

A note on hierarchical hubbing for a generalization of the VPN problem

Neil Olver^{a,b}

^a*Department of Econometrics & Operations Research, Vrije Universiteit Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, Netherlands*

^b*Centrum Wiskunde & Informatica (CWI), Science Park 123, 1098 XG Amsterdam, Netherlands*

Abstract

Robust network design refers to a class of optimization problems that occur when designing networks to efficiently handle variable demands. In this context, Fréchet et al. [1] recently explored *hierarchical hubbing*: a routing strategy involving a multiplicity of “hubs” connected to terminals and each other in a treelike fashion. For a natural generalization of the VPN problem, we prove a structural characterization implying that the optimal hierarchical hubbing solution can be found efficiently, and relate this to a “Generalized VPN Conjecture”.

Keywords: Robust network design, VPN Conjecture
2000 MSC: 68M10, 90B18, 90C27, 90C35

1. Introduction

1.1. Robust network design

Robust network design considers the problem of building networks under uncertainty in the pattern of utilization. Introduced by Ben-Ameur and Kerivin [2], the framework encompasses the important case of the “hose model” introduced by Fingerhut [3] and Duffield et al. [4]. It can itself be seen as falling under the broader umbrella of robust optimization [5].

We refer the reader to [6] for a more in-depth treatment; here we will give a brief self-contained exposition of the model. We are given an undirected graph $G = (V, E)$; this should be interpreted as an existing high-capacity network, in which we can reserve capacity. We assume there is an unlimited total capacity on any given link of the network, and that the cost to reserve capacity on any link is a linear function of the capacity required. Let $c : E \rightarrow \mathbb{R}_+$ denote the per-unit cost of capacity on each edge. A set $W \subseteq V$ of *terminals* needs to be adequately connected using the capacity reserved.

A *traffic pattern* (or *demand pattern*) describes the precise pairwise demand requirements at some moment in time. It can be specified by a traffic matrix D , indexed by pairs of terminals; for terminals i, j , the entry D_{ij} represents the bandwidth needed to send data from i to j . In our network, the traffic pattern is not fixed, but varying (and possibly uncertain). To deal with this, the robust network design framework allows for a *set* of traffic patterns to be prescribed. This (it turns out) can always be taken to be a convex set, and so we describe this set, or *demand universe*, as a convex body $\mathcal{U} \subset \mathbb{R}_+^{W \times W}$.

In this paper, we will be concerned only with the case of *symmetric* demands, meaning that demand from i to j is not distinguishable from demand from j to i . In this case, it is convenient to consider \mathcal{U} to be a subset of $\mathbb{R}_+^{\binom{W}{2}}$, where $\binom{W}{2}$ denotes the set of unordered pairs of terminals, so that $D_{ij} = D_{ji}$ refers to the same demand.

The robust network design (RND) problem asks for the cheapest capacity reservation $u : E \rightarrow \mathbb{R}_+$ that can support all traffic patterns in the specified universe \mathcal{U} . To fully specify the problem however, a further

Email address: n.olver@vu.nl (Neil Olver)

aspect must be considered: the *routing scheme*. The coarsest division is into *oblivious* or *dynamic* routing. In dynamic routing, the way in which traffic is routed may vary arbitrarily as a function of the current traffic pattern. This is typically infeasible, and we will be concerned here with the more practical oblivious routing, where the routing used for any given pair of terminals is specified in advance. We will also only consider *single-path routing*. The routing scheme in this case is described by a template $\mathcal{P} = \{P_{ij} : i, j \in W\}$, where P_{ij} is an i - j -path for each $i, j \in W$. (We do not require this path to be simple.) Since we consider symmetric demands, $P_{ij} = P_{ji}$ refers to the same path.

We may summarize the general robust network design problem (with oblivious, single-path routing and symmetric demands) as follows:

RND problem. *Given an undirected graph $G = (V, E)$ with edge costs $c(e)$, a terminal set $W \subseteq V$, and a convex demand universe $\mathcal{U} \subset \mathbb{R}_+^{\binom{W}{2}}$, a solution to the robust network design problem consists of a routing template $\mathcal{P} = \{P_{ij} : i, j \in W\}$, and a capacity allocation $u : E \rightarrow \mathbb{R}_+$, such that \mathcal{U} can be routed according to \mathcal{P} within the capacity u , i.e.,*

$$u(e) \geq \max_{D \in \mathcal{U}} \sum_{\{i,j\} \subseteq W} D_{ij} \ell(P_{ij}, e). \quad (1)$$

Here, $\ell(P, e)$ gives the number of times that edge e occurs on the (possibly non-simple) path P .

