i from r_i , and $\chi_i(U) = 0$ otherwise.*

Hu’s theorem immediately implies the following corollary.

Corollary 1. *Let $G = (V, E)$ be an undirected graph with edge capacity function $u : E \rightarrow \mathbb{R}_+$ and $s_1, s_2, r_1, r_2 \in V$. In the following, all demand values are equal to 1. If there exists a feasible 2-commodity flow for terminal pairs $\{s_1, r_1\}, \{s_2, r_2\}$ and for terminal pairs $\{s_1, r_2\}, \{s_2, r_1\}$, then there also exists a feasible 2-commodity flow for terminal pairs $\{s_1, s_2\}, \{r_1, r_2\}$.*

Proof. Hu’s 2-commodity flow theorem states that there exists a feasible 2-commodity flow if and only if the “cut condition” is satisfied. In the case of unit demands, the cut condition says that, for all $U \subseteq V$, the capacity $u(\delta(U))$ of the cut induced by U must be at least the number of terminal pairs which are separated by the cut.

It thus remains to show that the cut condition holds for terminal pairs $\{s_1, s_2\}, \{r_1, r_2\}$ if it holds for $\{s_1, r_1\}, \{s_2, r_2\}$ and for $\{s_1, r_2\}, \{s_2, r_1\}$. Consider an arbitrary $U \subseteq V$. If the corresponding cut separates neither $\{s_1, s_2\}$ nor $\{r_1, r_2\}$, nothing needs to be shown. If $\delta(U)$ separates one terminal pair, say $\{s_1, s_2\}$, then it separates either $\{s_1, r_1\}$ or $\{s_2, r_1\}$ since s_1 and s_2 lie on different sides of the cut. In particular, the capacity of the cut is at least 1. Finally, if $\delta(U)$ separates both terminal pairs $\{s_1, s_2\}, \{r_1, r_2\}$, then it must either separate $\{s_1, r_1\}$ and $\{s_2, r_2\}$ or $\{s_1, r_2\}$ and $\{s_2, r_1\}$. In both cases the capacity of the cut is at least 2. \square

We remark that Corollary 1 is no longer true if we replace “2-commodity flow” by “integral 2-commodity flow”. A counterexample is given in Figure 1. Even, Itai, and Shamir show that finding an integer 2-commodity flow is NP-hard [9]. On the other hand, Hu’s result states that there always exists a half-integral flow in this case. For a more detailed account of results we refer to Schrijver’s book [19, Chapter 71].

The next lemma shows that a partition of the senders and receivers into k -subsets each, and the “addition” of the optimal solutions of the k subproblems induced by the pairs of subsets, provides a lower bound on the optimal solution.

Lemma 2. Suppose that $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_k$ is a partition of \mathcal{S} and that $\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_k$ is a partition of \mathcal{R} . Let I_i be the VPND -instance on graph G with senders \mathcal{S}_i and receivers \mathcal{R}_i and let OPT_i be the value of an optimal solution to instance I_i . Then the following holds

$$\sum_{i=1}^k \text{OPT}_i \leq \text{OPT}.$$

Proof. Let \mathcal{P} be an optimal set of paths for the original VPND -instance with resulting capacity reservation $u : E \rightarrow \mathbb{R}_+$. The subset $\mathcal{P}_i \subseteq \mathcal{P}$ of paths with endpoints in $\mathcal{S}_i \cup \mathcal{R}_i$ defines a solution to sub-instances I_i with capacity reservation $u_i : E \rightarrow \mathbb{R}_+$. It suffices to show that $u(e) \geq \sum_{i=1}^k u_i(e)$ for each edge $e \in E$.

It follows from (1) that for each $i = 1, 2, \dots, k$

$$u_i(e) = \max_{M_i} |\{P_{rs} \in \mathcal{P}_i \mid e \in P_{rs} \text{ and } rs \in M_i\}|,$$

where M_i ranges over all bipartite matchings between senders \mathcal{S}_i and receivers \mathcal{R}_i . We denote the matching for which the maximum is attained by \tilde{M}_i . Then, the disjoint union $\tilde{M} := \bigcup_{i=1}^k \tilde{M}_i$ is a bipartite matching between senders \mathcal{S} and receivers \mathcal{R} . It thus follows from (1) that

$$u(e) \geq |\{P_{rs} \in \mathcal{P} \mid e \in P_{rs} \text{ and } rs \in \tilde{M}\}| = \sum_{i=1}^k |\{P_{rs} \in \mathcal{P}_i \mid e \in P_{rs} \text{ and } rs \in \tilde{M}_i\}| = \sum_{i=1}^k u_i(e)$$

for each edge $e \in E$. This concludes the proof. \square

We are now ready to prove the following theorem which gives a new lower bound on the cost of an optimal VPND solution.

