

A Short Proof of the VPN Tree Routing Conjecture on Ring Networks

Fabrizio Grandoni* Volker Kaibel† Gianpaolo Oriolo‡§ Martin Skutella¶

January 15, 2008

Abstract

Only recently, Hurkens, Keijsper, and Stougie proved the VPN Tree Routing Conjecture for the special case of ring networks. We present a short proof of a slightly stronger result which might also turn out to be useful for proving the VPN Tree Routing Conjecture for general networks.

Keywords: network design, virtual private networks, tree routing

1 Introduction

Consider a communication network which is represented by an undirected graph $G = (V, E)$ with edge costs $c : E \rightarrow \mathbb{R}_{\geq 0}$. Within this network there is a set of k terminals $W \subseteq V$ 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 $i \in W$ has an upper bound $b(i) \in \mathbb{Z}_+$ on the cumulative amount of traffic that terminal i can send or receive. The general aim is to install capacities on the edges of the graph supporting any possible communication scenario at minimum cost where the cost for installing one unit of capacity on edge e is $c(e)$.

A set of traffic demands $D = \{d_{ij} \mid i, j \in W\}$ specifies for each unordered pair of terminals $i, j \in W$ the amount $d_{ij} \in \mathbb{R}_{\geq 0}$ of traffic between i and j . A set D is *valid* if it respects the upper bounds on the traffic of the terminals. That is, (setting $d_{ii} = 0$ for all $i \in W$)

$$\sum_{j \in W} d_{ij} \leq b(i) \quad \text{for all terminals } i \in W.$$

A solution to the *symmetric Virtual Private Network Design* (sVPND) problem defined by G, c, W , and b consists of an i - j -path P_{ij} in G for each unordered pair $i, j \in W$, and edge capacities $u(e) \geq 0$, $e \in E$. Such a set of paths P_{ij} , $i, j \in W$, together with edge capacities $u(e)$, $e \in E$, is called a *virtual private network*. A virtual private network is *feasible* if all valid sets of traffic demands D can be routed without exceeding the installed capacities u where all traffic between terminals i and j is routed along path P_{ij} , that is, (with $\binom{W}{2}$ denoting the set of cardinality-two subsets of W)

$$u(e) \geq \sum_{\{i,j\} \in \binom{W}{2} : e \in P_{ij}} d_{ij} \quad \text{for all edges } e \in E.$$

*Dipartimento di Informatica, Sistemi e Produzione, Università di Roma “Tor Vergata”, Italy, grandoni@disp.uniroma2.it

†Fakultät für Mathematik, Otto-von-Guericke-Universität Magdeburg, Germany, kaibel@ovgu.de

‡Dipartimento di Ingegneria dell’Impresa, Università di Roma “Tor Vergata”, Italy, oriolodisp.uniroma2.it

§Corresponding author.

¶Institut für Mathematik, Technische Universität Berlin, Germany, skutella@math.tu-berlin.de

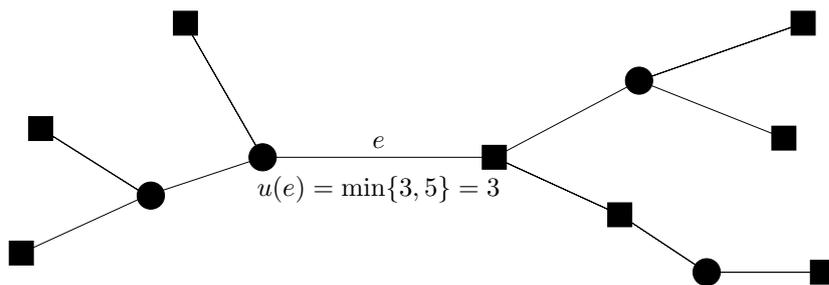


Figure 1: A tree solution for an $sVPND$ instance. The 8 rectangular nodes are the terminals in W , the remaining nodes are non-terminal nodes. For the sake of simplicity we assume in this instance that $b(j) = 1$ for all terminals $j \in W$. If some edge e is removed from the tree, two connected components remain. The smaller component on the left hand side contains 3 terminals while the larger component on the right hand side contains 5 terminals. Therefore the maximum amount of traffic on edge e is 3 which occurs when all terminals on the left hand side want to communicate with some terminal on the right hand side.

