

Virtual private network design: a proof of the tree routing conjecture on ring networks

C.A.J. Hurkens J.C.M. Keijsper
L. Stougie*

Technische Universiteit Eindhoven, *CWI, Amsterdam

November 18, 2004

Abstract

A basic question in Virtual Private Network (VPN) design is if the symmetric version of the problem always has an optimal solution which is a tree network. An affirmative answer would imply that the symmetric VPN problem is solvable in polynomial time. We give an affirmative answer in case the communication network within which to create the VPN is a circuit. This seems to be an important step towards an answer to the general question. The proof relies on a dual pair of linear programs and actually implies an even stronger property of VPNs. We show that this property also holds for some other special cases of the problem.

1 Introduction

In this paper, we consider a problem that is known as the *symmetric virtual private network* (VPN) problem, a problem emerging in telecommunication. Think of a large communication network represented by an undirected graph $G = (V, E)$, with a vertex for each user, and an edge for each link in the network. Within this network, a subgroup $W \subseteq V$ of the users wishes to reserve capacity on the links of the network for communication among themselves: they wish to establish a virtual private network. Vertices in W are also called *terminals*.

On each link, capacity (bandwidth) has a certain price per unit, $c : E \rightarrow \mathbb{R}_+$. The problem is to select one or more communication paths between every pair $\{i, j\}$ of users in W and to reserve enough capacity on the edges of the selected paths to accommodate any possible communication pattern amongst the users in W . Possible communication patterns are defined through an upper bound on the amount to be communicated (transmitted and received) for each node in W , specified by $b : W \rightarrow \mathbb{R}_+$. This so-called hose model for VPNs was proposed by Duffield et al. [2]. More precisely, a *communication scenario* for the symmetric VPN problem can be defined as a symmetric matrix $D = (d_{ij})_{\{i,j\} \subseteq W}$ with zeros on the diagonal, specifying for each *unordered* pair of distinct nodes $\{i, j\} \subseteq W$ the amount of communication $d_{ij} \geq 0$ between i and j . A communication scenario $D = (d_{ij})_{\{i,j\} \subseteq W}$ is said to be *valid* if $\sum_{j \in W \setminus \{i\}} d_{ij} \leq b(i)$, $\forall i \in W$. We denote the collection of valid communication scenarios by \mathcal{D} .

A network consisting of the selected communication paths with enough capacity reserved on the edges to accommodate every valid communication scenario we call a *feasible VPN*. The (symmetric) VPN problem is to find the cheapest feasible VPN.

There are several variants of the problem emerging from additional routing requirements.

SPR Single path routing: For each pair $\{i, j\} \subseteq W$, exactly one path $P_{ij} \subseteq E$ is to be selected to accommodate all traffic between i and j . The problem is then to choose the paths P_{ij} such as to minimize $\{\sum_{e \in E} c_e x_e \mid x_e \geq \sum_{\{i,j\}: e \in P_{ij}} d_{ij}, \forall e \in E \forall D = (d_{ij}) \in \mathcal{D}\}$.

TTR Terminal tree routing: This is single path routing with the additional restriction that $\cup_{j \in W} P_{ij}$ should form a tree in G for all $i \in W$.

TR *Tree routing*: This is SPR with the extra restriction that $\cup_{\{i,j\} \subseteq W} P_{ij}$ is a tree in G .

MPR *Multi-path routing*: For each pair $\{i, j\} \subseteq W$, and for each possible path between i and j , the fraction of communication between i and j to be routed along that path has to be specified.

The following lemma summarizes the rather obvious relations between the optimal solution values of these variants. By $OPT(\text{SPR})$ we denote the cost of an optimal solution for the SPR variant of the VPN problem. Similar notation is used for the other optimal values.

Lemma 1.1

$$OPT(\text{MPR}) \leq OPT(\text{SPR}) \leq OPT(\text{TTR}) \leq OPT(\text{TR}).$$

Proof. SPR is the MPR problem with the extra restriction that all fractions must be 0 or 1. The other inequalities are similarly trivial. \square

This paper shows that equality holds throughout in Lemma 1.1 for some special cases, in particular in the case that G is a circuit.

We mention some previous results. Gupta et al. [5] and Kumar et al. [8] have shown that $\text{TR} \in P$ (TR is polynomially solvable). In [5] it is also shown that $OPT(\text{TR}) = OPT(\text{TTR})$ and that $OPT(\text{TR}) \leq 2OPT(\text{MPR})$. Erlebach and Rüegg [3] proved that $\text{MPR} \in P$. This fact also follows from our LP-formulation in Section 2 (see also Section 6).

A prominent open question in VPN is if the optimal solution of the SPR problem is always a tree; i.e., $OPT(\text{SPR}) = OPT(\text{TR})$. An affirmative answer to this problem would imply that $\text{SPR} \in P$. Italiano et al. [7] mention this as an open problem. Erlebach and Rüegg [3] mention that no VPN instance has been found so far for which even $OPT(\text{MPR}) < OPT(\text{TR})$.

Our own conjecture is that $OPT(\text{MPR}) = OPT(\text{TR})$, from which $\text{SPR} \in P$ would follow. $OPT(\text{MPR}) = OPT(\text{TR})$ (or even $OPT(\text{SPR}) = OPT(\text{TR})$) was not known to be true for any class of graphs other than trees. It seems to be a crucial step forward to prove it for circuits, which is the main result in this paper.

Theorem 1.2 *Let $G = (V, E)$ be a circuit. Then $OPT(\text{MPR}) = OPT(\text{TR})$.*

This result is proved in Section 4 of this paper. Before that, an LP-formulation of MPR is given in Section 2, and a reformulation of the conjecture that $OPT(\text{MPR}) = OPT(\text{TR})$ in terms of this LP-formulation is given in Section 3, together with some preliminary results.

In Section 5, we proceed to prove our conjecture for some other special cases. We prove it for any graph G and any cost function c , if the communication bound of some terminal is larger than the sum of the bounds of the other terminals. We also prove it for any graph on at most 4 vertices, and for any complete graph if the cost function c is identical to 1. We also prove that the property $OPT(\text{TR}) = OPT(\text{MPR})$ is preserved under taking 1-sums of graphs, implying a common generalization of all the aforementioned results.

Finally, in Section 6, we discuss the more general *asymmetric* VPN problem. In the model of this problem, also proposed by Duffield et al. [2], there is a distinction between a bound $b^- : W \rightarrow \mathbb{R}_+$ for incoming communication and a bound $b^+ : W \rightarrow \mathbb{R}_+$ for outgoing communication. Gupta et al. [5] prove that even the TR problem is NP-hard for the VPN problem with asymmetric communication bounds. However, the TR problem is solvable in polynomial time if $b^-(v) = b^+(v)$ for all $v \in W$. Italiano et al. [7] show that this is true already if $\sum_{v \in W} b^-(v) = \sum_{v \in W} b^+(v)$. Gupta et al. [5] and [6] study approximation algorithms for NP-hard versions of the VPN problem.

For the MPR variant of this asymmetric problem, we give an LP-formulation of a size that is polynomial in the number of vertices of the graph. Polynomial time solvability of the asymmetric MPR VPN problem follows, using the ellipsoid method. This was already observed by Erlebach and Rüegg [3], but their formulation has exponential size.

The challenge remains to prove or disprove that SPR is polynomially solvable on any graph.

2 A linear programming formulation

Let $G = (V, E)$ be a graph, $W \subseteq V$ a set of terminals, $b \in \mathbb{R}_+^W$ a vector of communication upper bounds, and $c \in \mathbb{R}_+^E$ a vector of unit edge costs. We also use the notation b_v for $b(v)$ and similarly c_e for $c(e)$. We denote the set of all paths between vertices i and j in W by \mathcal{P}_{ij} . Let $\mathcal{P} = \cup_{\{i,j\} \subseteq W} \mathcal{P}_{ij}$. We introduce the variable x_p for each $p \in \mathcal{P}$. In the SPR VPN problem we are to select one path for each pair $\{i, j\}$ of distinct nodes in W ; i.e., we are to select values for the x -variables that satisfy $\sum_{p \in \mathcal{P}_{ij}} x_p = 1$, for all $\{i, j\} \subseteq W$ and $x_p \in \{0, 1\}$, for all $p \in \mathcal{P}$.

Once paths have been selected, i.e., values for x_p , $p \in \mathcal{P}$ have been set, the computation of the capacity that has to be reserved on the edges, z_e , $e \in E$, is straightforwardly formulated in the following linear program. Let $\alpha_p^e = 1$ if edge e is on path p , and 0 otherwise.

$$\begin{aligned} z_e = \max \quad & \sum_{\{i,j\} \subseteq W} \sum_{p \in \mathcal{P}_{ij}} \alpha_p^e x_p d_{ij}, \\ \text{s.t.} \quad & \sum_{j \in W} d_{ij} \leq b_i, \quad \forall i \in W, \\ & d_{ij} \geq 0, \quad \forall \{i, j\} \subseteq W, \end{aligned}$$

which, by strong duality, is equal to

$$\begin{aligned} z_e = \min \quad & \sum_{i \in W} b_i y_i^e, \\ \text{s.t.} \quad & y_i^e + y_j^e \geq \sum_{p \in \mathcal{P}_{ij}} \alpha_p^e x_p, \quad \forall \{i, j\} \subseteq W, \\ & y_i^e \geq 0, \quad \forall i \in W. \end{aligned}$$

The MPR problem is to make a feasible choice for the variables x_p such as to minimize total reservation costs $\sum_{e \in E} c_e z_e$, which by the above can be formulated as

$$\begin{aligned} \min \quad & \sum_{e \in E} c_e \sum_{i \in W} b_i y_i^e, \\ \text{s.t.} \quad & y_i^e + y_j^e - \sum_{p \in \mathcal{P}_{ij}} \alpha_p^e x_p \geq 0, \quad \forall \{i, j\} \subseteq W, \forall e \in E, \\ & \sum_{p \in \mathcal{P}_{ij}} x_p = 1, \quad \forall \{i, j\} \subseteq W, \\ & y_i^e \geq 0, \quad \forall i \in W, \forall e \in E, \\ & x_p \geq 0, \quad \forall p \in \mathcal{P}. \end{aligned} \tag{1}$$

The MPR problem, formulated above, is the LP-relaxation of the SPR problem, in which all variables x_p are restricted to be 0 or 1. The dual of MPR is given by

$$\begin{aligned} \max \quad & \sum_{\{i,j\} \subseteq W} \mu_{ij}, \\ \text{s.t.} \quad & \sum_{j \in W} \lambda_{ij}^e \leq c_e b_i, \quad \forall i \in W, \forall e \in E, \\ & \mu_{ij} - \sum_{e \in E} \alpha_p^e \lambda_{ij}^e \leq 0, \quad \forall \{i, j\} \subseteq W, \forall p \in \mathcal{P}_{ij} \\ & \lambda_{ij}^e \geq 0, \quad \forall \{i, j\} \subseteq W, \forall e \in E. \end{aligned} \tag{2}$$

At this point, let us note already that the separation problem over the dual polytope can be solved in polynomial time. There is only a polynomial number of constraints of the first type. For the second set of constraints, suppose we are given $\lambda_{ij}^e \forall \{i, j\} \subseteq W, \forall e \in E$, and $\mu_{ij} \forall \{i, j\} \subseteq W$. Take for any pair $\{i, j\} \subseteq W$ the constraints for all $p \in \mathcal{P}_{ij}$ together. To check if they are satisfied is a matter of computing a shortest path between i and j in W , where each edge $e \in E$ has weight λ_{ij}^e . There are only polynomially many i, j pairs in W .