Theorem 2. Consider an arbitrary subset $\mathcal{R}' \subseteq \mathcal{R}$ of S receivers. Let M be a matching in the complete graph on $\mathcal{S} \cup \mathcal{R}'$. Then

$$\sum_{vw \in M} \ell(v, w) \leq \text{OPT}.$$

Proof. Let $\mathcal{S} = \{s_1, s_2, \dots, s_S\}$ and $\mathcal{R}' = \{r_1, r_2, \dots, r_S\}$. It suffices to prove the claim for perfect matchings. Suppose that the matching consists of the edges

$$s_1 s_2, s_3 s_4, \dots, s_{2k-1} s_{2k}, \quad \text{and} \quad r_1 r_2, r_3 r_4, \dots, r_{2k-1} r_{2k}, \quad \text{and} \quad s_{2k+1} r_{2k+1}, s_{2k+2} r_{2k+2}, \dots, s_S r_S.$$

Consider the following partition of \mathcal{S} and \mathcal{R}' into $S - k$ subsets \mathcal{S}_i and \mathcal{R}'_i each

$$\begin{aligned} \mathcal{S}_i &= \{s_{2i-1}, s_{2i}\}, \mathcal{R}'_i = \{r_{2i-1}, r_{2i}\}, & 1 \leq i \leq k, \\ \mathcal{S}_i &= \{s_i\}, \mathcal{R}'_i = \{r_i\}, & 2k+1 \leq i \leq S. \end{aligned}$$

By Lemma 2, the sum of the optimum solutions OPT_i of the VPND sub-instances I_i with senders \mathcal{S}_i and receivers \mathcal{R}'_i is a lower bound on OPT . Thus we only need to prove that $\ell(s_{2i-1}, s_{2i}) + \ell(r_{2i-1}, r_{2i}) \leq \text{OPT}_i$ for each $1 \leq i \leq k$.

An optimum solution to the subproblem I_i , $1 \leq i \leq k$, is a reservation of capacities that supports 2-commodity flows with unit demands for the terminal pairs $\{s_{2i-1}, r_{2i-1}\}$, $\{s_{2i}, r_{2i}\}$, and for the terminal pairs $\{s_{2i-1}, r_{2i}\}$, $\{s_{2i}, r_{2i-1}\}$. By Corollary 1, it must also support a 2-commodity flow for the terminal pairs $\{s_{2i-1}, s_{2i}\}$, $\{r_{2i-1}, r_{2i}\}$. Therefore, the cost of this solution is at least $\ell(s_{2i-1}, s_{2i}) + \ell(r_{2i-1}, r_{2i})$. This concludes the proof. \square

3 The quality of a simple routing scheme

Consider the following simple VPND algorithm.

Algorithm 1: Simple Routing Scheme

- (1) Compute a vertex $v \in V$ such that $\sum_{s \in \mathcal{S}} \ell(s, v) + \sum_{r \in \mathcal{R}} \ell(r, v)$ is minimal.
- (2) Add one unit of capacity along the shortest path between each $u \in \mathcal{R} \cup \mathcal{S}$ and v .

Algorithm 1 selects the vertex $v \in V$ which minimizes the sum of the distances from v to the union of the senders and receivers, and reserves one unit of capacity along each of the shortest paths. Note that the effects of installing capacities along the shortest paths is cumulative. In other words, if k shortest paths share the same edge, the algorithm adds k units of capacity to that edge. Moreover, we assume the shortest paths are computed with a consistent tie-breaking rule. This way, the edges with nonzero capacity form a tree.

Algorithm 1 produces an optimal tree reservation in the symmetric case (*SymT*) [11] and in the balanced case (*BalT*) [15]. In the symmetric case, Gupta et al.[11] showed that the tree produced by the algorithm is a 2-approximate solution to the optimum graph reservation. Italiano et al. [15] show that, in the balanced case, the produced tree is a 3-approximate solution to the optimum graph reservation.

In this Section, we apply our new lower-bound result, to show that the algorithm produces a tree-solution, which is within $1 + R/S$ from the optimum graph reservation. Thus also (*BalG*) can be approximated within a factor of two. We first need the following simple lower bound on *OPT*. Let $\ell(E') := \sum_{uv \in E'} \ell(u, v)$ for any subset of edges E' .

Lemma 3. *The sum of the distances between each pair sender-receiver is at most R times the optimum:*

$$\sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}} \ell(s, r) \leq R \cdot OPT. \quad (2)$$

Proof. Let $B = (\mathcal{S} + \mathcal{R}, E^B)$ be the complete bipartite graph on the senders and receivers with edge weights $\ell : E^B \rightarrow \mathbb{R}_+$ given by the shortest path distances in the graph G . The edges of B can be partitioned into a set \mathcal{M} of R matchings. By Lemma 1, the cost $\ell(M)$ of each $M \in \mathcal{M}$ is at most *OPT*. Hence

$$\sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}} \ell(s, r) = \sum_{M \in \mathcal{M}} \ell(M) \leq R \cdot OPT.$$

□

We are now ready to bound the approximation ratio provided by Algorithm 1.