A feasible virtual private network is called *optimal* if the total cost of the capacity reservation $\sum_{e \in E} c(e) u(e)$ is minimal.

A well-known open question is whether the $sVPND$ problem can be solved efficiently (i.e., in polynomial time); see Erlebach and Rüegg [1] and Italiano, Leonardi, and Oriolo [5]. A feasible virtual private network is a *tree solution* if the subgraph of G induced by the support of u (i.e., edges $e \in E$ with $u(e) > 0$) is a tree. Gupta, Kleinberg, Kumar, Rastogi, and Yener [3] prove that a tree solution of minimum cost can be obtained in polynomial time by an all-pair shortest paths computation on the network G . In Figure 1 we explain how optimal (i.e., minimal) capacities can be determined for a fixed sub-tree of G spanning all terminals. In the following we assume that capacities are chosen accordingly in any tree solution.

Although many different (groups of) researchers are working on the $sVPND$ problem, there is no instance known where a tree solution of minimum cost is not simultaneously an optimal virtual private network. It is widely believed that the following conjecture holds.

Conjecture 1 (The VPN Tree Routing Conjecture). *For each $sVPND$ instance (G, c, W, b) there exists an optimal virtual private network which is a tree solution.*

The only progress in this direction is due to Hurkens, Keijsper, and Stougie [4] who prove that Conjecture 1 holds on ring networks and for some other special cases of the problem. Their results are based on a linear programming (LP) formulation of a relaxation of the $sVPND$ problem. In this relaxation several paths may be chosen between each pair of terminals but the fraction of traffic along each of these paths must be fixed, i.e., it may not depend on the actual set of traffic demands. Hurkens, Keijsper, and Stougie construct solutions to the dual linear program whose cost equals the cost of particular tree solutions. The details of their proof are somewhat involved.

In the following we present a simpler proof of the fact that Conjecture 1 holds on ring networks. It is based on a new and stronger conjecture which might be of independent interest. In Section 2 we argue that we can restrict to instances with unit communication bounds. Our new conjecture is presented in Section 3. Finally we give a short proof for the case of ring networks in Section 4.

2 Preliminaries

As already observed in [4], when proving Conjecture 1, we may assume that $b(i) = 1$ for each terminal $i \in W$. For the sake of self-containedness, we motivate this assumption in the following. For more

