

The VPN Conjecture is True

N. Goyal * N. Olver † F. B. Shepherd ‡

Abstract

We consider the following network design problem. We are given an undirected graph $G = (V, E)$ with edges costs $c(e)$ and a set of terminal nodes W . A *hose* matrix for W is any symmetric matrix (D_{ij}) such that for each i , $\sum_{j \neq i} D_{ij} \leq 1$. We must compute the minimum cost edge capacities that are able to support the oblivious routing of every hose matrix in the network. An *oblivious routing* template, in this context, is a simple path P_{ij} for each pair $i, j \in W$. Given such a template, if we are to route a demand matrix D , then for each i, j we send D_{ij} units of flow along each P_{ij} . Fingerhut et al. [13] and Gupta et al. [16] obtained a 2-approximation for this problem, using a routing template in the form of a tree. It has been widely asked (cf. [19]) if this solution actually results in the optimal capacity for the single path VPN design problem; this has become known as the *VPN conjecture*. The conjecture has previously been proven for some restricted classes of graphs [18, 15, 14]. Our main theorem establishes that this conjecture is true in general graphs. This also settles the complexity of the single path VPN problem. We also show that the version of the conjecture where fractional routing templates are allowed is false.

*Georgia Tech. Email: navin.goyal@gatech.edu

†McGill University. Email: olver@math.mcgill.ca

‡McGill University. Email: bruce.shepherd@mcgill.ca

1 Introduction

Robustness has become one of several paradigms used to address data uncertainty in optimization. While the input data may not be known perfectly, it assumes that each input comes from some known bounded region or *universe* of legal inputs. Robust optimization then seeks an optimal solution that is valid for all possible inputs in the universe; robustness is thus a worst-case approach for coping with uncertainty. The concept was initially developed in continuous optimization where a large class of robust linear programs were shown to be solvable via conic programming [5, 6].

Robust optimization has become increasingly popular in the discrete optimization community. One of the most widely studied examples arises in the context of network design. This has been motivated by the fact that traffic demands in IP networks are often hard to determine and/or are rapidly changing. In the most general model (cf [4]), the input consists of a graph (network topology) where each edge may come with a cost and possibly lower and upper bounds on its capacity. In addition, a *universe* of possible demand matrices is specified in polyhedral form. The problem is to find a minimum cost assignment of capacities to the edges such that each demand matrix in the polytope can be routed in the resulting capacitated network. Several versions result depending on the specifics of the particular application. For instance, routings may be fractional, integral, unsplittable, confluent and so on. While we only consider traffic routings that are computed offline, there is still an issue of whether and how rapidly one may “operationally” update a routing in response to a change in the traffic matrix. *Dynamic* (or adaptive) routing allows total freedom, and so the routing may be completely updated whenever incoming traffic demands change. This version, which is known to be NP-complete [10] but $O(\log n)$ approximable, has also been studied empirically [24, 1]. On the other extreme, *static* routing asks for an explicit reservation of capacity in the network for each demand pair ij . This potentially wastes a huge amount of network bandwidth.

A middle way, and the focus of this paper, is based on *oblivious routing* of the traffic matrices. This approach (inspired by routing in packet networks) asks for a *template* that defines ahead of time how any future demands will be routed. For each node pair i, j , the template f specifies a unit flow f_{ij} between i and j . The interpretation being that if there is a future demand of D_{ij} between nodes i, j , then along each ij path P , we should route $D_{ij}f_{ij}(P)$ flow on this path. Throughout this paper our focus is on undirected demands and so D_{ij}, D_{ji} represent the same demand.

There have been two main streams of recent work on robust optimization via oblivious routing, each based on a distinct universe of demand matrices. The first line of work asks for an oblivious fractional routing f in a graph G with the property that any *routable* demand matrix $[D_{ij}]$ has *low congestion* when routed according to the template f . A demand matrix is called routable if its corresponding multicommodity flow is feasible in G . In this context, the only network design element in the problem is to choose the routing so as to minimize the maximum capacity, i.e., the *congestion measure*. The seminal work of Leighton and Rao [23] forms the basis for work in this area; they showed that in graphs with good

expansion (relative to a set of terminal nodes W), one may fractionally route the uniform multicommodity flow for W with congestion $O(\log n)$. Hence, this multicommodity flow behaves as a crossbar for W , that is, one may fractionally route any matching on W with the same congestion. A more recent breakthrough was obtained by Räcke [27] who showed that in general graphs there exists such oblivious routings with polylogarithmic congestion. Note that for these problems, the target universe of demand matrices is $\mathcal{R} = \{[D_{ij}] : D \text{ is routable}\}$. It is worth noting that Räcke’s results actually hold for the larger polytope of demands that satisfy the cut condition: $D(\delta(S)) \leq |\delta(S)|$ for each subset of nodes S . Here $\delta(S)$ denotes the set of edges with exactly one endpoint in S (all our graphs will be undirected). Interestingly, the optimal (w.r.t. the congestion measure) oblivious routing can be computed in polytime either via the Ellipsoid Algorithm [3], or, as shown later, by a compact LP formulation [2].