Theorem 3. *Algorithm 1 is a $1 + R/S$ approximation algorithm for VPND.*

Proof. Let $G^m = (\mathcal{R} \cup \mathcal{S}, E^m)$ be the metric closure of $\mathcal{R} \cup \mathcal{S}$, i.e., the complete graph on $\mathcal{R} \cup \mathcal{S}$ with edge weight $\ell(u, v)$ given by the shortest path distance between u and v in G . We show that there exists a node $u \in \mathcal{R} \cup \mathcal{S}$ such that the cost of the complete star centered at u satisfies

$$\sum_{v \in \mathcal{R} \cup \mathcal{S}} \ell(u, v) \leq (1 + R/S) OPT.$$

If $R = S$, the edges of E^m can be partitioned into $2S - 1$ perfect matchings. By Theorem 2, the weight of each matching is a lower bound on *OPT*. Since each edge is contained in exactly two stars of G^m ,

$$\sum_{v \in \mathcal{R} \cup \mathcal{S}} \sum_{u \in \mathcal{R} \cup \mathcal{S}} \ell(v, u) = 2\ell(E^m) \leq 2(2S - 1) OPT.$$

Thus there must exist one star, whose weight is at most $2(2S - 1)OPT/(2S) < 2OPT$.

Suppose for the remainder of the proof that $R > S$. In the following we denote by \mathcal{M}_S and \mathcal{M}_R the set of matchings of cardinality at most $\lfloor S/2 \rfloor$ involving only senders and only receivers, respectively. Theorem 2 implies the inequality

$$\ell(M_S) + \ell(M_R) \leq OPT \quad \text{for each } M_S \in \mathcal{M}_S, M_R \in \mathcal{M}_R. \quad (3)$$

In consideration of (3), we distinguish two cases.

(A) $\ell(M_S) \leq OPT/2$ for each $M_S \in \mathcal{M}_S$.

Consider the subgraph G_S^m of G^m which is induced by the senders. The edges E_S^m of G_S^m can be partitioned into S matchings. Therefore

$$\sum_{s' \in S} \sum_{s \in S} \ell(s', s) = 2\ell(E_S^m) \leq 2S \frac{OPT}{2} = S OPT. \quad (4)$$

Combining (4) with (2), one obtains

$$\sum_{s' \in S} \left(\sum_{s \in S} \ell(s', s) + \sum_{r \in R} \ell(s', r) \right) \leq (S + R) OPT.$$

Therefore there is a sender s^* such that

$$\sum_{v \in S \cup R} \ell(s^*, v) \leq (1 + R/S) OPT.$$

(B) $\ell(M_S) > OPT/2$ for some maximum weight matching $M_S \in \mathcal{M}_S$.

Let G_R^m denote the subgraph of G^m which is induced by the receivers. We will show below that, for any maximal matching \tilde{M} in G_R^m ,

$$\ell(\tilde{M}) \leq (R/S) OPT/2. \quad (5)$$

Since the edges of G_R^m can be partitioned into at most R maximal matchings, we can then argue in a similar manner as in case (A) that

$$\sum_{r' \in R} \sum_{r \in R} \ell(r, r') \leq 2R \frac{R OPT}{S} = (R^2/S) OPT.$$

Together with (2) this implies

$$\sum_{r' \in R} \left(\sum_{s \in S} \ell(s, r') + \sum_{r \in R} \ell(r, r') \right) \leq (R + R^2/S) OPT,$$

from which we conclude that there is a receiver r^* satisfying

$$\sum_{v \in S \cup R} \ell(v, r^*) \leq (1 + R/S) OPT.$$

It remains to prove (5). We distinguish between the two subcases in which S is even and S is odd.

(B.1) S is even.

Theorem 2, together with the assumption $\ell(M_S) > OPT/2$, implies $\ell(M_R) \leq OPT/2$ for each matching

$M_{\mathcal{R}} \in \mathcal{M}_{\mathcal{R}}$. Note that $|\tilde{M}| = \lfloor R/2 \rfloor \geq S/2$. Consider the $S/2$ most expensive edges M' of \tilde{M} . Since $M' \in \mathcal{M}_{\mathcal{R}}$, $\ell(M') \leq OPT/2$. Hence the average cost of one edge of \tilde{M} is upper bounded by $(OPT/2)/(S/2) = OPT/S$. Since \tilde{M} has at most $R/2$ edges, we get

$$\ell(\tilde{M}) \leq |\tilde{M}| \frac{OPT}{S} \leq \frac{R}{2} \frac{OPT}{S} = (R/S)OPT/2.$$

(B.2) S is odd.