The difficulty in this optimization problem lies in choosing the routing template; once this is fixed, the optimal capacity allocation can be determined by solving a convex program described by (1), assuming we have access to at least a separation oracle for \mathcal{U} .

Note that there is always an optimum solution template whose paths P_{ij} are all simple, since any non-simple path can simply be replaced by a simple path within its support. The reason we allow non-simple paths is related to the specific type of routing templates we will be interested in.

The well-studied *symmetric hose model* [3, 4] is parameterized by a vector $b \in \mathbb{R}_+^W$, yielding the universe

$$\mathcal{H}(b) = \left\{ D \in \mathbb{R}_+^{\binom{W}{2}} : \sum_{\{i,j\} \subseteq W} D_{ij} \leq b_i \quad \forall i \in W \right\}.$$

This models the situation where terminals are connected to the network with “hoses” of known, fixed capacity, so that the total demand involving terminal i cannot exceed the capacity b_i of its associated hose link. Any demand pattern that fits through the hoses should be routable in the final network. These hoses may model real links, or chosen based on operational criteria; either way, the hose model gives a simple, useful and concise description of what the network must be able to handle, making it a very popular model in the literature.

A number of variations and generalizations of this model have been considered in the literature [7, 1, 8, 3]. For example, Fréchet et al. [1] consider the “capped” hose model, where in addition to the hose capacities b , point-to-point upper bounds on the demands are also given. In this paper, we will be interested in the following generalization introduced by Olver and Shepherd [8]. Let T^b be an arbitrary capacitated tree, with nonnegative edge capacities b and with leaf set in exact correspondence with the terminal set W . We will call any such capacitated tree a **demand tree**. We will use T^b to define a demand universe in a simple and natural way: let $\mathcal{U}(T^b)$ consist of all demand patterns that can be routed on T^b .

The case where T^b is a star corresponds precisely to the hose model; the capacity of the edge adjacent to terminal i precisely gives the marginal of i . This generalization allows the network operator more precise control over the demand universe, hopefully leading to more efficient solutions. In particular, if the terminals of the network can be logically divided into distinct groups (e.g., different branches of the company), with limited communication between groups, this information can be encoded via $\mathcal{U}(T^b)$.

We call the RND problem with oblivious routing for this class of demand universes the *generalized VPN problem*. It was shown in [8] that the generalized VPN problem is approximable to within a factor of 8.




Figure 1: An example of an embedding of T , with the resulting routing from 1 to 6 indicated.

1.2. Hierarchical hubbing

Fréchette et al. [1] define the following variant of the RND problem. Let T be a tree with leaf set W ; we will call such a tree a **hub tree**. A T -**embedding** is a mapping of the internal nodes of T into the network, and a mapping of each edge e of T to a “cable” that connects the images of the endpoints of e (see Figure 1). More than one node of T can be mapped to the same location, and multiple cables may run over the same edge of the network. More formally, we call a map $\varphi : V(T) \cup E(T) \rightarrow V(G) \cup E(G)$ a T -embedding if:

- (i) $\varphi(v) \in V(G)$ for all $v \in V(T)$,
- (ii) $\varphi(i) = i$ for all $i \in W$, and
- (iii) $\varphi(vw)$ is a simple $\varphi(v)$ - $\varphi(w)$ path in G for each $vw \in E(T)$.

The restriction to simple paths in the above definition is not necessary, but will be notationally convenient; in any case, there is no advantage to using non-simple paths. Such an embedding naturally defines a routing template: for each $\{i, j\} \subseteq W$, take the image of the unique i - j -path in the tree under the mapping (again, see Figure 1), yielding an i - j -path in G . Note that this path need *not* be simple.