A second area of active study is the work based on the class \mathcal{H} of *Hose demand matrices*, introduced in [13, 11]. These are determined by giving a *marginal* $b(v)$ to each node v in the graph $G = (V, E)$: $\mathcal{H} = \{D_{ij} = D_{ji} : \sum_{j \in V} D_{ij} \leq b(i) \forall i \in V\}$. The network design problem asks for the minimum cost (fractional) network capacity for which there is an oblivious routing supporting every hose matrix.

Work on the hose model [13, 11, 16, 12] is inspired by the need to design networks where only the capacity of nodes is known ahead of time, and not the precise pattern of how nodes will share traffic. This is the case for instance in switch design (which is one possible application of these techniques) where any possible matching of inputs to outputs must be supported. It is also the case for so-called *virtual private network (VPN)* design where a larger network operator carves out some of its capacity for sale to a smaller enterprise. Typically a request comes as a list of node locations and required capacity at each location. The VPN should be able to support whatever traffic the enterprise throws at it. We use the terms *multipath* (MPR) or *single path* (SPR) to refer to whether the flow template is allowed to be fractional or not. Motivated by the application to VPNs, prior to [3, 2, 12], work on the VPN problem focused primarily on the single path version. Note that even in the SPR model, the optimal VPN may still have fractional edge capacities.

Many of the foundations on the VPN design problem appear in Gupta et al. [16] and Fingerhut et al. [13]. In these papers, one restricts to single-path oblivious routings. In the undirected setting, both papers give a 2-approximation (for all versions). Their solution results in a VPN with integral capacities and has a particularly simple structure. Namely, they find a node t (not necessarily a terminal) such that the oblivious routing is defined by a shortest path tree T rooted at t . The optimal capacity is obtained as follows. For each node v , route $b(v)$ units from v on its (shortest) path in T to t . The optimal network design is just the aggregate of these capacities for each v . We call the resulting capacitated tree, the *VPN tree* for t . It had been discussed [17] and subsequently conjectured [19] that this solution actually results in the optimal capacity for the single path VPN design problem; this has become known as the *VPN conjecture*. The conjecture has previously been proven for the case of ring networks [18, 15] and just prior to submission we have learned of independent work establishing this for outerplanar graphs [14]. Our main theorem establishes that this

conjecture is true in general graphs.

Theorem 1.1. *For nonnegative edge costs, if there is a fractional VPN of cost $\beta \geq 0$ resulting from a single path oblivious routing, then there is some node t whose VPN tree costs β .*

This result also settles the complexity of the single path VPN problem.

In a tree routing template, every pair of nodes i, j routes on the unique simple path in the tree between them. In an optimal VPN-tree, there is actually enough capacity for nodes to route all of their traffic via the root node t . This is an oblivious routing template called *hub routing*. Hub routing has been proposed in the design of IP networks [30]. Instead of having to perform expensive IP switching at each node, cost-effective circuits (optical pipes) can be laid down from each node to the hub and so IP switching is only done at the node t . A downside is that the hub node becomes a critical point of failure. It is thus proposed to split traffic across a few hubs. Note that splitting across all possible hubs, one gets a well-known oblivious routing template: Valiant’s Randomized Load Balancing (RLB). This may have a practical impact, since it was empirically shown [30, 32], using realistic equipment costs, that by splitting across a small number of hubs (selective RLB), the network cost is lower than the standard practice of using IP switching at every node (full RLB was actually more expensive however).

Hub routing also plays an important role in the present paper, which builds on the recent elegant work of Grandoni et al. [15]. They introduce the notion of *Pyramidal routing* which corresponds to hub routing to a terminal t but where the edge costs are concave: the cost of a flow l on an edge is $\min\{l, k - l\}$, where k is the number of terminals. They show that the VPN conjecture is implied by what they call the *Pyramidal Routing Conjecture*: that a minimum cost pyramidal routing can be achieved using a tree. We present this work in detail in Section 2.2; we derive the VPN conjecture precisely as they prescribe, by establishing their conjecture on pyramidal routings.

A stronger version of the VPN conjecture asks whether the VPN tree is also optimal for the versions where we allow multipath routing templates [18]. In the final section, we give a family of counterexamples to this conjecture. Of course, the results of [13, 16] imply the cost difference is at most a factor 2. It is worth noting how this contrasts with classical results arising in the study of parallel computation. In the latter, the objective was to find oblivious routings that minimized congestion, and it was shown that fractional (multipath) routings are needed (i.e., randomization is required) in order to obtain near-optimal (polylogarithmic) congestion bounds [7, 20].

We do not have time to discuss many of the results on robustness and obliviousness. These include simplifications and numerous extensions of Räcke’s results (e.g. to directed graphs and node-based problems) as well as versions of the VPN problem using asymmetric demands. The reader is referred to [9] for a recent, more comprehensive overview. There has also been considerable work in the networking community on the application of RLB to switch design and networks with varying demand, especially in terms of throughput measures [8, 21, 22, 25, 26, 28, 31, 29].

