

New Approaches for Virtual Private Network Design

Friedrich Eisenbrand^{1*}, Fabrizio Grandoni², Gianpaolo Oriolo³, and Martin Skutella^{4**}

¹ Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, D-66123 Saarbrücken, Germany eisen@mpi-sb.mpg.de

² Università di Roma “La Sapienza”, Dipartimento di Informatica, Via Salaria 113, 00198 Roma, Italy grandoni@di.uniroma1.it

³ Università di Roma “Tor Vergata”, Dipartimento di Ingegneria dell’Impresa, Via del Politecnico 1, 00165, Roma, Italy oriolo@disp.uniroma2.it

⁴ Universität Dortmund, Fachbereich Mathematik, 44221 Dortmund, Germany, martin.skutella@uni-dortmund.de

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 lower bound on the cost of an optimum solution.
- Using this lower bound we reanalyze a simple routing scheme which has been described in the literature many times and provide a considerably stronger upper bound on its approximation ratio.
- We present a new randomized approximation algorithm for which, in contrast to earlier approaches from the literature, the resulting solution does not have tree structure.
- Finally we show that a combination of our new algorithm with the simple routing scheme yields a randomized algorithm with expected performance ratio 3.55 for virtual private network design. This is a considerable improvement of the previously best known approximation results (5.55 STOC’03, 4.74 SODA’05).

* Partially supported by the DFG as part of the Transregional Collaborative Research Center “Automatic Verification and Analysis of Complex Systems” (SFB/TR 14 AVACS)

** Partially supported by the DFG Focus Program 1126, “Algorithmic Aspects of Large and Complex Networks”, grant no. SK 58/4-1.

1 Introduction

Consider a communication network which is represented by an undirected graph $G = (V, E)$ with edge costs $c : E \rightarrow \mathbb{R}_+$. 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 , T and b 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 capacity 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 [9]. Given an instance of the Steiner tree problem, pick a terminal v which has to be connected with the other terminals in a Steiner tree. This terminal is assigned thresholds $b_{in}(v) := 0$ and $b_{out}(v) := 1$. All other terminals of the Steiner tree instance have input threshold one and output threshold zero. It is easy to see that a minimum cost Steiner tree yields an optimum reservation for this VPND-instance.

The virtual private network design problem has independently been defined by Fingerhut et al. [8] in the context of broadband ATM networks, and by Duffield et al. [5] (VPN “hose” model). Since then, it has been studied by various authors in several variations which we discuss next. In the following list, the last variant (*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. [9] 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. [9] 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. [13] show that this variant can be solved in polynomial time.
- (*AsymT*) Asymmetric thresholds, tree reservation: This problem is NP-hard [9]. Constant approximation algorithms are presented in [9,10,6]. Interestingly, while the algorithm in [9] is deterministic, the other algorithms are randomized and seem 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 approximation results presented in [10,6] in fact compare the computed tree solution to an optimal graph reservation. The current best approximation algorithm is the one in [6] which achieves an expected performance ratio 4.74.

Simplifying assumptions and a lower bound

Following [10], 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. Let \mathcal{S} and \mathcal{R} be the set of senders and the set of receivers, respectively. Let $S = |\mathcal{S}|$ and $R = |\mathcal{R}|$ denote the corresponding cardinalities. 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. Moreover, 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 analyses of all previous constant factor approximation algorithms for VPND.

Lemma 1 ([10]). *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 generalizes and thus 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 [11].

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 $v \in V$ whose shortest path tree to the union of senders and receivers is of minimal cost; cumulatively install a capacity of one on each shortest path. One interesting consequence of this result is that $(BalG)$, VPND with balanced thresholds and graph reservation, has a 2-approximation. Our result improves upon the 3-approximation of Italiano et al. [13] for this problem and generalizes the 2-approximation for $(SymG)$ by Gupta et al. [9].

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 receiver, the resulting solution has in general no tree structure. In contrast to our new approach, the previous algorithm by Gupta et al. [10] and its refinement in [6] construct only one ‘high-bandwidth’ core which is a Steiner tree with high capacity. In particular, all previous approximation algorithms for VPND produce tree solutions.