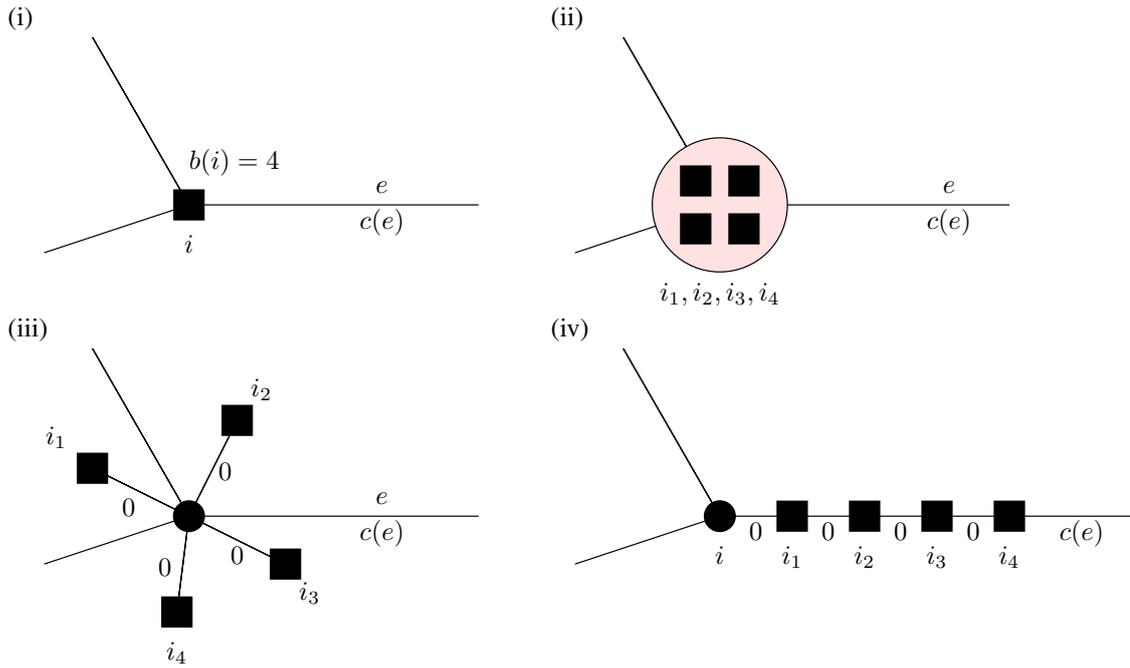


Figure 2: Illustration of the reduction described in Section 2. As an example, we consider a terminal $i \in W$ with $b(i) = 4$ and three incident edges; see part (i) of the figure. Part (ii) depicts the relaxed instance where the terminal at node i is split into 4 sub-terminals with unit communication bounds. In order to get an instance with at most one terminal at every node, we can introduce 4 new terminal nodes, one for each sub-terminal, and connect them to node i by edges of zero cost; see part (iii). Alternatively, in order to stay closer to the original topology of the network, one can subdivide an edge e incident to node i and place the terminals there with zero connection cost to i ; see part (iv).

details we refer to [4].

Consider a terminal $i \in W$ with $b(i) \geq 2$; see Figure 2 (i) for an example. We construct a new sVPND instance by replacing terminal i with $b(i)$ “sub-terminals” $i_1, \dots, i_{b(i)}$ with $b(i_j) = 1$ for $j = 1, \dots, b(i)$ that are all collocated at the same node $i \in V$; see Figure 2 (ii). Strictly speaking, the collocation of terminals at a node is a slight extension of the original definition of the sVPND problem. We show how to deal with this below.

Obviously, the new instance is a relaxation of the original instance. The additional degree of freedom of the new instance is that traffic from terminal i to some other terminal no longer needs to be routed along one fixed path but may be split into $b(i)$ packets of equal size that can be routed along different (but also fixed) paths.

In order to show that there exists an optimal virtual private network for the original instance which is a tree, it is obviously sufficient to find such an optimal solution to the new instance.

We finally argue that the new instance is itself equivalent to an sVPND instance (in the strict sense) with all $b(i)$ ’s equal to 1. Instead of having $b(i)$ sub-terminals located at node i , we add new terminal nodes $i_1, \dots, i_{b(i)}$ to the network and connect them to node i with edges of cost 0; see Figure 2 (iii). It is easy to observe that every feasible (tree) solution to the instance in Figure 2 (ii) naturally corresponds to a feasible (tree) solution to the instance in Figure 2 (iii) of the same cost and vice versa.

The only problem with the described approach is that the resulting network no longer has the same topology as the network of the original instance. For example, if we start with a ring network, the resulting network in Figure 2 (iii) is not a ring. This problem can be resolved by using the alternative construction illustrated in Figure 2 (iv). Here we subdivide an arbitrary edge e incident to node i

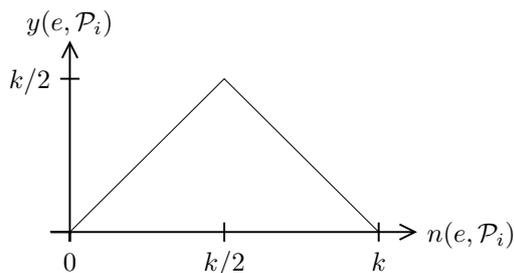


Figure 3: The concave cost function of the Pyramidal Routing Problem.

into $b(i) + 1$ parts by introducing $b(i)$ new terminal nodes $i_1, \dots, i_{b(i)}$. The first $b(i)$ edges connecting the new terminal nodes to node i have cost 0. To the last edge we assign cost $c(e)$. Again, it is easy to observe that every feasible solution to the instance in Figure 2 (ii) naturally corresponds to a feasible solution to the instance in Figure 2 (iv) of the same cost and vice versa.