There is a sender s^* which is missed by the maximum cost matching M_S of G_S^m . For each matching $M_{\mathcal{R}} \in \mathcal{M}_{\mathcal{R}}$ and $r_1^* \in \mathcal{R}$ which is not matched by $M_{\mathcal{R}}$, Theorem 2 yields

$$\ell(M_S) + \ell(M_{\mathcal{R}}) + \ell(s^*, r_1^*) \leq OPT,$$

and hence

$$\ell(M_{\mathcal{R}}) + \ell(s^*, r_1^*) \leq OPT/2.$$

Consider any other receiver r_2^* which is not matched by $M_{\mathcal{R}}$. This receiver must exist since $R > S$. By the triangle inequality one has $\ell(r_1^*, r_2^*) \leq \ell(s^*, r_1^*) + \ell(s^*, r_2^*)$. As a result we get

$$\ell(M_{\mathcal{R}}) + 1/2\ell(r_1^*, r_2^*) \leq OPT/2 \tag{6}$$

for each matching $M_{\mathcal{R}} \in \mathcal{M}_{\mathcal{R}}$ and receivers r_1^*, r_2^* which are missed by $M_{\mathcal{R}}$.

Now consider the $(S-1)/2$ most expensive edges M' of \tilde{M} , and let e' be the next most expensive edge of \tilde{M} . Since $M' \in \mathcal{M}_{\mathcal{R}}$, by (6),

$$\ell(M') + 1/2\ell(e') \leq OPT/2.$$

It follows that half the average cost of one edge of \tilde{M} is upper bounded by $(OPT/2)/(S-1+1) = OPT/(2S)$, from which we conclude

$$\ell(\tilde{M}) \leq 2|\tilde{M}| \frac{OPT}{2S} \leq 2 \frac{R}{2} \frac{OPT}{2S} = (R/S)OPT/2.$$

□

4 A new algorithm for VPND

In Section 3 we described an algorithm which guarantees a good approximation ratio for R close to S . Here we present a better approximation algorithm in case R is sufficiently larger than S .

The algorithm by Gupta et al. [12] constructs one random “high bandwidth core”, i.e. a small Steiner tree with high capacity, where the terminals are a random sender and a random subset of receivers. Such a Steiner tree collects and distributes the demands from outside, and routes them along its high capacity paths. Our algorithm is also based on Steiner-tree computations, but we proceed by connecting each sender to a previously sampled subset of receivers via S distinct Steiner trees of low capacity, and by connecting the other receivers along their shortest paths to the sampled subset (see Figure 2). More precisely, our algorithm works as follows.

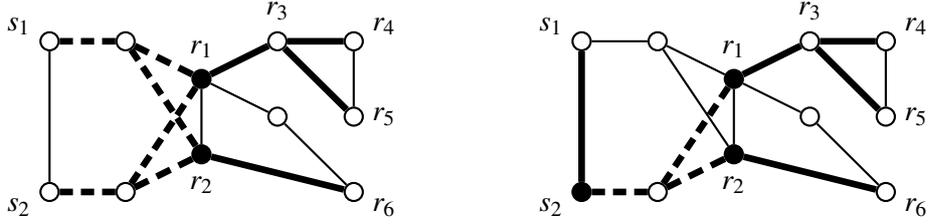


Figure 2: Intuitive comparison between Algorithm 2 (on the left) and the algorithm in [12]. Black nodes form the randomly sampled subsets. Positive capacity is reserved on thick edges only: dashed thick edges correspond to Steiner trees, while full ones to shortest paths.

Algorithm 2:

- (1) Partition \mathcal{R} into S subsets uniformly at random. Select one subset \mathcal{R}' uniformly at random, among the non-empty subsets in the partition.
- (2) For each sender $s \in \mathcal{S}$, compute a ρ_{st} -approximate Steiner tree $T(s)$ on $\{s\} \cup \mathcal{R}'$, and add one unit of capacity to each edge of $T(s)$.
- (3) Add one unit of capacity along the shortest path between each receiver $r \in \mathcal{R}$ and \mathcal{R}' .

It remains to specify the path between each sender-receiver pair (s, r) . Assume that the shortest paths are computed with a consistent tie-breaking rule. Let r^* be the receiver in \mathcal{R}' which is closest to r . The path P_{sr} between s and r is obtained by concatenating the (simple) path between s and r^* in $T(s)$, with the shortest path between r^* and r .

The thereby produced solution is not a tree solution. Though an optimal tree-solution is a constant-factor approximation to an optimal graph-solution, it is known [11] that an optimal solution to *AsymT* is not an optimal solution to *VPND*. All previous constant factor approximation algorithms for *VPND* however [12, 7], produce tree reservations.

Before we proceed with the analysis of Algorithm 2, we state a corollary of Lemma 2. Here, given a subset V' of nodes, we denote the cost of the optimum Steiner tree on V' by $st(V')$.