Finally, we can show that our new algorithm in combination with the simple algorithm from above yields a 3.55 randomized approximation algorithm. The previously best known algorithm [6] achieves performance ratio 4.74.

Related work

As discussed above, VPND can be seen as a generalization of the Steiner tree problem. The currently best known approximation ratio for the Steiner tree problem is $\rho < 1.55$ [15]. A related problem is *buy at bulk* network design (see, e.g. [1,2]). In this problem, there is a fixed demand d_{ij} 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. [10] consider the single source buy-at-bulk network design problem and present a constant factor approximation algorithm.

Another important issue in this context is to cope with edge failures [3]. Italiano et al. [14] consider the problem of restoring the network, when at most one edge in a tree-solution to VPND might fail and provide a constant factor approximation algorithm.

Recently, Hurkens, Keijsper and Stougie [12] considered the problem ($SymT$) in the special case when the given network is a ring. 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. [12] show that this conjecture holds for ring networks. 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 terminals in $S \cup \mathcal{R}'$ 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 [11].

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 $d_1, d_2 \in \mathbb{R}_+$. Then, there exists a (fractional) 2-commodity flow of value d_1, d_2 if and only if the cut condition is satisfied. Moreover, if all edge capacities are integral, then a half-integral flow exists.*

The cut condition requires that $u(\delta(X)) \geq d_1 \chi_1(X) + d_2 \chi_2(X)$ for each $X \subseteq V$. Here $\delta(X)$ denotes the cut induced by X . Moreover, for $i = 1, 2$, we set $\chi_i(X) = 1$ if the cut $\delta(X)$ separates s_i from r_i and $\chi_i(X) = 0$, otherwise.

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. In the case of unit demands, the cut condition requires that, for all $X \subseteq V$, the capacity $u(\delta(X))$ of the cut induced by X must be at least the number of terminal pairs which are separated by the cut. It thus remains to prove that

the cut condition holds for the 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 $X \subseteq V$. If the corresponding cut separates neither $\{s_1, s_2\}$ nor $\{r_1, r_2\}$, nothing needs to be shown. If $\delta(X)$ 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 $\delta(X)$ separates both terminal pairs $\{s_1, s_2\}, \{r_1, r_2\}$, then it either separates $\{s_1, r_1\}$ and $\{s_2, r_2\}$ or it separates $\{s_1, r_2\}$ and $\{s_2, r_1\}$. In both cases it follows that the capacity of the cut is at least 2.

We remark that Corollary 1 is no longer true if we replace “2-commodity flow” by “integral 2-commodity flow”. Even, Itai, and Shamir show that finding an integer 2-commodity flow is NP-hard [7]. 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 [16, Chapter 71].

Before we formulate and prove our new lower bound for VPND, we first state a general technique for deriving such lower bounds.

Lemma 2. *Let $\mathcal{S}_1, \dots, \mathcal{S}_k$ be a partition of \mathcal{S} and let $\mathcal{R}_1, \dots, \mathcal{R}_k$ be a partition of \mathcal{R} . We denote the VPND-instance on graph G with senders \mathcal{S}_i and receivers \mathcal{R}_i by I_i . Then,*

$$\sum_{i=1}^k OPT_i \leq OPT ,$$

where OPT_i is the cost of an optimal solution to instance I_i .

We are now ready to prove the main result of this section.

Theorem 2. *Let $\mathcal{R}' \subseteq \mathcal{R}$ be an arbitrary subset of cardinality $|\mathcal{R}'| = S$ and let M be a matching in the complete graph on $\mathcal{S} \cup \mathcal{R}'$. Then,*

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

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 M . Suppose that M consists of edges

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

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

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

By Lemma 2, the sum of OPT_i for the VPND-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

$$\ell(s_{2i-1}, s_{2i}) + \ell(r_{2i-1}, r_{2i}) \leq OPT_i \text{ for each } 1 \leq i \leq k.$$

For $1 \leq i \leq k$, any solution to instance I_i yields a reservation of capacities that supports 2-commodity flows with unit demands for terminal pairs $\{s_{2i-1}, r_{2i-1}\}$, $\{s_{2i}, r_{2i}\}$ and for $\{s_{2i-1}, r_{2i}\}$, $\{s_{2i}, r_{2i-1}\}$. By Corollary 1, it must also support a 2-commodity flow for terminal pairs $\{s_1, s_2\}$, $\{r_1, r_2\}$. Therefore, the cost of this solution is at least $\ell(s_1, s_2) + \ell(r_1, r_2)$. This concludes the proof.