Given a hub tree T , along with a T -embedding, a T -**hubbing** solution is defined as follows. A T -embedding φ is chosen; this defines the routing template. Moreover, each cable is given a capacity, determined by the maximum load that can be placed on the cable by a demand in \mathcal{U} . Stated differently, for any given $f \in E(T)$, the cable associated with f is given capacity $b(f) := \max_{D \in \mathcal{U}} \sum_{i \in S, j \notin S} D_{ij}$, where S is the terminal set of either component of $T \setminus f$. The capacity $u(e)$ allocated to an edge $e \in E(G)$ must be at least the sum of the capacities of the cables running over that edge:

$$u(e) \geq \sum_{f \in E(T): e \in \varphi(f)} b(f) \quad \text{for all } e \in E(G).$$

The embedding φ and the valid capacity allocation u together describe the T -hubbing solution. A **hierarchical hubbing** solution is simply a T -hubbing solution, for some choice of a hub tree T . The cost of a hierarchical hubbing solution is defined simply as the cost of its associated capacity allocation.

So we have the following hierarchical hubbing RND problem (again, in the case of symmetric demands).

RND_{HH} problem. *Given an undirected graph $G = (V, E)$ with edge costs $c(e)$, a terminal set $W \subseteq V$, and a convex demand universe $\mathcal{U} \subset \mathbb{R}_+^{\binom{W}{2}}$, the RND_{HH} problem is to find the tree T and an embedding of T that yields the cheapest hierarchical hubbing solution.*

Remark 1. *It would also be natural to instead choose capacities by considering the routing template induced by the hierarchical hubbing, and using (1). This alternative formulation is in general not the same as described above; there may be situations where not all cables on a given edge can be simultaneously saturated by a traffic pattern in \mathcal{U} , leading to a larger capacity requirement with the cable formulation. The formulation that we use in this paper, and which is also used in [1], seems overall easier to deal with (e.g., see Lemma 2 below). If the Generalized VPN Conjecture discussed in Section 3 is true, it follows immediately that for the universe $\mathcal{U}(T^b)$, both formulations have a common optimal solution.*

It is easy to confirm that any solution to the RND_{HH} problem is a solution to the RND problem, but not vice versa. So in general the optimal solution to RND_{HH} can be more expensive than the optimal RND solution; in fact, Fréchette et al. [1] demonstrate that the gap can be $\Omega(\log |V|)$, for some choices of the universe.

Fréchette et al. [1] are motivated to consider hierarchical hubbing for a few reasons. In *hub routing*, all traffic is routed via a single hub node; this has the advantage that routing decisions are localized at the hub. In order to address some practical shortfalls of hub routing, Shepherd and Winzer [9] ask for a “multihub” extension of this. Fréchette et al. argue that hierarchical hubbing provides a natural such extension (note that it is clearly a generalization; hub routing corresponds to taking the hub tree to be a star). They also show that it provides an effective heuristic for finding good solutions to the capped hose model mentioned earlier, which is APX-hard in the single-path oblivious routing model. They observe that for the capped hose model (and for other universes as well), hierarchical hubbing can yield much cheaper solutions than using a single hub. This is in contrast to the vanilla hose model, where a hub routing provides an optimal oblivious routing solution (see Section 3).

Finding the optimal hierarchical hubbing solution for a given universe is in general APX-hard, as observed in [10] (cf. [11] for tree routings); this can be seen by choosing the universe so that the resulting RND_{HH} problem is precisely the Steiner tree problem. It is thus natural to ask for which universes the problem is polynomially solvable. One case where the problem was previously known to be polynomial was for the hose model [12]; this will be discussed more in Section 3.

The difficulty of the RND_{HH} problem really comes from determining the optimal hub tree to use; once this has been decided, the optimal solution can be found easily. This is an immediate consequence of a dynamic programming algorithm of Olver and Shepherd [8].

Lemma 2. *For any universe $\mathcal{U} \subseteq \mathbb{R}_+^{\binom{W}{2}}$ over which we can optimize linear functions in polynomial time, and any hub tree T , the optimal T -hubbing for \mathcal{U} can be found in polynomial time.*

Proof. Observe that we can compute exactly the capacity $b(f)$ of any cable $f \in E(T)$, by solving the appropriate maximization problem. So this problem is just asking for the minimum cost embedding of a given capacitated tree. Such an algorithm was given in [8]. (Note that in any optimal solution, the path $\varphi(vw)$ used to route some $vw \in E(T)$ can always be taken to be a shortest path between $\varphi(v)$ and $\varphi(w)$. The algorithm uses dynamic programming to find the optimal placement of the hubs.) \square

Remark 3. *Any tree routing—meaning a routing template $\mathcal{P} = \{P_{ij} : i, j \in W\}$ such that $\bigcup_{i,j \in W} P_{ij}$ is a tree—can be described as a hierarchical hubbing. The hub tree is obtained from the support, adding additional dummy edges as needed so that all terminals are leaves in the hub tree.*

2. Hierarchical hubbing for the generalized VPN problem

Hierarchical hubbing was in fact introduced earlier in [8], but only for a specific algorithmic purpose: the factor 8 approximation algorithm to the generalized VPN problem that they propose produces a hierarchical hubbing solution. This of course also implies a factor 8 approximation to RND_{HH} with universe $\mathcal{U}(T^b)$.