Corollary 2 ([12]). *Let $\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_s$ be a partition of \mathcal{R} into S (disjoint) subsets. Consider an arbitrary perfect matching between \mathcal{S} and this family of subsets. Let $\mathcal{R}(s)$ be the subset matched with sender s . The sum over \mathcal{S} of the costs of the optimum Steiner trees on $\{s\} \cup \mathcal{R}(s)$ is a lower bound on *OPT*:*

$$\sum_{s \in \mathcal{S}} st(\{s\} \cup \mathcal{R}(s)) \leq OPT.$$

Proof. This follows from Lemma 2 since a solution of a sub-instance $\{s\}, \mathcal{R}(s)$ contains a Steiner tree with terminals $\{s\} \cup \mathcal{R}(s)$. □

Theorem 4. *Algorithm 2 is a $(2 + \rho_{st}) / (1 - e^{-R/S})$ approximation algorithm for *VPND*.*

Theorem 4 is a straightforward consequence of the following lemmas.

Lemma 4. *For a uniformly chosen random sender s' ,*

$$E[st(\{s'\} \cup \mathcal{R}')] \leq \frac{OPT}{S(1 - e^{-R/S})}.$$

Proof. Consider the following random process. For each receiver r , we *assign* r to a sender s chosen uniformly at random. Let $\mathcal{R}(s)$ be the subset of receivers assigned to s . Note that the subsets $\mathcal{R}(s)$ partition \mathcal{R} into S (possibly empty) subsets. Thus, by Corollary 2,

$$\sum_{s \in \mathcal{S}} st(\{s\} \cup \mathcal{R}(s)) \leq OPT.$$

This means that, for the random sender s' ,

$$E[st(\{s'\} \cup \mathcal{R}(s'))] \leq OPT/S.$$

Let A denote the event that $\mathcal{R}(s')$ is empty. By elementary probability theory,

$$E[st(\{s'\} \cup \mathcal{R}(s'))] = P(A) E[st(\{s'\} \cup \mathcal{R}(s')) | A] + P(\bar{A}) E[st(\{s'\} \cup \mathcal{R}(s')) | \bar{A}].$$

Now observe that

$$P(\bar{A}) = 1 - (1 - 1/S)^R \geq 1 - e^{-R/S}.$$

Moreover

$$E[st(\{s'\} \cup \mathcal{R}(s')) | A] = E[st(\{s'\})] = 0.$$

Thus

$$E[st(\{s'\} \cup \mathcal{R}(s')) | \bar{A}] = \frac{E[st(\{s'\} \cup \mathcal{R}(s'))]}{P(\bar{A})} \leq \frac{OPT}{S(1 - e^{-R/S})}.$$

The claim follows by observing that, given \bar{A} , $\mathcal{R}(s')$ and \mathcal{R}' are identically distributed. Thus

$$E[st(\{s'\} \cup \mathcal{R}')] = E[st(\{s'\} \cup \mathcal{R}(s')) | \bar{A}] \leq \frac{OPT}{S(1 - e^{-R/S})}.$$

□

Lemma 5. *The expected cost of the capacity installed in the second step of Algorithm 2 is at most*

$$\rho_{st} OPT / (1 - e^{-R/S}).$$

Proof. The expected cost considered is

$$E\left[\sum_{s \in \mathcal{S}} c(T(s))\right],$$

where $c(T(s))$ is the cost of the Steiner tree $T(s)$. Of course

$$c(T(s)) \leq \rho_{st} st(\{s\} \cup \mathcal{R}').$$

Let s' be a sender chosen uniformly at random. By Lemma 4,

$$E\left[\sum_{s \in \mathcal{S}} c(T(s))\right] \leq \rho_{st} E\left[\sum_{s \in \mathcal{S}} st(\{s\} \cup \mathcal{R}')\right] = \rho_{st} S E[st(\{s'\} \cup \mathcal{R}')] \leq \rho_{st} S \frac{OPT}{S(1 - e^{-R/S})} = \frac{\rho_{st} OPT}{1 - e^{-R/S}}.$$

□

Lemma 6. *The expected cost of the capacity installed in the third step of Algorithm 2 is at most*

$$2 OPT / (1 - e^{-R/S}).$$

Proof. Let r' be an arbitrary receiver in \mathcal{R}' . Observe that the probability of any other receiver r'' of being in \mathcal{R}' , given that r' is in \mathcal{R}' , is $1/S$. In fact, let \mathcal{R}^* be a random (possibly empty) partition element. Then

$$P(r' \in \mathcal{R}') = P(r' \in \mathcal{R}^* \mid \mathcal{R}^* \neq \emptyset) = \frac{P(r' \in \mathcal{R}^* \cap \mathcal{R}^* \neq \emptyset)}{P(\mathcal{R}^* \neq \emptyset)} = \frac{P(r' \in \mathcal{R}^*)}{P(\mathcal{R}^* \neq \emptyset)} = \frac{1/S}{1 - (1 - 1/S)^R}.$$

By basically the same argument

$$P(r' \in \mathcal{R}' \cap r'' \in \mathcal{R}') = \frac{(1/S)^2}{1 - (1 - 1/S)^R}.$$