In order to argue that also tree solutions correspond to each other, one has to be a little more careful. Notice that not every tree solution to the instance in Figure 2 (iv) induces a tree solution to the instance in Figure 2 (ii) (e.g., if terminals i_1 and i_2 are leaves of the tree). But there always exists an *optimal* tree solution to the instance in Figure 2 (iv) containing all edges $i_1i_2, \dots, i_{b(i)-1}i_{b(i)}$. This follows from the fact that these edges have cost 0 and an optimal tree solution is a shortest-paths tree for some source node $j \in V$; see [3].

We conclude this section with the following lemma resulting from our considerations above.

Lemma 1. *Consider a class of networks which is closed under the operation of subdividing edges by introducing new nodes of degree 2. Then Conjecture 1 holds for this class of networks if and only if it holds for the restricted class of instances on such networks where $b(i) = 1$ for all $i \in W$.*

3 The Pyramidal Routing Problem

Consider again a tree solution to an sVPND instance as depicted in Figure 1 with $k = |W|$ terminals and $b(i) = 1$ for all terminals $i \in W$. For a fixed terminal $i \in W$, let \mathcal{P}_i denote the set of simple paths P_{ij} in the tree that connect i to the remaining terminals $j \in W \setminus \{i\}$. Notice that the required capacity of edge e can be written as

$$u(e) = \min\{n(e, \mathcal{P}_i), k - n(e, \mathcal{P}_i)\}$$

where $n(e, \mathcal{P}_i)$ is the number of paths in \mathcal{P}_i containing edge e , that is,

$$n(e, \mathcal{P}_i) := |\{j \in W \setminus \{i\} \mid e \in P_{ij}\}| . \quad (1)$$

To simplify notation in the following we let

$$y(e, \mathcal{P}_i) := \min\{n(e, \mathcal{P}_i), k - n(e, \mathcal{P}_i)\} ; \quad (2)$$

see also Figure 3.

We now introduce another problem that we call *Pyramidal Routing* (PR) problem. The problem is defined by G, c, W —as above for the sVPND problem—and a *single source terminal* $i \in W$. A solution to the Pyramidal Routing Problem consists of a set \mathcal{P}_i of *simple i - j -paths* P_{ij} , one path for each terminal $j \in W \setminus \{i\}$. As above we denote the number of paths in \mathcal{P}_i containing a fixed edge $e \in E$ by $n(e, \mathcal{P}_i)$;

see (1). Moreover, we define $y(e, \mathcal{P}_i)$ accordingly as in (2). The PR problem is to find a solution \mathcal{P}_i that minimizes the objective function

$$\sum_{e \in E} c(e) y(e, \mathcal{P}_i) .$$

The PR problem can be seen as an unsplittable flow problem with concave cost functions on the edges (see Figure 3) where one unit of flow is sent from the source i to all destinations $j \in W \setminus \{i\}$. A *tree solution* to PR is a solution where the chosen paths $P_{ij}, j \in W \setminus \{i\}$, form a tree.

Lemma 2. *Consider an sVPND instance (G, c, W, b) with $b_j = 1$ for all $j \in W$ and a corresponding PR instance (G, c, W, i) for some terminal $i \in W$. Then any tree solution to the sVPND instance yields a tree solution of the same cost to the PR instance and vice versa.*

Proof. Consider any subtree of G containing all terminals. The induced set of i - j -paths, $j \in W \setminus \{i\}$, is \mathcal{P}_i . For the sVPND problem, the required capacity of some tree edge e is equal to $y(e, \mathcal{P}_i) = \min\{n(e, \mathcal{P}_i), k - n(e, \mathcal{P}_i)\}$. Therefore the cost of this tree solution is equal to $\sum_{e \in E} c(e) y(e, \mathcal{P}_i)$ for both problems. \square

We state an immediate corollary of the last lemma.