The main contribution of this note is the following structural characterization.

Theorem 4. *For a given demand tree T^b , the optimal hierarchical hubbing for $\mathcal{U}(T^b)$ is a T -hubbing.*

The following is then immediate by Lemma 2.

Theorem 5. *For any demand tree T^b , the optimal hierarchical hubbing solution with universe $\mathcal{U}(T^b)$ can be found in polynomial time.*

Thus RND_{HH} is polynomially solvable for a large and interesting class of demand universes.

Theorem 4 is certainly very natural, and it might seem even trivial at first glance. But consider, for example, the case where T is simply a star, and hence represents a hose model universe. Then recalling Remark 3—any tree routing is a hierarchical hubbing—this theorem includes the fact that the optimal tree

routing under the hose model is a hub routing, a result of Gupta et al. [12]. Theorem 4 may be seen as providing some evidence for a “Generalized VPN Conjecture” which would imply that the RND problem for $\mathcal{U}(T^b)$ is polynomially solvable. This is discussed in detail in Section 3.

2.1. Preliminaries

Given two disjoint subsets A, B of vertices in some graph $G = (V, E)$, an (A, B) -cut is any set $S \subseteq V$ with $A \subseteq S$ and $S \cap B = \emptyset$. We denote by $\delta(S)$ the set of edges crossing a cut S , and given capacities $u : E \rightarrow \mathbb{R}_+$, the capacity of S is simply $u(\delta(S)) := \sum_{e \in \delta(S)} u(e)$.

Given a tree T and edge capacities $b : E(T) \rightarrow \mathbb{R}_+$, T^b will denote the resulting edge-capacitated tree. It is possible that $\mathcal{U}(T^b) = \mathcal{U}(T^{b'})$ for distinct capacities b, b' . The following definition is convenient:

Definition 6. *Edge capacities b for a tree T are called **defining** if for every $f \in E(T)$, there exists a $D \in \mathcal{U}(T^b)$ that saturates edge f in T^b .*

Given $T^{b'}$, defining capacities $b \leq b'$ such that $\mathcal{U}(T^b) = \mathcal{U}(T^{b'})$ can easily be found, solving one maximum flow problem per edge of T . So we will assume for the remainder of the article that the capacities b are defining.

2.2. Proof of the main theorem

We will prove something stronger than Theorem 4: that given any hierarchical hubbing for $\mathcal{U}(T^b)$ with capacity allocation u , there is a T -hubbing for $\mathcal{U}(T^b)$ which requires no more than capacity $u(e)$ on each edge e . We begin by proving this result for the case where the network is itself a tree, with terminals forming the leaf set of the tree. We will then observe that the result for an arbitrary network follows easily.

Tree networks. In this section, we will denote the network, being a tree, by F ; the leaves of F form the terminal set W . Let $\text{fund}_F(e)$ denote the set of all pairs of leaves that are separated by edge $e \in E(F)$. Notice that there is an obviously optimal oblivious routing solution for the RND problem with universe $\mathcal{U}(T^b)$ in F : use the template $\mathcal{P} = \{P_{ij} : i, j \in W\}$ where P_{ij} is the unique simple path from i to j in F . The minimal required capacity $q^*(e)$ for any edge $e \in E(F)$ is then, from (1),

$$q^*(e) = \max_{D \in \mathcal{U}(T^b)} \sum_{\{i, j\} \in \text{fund}_F(e)} D_{ij}. \quad (2)$$

Note that *any* oblivious routing template will need capacity $q^*(e)$ at least on each $e \in E(F)$, since any i - j path in F contains P_{ij} .

Observe also that this oblivious routing template is an F -hubbing template, induced by the trivial F -embedding that maps F to itself. Since each cable in this embedding is a single edge, combining this F -embedding with the capacity allocation q^* yields an F -hubbing for $\mathcal{U}(T^b)$.