2 Proof of the VPN Conjecture

2.1 Definitions and Notation

We consider undirected graphs $G = (V, E)$ with *node* set V and *edge* set E . We may abuse notation and simply refer to an edge with endpoints u, v as uv . Each edge e has an associated nonnegative cost $c(e)$, representing the cost per unit of bandwidth required for that edge. For any subset S of nodes, we denote by $\delta_G(S)$ the set of edges with exactly one endpoint in S (if the context is clear we simply write $\delta(S)$). For a pair of nodes s, t an $s-t$ flow f (for our purposes) will be defined in the path formulation: for each simple path P with endpoints s, t it assigns a value $f(P)$. The *size* of the flow is $\sum_P f(P)$; if the value is one, we may refer to this as a *unit* flow.

A VPN problem instance consists of such an undirected graph G , as well as the marginals $b(v)$ for each node. As was noted in [18], it is sufficient to consider marginals from the set $\{0, 1\}$. Thus we define $W = \{v : b(v) = 1\}$, the set of *terminals* in our network. We use k to denote the number of terminals.

We will actually consider two distinct variants of the VPN problem:

- (MPR) In the *multipath* routing problem, the solution is specified by a unit $i-j$ flow f_{ij} , for each pair $i, j \in W$. The routing template \mathcal{P} is defined by $\mathcal{P} := \{f_{ij} : i, j \in W\}$. Given a hose matrix D satisfying the marginals, the flow from i to j will be routed proportionally according to this template, i.e. any $i-j$ -path P , carries a flow of $D_{ij}f_{ij}(P)$.
- (SPR) In the *single path* routing problem, the flows in the routing template are restricted to be unsplitable, i.e. each pair i, j routes along a single path P_{ij} .

We use \mathcal{R}_{ij} to denote the set of all simple $i-j$ -paths. A routing template induces a VPN solution as follows. A bandwidth requirement $u(e)$ on each edge is chosen as small as possible but allowing all hose demands to be met; thus

$$u(e) = \max_{D \in \mathcal{H}} \sum_{i, j \in W} f_{ij}(e) D_{ij},$$

where $f_{ij}(e) := \sum_{P \in \mathcal{R}_{ij}} f_{ij}(P)$.

The cost $C_{\text{vpn}}(\mathcal{P})$ of the solution is then given by $\sum_{e \in E} c(e)u(e)$.

A Pyramidal Routing (PR) problem instance, as introduced in [15], is also defined by an undirected graph G with costs, and a set of terminals W ; in addition, one node $t \in W$ is specified as the root. A routing template consists of a set \mathcal{P}_t of simple $i-t$ paths P_{it} , one for each $i \in W \setminus \{t\}$. Define $l(e, \mathcal{P}_t)$ to be the total flow through edge e , i.e. $l(e, \mathcal{P}_t) := |\{i \in W : e \in P_{it}\}|$. The bandwidth requirement $y(e, \mathcal{P}_t)$ is instead given by the function $y(e, \mathcal{P}_t) := \min\{l(e, \mathcal{P}_t), k - l(e, \mathcal{P}_t)\}$. Note that the function on the right hand side is concave in $l(e)$. The total cost is then $C_{\text{pyr}}(\mathcal{P}_t) := \sum_{e \in E} c(e)y(e, \mathcal{P}_t)$.

We can also define an analogous fractional version of the pyramidal problem, where instead of paths P_{it} , the routing template consists of a set of unit $i-t$ flows f_{it} .

2.2 Pyramidal Routing

Grandoni et al. [15] show that the VPN conjecture is implied by the following conjecture:

Conjecture 1 (The Pyramidal Routing Conjecture). *For every integral Pyramidal Routing instance, there exists a minimal cost solution that is a tree.*

A crucial part of their argument is the following lemma, the proof of which we reproduce below. Their proof is a slight extension of one given by Gupta et al cf. Theorem 3.1 in [16].

Lemma 2.1 ([15]). *Given an SPR instance, and a routing template $\mathcal{P} = \{P_{ij} : i \neq j \in W\}$, there exists a terminal $t \in W$ so that $C_{\text{vpn}}(\mathcal{P}) \geq C_{\text{pyr}}(\mathcal{P}_t)$, where $\mathcal{P}_t = \{P_{it} : i \in W \setminus \{t\}\}$.*

Proof. The strategy of the proof is to derive a lower bound for $u(e)$ for each e in the instance by judiciously selecting a demand matrix, which will then give us the desired lower bound on $C_{\text{vpn}}(\mathcal{P})$. Fix an edge e . The choice D^e of demand matrix for e is given by

$$d_{ij}^e := \begin{cases} \frac{1}{k} \left(\frac{y(e, \mathcal{P}_i)}{l(e, \mathcal{P}_i)} + \frac{y(e, \mathcal{P}_j)}{l(e, \mathcal{P}_j)} \right) & \text{if } e \in P_{ij}, \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

Claim 1. D^e is a valid hose demand matrix for all edges.