Thus

$$P(r'' \in \mathcal{R}' \mid r' \in \mathcal{R}') = \frac{P(r'' \in \mathcal{R}' \cap r' \in \mathcal{R}')}{P(r' \in \mathcal{R}')} = \frac{(1/S)^2 / (1 - (1 - 1/S)^R)}{(1/S) / (1 - (1 - 1/S)^R)} = \frac{1}{S}.$$

Now consider the following random process. In step t , let \mathcal{A}_t be the subset of receivers considered so far, and let \mathcal{B}_t be the bought receivers in \mathcal{A}_t . Initially $\mathcal{A}_1 = \mathcal{B}_1 = \{r'\}$, where r' is a random receiver. In step t , we consider the receiver $r_t \in \mathcal{R} \setminus \mathcal{A}_t$ which is closest to \mathcal{B}_t , and we set $\mathcal{A}_{t+1} = \mathcal{A}_t \cup \{r_t\}$. Then, with probability $1/S$, we *buy* r_t , that is

- we add S units of capacity along the shortest path from r_t to \mathcal{B}_t ;
- we set $\mathcal{B}_{t+1} = \mathcal{B}_t \cup \{r_t\}$.

Otherwise, we *rent* r_t , that is

- we add one unit of capacity along the shortest path from r_t to \mathcal{B}_t ;
- we set $\mathcal{B}_{t+1} = \mathcal{B}_t$.

Note that, at the end of the process, the set of bought receivers \mathcal{B}' has the same distribution as the set \mathcal{R}' of selected receivers. Let $\ell(v, V') = \min_{v' \in V'} \{\ell(v, v')\}$ denote the minimum distance between a node v and a subset of nodes V' . The expected cost of the third step of the algorithm is upper bounded by the expected cost c_{rent} of renting receivers

$$E\left[\sum_{r \in \mathcal{R} \setminus \mathcal{R}'} \ell(r, \mathcal{R}')\right] = E\left[\sum_{r_t \in \mathcal{R} \setminus \mathcal{B}'} \ell(r_t, \mathcal{B}')\right] \leq E\left[\sum_{r_t \in \mathcal{R} \setminus \mathcal{B}_t} \ell(r_t, \mathcal{B}_t)\right] = E\left[\sum_{r_t \in \mathcal{R}} \left(1 - \frac{1}{S}\right) \ell(r_t, \mathcal{B}_t)\right] = c_{rent},$$

where the inequality comes from the fact that $\mathcal{B}_t \subseteq \mathcal{B}'$ for any t .

The expected cost c_{buy} of buying receivers is an upper bound on c_{rent} :

$$c_{rent} = E\left[\sum_{r_t \in \mathcal{R}} \left(1 - \frac{1}{S}\right) \ell(r_t, \mathcal{B}_t)\right] \leq E\left[\sum_{r_t \in \mathcal{R}} \left(\frac{1}{S}\right) S \ell(r_t, \mathcal{B}_t)\right] = E\left[\sum_{r_t \in \mathcal{B}'} S \ell(r_t, \mathcal{B}_t)\right] = c_{buy}.$$

Note that c_{buy} is equal to S times the expected cost of the Steiner tree on \mathcal{B}' which is obtained via the minimum spanning tree heuristic (starting Prim's algorithm from node r'). It is well known that this heuristic provides a 2-approximate solution [17]. Thus:

$$c_{buy} \leq 2SE[st(\mathcal{B}')] = 2SE[st(\mathcal{R}')].$$

For a random sender s' , by Lemma 4,

$$E[st(\mathcal{R}')] \leq E[st(\{s'\} \cup \mathcal{R}')] \leq \frac{OPT}{S(1 - e^{-R/S})}.$$

Altogether,

$$E\left[\sum_{r \in \mathcal{R}} \ell(r, \mathcal{R}')\right] \leq c_{rent} \leq c_{buy} \leq 2S \frac{OPT}{S(1 - e^{-R/S})} = \frac{2OPT}{1 - e^{-R/S}}.$$

□

In Section 3 we described a $(1 + R/S)$ approximation algorithm. The factors $1 + R/S$ and $(2 + \rho_{st})/(1 - e^{-R/S})$ are equal for $R/S = 2.78\dots < 2.79$. Note that $1 + R/S$ is increasing in R/S and $(2 + \rho_{st})/(1 - e^{-R/S})$ is decreasing in R/S . It follows that a combination (taking the minimum cost solution) of Algorithms 1 and 2 has an expected approximation guarantee of 3.79, which is a considerable improvement compared to the 5.55 approximation ratio achieved by Gupta et al. [12].