It follows, using the ellipsoid method (see [4]), that MPR can be solved in polynomial time, which was also proved in [3]. The formulation above, as well as the formulation in [3] is not of polynomial size, however. In Section 6, we show how a linear program of polynomial size for MPR can be formulated using the ideas from this section. The formulation is extended there to a very general asymmetric version of the MPR VPN problem.

3 A reformulation of the conjecture, and some preliminary results

We start with some notation and terminology. Given a connected graph $G = (V, E)$, a set of terminals $W \subseteq V$, a demand function $b : W \rightarrow \mathbb{R}_+$ and a cost function $c : E \rightarrow \mathbb{R}_+$, we will regard b as a function on V rather than on W (with $b(v) = 0$ for $v \notin W$) and simply identify W with the set of vertices $\{v \in V \mid b(v) \neq 0\}$. The triple (G, b, c) is then called an *instance* (of the MPR or SPR or TR problem), and we will refer to a pair (λ, μ) feasible for the linear program (2) as a *feasible dual* or briefly *dual* for this instance. If moreover, $\sum_{\{u,v\}} \mu_{uv} = K$, we will say that (λ, μ) is a dual of *value* K . A *tree solution* for the instance (G, b, c) is a solution to the TR problem with these parameters. So a tree solution is a tree in G spanning the set of terminals $W = \{v \in V \mid b(v) > 0\}$, together with optimal capacity reservations on the edges of the tree.

By weak duality, and by $OPT(\text{MPR}) \leq OPT(\text{TR})$, we have for any feasible solution (λ, μ) to (2) that $\sum_{\{i,j\} \subseteq W} \mu_{ij}$ is at most the cost of any tree solution. It means that the following conjecture is equivalent to the conjecture that $OPT(\text{MPR}) = OPT(\text{TR})$.

Conjecture 3.1 *For any instance (G, b, c) , the cost of an optimal tree solution equals the value of an optimal dual.*

In this paper, we will show that Conjecture 3.1 holds in several special cases, the most important one being the case where G is a circuit, and b and c are arbitrary.

The next paragraph, which is essentially extracted from [5], summarizes how to compute the cost of a given tree solution. Let $G = (V, E)$, $b \in \mathbb{R}_+^V$, $c \in \mathbb{R}_+^E$ be given, and let W be the set of terminals. We will write $b(U)$ for $\sum_{v \in U} b(v)$, $U \subseteq V$. Given a tree $T \subseteq E$ spanning a vertex set $V(T) \supseteq W$ a directed tree can be constructed by directing the edges of T towards the ‘lighter’ side: if L_e and R_e are the components of $T - e$, and if $b(L_e) < b(R_e)$, direct e towards L_e , if $b(L_e) = b(R_e)$, direct e away from some fixed leaf l of the tree. (The italics indicate a correction of what is written in [5].) This directed tree has a unique vertex r of in-degree zero which is a ‘balance-point’ of the tree: every edge in the directed tree is directed away from r . The cost of the tree T is clearly equal to

$$\sum_e \min\{b(L_e), b(R_e)\}c(e). \quad (3)$$

Another expression for the cost of the tree is given in the following proposition from [5]. Here, we denote by $d_G^c(u, v)$ the distance from u to v in a graph G with respect to the length function c .

Proposition 3.2 ([5]) *Let $G = (V, E)$, $b : V \rightarrow \mathbb{R}_+$, $c : E \rightarrow \mathbb{R}_+$ be given. Then the cost of a tree solution T with ‘balance-point’ r equals*

$$\sum_{v \in W} b(v)d_T^c(r, v).$$

Moreover, this cost is bounded from below by $\sum_{v \in W} b(v)d_G^c(r, v)$, and bounded from above by $\sum_{v \in W} b(v)d_T^c(u, v)$ for any $u \in V(T)$.

As a consequence, we have that an optimal tree solution can be found by computing a shortest path tree T_u from every vertex $u \in V$ and taking the one with minimal cost $\sum_{v \in W} b(v)d_{T_u}^c(u, v) = \sum_{v \in W} b(v)d_G^c(u, v)$. Hence TR is solvable in polynomial time.

In the next section, we prove Conjecture 3.1 for circuits, that is we prove Theorem 1.2. We will restrict there to circuits $G = (V, E)$ with $|V|$ even and $b(v) = 1$ for all $v \in V$. To show that this is not an essential restriction, we prove a few preliminary lemmas. Some of the results here will be used in Section 5 as well.

First note that, for a fixed graph G , the optimum values of the various (integer) linear programs are continuous functions in b and c . Hence, for proving the conjecture, we may restrict ourselves to rational vectors b and c .

Lemma 3.3 *Let G be a fixed graph. If for any rational vectors b and c the cost of an optimal tree solution to (G, b, c) equals the value of an optimal dual, then the same is true for any instance (G, b, c) where b and c are real vectors. \square*

The next lemma claims that scaling of b or c is allowed when proving Conjecture 3.1.

Lemma 3.4 *For any $\beta \in \mathbb{R}_+$, the instance $(G, \beta b, c)$ has a feasible dual of value βK if and only if the instance (G, b, c) has a feasible dual of value K . Moreover, $(G, \beta b, c)$ has a tree solution of cost βK if and only if (G, b, c) has a tree solution of cost K . A similar statement holds if c is scaled instead of b .*

Proof. Obvious. Multiply all λ and μ values by β to obtain from a feasible dual of the original instance a feasible dual for the scaled instance. The cost of any tree changes by a factor β in the new situation as well (see for the cost of a tree solution Proposition 3.2). \square

The next lemma claims that edges of zero cost may be “contracted” or “decontracted” when proving Conjecture 3.1. Contraction of $e = \{u', v'\}$, by identifying the two vertices u' and v' with one new vertex w' , transforms G into $G/e = (V \setminus \{u', v'\} \cup \{w'\}, E')$, with

$$E' := \{\{u, v\} \in E \mid \{u, v\} \cap \{u', v'\} = \emptyset\} \cup \{\{w', v\} \mid \{u', v\} \in E, v \neq v'\} \\ \cup \{\{w', v\} \mid \{v', v\} \in E, v \neq u'\}.$$

By the contraction of e in the instance (G, b, c) we mean the instance (G', b', c') , where $G' = G/e$, $b'(v) = b(v)$ for $v \neq u', v'$, and $b'(w') = b(u') + b(v')$, and moreover $c'(\{u, v\}) = c(\{u, v\})$, if $w' \notin \{u, v\} \in E'$, and $c'(\{w', v\}) = c(\{u', v\})$ or $c(\{v', v\})$, or both values occur, in case parallel edges arise (edges in G' can be identified with those in $E \setminus \{e\}$). We will denote this contraction (G', b', c') by $(G, b, c)/e$.

Lemma 3.5 *Let $(G = (V, E), b, c)$ be an instance, where $e = \{x, y\} \in E$ has $c(e) = 0$. Then $(G = (V, E), b, c)$ has a dual of value K if and only if the contraction $(G, b, c)/e$ has a dual of value K . Moreover, (G, b, c) has an optimal tree solution of cost K if and only if $(G, b, c)/e$ has an optimal tree solution of cost K .*

Proof. For the proof of the first statement, consider a feasible dual (λ, μ) for $(G, b, c)/e$. Then the dual $(\hat{\lambda}, \hat{\mu})$ for (G, b, c) , defined by $\hat{\lambda}_{ij}^e := 0$, $\hat{\lambda}_{ij}^f := \lambda_{ij}^f$, $f \neq e$, $\hat{\mu}_{ij} := \mu_{ij}$, is evidently feasible too. Conversely, $\mu := \hat{\mu}$, $\lambda^f := \hat{\lambda}^f$, for $f \neq e$ also maintains feasibility, because each feasible dual $(\hat{\lambda}, \hat{\mu})$ for (G, b, c) has the property that $0 \leq \hat{\lambda}_{ij}^e \leq c(e)b_i = 0$. Solutions (λ, μ) and $(\hat{\lambda}, \hat{\mu})$ have the same value.

One implication of the second statement of the lemma is obvious: if we decontract the edge e , we obtain from a tree solution in $(G, b, c)/e$ of cost K a tree solution in (G, b, c) of the same cost (see Proposition 3.2).

The other implication is again obvious if edge e is in an optimal tree solution for (G, b, c) , or if at most one of its end points is covered by the tree. The remaining case is one in which the tree solution $T = (V', E')$ of cost K does not contain edge e , but $\{u', v'\} \subset V'$, so contraction would lead to a cycle. Let T have balance point r , and let T_r be a shortest path tree rooted at r for the graph $(V', E' \cup \{e\})$, that does contain e . T_r has tree cost $K_r \leq K$. As T is optimal, T_r must also have cost K . Contraction of e in T_r yields a tree solution for $(G, b, c)/e$ of cost K .

Combining these two results shows that $(G, b, c)/e$ and (G, b, c) have optimal tree solutions of the same cost. \square

The last lemma says that vertices with a communication bound of 0 and degree 2 in an instance $(G = (V, E), b, c)$ can be neglected when proving Conjecture 3.1. Suppose $v' \in V$ has degree 2 in V , and $e_1 = \{u, v'\}$, $e_2 = \{v', w\}$ are the two edges incident with v' . Let $e_3 := \{u, w\}$ be a new edge. Then *shortcutting* v' in (G, b, c) results in the instance (G', b', c') , where $G' = (V - v', E \setminus \{e_1, e_2\} \cup \{e_3\})$, $b'(v) = b(v)$, $\forall v \in V - v'$, and $c'(e_3) = c(e_1) + c(e_2)$, $c'(e) = c(e)$ for all $e \in E \setminus \{e_1, e_2\}$.

Lemma 3.6 *Let $(G = (V, E), b, c)$ be an instance, and $v \in V$ a vertex of degree 2 in G with $b(v) = 0$. Denote the instance obtained from (G, b, c) by shortcutting v by (G', b', c') . Then (G', b', c') has a dual of value K , if and only if (G, b, c) has a dual of value K . Moreover, (G', b', c') has a tree solution of cost K , if and only if the same holds for (G, b, c) .*

Proof. For the first statement, consider the dual (λ, μ) for (G, b, c) . Define $(\hat{\lambda}, \hat{\mu})$ by $\hat{\lambda}_{ij}^{e_3} := \lambda_{ij}^{e_1} + \lambda_{ij}^{e_2}$, $\hat{\lambda}_{ij}^e := \lambda_{ij}^e$, for $e \in E \setminus \{e_1, e_2\}$, and $\hat{\mu}_{ij} = \mu_{ij}$, for $\{i, j\} \subseteq W$. For feasible (λ, μ) this yields a feasible $(\hat{\lambda}, \hat{\mu})$. Similarly, for a feasible dual $(\hat{\lambda}, \hat{\mu})$ for (G', b', c') , define a feasible (λ, μ) by taking, in particular, $\lambda_{ij}^{e_1} := \hat{\lambda}_{ij}^{e_3} c(e_1)/c'(e_3)$, and $\lambda_{ij}^{e_2} := \hat{\lambda}_{ij}^{e_3} c(e_2)/c'(e_3)$.