Proof. We need to show that $\sum_{j \in W} d_{ij}^e \leq 1$ for all $i \in W$. We have

$$\begin{aligned} \sum_{j \in W} d_{ij}^e &= \sum_{j \in W: e \in P_{ij}} \frac{1}{k} \left(\frac{y(e, \mathcal{P}_i)}{l(e, \mathcal{P}_i)} + \frac{y(e, \mathcal{P}_j)}{l(e, \mathcal{P}_j)} \right) \\ &\leq \sum_{j \in W: e \in P_{ij}} \frac{1}{k} \left(\frac{k - l(e, \mathcal{P}_i)}{l(e, \mathcal{P}_i)} + \frac{l(e, \mathcal{P}_j)}{l(e, \mathcal{P}_j)} \right) \\ &= \sum_{j \in W: e \in P_{ij}} \frac{1}{k} \cdot \frac{k}{l(e, \mathcal{P}_i)} \\ &= 1. \end{aligned}$$

□

Claim 2. For every edge e we have

$$u(e) \geq \frac{1}{k} \sum_{i \in W} y(e, \mathcal{P}_i). \quad (2)$$

Proof. The claim follows from the definitions of d_{ij}^e and $l(e, \mathcal{P}_i) = |\{j \in W : e \in P_{ij}\}|$.

$$\begin{aligned} u(e) &\geq \sum_{i,j \in W: e \in P_{ij}} \frac{1}{k} \left(\frac{y(e, \mathcal{P}_i)}{l(e, \mathcal{P}_i)} + \frac{y(e, \mathcal{P}_j)}{l(e, \mathcal{P}_j)} \right) \\ &= \sum_{i \in W} \sum_{j \in W: e \in P_{ij}} \frac{1}{k} \cdot \frac{y(e, \mathcal{P}_i)}{l(e, \mathcal{P}_i)} \\ &= \frac{1}{k} \sum_{i \in W} y(e, \mathcal{P}_i). \end{aligned}$$

□

The theorem now follows by multiplying the inequality in (2) by $c(e)$ and summing over all $e \in E$:

$$\sum_{e \in E} c(e)u(e) \geq \sum_{e \in E} \frac{1}{k} \sum_{i \in W} y(e, \mathcal{P}_i) = \frac{1}{k} \sum_{i \in W} \sum_{e \in E} c(e)y(e, \mathcal{P}_i) \geq \min_{i \in W} \sum_{e \in E} c(e)y(e, \mathcal{P}_i).$$

□

The Pyramidal Routing problem has some interesting features that often make it more pleasant to work with. One such is the following.

Lemma 2.2. *There exists an optimal solution to a fractional Pyramidal Routing instance that is integral, i.e. f_{it} is an i - t -path for all $i \in W \setminus \{t\}$.*

Proof. (Sketch) We show this by proving that the problem consists of minimizing a concave function over a 0-1 polytope \mathcal{B} . \mathcal{B} encodes the flow problem solved by the template \mathcal{P} , namely, each terminal $i \in W$, we have a unit flow ($f_{it}(P)$), and so \mathcal{B} consists of the nonnegative vectors assigning weights to paths between terminal pairs such that the total flow from any terminal i to t is exactly one. By the max-flow min-cut theorem, the extreme points of \mathcal{B} are 0-1 vectors.

The objective function is $C_{\text{pyr}}(\mathcal{P}) = \sum_{e \in E} c(e)y(e, \mathcal{P}_t)$. As noted before, $y(e, \mathcal{P}_t)$ is a concave function in the load on e . In fact, one checks that it is actually concave over \mathcal{B} . Since the sum of concave functions remains concave, $C_{\text{pyr}}(\mathcal{P})$ is also concave over \mathcal{B} . It is well-known that an optimal solution always exists at a vertex of the polytope, which corresponds to an integral routing template. □

We have seen that the VPN cost is lower bounded by some Pyramidal Routing cost. In fact, a converse result holds too; the cost of optimal SPR solutions can also be upper bounded by the cost of an associated Pyramidal Routing problem. To do this, for each solution \mathcal{P}_t to a PR problem instance with root t , we define an oblivious routing template, called the *truncated hub template* associated with \mathcal{P}_t . This is defined as the template $\mathcal{Q} = \{Q_{ij} : i, j \in W\}$, where Q_{ij} is any i - j -path in the component of $P_i \Delta P_j$ (symmetric difference) containing i and j . Note that since i and j are the only odd-degree nodes in $P_i \Delta P_j$, they will indeed be in the same component.

Lemma 2.3. *Given a solution \mathcal{P}_t to a PR problem instance with root t , the capacity on any edge e required by its truncated hub template \mathcal{Q} is at most $y(e, \mathcal{P}_t)$. In particular, $C_{\text{vpn}}(\mathcal{Q}) \leq C_{\text{pyr}}(\mathcal{P}_t)$.*

Proof. Let $D \in \mathcal{H}$ be any valid demand matrix. Consider any edge e and define the set of nodes which route through e by $R_e := \{i \in W : e \in P_{it}\}$.