Corollary 1. *A tree solution of minimum cost to a PR instance (G, c, W, i) yields simultaneously a tree solution of (the same) minimum cost to the PR instances (G, c, W, j) for all $j \in W$. In particular, $y(e, \mathcal{P}_i) = y(e, \mathcal{P}_j)$ where \mathcal{P}_i and \mathcal{P}_j denote the two sets of paths connecting terminal i and j , respectively, with all other terminals.*

We state the following conjecture for the Pyramidal Routing problem.

Conjecture 2 (The Pyramidal Routing Conjecture). *For each PR instance (G, c, W, i) there exists an optimal solution which is a tree solution.*

As we show in the remainder of this section, Conjecture 2 is strongly related to Conjecture 1.

Theorem 1. *If Conjecture 2 holds on some class of networks, which is closed under the operation of subdividing edges by introducing new nodes of degree 2, then Conjecture 1 holds on the same class.*

Before we can prove this theorem we need one further result. The following lemma and its proof are a slight extension of Theorem 3.2 in [3].

Lemma 3. *Consider an sVPND instance (G, c, W, b) with $b(j) = 1$ for all $j \in W$ and some feasible virtual private network given by simple i - j -paths $P_{ij}, i, j \in W$, and capacities $u(e), e \in E$. There exists a terminal $i \in W$ such that $\sum_{e \in E} c(e) u(e) \geq \sum_{e \in E} c(e) y(e, \mathcal{P}_i)$, where $\mathcal{P}_i = \{P_{ij} \mid j \in W \setminus \{i\}\}$.*

Proof. We first consider some fixed edge $e \in E$. In order to derive a lower bound on $u(e)$, we define a set of traffic demands D^e by setting

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

Claim 1. D^e is a set of valid traffic demands.

Proof. We need to show that $\sum_{j \in W} d_{ij}^e \leq 1$ for each $i \in W$. Remember that $y(e, \mathcal{P}_i) = \min\{n(e, \mathcal{P}_i), k - n(e, \mathcal{P}_i)\}$. By definition of D^e we thus get

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

The last equation follows since $|\{j \in W \mid e \in P_{ij}\}| = n(e, \mathcal{P}_i)$. \square

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

Proof. We know from the previous claim that D^e is a valid set of traffic demands. Moreover, by definition, $d_{ij}^e > 0$ implies $e \in P_{ij}$. Therefore

$$u(e) \geq \frac{1}{k} \sum_{\{i,j\} \in \binom{W}{2}: e \in P_{ij}} \left(\frac{y(e, \mathcal{P}_i)}{n(e, \mathcal{P}_i)} + \frac{y(e, \mathcal{P}_j)}{n(e, \mathcal{P}_j)} \right) = \frac{1}{k} \sum_{i \in W: n(e, \mathcal{P}_i) > 0} y(e, \mathcal{P}_i) = \frac{1}{k} \sum_{i \in W} y(e, \mathcal{P}_i)$$

since $y(e, \mathcal{P}_i) = 0$ if $n(e, \mathcal{P}_i) = 0$. \square

Since Claim 2 holds for all $e \in E$, it follows that

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

This concludes the proof of the lemma. \square

We can now prove Theorem 1.

Proof of Theorem 1. Consider an sVPND instance (G, c, W, b) with $b_j = 1$ for all $j \in W$; this can be assumed without loss of generality due to Lemma 1. Consider an optimal solution to the instance. By Lemma 3 there exists a terminal $i \in W$ such that the optimal solution value of the PR instance (G, c, W, i) is a lower bound on the optimal solution value of the sVPND instance. If Conjecture 2 holds for the PR instance, there exists an optimal tree solution to the PR instance which, by Lemma 2, also yields an optimal virtual private network. \square

4 The Case of Ring Networks

In this section we prove Conjecture 2 for the case that G is an arbitrary ring network, i.e., a cycle.