A tree solution for (G', b', c') (not) using edge e_3 can be translated into a tree solution of the same cost for (G, b, c) (not) using both e_1 and e_2 , and vice versa. \square

4 The circuit

In this section, we prove $OPT(\text{TR}) = OPT(\text{SPR}) = OPT(\text{MPR})$ for circuits, by proving Conjecture 3.1 for circuits. We use the lemmas from the previous section to prove that we may restrict ourselves to *even* circuits in which each vertex has communication bound 1 (an even circuit is a circuit with an even number of vertices).

Lemma 4.1 *If Conjecture 3.1 holds for every instance (G, b, c) where G is an even circuit and $b \equiv 1$, then it holds for every instance (G, b, c) where G is a circuit.*

Proof. Consider a general circuit instance $(G = C_n, b, c)$. From Lemma 3.6 we know that without loss of generality $b(v) > 0$ for all nodes v . Lemma 3.3 implies that we may assume that b is rational. Dividing b by $\gcd\{\frac{1}{2}b(v) \mid v \in V\}$ is allowed by Lemma 3.4. So we may assume that each $b(v)$ is a positive even integer. Finally, by Lemma 3.5, any path u, v, w , for v with $b(v) \neq 1$, may be substituted by a path $u, v_1, v_2, \dots, v_N, w$, where $N = b(v)$, setting $b(v_1) = \dots = b(v_N) = 1$, and $c(\{u, v_1\}) = c(\{u, v\})$, $c(\{v_k, v_{k+1}\}) = 0$, $c(\{v_N, w\}) = c(\{v, w\})$. Thus, we arrive at an even circuit with $b \equiv 1$. \square

4.1 Properties of an even circuit with unit bounds

Given an even circuit $G = C_{2n} = (V, E)$ on which all vertices have communication bound 1, we number the vertices starting at some vertex and following the circuit in counterclockwise direction $0, 1, 2, \dots, 2n-1$. All vertex and edge labels are taken modulo $2n$. The edge $\{i-1, i\}$ is denoted by e_i , or by i in case we use it as an index and no confusion with vertices is possible, $i = 0, \dots, 2n-1$. E.g., we will write c_k for the unit cost $c(e_k)$ of edge e_k . On even circuits each edge e_k has an *opposite edge* e_{k+n} .

The cost of the tree solution obtained from the circuit by deleting the edge e_k is denoted by $C(e_k; c)$; we explicitly indicate dependence on c as the unit cost function of the edges, since we will use other unit cost functions later. By (3), we have

$$C(e_k; c) = \sum_{i=1}^{n-1} i(c_{k-i} + c_{k+i}) + nc_{k+n}.$$

Define for each vertex j the *half-sum*,

$$H(j; c) := \sum_{i=j+1}^{j+n} c_i,$$

i.e., the sum of the unit costs of the n edges on the path starting in j and ending with the edge opposite to e_j . This allows us to rewrite $C(e_k; c)$ in what turns out to be a more convenient way.

$$C(e_k; c) = \sum_{j=k}^{k+n-1} H(j; c). \quad (4)$$

From this expression the following equations are easily derived.

$$C(e_k; c) - C(e_{k+1}; c) = H(k; c) - H(k+n; c), \quad \forall k = 0, 1, \dots, 2n-1; \quad (5)$$

$$H(j; c) - H(j-1; c) = c_{j+n} - c_j, \quad \forall j = 0, 1, \dots, 2n-1. \quad (6)$$

Now suppose that the tree obtained by deleting edge e_k has minimum cost among all spanning trees of the circuit, in other words $C(e_k; c) = \min_{e \in E} C(e; c)$. Then, using (5), we have

$$H(k; c) - H(k+n; c) = C(e_k; c) - C(e_{k+1}; c) \leq 0, \quad \text{and} \quad (7)$$

$$H(k-1; c) - H(k+n-1; c) = C(e_{k-1}; c) - C(e_k; c) \geq 0. \quad (8)$$

Subtracting (8) from (7) and applying (6) yields

$$\begin{aligned} 2c_{k+n} - 2c_k &= H(k; c) - H(k+n; c) - H(k-1; c) + H(k+n-1; c) \\ &= 2C(e_k; c) - C(e_{k+1}; c) - C(e_{k-1}; c) \leq 0. \end{aligned} \quad (9)$$

Consequently, if $c_k = 0$ then $c_{k+n} = 0$ and $2C(e_k; c) - C(e_{k+1}; c) - C(e_{k-1}; c) = 0$. Hence,

$$c_k = 0 \wedge C(e_k; c) = \min_{e \in E} C(e; c) \Rightarrow C(e_{k+1}; c) = C(e_{k-1}; c) = C(e_k; c) = \min_{e \in E} C(e; c), \quad (10)$$

i.e., if e_k is a minimizer of $C(e; c)$ and $c_k = 0$, then the adjacent edges e_{k-1} and e_{k+1} are also minimizers of $C(e; c)$.

In the case of an even circuit with $b \equiv 1$, the constraints of the dual linear program (2) reduce to the following. (In a circuit there are only two possible paths between any pair of vertices.)

$$\begin{aligned} \sum_{j \in V \setminus \{i\}} \lambda_{ij}^e &\leq c_e, \quad \forall i \in V, \forall e \in E, \\ \mu_{ij} &\leq \sum_{l=i+1}^j \lambda_{ij}^{e_l}, \quad \forall \{i, j\} \subseteq V, \\ \mu_{ij} &\leq \sum_{l=j+1}^i \lambda_{ij}^{e_l}, \quad \forall \{i, j\} \subseteq V, \\ \lambda_{ij}^e &\geq 0, \quad \forall \{i, j\} \subseteq V, \forall e \in E. \end{aligned} \quad (11)$$

4.2 Proof for the even circuit with unit bounds

Given a cost function $c : E \rightarrow \mathbb{R}_+$, we call the set of edges with non-zero unit cost the *support* of c , i.e. $\text{supp}(c) = \{e \in E \mid c_e > 0\}$. The following lemma is crucial to the main result.

Lemma 4.2 *Let $G = (V, E) = C_{2n}$ be an even circuit, and $b \equiv 1$. Let F be a non-empty subset of E . Then there exist a nonnegative cost function $\hat{c} : E \rightarrow \mathbb{R}_+$, not identical to 0, with $\text{supp}(\hat{c}) \subseteq F$, and a constant K , such that for all $f \in F$: $K = C(f; \hat{c}) = \min_{e \in E} C(e; \hat{c})$. Moreover, there is a dual solution $(\hat{\lambda}, \hat{\mu})$ for the problem with cost function \hat{c} , with value K .*

Proof. The proof is by induction on $|F|$. The theorem is clearly true if $|F| = 1$. For suppose $F = \{e_k\}$. Then we can take $\hat{c}_k = 1$ and $\hat{c}_i = 0$ for other i . Clearly, $\min_{e \in E} C(e; \hat{c}) = C(e_k; \hat{c}) = 0$. A feasible dual solution with objective value 0 is $\hat{\lambda}_{ij}^e = 0$, $\hat{\mu}_{ij} = 0$, $\forall e \in E$, $\forall \{i, j\} \subseteq V$.

Suppose now that $|F| > 1$. We distinguish three cases.

Case 1. There exists a k such that $e_k \in F$ and $e_{k+n} \in F$. In this case we call the edge set F *singular*. See Figure 1 for an example.

Consider the cost function $\hat{c} : E \rightarrow \mathbb{R}_+$ defined by $\hat{c}_k = \hat{c}_{k+n} = 1$ and $\hat{c}_i = 0$ otherwise. It satisfies $C(e; \hat{c}) = n$, for all $e \in E$. The following dual solution is feasible with respect to this cost function \hat{c} (see (11)) and has objective value $\sum_{i < j} \hat{\mu}_{ij} = n$:

$$\begin{aligned} \hat{\lambda}_{i, i+n}^{e_k} &:= 1, \quad \forall i = 0, 1, \dots, n-1, \\ \hat{\lambda}_{i, i+n}^{e_{k+n}} &:= 1, \quad \forall i = 0, 1, \dots, n-1, \\ \hat{\lambda}_{ij}^e &:= 0, \quad \text{otherwise,} \\ \hat{\mu}_{i, i+n} &:= 1, \quad i = 0, \dots, n-1 \\ \hat{\mu}_{ij} &:= 0, \quad \text{otherwise.} \end{aligned}$$




Figure 1: Example of a singular subset, denoted by bold lines.




Figure 2: Contracting the circuit to a smaller one.

Case 2. F is not singular, and there exist k and m , $k < m < k+n$, such that $e_k \in F$, $e_m \in F$, $e_l \notin F$ for all $k < l < m$, and $e_l \notin F$ for all $k+n \leq l \leq m+n$. See Figure 2 for an example, with $e = e_k$ and $f = e_m$.

We ‘contract’ the edges e_k, \dots, e_m to a new edge $e' = (k-1, m)$, and the edges e_{k+n}, \dots, e_{m+n} to an edge $\bar{e}' = (k-1+n, m+n)$, to arrive at a new even cycle (V', E') , with $|V'| = |E'| = 2(n-m-k)$. We maintain the vertex labels $V' = \{m, m+1, \dots, k-1+n, m+n, \dots, k-1\}$. The new edge set is $E' = \{e_{m+1}, \dots, e_{k-1+n}, e', e_{m+n+1}, \dots, e_{k-1}, e'\}$.

Note that edges that were opposite before contraction remain opposite after contraction. Also, the new edges e' and \bar{e}' are opposite.

Consider the subset of edges $F' = F \setminus \{e_k, e_m\} \cup \{e'\}$. As $0 < |F'| < |F|$, we can apply the induction hypothesis. Thus, there exist a cost function $c' : E' \rightarrow \mathbb{R}_+$, not identical to 0, with $\text{supp}(c') \subseteq F'$, a constant $K' = \min_{e \in E'} C(e; c') = C(f; c'), \forall f \in F'$, and a dual solution (λ, μ) , with value K' .