Now notice that if we have a pair $i, j \in W$ where $e \in Q_{ij}$, then exactly one of i and j is in R_e , because of the symmetric difference construction. So we have

$$\begin{aligned} \sum_{i,j:e \in Q_{ij}} D_{ij} &= \sum_{i \in R_e} \sum_{j \notin R_e} D_{ij} \\ &\leq \sum_{i \in R_e} \sum_{j \in W} D_{ij} \\ &\leq \sum_{i \in R_e} 1 \\ &= |R_e|. \end{aligned}$$

Similarly,

$$\sum_{i,j:e \in Q_{ij}} D_{ij} = \sum_{j \notin R_e} \sum_{i \in R_e} D_{ij} \leq \sum_{j \notin R_e} 1 = |W \setminus R_e|.$$

Thus we have that $\sum_{i,j:e \in Q_{ij}} D_{ij} \leq \min\{l_e, k - l_e\} = y(e, \mathcal{P}_t)$. But then the required capacity on edge e is

$$u(e) = \max_{D \in \mathcal{H}} \sum_{i,j:e \in Q_{ij}} D_{ij} \leq y(e, \mathcal{P}_t),$$

as required. □

Note that by Lemma 2.1 the optimal SPR cost is at least a convex combination of costs $C_{\text{pyr}}(\mathcal{P}_t)$. The preceding lemma shows that it is also at most the cost of any given optimal PR solution. Thus we have the following.

Theorem 2.4. *For any pair $t, t' \in W$, the optimal solutions to the PR problems with root t and t' are the same, and have the same value as the optimal solution of the associated SPR problem. Hence, the PR conjecture and the VPN conjecture are equivalent.*

2.3 A Reduction to T -joins

Begin with an instance of the Pyramidal Routing problem, with root t . Let \mathcal{P}_t be a routing template for this instance.

Definition 1. *Call an edge $e \in E$ **heavy** if $l(e) \geq k/2$.*

Let H be the set of heavy edges determined by \mathcal{P}_t . Note that

$$y(e) = \begin{cases} l(e) & \text{if } e \notin H \\ k - l(e) & \text{if } e \in H. \end{cases}$$

Let T' be the set of odd degree vertices in the subgraph induced by H . Now define $T = T' \Delta \{t\}$, and $T_u := T \Delta \{u\}$ for all $u \in W$. Note that $|T'|$ is even, so $|T|$ is odd, and so $|T_u|$ is even for all $u \in W$.

Let M_u be the minimum cost T_u -join on G . $C(M_u) := \sum_{e \in M_u} c(e)$ is defined to be the cost of M_u .

Define $C'(u) := C(P_{ut} \Delta H)$. This we think of as being u 's contribution to the total cost of the pyramidal routing. Notice that u pays for light edges on its path, and heavy edges not on its path. We also have

$$\begin{aligned} \sum_{u \in W} C'(u) &= \sum_{u \in W} \sum_{e \in P_{ut} \Delta H} c(e) \\ &= \sum_{e \in E \setminus H} l(e) \cdot c(e) + \sum_{e \in H} (k - l(e))c(e) \\ &= C_{\text{pyr}}(\mathcal{P}_t). \end{aligned}$$

So this really is a division of the total cost between the terminals.

Theorem 2.5. *A lowerbound for the cost of solution \mathcal{P}_t is $\sum_{u \in W} C(M_u)$.*

Proof. We want to show $C'(u) \geq C(M_u)$. Consider the symmetric difference $H_u = P_{ut} \Delta H$. Since P_{ut} has even degree at every node except u and t , and H is a T' -join, it follows that H_u is a T_u -join. So $C(H_u) \geq C(M_u)$. But by definition, $C'(u) = C(H_u)$. Thus

$$C_{\text{pyr}}(\mathcal{P}_t) = \sum_{u \in W} C'(u) = \sum_{u \in W} C(H_u) \geq \sum_{u \in W} C(M_u).$$

□

Note that the right hand side of this inequality is independent of H . We also have the following pleasant result. Define the truncated template $\mathcal{Q} = \{Q_{uv} : u, v \in W\}$, where Q_{uv} is any u - v path contained within the component of $M_u \Delta M_v$ containing u and v (both must be in the same component by considering the degree parities). Then:

Theorem 2.6. *The truncated template \mathcal{Q} satisfies*

$$C_{\text{vpn}}(\mathcal{Q}) \leq \sum_{u \in W} C(M_u).$$

The proof is very similar to the proof of Lemma 2.3, and so we omit it.

Note that if we had $T = \{v\}$, then M_u is a shortest path from u to v . So if v is the centre of the minimum cost VPN tree, $\sum_{u \in W} C(M_u)$ is exactly the cost of the optimal tree solution. In the next section, we show that no other choice of T improves upon this.