Theorem 2. *Conjecture 2 holds true when G is a ring network.*

Proof. Consider a PR instance (G, c, W, i) where G is a cycle. We can assume without loss of generality that each node $v \in V$ is a terminal, i.e., $W = V$: If there is a non-terminal node v with two neighbors $x, y \in V$, we can remove v and replace edges xv and vy by a new edge xy with $c(xy) := c(xv) + c(vy)$.

Let \mathcal{P}_i be an arbitrary optimal solution to this PR instance.

Claim 3. *Let $e, f \in E$ be a pair of incident edges with $i \notin e \cap f$ (i.e., not both edges are incident to terminal i). Then $n(e, \mathcal{P}_i) = n(f, \mathcal{P}_i) \pm 1$.*

Proof. Let $j \in e \cap f$ be the terminal that is incident to both e and f . Since $j \neq i$, exactly one path in \mathcal{P}_i (namely the i - j -path) ends in j . All other paths either contain both edges e and f or none of the two. \square

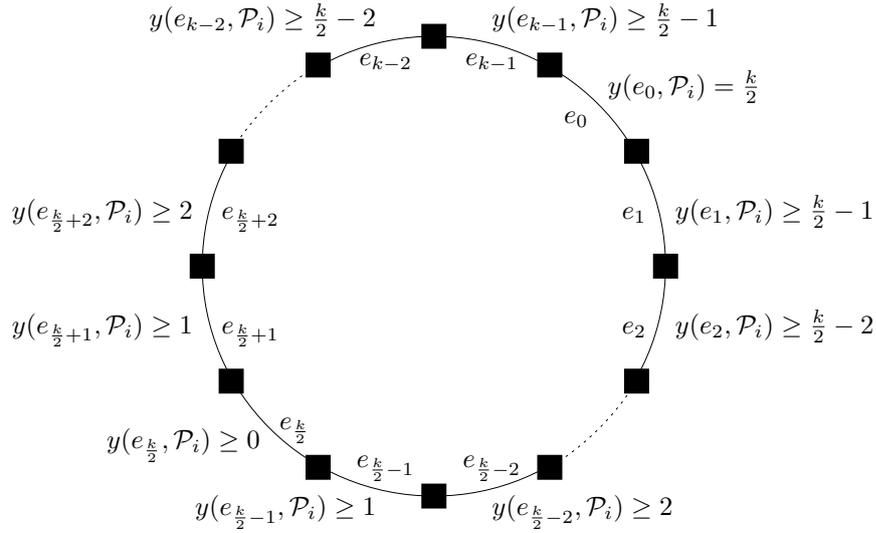


Figure 4: Lower bounding the y -values along the cycle.

Claim 4. Let $e, f \in E$ be an arbitrary pair of incident edges. Then $y(e, \mathcal{P}_i) = y(f, \mathcal{P}_i) \pm 1$.

Proof. If $i \notin e \cap f$, the claim follows immediately from Claim 3. We can therefore assume that e and f are the two edges incident to terminal i . Since there are exactly $k - 1$ simple paths in \mathcal{P}_i starting at i , we get $n(e, \mathcal{P}_i) + n(f, \mathcal{P}_i) = k - 1$. The result follows by definition of $y(e, \mathcal{P}_i)$ and $y(f, \mathcal{P}_i)$; see (2). \square

We distinguish two cases.

First case: k is even. Since $n(e, \mathcal{P}_i) + n(f, \mathcal{P}_i) = k - 1$ for the two edges e and f incident to i , we can assume that $n(e, \mathcal{P}_i) \geq \frac{k}{2}$ and $n(f, \mathcal{P}_i) < \frac{k}{2}$. Going along the cycle, it follows from Claim 3 that there exists an edge e_0 with $n(e_0, \mathcal{P}_i) = \frac{k}{2}$. We number the edges along the cycle consecutively e_0, e_1, \dots, e_{k-1} . Using Claim 4 inductively, it follows that

$$y(e_\ell, \mathcal{P}_i) \geq \left\lfloor \frac{k}{2} - \ell \right\rfloor \quad \text{for } \ell = 0, \dots, k - 1;$$

see Figure 4.