Since e' is a minimizer of $C(e; c')$, using (9) and the fact that $c'(\bar{e}') = 0$, we have

$$c'(e') = \frac{1}{2}(H'(m+n; c') - H'(m; c')) + \frac{1}{2}(H'(k-1; c') - H'(k+n-1; c')),$$

where $H'(j, c')$ is the half-sum for vertex j on the smaller, contracted circuit, with its corresponding cost function c' . Now we define the cost $\hat{c} : E \rightarrow \mathbb{R}_+$ from c' as follows:

$$\begin{aligned}\hat{c}_m &:= \frac{1}{2}(H'(m+n; c') - H'(m; c')), \\ \hat{c}_k &:= \frac{1}{2}(H'(k-1; c') - H'(k+n-1; c')), \\ \hat{c}_i &:= c'_i, & \forall \{i-1, i\} \in F, i \neq k, m, \\ \hat{c}_i &:= 0, & \forall \{i-1, i\} \in E \setminus F.\end{aligned}$$

It follows from (7) and (8) that $\hat{c}_m \geq 0$ and $\hat{c}_k \geq 0$. Note that $\hat{c}_k + \hat{c}_m = c'(e')$ and

$$\begin{aligned}\hat{c}_k + H'(k+n-1; c') &= \hat{c}_m + H'(m; c') = \frac{1}{2} \sum_{e \in E'} c'_e = \frac{1}{2} \sum_{f \in F'} c'_f \\ &= \frac{1}{2} \sum_{f \in F} \hat{c}(f) = \frac{1}{2} \sum_{e \in E} \hat{c}(e).\end{aligned}$$

Hence, for the half-sums in the larger circuit, we have

$$\begin{aligned}H(j; \hat{c}) &= \hat{c}_m + H'(m; c') = \frac{1}{2} \sum_{e \in E} \hat{c}(e), & j = k, \dots, m-1, \\ H(j; \hat{c}) &= \hat{c}_k + H'(k+n-1; c') = \frac{1}{2} \sum_{e \in E} \hat{c}(e), & j = k+n, \dots, m-1+n, \\ H(j; \hat{c}) &= H'(j; c'), & \text{otherwise.}\end{aligned}$$

Using this in (4) yields $C(e_k; \hat{c}) = C(e_m; \hat{c})$, and $C(e_k; \hat{c}) - C(f; \hat{c}) = C(e'; c') - C(f; c')$, for all $f \in F'$. Hence, $C(f; \hat{c}) = \min_{e \in E} C(e; \hat{c}) = K, \forall f \in F$, with $K = K' + \frac{1}{2}(m-k) \sum_{e \in E} \hat{c}(e)$.

To define the dual feasible solution with value K we represent V as $V = V' \cup V_{e'} \cup V_{\bar{e}'}$, with $V_{e'} = \{k, k+1, \dots, m-1\}$ and $V_{\bar{e}'} = \{k+n, k+n+1, \dots, m+n-1\}$. Again we construct the dual solution $(\hat{\lambda}, \hat{\mu})$ for (V, E, \hat{c}) from the dual (λ, μ) associated with (V', E', c') :

$$\begin{aligned}\hat{\lambda}_{ij}^e &:= \lambda_{ij}^e, & \text{for } e \in F \setminus \{e_k, e_m\}, \{i, j\} \subseteq V', \\ \hat{\lambda}_{ij}^{e_k} &:= \frac{\lambda_{ij}^{e'} \hat{c}_k}{c'(e')}, & \text{for } \{i, j\} \subseteq V', \\ \hat{\lambda}_{ij}^{e_m} &:= \frac{\lambda_{ij}^{e'} \hat{c}_m}{c'(e')}, & \text{for } \{i, j\} \subseteq V', \\ \hat{\lambda}_{ij}^e &:= \hat{c}(e), & \text{for } e \in E, i \in V_{e'}, j = i+n, \\ \hat{\lambda}_{ij}^e &:= 0, & \text{otherwise,} \\ \hat{\mu}_{ij} &:= \mu_{ij}, & \text{for } \{i, j\} \subseteq V', \\ \hat{\mu}_{ij} &:= \frac{1}{2} \sum_{e \in E} \hat{c}(e), & \text{for } i \in V_{e'}, j = i+n, \\ \hat{\mu}_{ij} &:= 0, & \text{otherwise.}\end{aligned}$$

Case 2 is settled by an easy verification that $(\hat{\lambda}, \hat{\mu})$ satisfies the dual constraints (11), and has value

$$\sum_{i, j \in V, i < j} \hat{\mu}_{ij} = \sum_{i, j \in V', i < j} \mu_{ij} + \sum_{i \in V_{e'}} \hat{\mu}_{i, i+n} = K' + \frac{1}{2}(m-k) \sum_{e \in E} \hat{c}(e) = K.$$

Case 3. F is not singular, and Case 2 does not apply. It means that between each pair of consecutive edges from F , e and f say, there is exactly one opposite edge \bar{g} of some edge $g \in F$. See Figure 3 for an example.

Hence F consists of an odd number of edges, $2k+1$ say, such that these edges and their opposites are perfectly alternating. As a consequence, an edge $e \in F$ and its opposite \bar{e} , split the other edges of F into two groups of size k . Therefore, we call this a *uniform* configuration. Let $F = \{f_0, f_1, \dots, f_{2k}\}$, with $f_i = (m_i - 1, m_i)$, and denote the number of nodes between




Figure 3: Example of an edge set with alternating edges and opposites.

edge f_i and f_{i+1} by ν_i , i.e., $\nu_i = m_{i+1} - m_i$ modulo $2n$. We have that $\sum_{i=0}^{2k} \nu_i = 2n$. Please note that from now on subscripts i for ν, f etcetera are taken modulo $2k+1$. Define

$$\tau_i := \sum_{j=0}^k \nu_{i+j} - \sum_{j=k+1}^{2k} \nu_{i+j} = 2(m_{i+k+1} - m_i) \pmod{2n}, \quad i = 0, 1, \dots, 2k,$$

By uniformity, $\sum_{j=0}^k \nu_{i+j} > \sum_{j=k+1}^{2k} \nu_{i+j}$, whence $\tau_i > 0$, $i = 0, \dots, 2k$. Notice that $\frac{1}{2}\tau_i$ equals the number of nodes between \bar{f}_i and f_{i+k+1} (or equivalently between f_i and \bar{f}_{i+k+1}). Note that, by definition, $\nu_{i+k} = \frac{1}{2}(\tau_i + \tau_{i+k})$. This justifies the equality in the next definition of the cost function:

$$\hat{c}(f_i) := \frac{\nu_{i+k}}{\tau_i \tau_{i+k}} = \frac{1}{2} \left(\frac{1}{\tau_i} + \frac{1}{\tau_{i+k}} \right), \quad i = 0, 1, \dots, 2k; \quad \hat{c}(e) := 0, \quad e \notin F.$$

Observe the following identity, for $i = 0, 1, \dots, 2k$,

$$\begin{aligned} & \sum_{j=i}^{i+k} \hat{c}(f_j) - \sum_{j=i+k+1}^{i+2k} \hat{c}(f_j) = \sum_{j=i}^{i+k-1} (\hat{c}(f_j) - \hat{c}(f_{j+k+1})) + \hat{c}(f_{i+k}) \\ &= \frac{1}{2} \left(\sum_{j=i}^{i+k-1} \left(\left(\frac{1}{\tau_j} + \frac{1}{\tau_{j+k}} \right) - \left(\frac{1}{\tau_{j+k+1}} + \frac{1}{\tau_j} \right) \right) + \left(\frac{1}{\tau_{i+k}} + \frac{1}{\tau_{i+2k}} \right) \right) \\ &= \frac{1}{\tau_{i+k}} > 0 \end{aligned} \tag{12}$$

As a consequence, we have that consecutive edges from F yield the same cost, as

$$\begin{aligned} & C(f_i; \hat{c}) - C(f_{i+1}; \hat{c}) = \\ & C(f_i; \hat{c}) - C(\bar{f}_{i+k+1}; \hat{c}) + C(\bar{f}_{i+k+1}; \hat{c}) - C(f_{i+1}; \hat{c}) = \\ & \frac{1}{2} \tau_i (\hat{c}(f_{i+1}) + \dots + \hat{c}(f_{i+k}) - (\hat{c}(f_{i+k+1}) + \dots + \hat{c}(f_{i+2k+1}))) + \\ & \frac{1}{2} \tau_{i+k+1} (\hat{c}(f_{i+1}) + \dots + \hat{c}(f_{i+k+1}) - (\hat{c}(f_{i+k+2}) + \dots + \hat{c}(f_{i+2k+1}))) = \\ & \frac{1}{2} \tau_i \left(-\frac{1}{\tau_i} \right) + \frac{1}{2} \tau_{i+k+1} \frac{1}{\tau_{i+k+1}} = 0 \end{aligned}$$

Using (10), we conclude that $C(f_i; \hat{c}) = \min_{e \in E} C(e; \hat{c})$, for all $f_i \in F$.

We introduce the following entities to facilitate the exposition of the dual solution that we propose:

$$\Lambda_i^e := \begin{cases} \frac{1}{2} \frac{\tau_i}{\nu_i \nu_{i+k}} \hat{c}(e) \left(1 + \frac{1}{\tau_i} \frac{1}{\sum_j \hat{c}(f_j)} \right), & \text{if } f_i < e \leq f_{i+k}, \\ \frac{1}{2} \frac{\tau_i}{\nu_i \nu_{i+k}} \hat{c}(e) \left(1 - \frac{1}{\tau_i} \frac{1}{\sum_j \hat{c}(f_j)} \right), & \text{if } f_{i+k} < e \leq f_{i+2k+1}. \end{cases}$$

Now we define the dual solution as

$$\hat{\lambda}_{st}^e := \begin{cases} \Lambda_i^e, & \forall e \in E \quad \forall i = 0, 1, \dots, 2k \quad \forall s, t : m_i \leq s < m_{i+1}, m_{i+k} \leq t < m_{i+k+1}, \\ 0, & \text{otherwise,} \end{cases}$$

$$\hat{\mu}_{st} := \sum_{e: s < e \leq t} \hat{\lambda}_{st}^e, \quad \forall \{s, t\} \subseteq V.$$

To verify feasibility of this solution, first notice that from (12) we have

$$0 < \frac{1}{\tau_{i+k}} \frac{1}{\sum_j \hat{c}(f_j)} \leq \frac{\sum_{j=i}^{i+k} \hat{c}(f_j)}{\sum_j \hat{c}(f_j)} \leq 1, \quad \forall i,$$

whence $\lambda_{uv}^e \geq 0$ is satisfied for all $u, v \in V$ and all $e \in E$. Simple algebraic computations show that, for all $e \in E$, for all $i = 0, 1, \dots, 2k$, and for all s with $m_i \leq s < m_{i+1}$,

$$\sum_{t=m_{i+k}}^{m_{i+k+1}-1} \hat{\lambda}_{st}^e + \sum_{t=m_{i+k+1}}^{m_{i+k+2}-1} \hat{\lambda}_{st}^e = \nu_{i+k} \Lambda_i^e + \nu_{i+k+1} \Lambda_{i+k+1}^e \leq \hat{c}(e)$$

Actually, the inequality is tight except for $e = f_{i+k+1}$.

To show that $(\hat{\lambda}, \hat{\mu})$ satisfies the second and third type of dual constraints in (11), it suffices to show that the sum of $\hat{\lambda}$ -values over edges along the $s - t$ -path is the same as the sum of $\hat{\lambda}$ -values over edges along the $t - s$ -path; i.e., $\sum_{e: s < e \leq t} \hat{\lambda}_{st}^e = \sum_{e: t < e \leq s} \hat{\lambda}_{st}^e$, for all $s, t \in V$. Clearly we need to show this only for pairs $\{s, t\}$ with $m_i \leq s < m_{i+1}, m_{i+k} \leq t < m_{i+k+1}$ for some i . The claim follows straightforwardly from the fact that the equality

$$(\hat{c}(f_{i+1}) + \dots + \hat{c}(f_{i+k})) \left(1 + \frac{1}{\tau_i} \frac{1}{\sum_f \hat{c}(f)}\right) = (\hat{c}(f_{i+k+1}) + \dots + \hat{c}(f_{i+2k+1})) \left(1 - \frac{1}{\tau_i} \frac{1}{\sum_f \hat{c}(f)}\right)$$

is equivalent to the equality

$$\frac{1}{\tau_i} \frac{1}{\sum_f \hat{c}(f)} (\hat{c}(f_{i+1}) + \dots + \hat{c}(f_{i+k}) + \hat{c}(f_{i+k+1}) + \dots + \hat{c}(f_{i+2k+1})) = (\hat{c}(f_{i+k+1}) + \dots + \hat{c}(f_{i+2k+1})) - (\hat{c}(f_{i+1}) + \dots + \hat{c}(f_{i+k})),$$

which is evident, since the right hand side equals $\frac{1}{\tau_i}$ by (12). This completes the feasibility check.