2.4 A T -join inequality

For each node t in G , define $C_{\text{SP}}(v)$ to be the cost of the VPN tree from the terminals to t , i.e.

$$C_{\text{SP}}(v) = \sum_{u \in W} SP(u, v),$$

where $SP(u, v)$ is the cost of the shortest path between u and v in G . Call a set $S \subseteq V$ T -even (respectively T -odd) if $|S \cap T|$ is even (respectively odd). Note that since $|T|$ is odd, exactly one of S and $V \setminus S$ is T -even for any $S \subseteq V$.

We prove the following inequality, which in turn proves the Pyramidal Routing conjecture by the reduction in the previous section:

Theorem 2.7. *There exists a node $v \in V$ so that*

$$\sum_{u \in W} C(M_u) \geq C_{\text{SP}}(v).$$

In fact, we prove the following slightly stronger theorem:

Theorem 2.8. *Let F be the multigraph obtained by taking the disjoint union of the M_u 's. Then there exists a node $v \in V$ so that there are edge-disjoint paths from all the vertices in W to v .*

We will need the following lemma:

Lemma 2.9. *For any set $S \subseteq V$ which is T -even, $|\delta_F(S)| \geq |S \cap W|$.*

Proof. Consider any $u \in S \cap W$. Since S is T -even, it is T_u -odd, and so M_u must intersect $\delta_F(S)$. \square

Proof of 2.8. Construct the graph F' from F by adding an extra node s , and edges su for all $u \in W$. The statement of the theorem is equivalent to showing that there exists a node v so that there is an s - v flow of size k on F' , taking all the edges to have unit capacity.

Define D_z to be the side of the minimum s - z cut containing $z \in V$. Suppose for a contradiction that $|\delta_{F'}(D_z)| < k$ for all $z \in V$, since otherwise we have a valid routing by the max-flow min-cut theorem. So then $|\delta_F(D_z)| + |D_z \cap W| < k$, i.e.

$$|\delta_F(D_z)| < k - |D_z \cap W| \quad \forall z \in V. \quad (3)$$

Note that D_z is T -even, since otherwise $V \setminus D_z$ would be T -even, which would imply

$$|\delta_F(D_z)| = |\delta_F(V \setminus D_z)| \geq |W \setminus D_z| = k - |D_z \cap W|,$$

contradicting (3).

Suppose inductively that all intersections of less than l sets are T -even. This is true for $l = 2$, since D_z is T -even for all $z \in V$. Now suppose for a contradiction that for some arbitrary collection D_1, \dots, D_l , the set $D_1 \cap D_2 \cdots \cap D_l$ is T -odd. Let $D'_i = D_i \setminus (\cup_{j \neq i} D_j)$. It follows from the inclusion-exclusion principle, and our inductive assumption, that

Claim 3. D'_i is T -odd for all i .

Proof. Note that $D'_i = D_i \setminus (\cup_{j \neq i} (D_j \cap D_i))$. Now by the inclusion-exclusion principle, we have

$$|\cup_{j \neq i} (D_j \cap D_i) \cap T| = \sum_{j \neq i} |(D_j \cap D_i) \cap T| - \sum_{j_1 < j_2: j_1, j_2 \neq i} |(D_{j_1} \cap D_{j_2} \cap D_i) \cap T| + \dots \\ + (-1)^{l-1} |(D_1 \cap D_2 \cap \dots \cap D_l) \cap T|.$$

The last term on the right is odd and the rest are even by our assumptions, and thus $\cup_{j \neq i} (D_j \cap D_i)$ is T -odd; since D_i is T -even, it follows that D'_i is T -odd. \square

Claim 4.

$$\sum_{i=1}^l |\delta_F(D_i)| \geq \sum_{i=1}^l |\delta_F(D'_i)|$$

Proof. Consider any edge e that contributes to the right hand side. If it has endpoints in D'_i and D'_j for some $i \neq j$, then it will contribute twice to the right hand side; such an edge will also contribute at least twice to the left hand side, in $\delta_F(D_i)$ and $\delta_F(D_j)$. If e has an endpoint in D'_i only, and not in any other D'_j , then it is counted once on the right hand side, and at least once on the left hand side. \square

Now we have

$$\begin{aligned} \sum_{i=1}^l |\delta_F(D_i)| &\geq \sum_{i=1}^l |\delta_F(D'_i)| && \text{by Claim 4} \\ &\geq \sum_{i=1}^l (k - |D'_i \cap W|) && \text{as } D'_i \text{ is } T\text{-odd} \\ &\geq \sum_{i=1}^l (k - |D_i \cap W|) && \text{as } D'_i \subseteq D_i. \end{aligned}$$