Removing edge $e_{\frac{k}{2}}$ from the cycle yields a tree and a corresponding tree solution \mathcal{P}'_i to the PR instance. Moreover, we argue that

$$y(e_\ell, \mathcal{P}'_i) = \left\lfloor \frac{k}{2} - \ell \right\rfloor \quad \text{for } \ell = 0, \dots, k - 1.$$

By Corollary 1 we can assume without loss of generality that terminal i is incident to edge e_0 . But in this case $n(e_\ell, \mathcal{P}'_i) = \left\lfloor \frac{k}{2} - \ell \right\rfloor$ and thus $y(e_\ell, \mathcal{P}'_i) = \left\lfloor \frac{k}{2} - \ell \right\rfloor$ for $\ell = 0, \dots, k - 1$.

As a result, $\sum_{e \in E} c(e)y(e, \mathcal{P}_i) \geq \sum_{e \in E} c(e)y(e, \mathcal{P}'_i)$. Since \mathcal{P}_i is an optimal solution, the tree solution \mathcal{P}'_i is optimal as well.

Second case: k is odd. In this case it follows from Claims 3 and 4 that there exist two incident edges that we call $e_{\frac{1}{2}}$ and $e_{k-\frac{1}{2}}$ with

$$y\left(e_{\frac{1}{2}}, \mathcal{P}_i\right) = y\left(e_{k-\frac{1}{2}}, \mathcal{P}_i\right) = \frac{k-1}{2}.$$

We number the edges along the cycle consecutively $e_{\frac{1}{2}}, e_{\frac{1}{2}+1}, \dots, e_{k-\frac{1}{2}}$. Using Claim 4 inductively, it follows that

$$y(e_\ell, \mathcal{P}_i) \geq \left\lfloor \frac{k}{2} - \ell \right\rfloor \quad \text{for } \ell = \frac{1}{2}, \dots, k - \frac{1}{2}.$$

The remainder of the proof is identical to the first case. \square

Theorem 1 and Theorem 2 imply that the VPN Tree Routing Conjecture is true for ring networks.

Corollary 2 ([4]). *Conjecture 1 holds true if G is a ring network.*

We conclude the paper with some remarks. First, in the definition of the Pyramidal Routing problem we might relax the constraints that the paths have to be simple and allow *trails*; a trail may visit nodes (but not edges) more than once. It is easy to see that Lemma 2, Lemma 3, and Theorem 1 still hold. (Observe that on a ring a path is a trail and vice versa.)

Then, in their paper [4], Hurkens, Keijsper, and Stougie prove a result that is slightly stronger than Corollary 2. Consider the relaxation of the $sVPND$ problem where traffic between each pair of terminals may be routed along several different paths but the fraction of traffic along each of these paths must be fixed. Hurkens et al. show that, for the case of ring networks, even the relaxed problem has always an optimal solution which is a tree solution.

Acknowledgements: The authors wish to thank Yuri Faenza for a useful discussion concerning Lemma 1. The third author has been supported by the Italian Research Project Cofin 2006 “Modelli ed algoritmi per l’ottimizzazione robusta delle reti”. The last author has been supported by the DFG Focus Program 1126, “Algorithmic Aspects of Large and Complex Networks”, grant SK 58/5-3.

Note added in proof: Conjecture 2 has been recently proved in [2].

References

- [1] T. Erlebach and M. Rüegg. Optimal Bandwidth Reservation in Hose-Model VPNs with Multi-Path Routing. In *Proceedings of the Twenty-Third Annual Joint Conference of the IEEE Computer and Communications Societies INFOCOM*, 4:2275-2282, 2004.
- [2] N. Goyal, N. Olver, F.B. Shepherd. The VPN Conjecture is true. Submitted, 2007.
- [3] 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 the 33rd ACM Symposium on Theory of Computing*, 389–398, ACM Press, 2001.
- [4] 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. *SIAM Journal on Discrete Mathematics*, 21:482–503, 2007.
- [5] G. Italiano, S. Leonardi, and G. Oriolo. Design of trees in the hose model: the balanced case. *Operation Research Letters*, 34(6): 601–606, 2006.