It remains to verify that $\sum_{\{s,t\} \subseteq V} \hat{\mu}_{st} = C(f_0; \hat{c})$. We use that $C(f_0; \hat{c}) = C(f_i; \hat{c})$ for all $i = 0, 1, \dots, 2k$ implies that $C(f_0; \hat{c}) = \frac{1}{2k+1} \sum_{i=0}^{2k} C(f_i; \hat{c})$. The half-sums that constitute the cost $C(f_i; \hat{c})$ start in nodes $j = m_i, \dots, m_{i+n}-1$ (4). These starting nodes can be subdivided into $2k+1$ subsets: for $j = i, \dots, i+k$ the half-sums starting in the $\frac{1}{2}\tau_j$ nodes between f_j and \bar{f}_{j+k+1} all have value $\sum_{t=j+1}^{j+k} \hat{c}(f_t)$, whereas for $j = i, \dots, i+k-1$ the half-sums starting in the $\frac{1}{2}\tau_{j+k+1}$ nodes between \bar{f}_{j+k+1} and f_{j+1} all have value $\sum_{t=j+1}^{j+k+1} \hat{c}(f_t)$. Altogether, the total cost is

$$\begin{aligned} \frac{1}{2k+1} \sum_{i=0}^{2k} C(f_i; \hat{c}) &= \frac{1}{2k+1} \sum_{i=0}^{2k} \left[\sum_{j=i}^{i+k-1} \left(\frac{1}{2}\tau_j \sum_{t=j+1}^{j+k} \hat{c}(f_t) + \frac{1}{2}\tau_{j+k+1} \sum_{t=j+1}^{j+k+1} \hat{c}(f_t) \right) \right. \\ &\quad \left. + \frac{1}{2}\tau_{i+k} \sum_{t=i+k+1}^{i+2k} \hat{c}(f_t) \right] \\ &= \sum_{i=0}^{2k} \left[\frac{1}{2}\tau_i \frac{k+1}{2k+1} \sum_{j=i+1}^{i+k} \hat{c}(f_j) + \frac{1}{2}\tau_{i+k+1} \frac{k}{2k+1} \sum_{j=i+1}^{i+k+1} \hat{c}(f_j) \right] \quad (13) \end{aligned}$$

In turn, the value of the dual solution can be rewritten as

$$\begin{aligned}
\sum_{\{s,t\} \subset V} \hat{\mu}_{st} &= \sum_{i=0}^{2k} \nu_i \nu_{i+k} \sum_{e=f_{i+1}, \dots, f_{i+k}} \Lambda_i^e \\
&= \sum_{i=0}^{2k} \frac{1}{2} \tau_i \left[\sum_{j=i+1}^{i+k} \hat{c}(f_j) + \frac{1}{\tau_i} \frac{\sum_{j=i+1}^{i+k} \hat{c}(f_j)}{\sum_f \hat{c}(f)} \right] \\
&= \sum_{i=0}^{2k} \frac{1}{2} \tau_i \left[\sum_{j=i+1}^{i+k} \hat{c}(f_j) + \frac{1}{\tau_i} \frac{k}{2k+1} \right] \\
&= \sum_{i=0}^{2k} \frac{1}{2} \tau_i \left[\sum_{j=i+1}^{i+k} \hat{c}(f_j) + \left(\sum_{j=i+k+1}^{i+2k+1} \hat{c}(f_j) - \sum_{j=i=1}^{i+k} \hat{c}(f_j) \right) \frac{k}{2k+1} \right] \\
&= \sum_{i=0}^{2k} \frac{1}{2} \tau_i \left[\frac{k+1}{2k+1} \sum_{j=i+1}^{i+k} \hat{c}(f_j) + \frac{k}{2k+1} \sum_{j=i+k+1}^{i+2k+1} \hat{c}(f_j) \right],
\end{aligned}$$

which equals (13). This settles the proof of Case 3, and thereby the proof of the lemma. \square

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

Theorem 4.3 *Given an even circuit $G = (V, E)$ with $b(i) = 1, \forall i \in V$, and a nonnegative cost function $c : E \rightarrow \mathbb{R}_+$, the cost of an optimal tree solution equals the value of an optimal dual.*

Proof. The proof is by induction on $|\text{supp}(c)|$. The theorem is clearly true if $|\text{supp}(c)| = 1$, when deleting the only edge with positive unit cost yields a tree solution with total cost 0. Setting all dual variables to 0 is feasible and yields value 0.

Now suppose $|\text{supp}(c)| > 1$. Lemma 4.2 (applied to $F := \text{supp}(c)$) tells us that there exist a non-negative non-zero cost function \hat{c} , such that $C(f; \hat{c}) = \min_{e \in E} C(e; \hat{c}), \forall f \in \text{supp}(c)$, and a dual solution $(\hat{\lambda}, \hat{\mu})$ w.r.t. \hat{c} with the same objective value. Define cost vector c' as $c' := c - \sigma \hat{c}$, where σ is a scalar chosen such that c' is non-negative and at least one $f \in \text{supp}(c)$ has $c'(f) = 0$. Such a scalar exists, since $\text{supp}(\hat{c}) \subseteq \text{supp}(c)$ and $\text{supp}(\hat{c}) \neq \emptyset$.

Let $C(e^*; c') = \min_{e \in E} C(e; c')$. By (10), we can assume that $e^* \in \text{supp}(c') \subset \text{supp}(c)$. Since $|\text{supp}(c')| < |\text{supp}(c)|$, the induction hypothesis may be applied to c' , giving a feasible dual (λ', μ') w.r.t. c' of value $\sum \mu'_{vw} = C(e^*; c')$. The solution $(\lambda, \mu) := (\lambda', \mu') + \sigma(\hat{\lambda}, \hat{\mu})$ is feasible w.r.t. c , as $c = c' + \sigma \hat{c}$. Its value is equal to $C(e^*; c') + \sigma C(e^*; \hat{c}) = \min_{e \in E} C(e; c)$. \square

This theorem together with Lemma 4.1 implies our main result, Theorem 1.2.

5 Other cases where the conjecture holds

In this section, we prove the other results announced in the introduction. We start with the observation that Conjecture 3.1 holds for trees, since it is equivalent to the statement that $OPT(\text{TR}) = OPT(\text{MPR})$, which is obviously true for trees. In fact, it is possible to construct, for any instance (G, b, c) where G is a tree, an explicit dual (λ, μ) with value equal to the cost of the tree.

Indeed, denote the unique path in the tree G between two distinct vertices $i, j \in V$ by P_{ij} . For any edge e of the tree, denote the two components of $G - e$ by L_e and R_e . Note that the cost of the tree G is given by $K = \sum_e \min\{b(L_e), b(R_e)\}c(e)$. Define the dual as follows:

$$\begin{aligned}
\lambda_{ij}^e &:= \frac{b(i)b(j)}{b(L_e)b(R_e)} \min\{b(L_e), b(R_e)\}c(e), & \text{if } |\{i, j\} \cap L_e| = 1, \\
\lambda_{ij}^e &:= 0, & \text{otherwise,} \\
\mu_{ij} &:= \sum_{e \in P_{ij}} \lambda_{ij}^e, & \text{for } \{i, j\} \subseteq V.
\end{aligned}$$

Then λ, μ is a feasible dual: for $i \in L_e$ we have

$$\begin{aligned} \sum_{j \neq i} \lambda_{ij}^e &= \sum_{j \in R_e} \frac{b(i)b(j)}{b(L_e)b(R_e)} \min\{b(L_e), b(R_e)\}c(e) \\ &= \frac{b(i)}{b(L_e)} \min\{b(L_e), b(R_e)\}c(e) \leq c(e)b(i), \end{aligned}$$

and similarly for $i \in R_e$. Since there is only one path between any two vertices i and j , the constraint for μ holds by definition. The value of this dual is, as required,

$$\sum_{i,j} \mu_{ij} = \sum_e \sum_{i,j:e \in P_{ij}} \lambda_{ij}^e = \sum_e \min\{b(L_e), b(R_e)\}c(e) = K.$$

As we gather from the next lemma, for proving Conjecture 3.1 we may assume that the graph G is complete, and that the cost function c is a metric (satisfies the triangle inequality), i.e. $c(uw) \leq c(uv) + c(vw)$ for any three vertices u, v , and w . From now on, we use labels uv or vu to denote the edge $\{u, v\}$. The ‘metric’-assumption will be used frequently in the remainder.

Lemma 5.1 *Let $G = (V, E), b \in \mathbb{R}_+^V, c \in \mathbb{R}_+^E$ be given. Let $H = (V, F)$ be the complete graph on V . Define $c'(uv)$ for $\{u, v\} \in F$ as the length of a shortest path between u and v with respect to the length function c (so that the new cost function $c' \in \mathbb{R}_+^F$ satisfies the triangle inequality). If an optimal tree solution for the instance (H, b, c') has cost K equal to the value of an optimal dual, then K is also the cost of an optimal tree solution and the value of an optimal dual for the instance (G, b, c) .*

Proof. A dual (λ, μ) which is feasible for c' and F is also feasible for c and E , since $c' \leq c$ on E . Moreover, for any optimal tree solution in F of cost K with respect to c' (which may be taken equal to a star S centered in its ‘balance-point’ s) there exists a tree solution in E of cost at most K with respect to c : replace the star with center s by a shortest path tree T rooted at s , then $\sum_v b(v)d_T^c(s, v) = \sum_v b(v)d_S^{c'}(s, v) = K$, and the result follows by Proposition 3.2. \square

The next lemma says that if one of the vertices is “dominant” (in the sense that its b -value is greater than the sum of the b -values of the other vertices), then Conjecture 3.1 holds.

Lemma 5.2 *Let (G, b, c) be an instance such that $b(s) \geq \sum_{v \neq s} b(v)$ for some $s \in W$. Then the shortest path tree T rooted at s is an optimal tree solution for this instance. Moreover, the cost $\sum_v b(v)d_T^c(s, v)$ of this tree solution equals the value of an optimal dual.*

Proof. By Lemma 5.1, we may assume that G is complete, and that c satisfies the triangle inequality. Then the optimal tree solution can be taken equal to a star. But the cost of the star centered at s is at most the cost of any other star. Indeed, by Proposition 3.2, the star centered at $x \neq s$ has cost

$$\begin{aligned} \sum_{v \neq x} b(v)c(xv) &= b(s)c(xs) + \sum_{x \neq v \neq s} b(v)c(xv) \\ &\geq \sum_{v \neq s} b(v)c(xs) + \sum_{x \neq v \neq s} b(v)c(xv) \text{ (assumption)} \\ &= b(x)c(xs) + \sum_{x \neq v \neq s} b(v)(c(xs) + c(xv)) \\ &\geq b(x)c(xs) + \sum_{x \neq v \neq s} b(v)c(sv) \text{ (triangle inequality)} \\ &= \sum_{v \neq s} b(v)c(sv). \end{aligned}$$

Therefore, the cost of the optimal tree solution is $\sum_{v \neq s} b(v)c(sv)$.