But this is a contradiction because (3) implies

$$\sum_{i=1}^l |\delta_F(D_i)| < \sum_{i=1}^l (k - |D_i \cap W|).$$

So our assumption that $D_1 \cap \dots \cap D_l$ is T -odd was incorrect. Thus inductively, arbitrary intersections of the D_v 's are T -even. It follows that $\cup_{u \in V} D_u$ is T -even, again by the inclusion-exclusion principle. But $u \in D_u$, so $\cup_{u \in V} D_u = V$, which is T -odd. This contradiction implies the result. \square

This argument actually proves a more general polyhedral result for any odd subset T of nodes. In any vertex x^* of the polyhedron

$$\{x \in \mathbf{R}^{E(G)} : x \geq 0, x(\delta(S)) \geq |S| \text{ for any } T\text{-odd set } S\},$$

there is some node t such that x^* assigns enough capacity for all nodes to simultaneously route one unit of flow to t . One argues that x^* is defined by a laminar family, and uses uncrossing and parity as in the proof above. Further details will appear in a full version.

3 The Multipath VPN Conjecture

It has been conjectured [18] (cf. [12]) that the multipath version of the VPN problem, where routing templates may be fractional, also has an optimal solution in the form of a tree. In this section, we show that this conjecture is false.

Recall that a *projective plane* of order n is a $(n^2 + n + 1, n + 1, 1)$ block design. It consists of a set W of $n^2 + n + 1$ *points*, and a collection \mathcal{L} of $n^2 + n + 1$ subsets of points (the *lines*). Every line contains exactly $n + 1$ points, every point is in exactly $n + 1$ lines, every pair of lines determines a unique point, and every pair of points lies on a unique line.

The complement (W, \mathcal{B}) of a projective plane (W, \mathcal{L}) is obtained by replacing each line with its complement: $\mathcal{B} = \{W \setminus L : L \in \mathcal{L}\}$. Call the sets in \mathcal{B} *blocks*. This is an $(n^2 + n + 1, n^2, n^2 - n)$ design: every block contains n^2 points, every point is in n^2 blocks, every pair of blocks have exactly $n^2 - n$ points in common, and every pair of points is contained in exactly $n^2 - n$ blocks. It is well known that projective planes (and hence their complements) exist for all orders that are powers of primes.

We now construct the bipartite graph $G = (W \cup U, E)$ as follows from the complement (W, \mathcal{B}) of a projective plane of order n as follows. The nodes in W correspond to points, and the nodes in U correspond to the blocks. An edge (w, B) exists iff the w is contained in the block B . Our construction implies that $|\delta(v)| = n^2$ for all $v \in W \cup U$, and $|\delta(v) \cap \delta(w)| = n^2 - n$ for $v, w \in W$ and $v, w \in U$. We also have $|W| = |U| = n^2 + n + 1$.

We now consider the MPR instance on this graph, where all edges have unit cost, and W is the set of terminals. Notice that for $n \geq 2$, the optimal VPN tree has cost

$$C_t = n^2 + 3(n + 1),$$

by rooting at any node in U ; a VPN tree routed on a node in W has cost $2(n^2 + n)$, which is larger. But consider the following multipath routing: for any pair $i, j \in W$, route a fraction $1/(n^2 - n)$ through each of the $n^2 - n$ common neighbours of i and j . Then it is easily seen that the required capacity on any edge e is just $u(e) = 1/(n^2 - n)$. This gives a total cost of

$$C_f = \frac{1}{n^2 - n} \cdot (n^2 + n + 1) \cdot n^2 = \frac{n(n^2 + n + 1)}{n - 1}.$$

But a quick calculation shows that $C_f < C_t$ for $n \geq 2$, thus giving a class of counterexamples.

In fact, the same technique yields a counterexample for any (v, l, λ) block design satisfying

$$\frac{v(v-1)}{l-1} < 3v - 2l.$$

Some examples are the complements of Steiner triple systems, and the unique $(9, 6, 5)$ design.

Acknowledgements: We are grateful to the authors of [15] who provided a preprint of their paper. Their ideas substantially altered our approach to the VPN problem.

References

- [1] A. Altin, E. Amaldi, P. Belotti, and M. C. Pinar. Provisioning virtual private networks under traffic uncertainty. *Networks*, 49(1):100–115, January 2007.
- [2] D. Applegate and E. Cohen. Making intra-domain routing robust to changing and uncertain traffic demands: understanding fundamental tradeoffs. In *Proceedings of ACM SIGCOMM*, pages 313–324, New York, NY, USA, 2003. ACM.
- [3] Y. Azar, E. Cohen, A. Fiat, H. Kaplan, and H. Räcke. Optimal oblivious routing in polynomial time. *Journal of Computer and System Sciences*, 69(3):383–394, November 2004.
- [4] W. Ben-Ameur and H. Kerivin. New economical virtual private networks. *Commun. ACM*, 46(6):69–73, June 2003.
- [5] A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Mathematics of Operations Research*, 23(4):769–805, November 1998.
- [6] A. Ben-Tal and A. Nemirovski. Robust solutions of uncertain linear programs. *Operations Research Letters*, 25(1):1–13, August 1999.
- [7] A. Borodin and J. E. Hopcroft. Routing, merging and sorting on parallel models of computation. In *Proceedings of ACM STOC*, pages 338–344, New York, NY, USA, 1982. ACM.
- [8] C.-S. Chang, D.-S. Lee, and Y.-S. Jou. Load balanced Birkhoff-von Neumann switches, part I: one-stage buffering. *Journal of Computer Communications*, 25(6):611–622, April 2002.
- [9] C. Chekuri. Routing and network design with robustness to changing or uncertain traffic demands. *SIGACT News*, 38(3):106–128, September 2007.
- [10] C. Chekuri, G. Oriolo, M. G. Scutella, and F. B. Shepherd. Hardness of robust network design. *Networks*, 50(1):50–54, 2007.