The next theorem, which is the key technical theorem of this paper, shows that there is a T -hubbing for $\mathcal{U}(T^b)$ with capacity allocation q^* . In general, the T -embedding that certifies this will induce *non-simple* paths; despite this, extra capacity will not be needed.

Theorem 7. *Let F be any tree with leaf set W , and let $q^* : E(F) \rightarrow \mathbb{R}_+$ be as in (2). Then there is a T -hubbing for $\mathcal{U}(T^b)$ into the network F with capacity allocation q^* .*

Proof. We will need the following lemma, which follows easily from standard uncrossing techniques.

Lemma 8. *Let $W_1 \subseteq W_2 \subseteq W$, and for $i \in \{1, 2\}$, let S_i be a minimum capacity $(W_i, W \setminus W_i)$ -cut in T^b (with $W_i \subseteq S_i$). Then $S_1 \cap S_2$ is a minimum capacity $(W_1, W \setminus W_1)$ -cut, and $S_1 \cup S_2$ is a minimum capacity $(W_2, W \setminus W_2)$ -cut.*

Proof. By submodularity of the cut function,

$$b(\delta(S_1 \cap S_2)) + b(\delta(S_1 \cup S_2)) \leq b(\delta(S_1)) + b(\delta(S_2)). \quad (3)$$

But $S_1 \cap S_2$ is a $(W_1, W \setminus W_1)$ -cut, and so by the definition of S_1 , $b(\delta(S_1 \cap S_2)) \geq b(\delta(S_1))$. Moreover, $S_1 \cup S_2$ is a $(W_2, W \setminus W_2)$ -cut, and hence $b(\delta(S_1 \cup S_2)) \geq b(\delta(S_2))$. We deduce that (3) holds with equality, and hence that $S_1 \cap S_2$ is a minimum $(W_1, W \setminus W_1)$ -cut, and $S_1 \cup S_2$ a minimum $(W_2, W \setminus W_2)$ -cut. \square

Pick an arbitrary leaf $r \in W$, and call it the root. An edge e of F divides the terminal set into two, W_e and $W \setminus W_e$, where we choose W_e to *not* contain the root. Let S_e be a minimum $(W_e, W \setminus W_e)$ cut containing W_e in the tree T^b , breaking ties by choosing S_e to have minimum cardinality.

We now describe the T -embedding φ . We will define, for each internal node $w \in V(T) \setminus W$, an orientation \vec{F}_w of F . For each edge $e \in E(F)$, orient e away from the root if $w \in S_e$, and otherwise orient e towards the root. We then define $\varphi(w)$ to be the unique sink node of \vec{F}_w , whose existence we guarantee as follows.

Lemma 9. *There is a unique sink node in \vec{F}_w .*

Proof. We begin by showing that every node has outdegree at most 1 in \vec{F}_w . Suppose for a contradiction that some node $u \in V(F)$ has two outgoing arcs e, e' in \vec{F}_w . There are two cases to consider:

- (i) Both arcs e and e' are oriented away from r . Then $w \in S_e$ and $w \in S_{e'}$. Now $W_e \subseteq (W \setminus W_{e'})$, and hence by Lemma 8, $\hat{S}_e := S_e \cap (V(T) \setminus S_{e'})$ is a minimum $(W_e, W \setminus W_{e'})$ -cut in T^b . But since $w \notin \hat{S}_e$, $\hat{S}_e \subsetneq S_e$, contradicting the size minimality of S_e .
- (ii) One of e, e' is oriented towards r , say e' . Then $w \in S_e$ and $w \notin S_{e'}$. Then simply note that $W_e \subseteq W_{e'}$, and hence (by Lemma 8) $S_e \cap S_{e'}$ is a minimum $(W_e, W \setminus W_{e'})$ -cut in T^b . Again since $w \notin S_e \cap S_{e'}$, this contradicts the size minimality of S_e .

By starting at an arbitrary node, and following the unique outgoing arc until we reach a node with no outgoing arcs (this may be a leaf, or not), the existence of some sink v follows. To see that this sink is unique, observe that on any path in \vec{F}_w terminating at v , all arcs must be oriented towards v by the condition on the outdegree, and hence none of the nodes on the path aside from v can possibly be a sink. \square

We complete the definition of the T -embedding φ in the obvious way, by taking $\varphi(vw)$ to be the unique simple path between $\varphi(v)$ and $\varphi(w)$ in F , for each $vw \in E(T)$.