3 The quality of a simple routing scheme

Consider the following simple VPND algorithm: Select the node with the cheapest shortest path tree to the union of senders and receivers and reserve one unit of capacity along each shortest path. The effect of installing capacities along shortest paths is cumulative. In other words, if k shortest paths share the same edge, the algorithm assigns k units of capacity to that edge. Moreover, the shortest paths can be computed with a consistent tie-breaking rule such that the edges with nonzero capacity form a tree.

This algorithm produces an optimal tree reservation in the symmetric case (*SymT*) [9] and in the balanced case (*BalT*) [13]. In the symmetric case, Gupta et al. [9] show that the tree produced by the algorithm is a 2-approximate solution to the optimum graph reservation. Italiano et al. [13] 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 this algorithm produces a tree-solution whose cost is within a factor of $1 + R/S$ of the optimum graph reservation cost. As a consequence, also (*BalG*) can be approximated within a factor of two.

In [6] the following inequality is shown which follows from Lemma 1 and which we will use in the proof of the theorem below:

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

We are now ready to bound the approximation ratio provided by this simple routing scheme.

Theorem 3. *The above described algorithm achieves a performance ratio $1 + R/S$.*

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 its star satisfies

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

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

Suppose for the remainder of the proof that $R > S$. In the following we denote by \mathcal{M}_S the set of (possibly not perfect) matchings of the senders in G^m and by \mathcal{M}_R the matchings of at most S receivers. Theorem 2 implies the inequality

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

where $\ell(E') := \sum_{uv \in E'} \ell(u, v)$ for any $E' \subseteq E^m$. In consideration of (3), we distinguish two cases.

First case: $\ell(M_S) \leq OPT/2$ 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 of G_S^m can be partitioned into S matchings. On the other hand one has

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

and thus $\sum_{s' \in \mathcal{S}} \sum_{s \in \mathcal{S}} \ell(s', s) \leq S \cdot OPT$. This means that the average weight of a complete star in G_S^m is at most OPT . Let s' be a random sender. Inequality (2) implies that $E[\sum_{r \in \mathcal{R}} \ell(s', r)] \leq (R/S)OPT$. Together with the above discussion this implies for a random sender s'

$$E \left[\sum_{s \in \mathcal{S}} \ell(s', s) + \sum_{r \in \mathcal{R}} \ell(s', r) \right] \leq (1 + R/S)OPT. \quad (5)$$

Second case: $\ell(M_S) > OPT/2$ 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 the cost of any matching \widetilde{M} in G_R^m is bounded by $(R/S)OPT/2$. Since the edges of G_R^m can be partitioned into R matchings, we can then argue in a similar manner as above that $E[\sum_{r \in \mathcal{R}} \ell(r', r)] \leq (R/S)OPT$ for a random receiver r' . Together with (2) this implies

$$E \left[\sum_{s \in \mathcal{S}} \ell(s, r') + \sum_{r \in \mathcal{R}} \ell(r, r') \right] \leq (1 + R/S)OPT. \quad (6)$$

It remains to bound the cost of any matching \widetilde{M} in G_R^m . First assume that S is even. Theorem 2 implies $\ell(M_R) \leq OPT/2$ for each matching $M_R \in \mathcal{M}_R$ of at most S receivers. As a consequence, the average cost of an edge in a matching \widetilde{M} is at most $(OPT/2)/(S/2) = OPT/S$. Since \widetilde{M} has at most $R/2$ edges, we get $\ell(\widetilde{M}) \leq (R/S)OPT/2$ for any matching \widetilde{M} in G_R^m .