Now, define λ and μ as follows:

$$\begin{aligned}\lambda_{sv}^e &:= c(e)b(v), & \text{for } v \neq s \text{ and } e \in E, \\ \lambda_{uv}^e &:= 0, & \text{for } u \neq s \neq v \text{ and } e \in E, \\ \mu_{sv} &:= c(sv)b(v), & \text{for } v \neq s, \\ \mu_{uv} &:= 0, & \text{for } u \neq s \neq v.\end{aligned}$$

Then it is not hard to check (using the assumption and the triangle inequality) that (λ, μ) is feasible for the dual of MPR. Moreover, the value $\sum_{uv \in F} \mu_{uv} = \sum_{v \neq s} c(sv)b(v)$ of this solution equals the cost of the optimal tree solution. \square

The next lemma shows that the property that Conjecture 3.1 holds is preserved under taking 1-sums. A *1-sum* of two graphs is the graph obtained by identifying a vertex of one graph with a vertex of the other graph. More precisely, let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be disjoint graphs, take any $v_1 \in V_1$ and $v_2 \in V_2$ and identify them, creating a vertex z , which is then the only vertex common to V_1 and V_2 , i.e., $V_1 \cap V_2 = \{z\}$. The *1-sum of G_1 and G_2 in z* is then the graph $G = (V_1 \cup V_2, E_1 \cup E_2)$.

Lemma 5.3 *Let $G = (V, E)$ be the 1-sum of $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ in a vertex z . Let $b \in \mathbb{R}_+^V$, and $c \in \mathbb{R}_+^E$ be given. Then the cost of an optimal tree solution equals the value of an optimal dual for the instance (G, b, c) if the same holds for every instance of the form $(G_1, b_1, c|_{E_1})$ (where $b_1 \in \mathbb{R}_+^{V_1}$ is arbitrary) and every instance of the form $(G_2, b_2, c|_{E_2})$ (where $b_2 \in \mathbb{R}_+^{V_2}$ is arbitrary). Here, $c|_{E_i}$ denotes the restriction of c to E_i , $i = 1, 2$.*

Proof. Define $B_i := b(V_i)$ for $i = 1, 2$. Then without loss of generality, $B_1 \leq B_2$. If we look at the instance (G_1, b_1, c_1) , where c_1 is the restriction of c to E_1 , and $b_1(z) = B_2$, $b_1(v) = b(v)$ for all other $v \in V_1$, then we see that z is a “dominant” vertex for that instance, in the sense of Lemma 5.2. This means that the shortest path tree T_1 rooted at z is a tree solution of minimum cost $K_1 := \sum_{v \in V_1} b_1(v)d_{T_1}^{c_1}(z, v)$ for this instance. If we define the instance (G_2, b_2, c_2) in a similar way, there is a tree solution T_2 of minimum cost K_2 for this instance, which is a shortest path tree rooted at some $r \in V_2$ (not necessarily $r = z$ this time). Now $T := T_1 \cup T_2$ is a tree solution of (G, b, c) of cost at most $K_1 + K_2$: if we compute the cost of this tree from the root r , we get

$$\begin{aligned}\sum_{v \in V} b(v)d_T^c(r, v) &= \sum_{v \in V_2 - z} b_2(v)d_{T_2}^{c_2}(r, v) + b(z)d_{T_2}^{c_2}(r, z) + \sum_{v \in V_1 - z} b_1(v)(d_{T_2}^{c_2}(r, z) + d_{T_1}^{c_1}(z, v)) \\ &= \sum_{v \in V_2 - z} b_2(v)d_{T_2}^{c_2}(r, v) + B_1 d_{T_2}^{c_2}(r, z) + \sum_{v \in V_1 - z} b_1(v)d_{T_1}^{c_1}(z, v) \\ &= K_2 + K_1.\end{aligned}$$

Next, we show that a feasible dual for (G, b, c) with value $K_1 + K_2$ (as defined above) exists. We use the fact that a feasible dual λ^i, μ^i of value K_i exists for (G_i, b_i, c_i) , $i = 1, 2$. For convenience we write $\lambda_{uv}^i(e)$ instead of $(\lambda^i)_{uv}^e$ in the remainder of this proof. Define:

$$\begin{aligned}
\lambda_{st}^e &:= \frac{b(t)}{B_2} \lambda_{sz}^1(e), & \text{for } s \in V_1 - z, t \in V_2, \text{ and } e \in E_1, \\
\lambda_{st}^e &:= \frac{b(s)}{B_1} \lambda_{zt}^2(e), & \text{for } s \in V_1, t \in V_2 - z, \text{ and } e \in E_2, \\
\lambda_{st}^e &:= \lambda_{st}^1(e), & \text{for } s, t \in V_1 - z, e \in E_1, \\
\lambda_{st}^e &:= \lambda_{st}^2(e), & \text{for } s, t \in V_2 - z, e \in E_2, \\
\lambda_{st}^e &:= 0, & \text{otherwise,} \\
\mu_{st} &:= \mu_{st}^1, & \text{for } s \in V_1 - z, t \in V_1 - z, \\
\mu_{st} &:= \mu_{st}^2, & \text{for } s \in V_2 - z, t \in V_2 - z, \\
\mu_{st} &:= \frac{b(t)}{B_2} \mu_{sz}^1 + \frac{b(s)}{B_1} \mu_{zt}^2, & \text{for } s \in V_1, t \in V_2 \text{ (where } \mu_{zz}^1 := 0 =: \mu_{zz}^2).
\end{aligned}$$

Then for $s \in V_1 - z, e \in E_1$:

$$\sum_{t \neq s} \lambda_{st}^e = \sum_{t \in V_1 - z - s} \lambda_{st}^1(e) + \sum_{t \in V_2} \frac{b(t)}{B_2} \lambda_{sz}^1(e) = \sum_{t \in V_1 - s} \lambda_{st}^1(e) \leq c_1(e) b_1(s),$$

and similarly for $s \in V_2 - z, e \in E_2$. For $s \in V_1, e \in E_2$ we have

$$\sum_{t \neq s} \lambda_{st}^e = \sum_{t \in V_2 - z} \frac{b(s)}{B_1} \lambda_{zt}^2(e) \leq \frac{b(s)}{B_1} b_2(z) c_2(e) = b(s) c(e),$$

and similarly for $s \in V_2, e \in E_1$. The definition of μ is exactly such that it satisfies the constraints for μ in (2). Finally, we have

$$\begin{aligned}
\sum_{s,t} \mu_{st} &= \sum_{s,t \in V_1 - z} \mu_{st}^1 + \sum_{s,t \in V_2 - z} \mu_{st}^2 + \sum_{s \in V_1 - z} \sum_{t \in V_2} \frac{b(t)}{B_2} \mu_{sz}^1 + \sum_{s \in V_1} \sum_{t \in V_2 - z} \frac{b(s)}{B_1} \mu_{zt}^2 \\
&= \sum_{s,t \in V_1 - z} \mu_{st}^1 + \sum_{s,t \in V_2 - z} \mu_{st}^2 + \sum_{s \in V_1 - z} \mu_{sz}^1 + \sum_{t \in V_2 - z} \mu_{zt}^2 \\
&= \sum_{s,t \in V_1} \mu_{st}^1 + \sum_{s,t \in V_2} \mu_{st}^2 = K_1 + K_2.
\end{aligned}$$

□

We are now ready to prove that Conjecture 3.1 holds for graphs in which any two circuits are separated by a vertex-cut of size 1, which is a slight extension of Theorem 1.2.

Theorem 5.4 *If G is a graph in which any two circuits are separated by a vertex-cut of size 1, then the cost of an optimal tree solution equals the value of an optimal dual for any instance (G, b, c) , where $b \in \mathbb{R}_+^V, c \in \mathbb{R}_+^E$.*

Proof. Directly from Theorem 1.2, from the fact that Conjecture 3.1 holds for trees, and from Lemma 5.3, since G is a 1-sum of trees and circuits. □

The next lemma says that if Conjecture 3.1 holds for an instance, it still holds if we add edges to the graph with cost equal to the length of a shortest path between their endpoints. It provides a kind of converse to Lemma 5.1.

Lemma 5.5 *Suppose for the instance $(G = (V, E), b, c)$ a tree solution of cost K and a feasible dual of value K exist, and suppose $x, y \in V$, but $e := \{x, y\} \notin E$. Then for the instance (G', b, c') a tree solution of cost K and a feasible dual of value K exist, where $G' := (V, E \cup \{e\})$, $c'|_E = c$, and $c'(e)$ is equal to $d_G^c(x, y)$, the distance in G between x and y with respect to the cost function c .*

Proof. A tree solution for (G, b, c) is also a tree solution for (G', b, c') , of the same cost. Moreover, suppose (λ, μ) is a feasible dual for (G, b, c) of value $K = \sum_{u,v} \mu_{uv}$. Define $\hat{\lambda}_{uv}^e$ as the length of a shortest path from x to y in G with respect to the length function λ_{uv} , let $\hat{\lambda}_{uv}^f = \lambda_{uv}^f$ for $f \in E$, and $\hat{\mu} := \mu$. Then $(\hat{\lambda}, \hat{\mu})$ is a feasible dual for (G', b, c') of value K . Indeed, if P_{xy} is a shortest x - y path in G with respect to the length function λ_{uv} , then we have

$$\begin{aligned} \sum_{j \neq i} \hat{\lambda}_{ij}^f &= \sum_{j \neq i} \lambda_{ij}^f \leq c(f)b(i) = c'(f)b(i) \quad \forall i \in W, \forall f \in E \\ \sum_{j \neq i} \hat{\lambda}_{ij}^e &= \sum_{j \neq i} \sum_{f \in P_{xy}} \lambda_{ij}^f = \sum_{f \in P_{xy}} \sum_{j \neq i} \lambda_{ij}^f \leq \sum_{f \in P_{xy}} c(f)b(i) = d_G^c(x, y)b(i) = c'(e)b(i). \end{aligned}$$

□

From the above lemma together with Theorem 1.2, it follows that if the cost function c on the graph is a “circuit metric”, Conjecture 3.1 holds. (c is a circuit metric if it satisfies the triangle inequality, and every edge outside some Hamilton circuit has cost equal to the length of a shortest path, along the circuit, between its endpoints).

For the remaining results in this section, let us rephrase our conjecture as follows. Let $\Phi(b, c)$ denote the minimum value of the LP (1) as a function of b and c . Note that for fixed b , the function Φ is concave in c . Similarly, Φ is concave in b for fixed c . If the optimal tree solution for a given instance has nonzero cost, then by scaling the cost function c (Lemma 3.4), which we assume to be a metric (Lemma 5.1), it is always possible to arrive at the situation where the optimal star solution (which is an optimal tree solution) has cost 1, without losing the property that c is a metric. We may also assume that the graph is complete (Lemma 5.1). Therefore, Conjecture 3.1 follows from the statement that for every complete graph $G = (V, E)$ the minimum over all b and c of $\Phi(b, c)$ subject to the following restrictions is at least 1:

$$\begin{aligned} \min \quad & \Phi(b, c), \\ \text{s.t.} \quad & b(v) \geq 0, \quad \forall v \in V, \\ & \sum_{v \neq s} b(v)c(sv) \geq 1, \quad \forall s \in V, \\ & c(e) \geq 0, \quad \forall e \in E, \\ & c(uw) \leq c(uv) + c(vw), \quad \forall u, v, w \in V. \end{aligned} \tag{14}$$

For fixed c , the constraints are linear, and the area over which we minimize the concave function $\Phi(b, c)$ is therefore convex, and hence the minimum is attained in a vertex of the polyhedron determined by (14). Similarly for fixed b .