Now let us consider $q^*(e)$ for some edge $e \in E(F)$. Rewriting (2), we have

$$q^*(e) = \max_{D \in \mathcal{U}(T^b)} \sum_{i \in W_e, j \notin W_e} D_{ij}. \quad (4)$$

The right hand side of (4) can be seen as a maximum flow problem; send as much flow as possible in T^b from W_e to $W \setminus W_e$. Invoking the max-flow min-cut theorem, we obtain

$$q^*(e) = b(\delta(S_e)). \quad (5)$$

We show now that the T -embedding φ along with the capacity allocation q^* defines a valid hierarchical hubbing solution. We will use the Iverson bracket $[A]$ to denote the indicator function of a predicate A . Consider any edge e in F . The capacity required on e by the hierarchical hubbing solution induced by φ is

$$\begin{aligned} & \sum_{f=uv \in E(T)} [\varphi(u) \text{ and } \varphi(v) \text{ are separated by } e] b(f) \\ &= \sum_{f=uv \in E(T)} [\vec{F}_u \text{ and } \vec{F}_v \text{ orient } e \text{ in opposite directions}] b(f) \\ &= \sum_{f=uv \in E(T)} [\text{exactly one of } u \text{ and } v \text{ is in } S_e] b(f) \\ &= b(\delta(S_e)). \end{aligned}$$

Combined with (5), this completes the proof. \square

General networks. We now show how this result for tree networks can be leveraged to demonstrate the general case.

Since multiple graphs will play the role of the network in the following arguments, we will take care here to be explicit regarding what network any given embedding is mapping into. Suppose now that F is a hub tree, and let (φ, u) be any F -hubbing for $\mathcal{U}(T^b)$ into G . By Theorem 7, there exists a T -hubbing (η, q^*) for $\mathcal{U}(T^b)$ into F , where q^* is the minimal possible capacity allocation discussed earlier. We will essentially compose this hierarchical hubbing with φ to obtain a T -hubbing (ρ, u) into G .

We define the T -embedding ρ as follows. Let $\rho(v) = \varphi(\eta(v))$ for all $v \in V(T)$. For any edge $vw \in E(T)$, consider the path $\eta(vw)$ in F , and write it in terms of its edges: $\eta(vw) = e_1 e_2 \dots e_t$. Take $\rho(vw)$ to be any simple $\rho(v)$ - $\rho(w)$ path in G contained in the concatenation of the paths $\varphi(e_1), \varphi(e_2), \dots, \varphi(e_t)$. Clearly ρ does define a T -embedding into G .

We have that for any $e \in E(G)$,

$$u(e) \geq \sum_{f \in E(F)} [e \in \varphi(f)] q^*(f).$$

This follows from the definition of an F -hubbing for $\mathcal{U}(T^b)$, combined with (2). Now since (η, q^*) is a T -hubbing for $\mathcal{U}(T^b)$ into F , we have that

$$q^*(f) \geq \sum_{f' \in E(T)} [f \in \eta(f')] b(f').$$

Hence

$$u(e) \geq \sum_{f' \in E(T)} \sum_{f \in E(F)} [e \in \varphi(f)] \cdot [f \in \eta(f')] b(f') \geq \sum_{f' \in E(T)} [e \in \rho(f')] b(f').$$

So (ρ, u) is indeed a T -hubbing for $\mathcal{U}(T^b)$ into G . The proof of Theorem 4, and hence Theorem 5, is complete.

3. Connection to the VPN Conjecture

A well-known conjecture in the area was the *VPN Conjecture*. This was resolved by Goyal et al. [13], who proved

Theorem 10 ([13]). *There is an optimal solution to the RND problem with oblivious routing for the hose model whose support is a tree.*

This result leads to a polynomial time algorithm to solve the RND problem for the hose model exactly: indeed, Gupta et al. [12] had previously provided a polynomial time algorithm that computes the optimal tree solution. The algorithm simply finds a hub routing solution. In other words, for each node v , compute the sum (weighted by the hose capacities b_i) of the lengths of the shortest paths from each terminal to v ; then choosing v^* that minimizes this quantity, allocate b_i units of capacity along the i - v^* shortest path for each $i \in W$, additively. (A technicality: the marginals should be “defining” in the sense of Definition 6, meaning that $b_i \leq \frac{1}{2} \sum_j b_j$ for all $i \in W$.) The routing is simply a hub routing centered at v^* : for each pair i, j of terminals, the path from i to j is obtained by appending the shortest path from i to v^* to that from v^* to j . This is of course an $R^{b'}$ -hubbing solution, where R is a star with leaf set W , and the capacity of edge ir (with r being the internal node of the star) is $b'(ir) = b_i$.