It remains to consider the case that S is odd. Then there is a sender s^* which is missed by the maximum cost matching M_S of G_S^m . Theorem 2 yields

$$\ell(M_R) + \ell(s^*, r_1^*) \leq OPT/2 \quad (7)$$

for each matching $M_{\mathcal{R}} \in \mathcal{M}_{\mathcal{R}}$ and $r_1^* \in \mathcal{R}$ which is not matched by $M_{\mathcal{R}}$. Since $R > S$, there exists another receiver r_2^* which is missed by $M_{\mathcal{R}}$. 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

$$\ell(M_{\mathcal{R}}) + 1/2 \ell(r_1^*, r_2^*) \leq OPT/2 \tag{8}$$

for each matching $M_{\mathcal{R}} \in \mathcal{M}_{\mathcal{R}}$ and receivers r_1^*, r_2^* which are missed by $M_{\mathcal{R}}$. This implies that the average weight of an edge of \widetilde{M} is bounded by OPT/S and thus $\ell(\widetilde{M}) \leq (R/S)OPT/2$.

4 A new algorithm for VPND

In Section 3 we described an algorithm which guarantees a good approximation ratio for R close to S . In this section we present a better approximation algorithm for the case that R is sufficiently larger than S . As the algorithm by Gupta et al. [10], this algorithm is based on Steiner-tree computations. However, in contrast to the algorithm in [10], it does not construct a “high bandwidth core”, which is a small Steiner tree with high capacity, which collects and distributes the demands from outside, and routes them along its high capacity paths. Instead, we proceed by constructing Steiner trees for each sender to a previously sampled subset of receivers, and by connecting the other receivers along their shortest paths to the sampled subset.

Algorithm 1.

- (1) Partition \mathcal{R} into S subsets uniformly at random. Among the non-empty subsets in the partition, select one subset \mathcal{R}' uniformly at random.
- (2) For each sender $s \in \mathcal{S}$, compute a ρ -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 a shortest path between each receiver $r \in \mathcal{R}$ and \mathcal{R}' .

It is interesting to note here that 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 [9] that an optimal solution to (*AsymT*) is in general not an optimal solution to VPND. All previous constant factor approximation algorithms for VPND produce tree reservations [10,6].

It remains to specify the path between each sender-receiver pair (s, r) . Assume that the shortest paths are computed in a consistent way. 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 .

Before we proceed with the analysis of Algorithm 1, we state a corollary of Lemma 2. Here, given a subset V' of nodes, we denote the cost of an optimum Steiner tree on terminal set V' by $st(V')$.

Corollary 2 ([10]). Let $\mathcal{R}_1, \dots, \mathcal{R}_s$ 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 . Then,

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

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

Theorem 4 is a straightforward consequence of the following lemmas.

Lemma 3. 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,

$$\begin{aligned} E[st(\{s'\} \cup \mathcal{R}(s'))] &= P(A) E[st(\{s'\} \cup \mathcal{R}(s')) \mid A] \\ &\quad + P(\bar{A}) E[st(\{s'\} \cup \mathcal{R}(s')) \mid \bar{A}]. \end{aligned}$$

Now observe that $P(\bar{A}) = 1 - (1 - 1/S)^R \geq 1 - e^{-R/S}$. Moreover $E[st(\{s'\} \cup \mathcal{R}(s')) \mid A] = E[st(\{s'\})] = 0$. Thus

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

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

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

This implies the following upper bound on the expected cost of capacity which is installed in the second step of the algorithm.

Lemma 4. The expected cost of the capacity installed in the second step of Algorithm 1 is at most

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

The cost of the third step can be bounded as in the following lemma. A proof can be found in the full version of this paper.

Lemma 5. *The expected cost of the capacity installed in the third step of Algorithm 1 is at most*

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

In Section 3 we described a $(1 + R/S)$ approximation algorithm. The factors $1 + R/S$ and $(2 + \rho)/(1 - e^{-R/S})$ are equal for $R/S = 2.78\dots < 2.79$. Note that $1 + R/S$ is increasing in R/S and $(2 + \rho)/(1 - e^{-R/S})$ is decreasing in R/S . It follows that a combination (taking the minimum cost solution) of the simple routing scheme of section 3 and Algorithm 1 has an expected approximation guarantee of 3.79.

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

4.1 Computing the Optimum Steiner Tree

The performance ratios of the simple scheme and Algorithm 1 meet roughly at $R/S = 2.78$. 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 [4]. This suggests the following variant of Algorithm 1.