This allows us to prove that Conjecture 3.1 holds for graphs on at most 4 vertices, by proving that it holds, when b is fixed, for all vertices c of (14) for such a graph.

Theorem 5.6 *Let G be a graph graphs on at most 4 vertices. Then the cost of an optimal tree solution equals the value of an optimal dual for any instance (G, b, c) , where $b \in \mathbb{R}_+^V$, $c \in \mathbb{R}_+^E$.*

Proof. We assume that c is a metric. For a graph on 3 or fewer vertices, c is necessarily a “circuit” or “tree” metric, and the conjecture holds (by Theorem 1.2, and Lemma 5.5). Since we may assume that the graph is complete, it suffices to prove that the conjecture holds for K_4 .

So $G = K_4$, and c is a metric on G . First of all, we may assume that $c > 0$, since otherwise, by Lemma 3.5, we are back in the case where the graph has 3 vertices. In the case of K_4 , there are in (14) four constraints saying that every star solution has cost at least 1, and 12 triangle inequalities. For fixed $b \geq 0$ (chosen such that there exists at least one nonzero metric c such that the optimal star solution has cost 1), the polyhedron determined by (14) is contained in \mathbb{R}^6 . So in a vertex of the polyhedron, 6 linearly independent constraints for c should be tight. Since we may assume that $c > 0$, this means that in any case, at least 2 triangle inequalities are tight for a cost function c minimizing $\Phi(b, c)$ over (14). We discern 4 cases.

Case 1. There are two tight triangle inequalities on the same triangle. That is, after suitably renaming the vertices 1, 2, 3, 4, we have $c(12) = c(13) + c(23)$ and $c(13) = c(12) + c(23)$. It follows that $c(23) = 0$, contradicting our assumption that $c > 0$.

So this case does not occur; only distinct triangles can be tight for the triangle inequality. Note that in K_4 , any two distinct triangles intersect in precisely one edge.

Case 2. The edge on which two distinct tight triangles intersect is on the left hand side in one of the tight inequalities, on the right hand side in the other. That is, possibly after renaming vertices, we have $c(12) = c(13) + c(23)$ and $c(23) = c(24) + c(34)$. It can be derived that $c(12) = c(13) + c(24) + c(34)$ and $c(14) = c(13) + c(34)$. This means that c is a “tree-metric” completely determined by its value on only the edges 13, 24 and 34. Since the conjecture is true for trees, by Lemma 5.5 it is also true for K_4 in this case.

Case 3. The edge on which two distinct tight triangles intersect is on the right hand side in both tight inequalities. That is, possibly after renaming vertices, we have $c(12) = c(13) + c(23)$ and $c(14) = c(13) + c(34)$, or $c(12) = c(13) + c(23)$ and $c(34) = c(13) + c(14)$. In the latter case, c is a circuit metric, completely determined by its value on the edges 13, 23, 24 and 14; the result follows from Theorem 1.2, using Lemma 5.5. In the former case, c is determined by its value on the edges 13, 23, 24, 34, and hence is a “tree plus edge” metric; the result follows from Theorem 5.4, using Lemma 5.5.

Case 4. The edge on which two distinct tight triangles intersect is on the left hand side in both tight inequalities. That is, possibly after renaming vertices, we have

$$\begin{aligned} c(12) &= c(13) + c(23), \\ c(12) &= c(14) + c(24). \end{aligned}$$

If any other triangle inequality is tight, we are back in one of the previous cases. So we may assume that the set of six linearly independent tight constraints consists of the above two triangle inequalities together with all four star-inequalities. Hence, the stars centered at all four vertices have the same cost of 1:

$$\begin{aligned} c(12)b_2 + c(13)b_3 + c(14)b_4 &= 1, \\ c(12)b_1 + c(23)b_3 + c(24)b_4 &= 1, \\ c(13)b_1 + c(23)b_2 + c(34)b_4 &= 1, \\ c(14)b_1 + c(24)b_2 + c(34)b_3 &= 1. \end{aligned}$$

The determinant of the matrix of coefficients of the above set of 6 equations is $2b_3b_4(b_1 - b_2)(b_1 + b_2 + b_3 + b_4)$. As the equations are linearly independent, this determinant is nonzero. This implies that the system has a unique solution for c , where every $c(ij)$ is an expression in the b_v 's. For this unique solution c the following can be calculated:

$$c(13) + c(14) - c(34) = \frac{(b_1 - b_2 + b_3 - b_4)(b_1 - b_2 - b_3 + b_4)(b_1 - b_2 + b_3 + b_4)}{2(b_1 - b_2)b_3b_4(b_1 + b_2 + b_3 + b_4)},$$

and

$$c(23) + c(24) - c(34) = \frac{(b_1 - b_2 - b_3 - b_4)(b_1 - b_2 + b_3 - b_4)(b_1 - b_2 - b_3 + b_4)}{2(b_1 - b_2)b_3b_4(b_1 + b_2 + b_3 + b_4)}.$$

Since both expressions are positive (no more triangle inequalities are tight), their product is also positive. It follows that

$$(b_1 - b_2 - b_3 - b_4)(b_1 - b_2 + b_3 + b_4) > 0,$$

since the other factors of the product are squares. So either the factors $b_1 - b_2 - b_3 - b_4$ and $b_1 - b_2 + b_3 + b_4$ are both positive, or they are both negative. In the first case, we have that $b_1 > b_2 + b_3 + b_4$; in the second case, $b_2 > b_1 + b_3 + b_4$ holds. So in either case, b is such that there is a “dominant” vertex, and the conjecture follows from Lemma 5.2. \square

Conjecture 3.1 also holds if G is a complete graph, and the unit cost of all edges is the same.

Theorem 5.7 *If $G = K_n$ ($n \geq 1$) and $c(e) = 1$ for every $e \in E$, then the cost of an optimal tree solution equals the value of an optimal dual for any instance (G, b, c) , where $b \in \mathbb{R}_+^V$.*

Proof. Since $c \equiv 1$, c satisfies the triangle inequality, but no triangle inequality is tight. We may assume that b is such that the cost of an optimal tree solution is greater than zero. Then, by scaling b instead of c (Lemma 3.4), it is still possible to arrive at a situation where the star-solution of minimal cost has a cost of 1. So we are now studying the minimization problem (14) for fixed c . The minimum is attained in a vertex of the polyhedron in \mathbb{R}^n determined by

$$\begin{aligned} b_i &\geq 0, \quad \forall i \in \{1, \dots, n\}, \\ \sum_{j \neq i} b_j &\geq 1, \quad \forall i \in \{1, \dots, n\}. \end{aligned} \tag{15}$$

Here, we identify the vertex set of G with $\{1, \dots, n\}$.

In a vertex of this n -dimensional polyhedron, n independent inequalities should be tight. Choose a vertex minimizing $\Phi(b, 1)$ over (15). Then for some index-sets I and J from $\{1, \dots, n\}$, with $|I| + |J| = n$,

$$\begin{aligned} b_i &= 0, \quad \forall i \in I, \\ \sum_{j=1}^n b_j - b_i &= 1, \quad \forall i \in J, \end{aligned}$$

is a set of n independent equalities. We have that $b_i = b_j$ for every $i, j \in J$, since $\sum_{j=1}^n b_j - 1$ is a constant. We will argue that we may assume that $I \cap J = \emptyset$.

First, note that $|I \cap J| \leq 1$, since if $\{i, j\} \subseteq I \cap J$, $i \neq j$, then the four equalities $b_i = 0$, $b_j = 0$, $\sum_{k=1}^n b_k - b_i = 1$, $\sum_{k=1}^n b_k - b_j = 1$ are not independent. Suppose next that i is the unique index with $i \in I \cap J$. Then $b_j = b_i = 0$ for every $j \in J$. As $|I \cup J| = n - 1$, there is exactly one vertex $k \notin I \cup J$, and from the equation for $i \in J$ we know that $b_k = 1$. But then $\sum_{j \neq k} b_j = 0 \not\geq 1$, so this solution is not feasible for (15).

So we may assume that $I \cap J = \emptyset$, or in other words, that $b_i = b > 0$ for k vertices of the graph, and $b_i = 0$, for $n - k$ vertices of the graph. Feasibility for (15) requires $k > 1$. We will explicitly construct a tree solution of minimum cost and a dual of the same positive value for this special case.

First, note that since c satisfies the triangle inequality, the optimal tree solution may be taken equal to a star. The star centered at a vertex i with $b_i = 0$ has cost kb , whereas the star centered at a vertex i with $b_i = b$ has cost $(k - 1)b$. Therefore, an optimal tree solution has cost $(k - 1)b$, and we have to construct a dual of this value. Define:

$$\begin{aligned} \lambda_{st}^e &:= \frac{2b}{k}, \quad \text{if } b_s = b_t = b, \text{ and } e = \{s, t\}, \\ \lambda_{st}^e &:= \frac{b}{k}, \quad \text{if } b_s = b_t = b, \text{ and } e \neq \{s, t\}, \\ \lambda_{st}^e &:= 0, \quad \text{if } b_s = 0 \text{ or } b_t = 0, \\ \mu_{st} &:= \frac{2b}{k}, \quad \text{if } b_s = b_t = b, \\ \mu_{st} &:= 0, \quad \text{if } b_s = 0 \text{ or } b_t = 0. \end{aligned}$$

Then, it is not hard to check that (λ, μ) is a feasible dual, with value

$$\sum_{s,t} \mu_{st} = \frac{k(k-1)}{2} \cdot \frac{2b}{k} = (k-1)b.$$

□

For completeness, we now formulate the most general statement we can obtain by combining all the results in this section. We use the following definitions. The *connectivity* of a graph $G = (V, E)$ is the minimum size of a subset U of V for which $G - U$ is not connected. If no such U exists (or equivalently, if G is complete), then the connectivity is ∞ . A graph is k -connected if its connectivity is at least k . Now, a k -connected component of a graph

$G = (V, E)$ is an inclusionwise maximal subset U of V for which $G[U]$ (the subgraph of G induced by the vertices in U) is k -connected. A *block* is a 2-connected component U with $|U| \geq 2$. We identify the blocks of a graph with the subgraphs they induce. Then distinct blocks of a graph can be seen to have at most one vertex in common. A connected graph may therefore be obtained from its blocks by taking repeated 1-sums.