- [11] 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 ACM SIGCOMM*, pages 95–108, New York, NY, USA, 1999. ACM.
- [12] T. Erlebach and M. Ruegg. Optimal bandwidth reservation in hose-model vpns with multi-path routing. In *Proceedings of INFOCOM*, volume 4, pages 2275–2282 vol.4, 2004.
- [13] A. J. Fingerhut, S. Suri, and J. S. Turner. Designing least-cost nonblocking broadband networks. *Journal of Algorithms*, 24(2):287–309, August 1997.
- [14] S. Fiorini, G. Oriolo, L. Sanità, and D. O. Theis. The VPN tree routing conjecture for outerplanar networks. Submitted manuscript, 2007.
- [15] F. Grandoni, V. Kaibel, G. Oriolo, and M. Skutella. A short proof of the VPN tree routing conjecture on ring networks. Accepted to *Operations Research Letters*, Oct 2007.
- [16] A. Gupta, J. Kleinberg, A. Kumar, R. Rastogi, and B. Yener. Provisioning a virtual private network: a network design problem for multicommodity flow. In *Proceedings of ACM STOC*, pages 389–398, New York, NY, USA, 2001. ACM Press.
- [17] A. Gupta and A. Kumar. Private communication, 2001.
- [18] C. A. J. Hurkens, J. C. M. Keijsper, and L. Stougie. Virtual private network design: A proof of the tree routing conjecture on ring networks. *Proceedings of Integer Programming and Combinatorial Optimization*, pages 407–421, 2005.
- [19] G. Italiano, S. Leonardi, and G. Oriolo. Design of trees in the hose model: the balanced case. *Operations Research Letters*, 34(6):601–606, 2006. A preliminary version appeared in *Proceedings of ARACNE*, 65-76, 2002.
- [20] C. Kaklamanis, D. Krizanc, and T. Tsantilas. Tight bounds for oblivious routing in the hypercube. In *Proceedings of ACM SPAA*, pages 31–36, New York, NY, USA, 1990. ACM.
- [21] I. Keslassy, S.-T. Chuang, K. Yu, D. Miller, M. Horowitz, O. Solgaard, and N. McKeown. Scaling internet routers using optics. In *Proceedings of ACM SIGCOMM*, pages 189–200, New York, NY, USA, 2003. ACM.
- [22] M. Kodialam, T. V. Lakshman, and S. Sengupta. Efficient and robust routing of highly variable traffic. In *Proceedings of HotNets III*, 2004.
- [23] T. Leighton and S. Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM*, 46(6):787–832, November 1999.

- [24] C. Lemaréchal, A. Ouorou, and G. Petrou. Robust network design in telecommunications. In *Proceedings of INOC*, 2005.
- [25] D. Mitra and R. A. Cieslak. Randomized parallel communications on an extension of the omega network. *Journal of the ACM*, 34(4):802–824, October 1987.
- [26] H. Nagesh, V. Poosala, S. Sengupta, M. Alicherry, and V. Kumar. Netswitch: Load-balanced data-over-optical architecture for mesh networks. Technical report, Lucent Technical Memorandum ITD-04-45867F, 2004.
- [27] H. Räcke. Minimizing congestion in general networks. In *Proceedings of IEEE FOCS*, pages 43–52, 2002.
- [28] S. Sengupta, V. Kumar, and D. Saha. Switched optical backbone for cost-effective scalable core IP networks. *IEEE Communications Magazine*, 41(6):60–70, 2003.
- [29] R. Z. Shen and N. Mckeown. Designing a predictable internet backbone network. In *Proceedings of HotNets III*, 2004.
- [30] F. B. Shepherd and P. J. Winzer. Selective randomized load balancing and mesh networks with changing demands. *Journal of Optical Networking*, 5(5):320–339, May 2006.
- [31] I. Widjaja and A. I. Elwalid. Exploiting parallelism to boost data-path rate in high-speed IP/MPLS networking. In *Proceedings of IEEE INFOCOM*, volume 1, pages 566–575, 2003.
- [32] P. J. Winzer, F. B. Shepherd, P. Oswald, and M. Zirngibl. Robust network design and selective randomized load balancing. In *Proceedings of the 31st European Conference on Optical Communications (ECOC 2005)*, volume 1, pages 23–24, 2005.