In the second step of Algorithm 1 compute an *optimal* Steiner tree whenever $|\mathcal{R}'| \leq \log n$, where n is the number of nodes in (the original graph) G . A similar technique was already applied in [6].

Clearly, this modification is a polynomial time algorithm whose expected approximation guarantee is not worse than the one of Algorithm 1. In particular, if $R/S \geq \log \log n$, the approximation achieved is

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

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, by Markov's inequality, at most $(1 + \log \log n)/\log n$. In this unlikely event however, we can estimate the outcome of the combination of the thereby computed solution with the solution computed by the simple routing scheme described in section 3.

The next theorem is proved in a similar way as the main theorem in [6]. The proof can be found in the full version of this paper.

Theorem 6. *The combination (taking the cheaper solution) of the above described modification of Algorithm 1 and the simple routing scheme described in section 3 is an expected 3.55-approximation algorithm for VPND.*

References

1. M. Andrews and L. Zhang. The access network design problem. In *Proceedings of the 39th Annual Symposium on Foundations of Computer Science(FOCS-98)*, pages 40–49, Los Alamitos, CA, Nov.8–11 1998. IEEE Computer Society.
2. B. Awerbuch and Y. Azar. Buy-at-bulk network design. In *38th Annual Symposium on Foundations of Computer Science*, pages 542–547, Miami Beach, Florida, 20–22 Oct. 1997. IEEE.
3. G. Brightwell, G. Oriolo, and F. B. Shepherd. Reserving resilient capacity in a network. *SIAM J. Discrete Math.*, 14(4):524–539, 2001.
4. S. E. Dreyfus and R. A. Wagner. The Steiner problem in graphs. *Networks*, 1:195–207, 1971/72.
5. N. G. Duffield, P. Goyal, A. Greenberg, P. Mishra, K. K. Ramakrishnan, and J. E. van der Merive. 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. ACM Press, 1999.
6. F. Eisenbrand and F. Grandoni. An improved approximation algorithm for virtual private network design. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms, SODA 05*, pages 928–932, 2005.
7. S. Even, A. Itai, and A. Shamir. On the complexity of timetable and multicommodity flow problems. *SIAM Journal on Computing*, 5(4):691–703, 1976.
8. J. A. Fingerhut, S. Suri, and J. S. Turner. Designing least-cost nonblocking broadband networks. *Journal of Algorithms*, 24(2):287–309, 1997.
9. 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*, pages 389–398, 2001.
10. A. Gupta, A. Kumar, and T. Roughgarden. Simpler and better approximation algorithms for network design. In ACM, editor, *Proceedings of the Thirty-Fifth ACM Symposium on Theory of Computing, San Diego, CA, USA, June 9–11, 2003*, pages 365–372, New York, NY, USA, 2003. ACM Press.
11. T. Hu. Multi-commodity network flows. *Operations Research*, 11:344–360, 1963.
12. C. Hurkens, J. Keijsper, and L. Stougie. Virtual private network design: A proof of the tree routing conjecture on ring networks. In *Proceedings of the eleventh Conference on Integer Programming and Combinatorial Optimization, IPCO XI*, 2005. to appear.
13. G. Italiano, S. Leonardi, and G. Oriolo. Design of networks in the hose model. In *Proceedings of ARACNE 2002*, pages pp 65–76, 2002.
14. G. F. Italiano, R. Rastogi, and B. Yener. Restoration algorithms for virtual private networks in the hose model. In *Proceedings of the 21st Annual Joint Conference of the IEEE Computer and Communications Society (INFOCOM-02)*, volume 1 of *Proceedings IEEE INFOCOM 2002*, pages 131–139. IEEE Computer Society, 2002.
15. G. Robins and A. Zelikovsky. Improved steiner tree approximation in graphs. In *Proceedings of the eleventh annual ACM-SIAM Symposium on Discrete Algorithms, SODA'00*, pages 770–779, 2000.
16. A. Schrijver. *Combinatorial optimization. Polyhedra and efficiency (3 volumes)*. Algorithms and Combinatorics 24. Berlin: Springer., 2003.