The result of Gupta et al. [12] can be seen as a precursor to, and evidence for, the VPN Conjecture. Similarly, the result of this paper, which is a generalization of the result of Gupta et al., can be seen as evidence for the following conjecture proposed in [8].

Generalized VPN Conjecture. *There is an optimal solution to the RND problem with universe $\mathcal{U}(T^b)$ that is a T -hubbing solution.*

Theorem 4 shows that the following version of the conjecture, while seemingly weaker, is equivalent:

Generalized VPN Conjecture (II). *There is an optimal solution to the RND problem with universe $\mathcal{U}(T^b)$ that is a hierarchical hubbing solution.*

It is known that the optimal T -hubbing solution for $\mathcal{U}(T^b)$ is always within a factor 8 of the optimal oblivious routing solution [8]. This in particular implies that the RND problem for the generalized hose model is constant approximable. A positive resolution to the Generalized VPN Conjecture would imply that the RND problem for $\mathcal{U}(T^b)$ can be solved optimally in polynomial time.

Acknowledgements. The author is very grateful to Alexandre Fréchet and Bruce Shepherd, both for asking the question that led to this note, and for their help in the exposition of the main proof. The author also thanks an anonymous referee for a careful reading. This research was supported by a Veni grant of the Netherlands Organisation for Scientific Research (NWO).

References

- [1] A. Fréchet, F. B. Shepherd, M. K. Thottan, P. J. Winzer, Shortest path versus multi-hub routing in networks with uncertain demand, in: Proc. INFOCOM, 2013, pp. 710–718. doi:10.1109/INFOCOM.2013.6566857.
- [2] W. Ben-Ameur, H. Kerivin, New economical virtual private networks, Commun. ACM 46 (6) (2003) 69–73. doi:10.1145/777313.777314.
- [3] J. A. Fingerhut, S. Suri, J. S. Turner, Designing least-cost nonblocking broadband networks, J. Algorithms 24 (2) (1997) 287–309. doi:10.1006/jagm.1997.0866.
- [4] N. G. Duffield, P. Goyal, A. Greenberg, P. Mishra, K. K. Ramakrishnan, J. E. van der Merwe, A flexible model for resource management in virtual private networks, in: Proc. SIGCOMM, 1999, pp. 95–108. doi:10.1145/316188.316209.
- [5] A. Ben-Tal, L. El Ghaoui, A. Nemirovski, Robust Optimization, Princeton Series in Applied Mathematics, Princeton University Press, 2009.
- [6] N. Olver, Robust network design, Ph.D. thesis, McGill University (2010).
- [7] F. Eisenbrand, E. Happ, Provisioning a virtual private network under the presence of non-communicating groups, in: Proc. CIAC, Vol. 3998 of Lect. Notes Comput. Sci., Springer, 2006, pp. 105–114. doi:10.1007/11758471_13.
- [8] N. Olver, F. B. Shepherd, Approximability of robust network design, in: Proc. SODA, 2010, pp. 1097–1105. doi:10.1137/1.9781611973075.89.
- [9] F. B. Shepherd, P. J. Winzer, Selective randomized load balancing and mesh networks with changing demands, J. Opt. Netw. 5 (5) (2006) 320–339. doi:10.1364/JON.5.000320.
- [10] A. Fréchet, Hub routing for the robust network design problem, Master’s thesis, McGill University (2012).
- [11] A. Kumar, R. Rastogi, A. Silberschatz, B. Yener, Algorithms for provisioning virtual private networks in the hose model, IEEE ACM T. Network. 10 (4) (2002) 565–578. doi:10.1109/TNET.2002.802141.
- [12] A. Gupta, J. Kleinberg, A. Kumar, R. Rastogi, B. Yener, Provisioning a virtual private network: a network design problem for multicommodity flow, in: Proc. STOC, 2001, pp. 389–398. doi:10.1145/380752.380830.
- [13] N. Goyal, N. Olver, F. B. Shepherd, The VPN Conjecture is true, J. ACM 60 (3) (2013) 17:1–17:17. doi:10.1145/2487241.2487243.