Theorem 5. *The combination (taking the cheaper solution) of Algorithm 1 and Algorithm 2 is an expected 3.79 approximation algorithm for VPND.*

5 Computing optimum Steiner trees

The performance ratio of Algorithm 1 and Algorithm 2 meet roughly at $R/S \simeq 2.78$. In this case, the sampled set \mathcal{R}' of receivers has expected constant size. An optimum Steiner tree on a graph with n nodes and t terminals can be computed in $O(3^t n + 2^t n^2 + n^3)$ time with the Dreyfus-Wagner algorithm [5]. This suggests the following variant of Algorithm 2, which computes optimal Steiner trees in step (2), instead of ρ_{st} -approximate ones, whenever $|\mathcal{R}'| \leq \log n$. Here n is the number of nodes in the original graph G . Without loss of generality, we assume n is sufficiently large.

Algorithm 3:

- (1) Partition \mathcal{R} into S subsets uniformly at random. Select one subset \mathcal{R}' uniformly at random, among the non-empty subsets in the partition.
- (2) For each sender $s \in \mathcal{S}$, compute a ρ_{st} -approximate Steiner tree $T(s)$ on $\{s\} \cup \mathcal{R}'$ if $|\mathcal{R}'| > \log n$. Otherwise, compute an optimum Steiner tree $T(s)$ on $\{s\} \cup \mathcal{R}'$. In both cases, add one unit of capacity to each edge of $T(s)$.
- (3) Add one unit of capacity along the shortest path between each receiver $r \in \mathcal{R}$ and \mathcal{R}' .

Clearly, Algorithm 3 is a polynomial time algorithm whose expected approximation guarantee is not worse than the one of Algorithm 2. In particular, if R/S is very large, say $R/S \geq \log \log n$, the approximation achieved is

$$\frac{2 + \rho_{st}}{1 - e^{-R/S}} \leq \frac{2 + \rho_{st}}{1 - 1/\log n} = 2 + \rho_{st} + o(1) < 3.55.$$

What can be said about the approximation guarantee if $R/S \leq \log \log n$? In that case, the expected size of \mathcal{R}' is $1 + (R - 1)/S < 1 + \log \log n$. The probability that the size of \mathcal{R}' exceeds $\log n$ is at most $(1 + \log \log n)/\log n$ by Markov's inequality. Thus with high probability Algorithm 3 computes optimum Steiner trees. When this happens, the approximation ratio is bounded by $3.325\dots < 3.326$. In the very unlikely event that Algorithm 3 computes ρ_{st} -approximate Steiner trees, the approximation guaranteed by Algorithm 1 alone is $O(\log \log n)$. Hence this event contributes to the expected approximation ratio with $o(1)$ only. This is the intuition behind the following theorem.

Theorem 6. *For a sufficiently large n , the combination (taking the cheaper solution) of Algorithm 1 and Algorithm 3 is an expected 3.55-approximation algorithm for VPND.*

Proof. Following the discussion above, if $R/S > \log \log n$, the expected approximation ratio achieved is $2 + \rho_{st} + o(1) < 3.55$. Then we can restrict our analysis to the case $R/S \leq \log \log n$. Let APX denote the expected cost of the solution computed. One has,

$$APX \leq E \left[\min \left\{ (1 + R/S) OPT, \sum_{s \in \mathcal{S}} c(T(s)) + \sum_{r \in \mathcal{R}} \ell(r, \mathcal{R}') \right\} \right].$$

Of course

$$APX \leq (1 + R/S) OPT, \quad (7)$$

and

$$APX \leq E \left[\sum_{r \in \mathcal{R}} \ell(r, \mathcal{R}') \right] + E \left[\min \left\{ (1 + R/S) OPT, \sum_{s \in \mathcal{S}} c(T(s)) \right\} \right].$$

By Lemma 6,

$$E \left[\sum_{r \in \mathcal{R}} \ell(r, \mathcal{R}') \right] \leq \frac{2 OPT}{1 - e^{-R/S}}. \quad (8)$$

Let A denote the event that $R' = |\mathcal{R}'| \leq \log n$. By elementary probability theory one has

$$\begin{aligned} E \left[\min \left\{ (1 + R/S) OPT, \sum_{s \in \mathcal{S}} c(T(s)) \right\} \right] &\leq P(\bar{A}) E \left[(1 + R/S) OPT \mid \bar{A} \right] \\ &\quad + P(A) E \left[\sum_{s \in \mathcal{S}} c(T(s)) \mid A \right]. \end{aligned} \quad (9)$$

We now consider both terms separately. By Markov's inequality one has $P(\bar{A}) \leq (1 + \log \log n) / \log n$. Thus

$$P(\bar{A}) E \left[(1 + R/S) OPT \mid \bar{A} \right] \leq P(\bar{A}) E \left[(1 + \log \log n) OPT \mid \bar{A} \right] \leq \frac{(1 + \log \log n)^2}{\log n} OPT.$$

Given A , Algorithm 3 computes optimal Steiner trees $T(s)$ of cost $st(\{s\} \cup \mathcal{R}')$. Also $E[st(\{s\} \cup \mathcal{R}') \mid R' \leq h]$ is a non-decreasing function of h . Thus, from the proof of Lemma 5, the second term on the right of (9) can be bounded by

$$P(A) E \left[\sum_{s \in \mathcal{S}} c(T(s)) \mid A \right] \leq E \left[\sum_{s \in \mathcal{S}} st(\{s\} \cup \mathcal{R}') \right] \leq \frac{OPT}{1 - e^{-R/S}}.$$

One therefore has

$$\begin{aligned} E \left[\min \left\{ (1 + R/S) OPT, \sum_{s \in \mathcal{S}} c(T(s)) \right\} \right] &\leq \frac{(1 + \log \log n)^2}{\log n} OPT + \frac{OPT}{1 - e^{-R/S}} \\ &\leq \left(1 + \frac{(1 + \log \log n)^2}{\log n} \right) \frac{OPT}{1 - e^{-R/S}}. \end{aligned} \quad (10)$$

Combining (7), (8), and (10), we conclude that

$$APX \leq \min \left\{ (1 + R/S) OPT, \left(3 + \frac{(1 + \log \log n)^2}{\log n} \right) \frac{OPT}{1 - e^{-R/S}} \right\}.$$

Thus, for a sufficiently large n , the expected approximation ratio for $R/S \leq \log \log n$ is upper bounded by $3.325 \dots < 3.326$. Altogether, the approximation ratio is

$$\min \{ 3.326, 2 + \rho_{st} + o(1) \} < 3.55.$$

□

Corollary 3. *There is an expected 3.55-approximation algorithm for VPND.*

Proof. When n is upper bounded by a constant, the optimum solution can be computed in polynomial time by trivial enumeration. The claim follows from Theorem 6. \square

References

- [1] M. Andrews and L. Zhang. The access network design problem. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 40–49, 1998.
- [2] B. Awerbuch and Y. Azar. Buy-at-bulk network design. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 542–547, 1997.
- [3] G. Brightwell, G. Oriolo, and F. B. Shepherd. Reserving resilient capacity in a network. *SIAM Journal on Discrete Mathematics*, 14(4):524–539, 2001.
- [4] G. Brightwell, G. Oriolo, and F. B. Shepherd. Reserving resilient capacity for a single commodity with upper-bound constraints. *Networks*, 41(2):87–96, 2003.
- [5] S. E. Dreyfus and R. A. Wagner. The Steiner problem in graphs. *Networks*, 1:195–207, 1971/72.
- [6] 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 *Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication*, pages 95–108, 1999.
- [7] F. Eisenbrand and F. Grandoni. An improved approximation algorithm for virtual private network design. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 928–932, 2005.
- [8] F. Eisenbrand, F. Grandoni, G. Oriolo, and M. Skutella. New approaches for virtual private network design. In *International Colloquium on Automata, Languages and Programming (ICALP)*, pages 1152–1162, 2005.
- [9] S. Even, A. Itai, and A. Shamir. On the complexity of timetable and multi-commodity flow problems. *SIAM Journal on Computing*, 5(4):691–703, 1976.
- [10] J. A. Fingerhut, S. Suri, and J. S. Turner. Designing least-cost nonblocking broadband networks. *Journal of Algorithms*, 24(2):287–309, 1997.
- [11] A. Gupta, J. Kleinberg, A. Kumar, R. Rastogi, and B. Yener. Provisioning a virtual private network: a network design problem for multicommodity flow. In *ACM Symposium on the Theory of Computing (STOC)*, pages 389–398, 2001.
- [12] A. Gupta, A. Kumar, and T. Roughgarden. Simpler and better approximation algorithms for network design. In *ACM Symposium on the Theory of Computing (STOC)*, pages 365–372, 2003.
- [13] T. C. Hu. Multi-commodity network flows. *Operations Research*, 11:344–360, 1963.
- [14] C. Hurkens, J. Keijsper, and L. Stougie. Virtual private network design: A proof of the tree routing conjecture on ring networks. In *International Conference on Integer Programming and Combinatorial Optimization (IPCO)*, pages 407–421, 2005.

- [15] G. Italiano, S. Leonardi, and G. Oriolo. Design of networks in the hose model. In *Proceedings of ARACNE*, pages 65–76, 2002.
- [16] G. F. Italiano, R. Rastogi, and B. Yener. Restoration algorithms for virtual private networks in the hose model. In *Annual Joint Conference of the IEEE Computer and Communications Society (INFOCOM)*, pages 131–139, 2002.
- [17] L. Kou, G. Markowsky, and L. Berman. A fast algorithm for Steiner trees. *Acta Informatica*, 15(2):141–145, 1981.
- [18] G. Robins and A. Zelikovsky. Improved steiner tree approximation in graphs. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 770–779, 2000.
- [19] A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*, volume 24 of *Algorithms and Combinatorics*. Springer, Berlin, 2003.