Theorem 5.8 *Suppose $G = (V, E)$ is a connected graph and $c \in \mathbb{R}_+^E$ is a cost function such that every block $H = (V', E')$ of G endowed with the cost function $c|_{E'}$ is either a circuit, or a graph on at most 4 vertices, or a complete graph with uniform edge costs. Then the cost of an optimal tree solution equals the value of an optimal dual for the instance (G, b, c) , for any $b \in \mathbb{R}_+^V$.*

Proof. This follows directly from Theorem 1.2, Theorem 5.6, Theorem 5.7 and Lemma 5.3, since G can be obtained from its blocks by taking repeated 1-sums. \square

Note that Theorem 5.8 extends Theorem 5.4, since every block of a tree on two or more vertices is an edge (a complete graph on two vertices).

6 The asymmetric VPN

In this section, we give a compact LP-formulation for the MPR variant of the asymmetric VPN. An LP-formulation of exponential size for this problem (which did prove polynomial time solvability of MPR by the ellipsoid method) has been given by [3]. By using the dual variables as we did in Section 2, below we are able to obtain a formulation of polynomial size. Something similar is done (independently) in [1].

We now describe the most general asymmetric setting of the VPN problem. We are given a graph $G = (V, E)$ and a subset $W \subseteq V$ of the vertices of users. We are given a unit edge cost vector c , with possibly $c_{uv} \neq c_{vu}$ for edges $\{u, v\} \in E$. For communication from $i \in W$ to $j \in W$, a *routing* is a vector describing a number of directed paths from i to j , together with fractions that tell how communication from i to j is split over these paths. Let variable ρ_{ij}^{uv} denote the fraction of traffic from i to j , making use of edge $\{u, v\}$ in the direction from u to v . In graph terminology, ρ_{ij} describes a flow of value 1 from i to j , in the directed graph obtained from G by replacing each edge $\{u, v\}$ by two arcs (u, v) and (v, u) . It satisfies the following constraints:

$$\sum_{v:\{u,v\} \in E} (\rho_{ij}^{uv} - \rho_{ij}^{vu}) = \begin{cases} 1 & \text{if } u = i; \\ -1 & \text{if } u = j; \\ 0 & \text{otherwise,} \end{cases} \quad \text{for all } u \in V$$

We are furthermore given bounds on the amount of traffic leaving i : b_i^+ , and on the amount of traffic towards i : b_i^- . Possibly, $b_i^+ \neq b_i^-$. Given the routing ρ , we have to find capacities x_{uv} for the directed edge (u, v) such that for any communication scenario D (possibly asymmetric) that satisfies $\sum_{j \in W} d_{ij} \leq b_i^+$, and $\sum_{j \in W} d_{ji} \leq b_i^-$, the edge (u, v) is used within its capacity. Evidently, we choose x_{uv} to be as small as possible, which means that

$$\begin{aligned} x_{uv} &= \max \sum_{i \in W} \sum_{j \in W, j \neq i} d_{ij} \rho_{ij}^{uv} = \min \sum_{i \in W} (y_i^{uv} b_i^+ + z_i^{uv} b_i^-) \\ \text{s.t. } \sum_{j \in W} d_{ij} &\leq b_i^+ & \text{s.t. } y_i^{uv} + z_j^{uv} &\geq \rho_{ij}^{uv}, \quad i, j \in W, j \neq i \\ \sum_{i \in W} d_{ij} &\leq b_j^- & y_i^{uv}, z_j^{uv} &\geq 0 \\ d_{ij} &\geq 0 \end{aligned}$$

Here, the second equation follows from strong duality. The purpose is to minimize total cost $\sum_u \sum_{v:\{u,v\} \in E} c_{uv} x_{uv}$, hence the general problem is described by the following linear program

$$\begin{aligned} \min \sum_u \sum_{v:\{u,v\} \in E} \sum_{i \in W} (y_i^{uv} b_i^+ + z_i^{uv} b_i^-) c_{uv}, \\ \text{s.t. } \sum_{v:\{u,v\} \in E} (\rho_{ij}^{uv} - \rho_{ij}^{vu}) &= \begin{cases} 1 & \text{if } u = i; \\ -1 & \text{if } u = j; \\ 0 & \text{otherwise,} \end{cases} \quad u \in V, i, j \in W, j \neq i, \\ y_i^{uv} + z_j^{uv} - \rho_{ij}^{uv} &\geq 0, & i, j \in W, j \neq i; u, v \in V, \{u, v\} \in E, \\ y_i^{uv}, z_i^{uv}, \rho_{ij}^{uv} &\geq 0, & i, j \in W, j \neq i; u, v \in V, \{u, v\} \in E. \end{aligned} \tag{16}$$

The routing can be made unsplitable, by adding the constraint $\rho_{ij}^{uv} \in \{0, 1\}$. It can be made symmetric by adding $\rho_{ij}^{uv} = \rho_{ji}^{vu}$. If we do both, we force a single path routing. Note that if, in addition, the data has the property that communication bounds are symmetric, that is $b_i^+ = b_i^-$, for all $i \in W$, it follows that $x_{uv} = x_{vu}$. This can be seen by observing that if (y, z) is feasible for the dual of x_{uv} , then (z, y) is feasible for the dual of x_{vu} , and has the same objective value. Hence $x_{vu} \leq x_{uv}$, and by symmetry the converse also holds. Then the total cost equals $\sum_{u,v} c_{uv} x_{uv} = \sum_{u,v, u < v} x_{uv} (c_{uv} + c_{vu}) = \sum_{u,v} x_{uv} (c_{uv} + c_{vu})/2$. Hence, for single path routing with symmetric bounds, we may assume without loss of generality, that edge costs are symmetric, $c_{uv} = c_{vu}$.

Polynomial time solvability of the asymmetric MPR problem is all the more interesting, as the asymmetric TR problem has been shown to be NP-complete [5]. The asymmetric TR problem becomes solvable in polynomial time however if $\sum_{i \in W} b_i^+ = \sum_{i \in W} b_i^-$ [7]. This is no longer true if one considers another variant of the VPN problem where upper bounds are given on the capacities that can be reserved on the edges. In fact, as was claimed (without proof) in [5], in this last mentioned VPN variant it is already NP-hard to decide whether a feasible solution to TR or SPR exists (for both the symmetric and the asymmetric case), even if all capacities are at most 2.

Our compact LP-formulation for the asymmetric MPR problem however can be easily extended to the case where upper bounds on the available edge capacities are given (just add constraints $\sum_i y_i^{uv} b_i^+ + z_i^{uv} b_i^- \leq C_{uv}$ to the program (16), where C_{uv} is the capacity upper bound on edge uv). So the MPR problem remains solvable in polynomial time for this variant, as was also observed in [3].

In [5] and [6] approximation algorithms are studied for NP-hard versions of the VPN problem.

Note the difference between the asymmetric setting with the bidirected graph, and the symmetric problem discussed in the paper. In the symmetric case, edge traffic is counted regardless of its direction, and the bounds on incoming and outgoing traffic are also aggregated. For a given routing vector ρ the required edge capacities are computed by

$$\begin{aligned} x_{\{u,v\}} &= \max \sum_{i \in W} \sum_{j \in W-i} d_{ij} (\rho_{ij}^{uv} + \rho_{ij}^{vu}) \\ \text{s.t. } \sum_{j \in W} (d_{ij} + d_{ji}) &\leq b_i, \\ d_{ij} &\geq 0. \end{aligned}$$

Note that, under symmetric routing conditions ($\rho_{ij}^{uv} = \rho_{ji}^{vu}$), variables d_{ij} and d_{ji} have the same cost coefficients, and hence the optimum is achieved by symmetric matrix (d_{ij}) . As a consequence, the value $x_{\{u,v\}}$ above then equals

$$\begin{aligned} x_{uv} + x_{vu} &:= \max \sum_{i \in W} \sum_{j \in W-i} d_{ij} \rho_{ij}^{uv} + \max \sum_{i \in W} \sum_{j \in W-i} d_{ij} \rho_{ij}^{vu} \\ \text{s.t. } \sum_{j \in W} d_{ij} &\leq b_i/2 & \text{s.t. } \sum_{j \in W} d_{ij} &\leq b_i/2 \\ \sum_{i \in W} d_{ij} &\leq b_j/2 & \sum_{i \in W} d_{ij} &\leq b_j/2 \\ d_{ij} - d_{ji} &= 0 & d_{ij} - d_{ji} &= 0 \\ d_{ij} &\geq 0 & d_{ij} &\geq 0 \end{aligned}$$

As the optimal values for $x_{\{u,v\}}$ are attained by symmetric communication patterns d_{ij} , we may write for $e = \{u, v\}$,

$$\begin{aligned} x_e &= \max \sum_{i,j \in W, i < j} d_{\{i,j\}} (\rho_{ij}^{uv} + \rho_{ij}^{vu}) = \min \sum_{i \in W} y_i^e b_i \\ \text{s.t. } \sum_{j \in W, j \neq i} d_{\{i,j\}} &\leq b_i & \text{s.t. } y_i^e + y_j^e &\geq \rho_{ij}^{uv} + \rho_{ij}^{vu} \\ d_{\{i,j\}} &\geq 0 & y_i^e &\geq 0 \end{aligned}$$

and we have the following, polynomially sized reformulation of MPR for the symmetric VPN:

$$\begin{aligned}
& \min \sum_{e=\{u,v\} \in E} \sum_{i \in W} y_i^e b_i c_e, \\
& \text{s.t.} \quad \sum_{v:\{u,v\} \in E} (\rho_{ij}^{uv} - \rho_{ij}^{vu}) = \begin{cases} 1 & \text{if } u = i, \\ -1 & \text{if } u = j, \\ 0 & \text{otherwise,} \end{cases} \quad u \in V, i, j \in W, i < j, \\
& \quad y_i^e + y_j^e - \rho_{ij}^{uv} - \rho_{ij}^{vu} \geq 0, \quad i, j \in W, i < j, e = \{u, v\} \in E, \\
& \quad y_i^e, \rho_{ij}^{uv}, \rho_{ij}^{vu} \geq 0, \quad i, j \in W, i < j, e = \{u, v\} \in E.
\end{aligned}$$

References

- [1] A. Altin, E. Amaldi, P. Belotti, and M.Ç. Pinar, *Virtual private network design under traffic uncertainty*, Proceedings of CTW04, 2004, pp. 24–27.
- [2] N.G. Duffield, P. Goyal, and A Greenberg, *A flexible model for resource management in virtual private networks*, ACM SIGCOMM Computer Communication Review **29** (1999), no. 4, 95–108.
- [3] T. Erlebach and M. Rüegg, *Optimal bandwidth reservation in hose-model vpns with multi-path routing*, preprint (2003).
- [4] M. Grötschel, L. Lovász, and A. Schrijver, *The ellipsoid method and its consequences in combinatorial optimization*, Combinatorica **1** (1981), no. 2, 169–197.
- [5] R.P. Gupta, J. Kleinberg, A. Kumar, R. Rastogi, and B. Yener, *Provisioning a virtual private network: A network design problem for multicommodity flow*, Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC), 2001, pp. 389–398.
- [6] R.P. Gupta, A. Kumar, and T. Roughgarden, *Simpler and better approximation algorithms for network design*, Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC), 2003, pp. 365–372.
- [7] G. Italiano, S. Leonardi, and G. Oriolo, *Design of networks in the hose model*, Proceedings of the 3rd Workshop on Approximation and Randomization Algorithms in Communication Networks (ARACNE), Carleton Scientific, 2002, pp. 65–76.
- [8] A. Kumar, R. Rastogi, A. Silberschatz, and B. Yener, *Algorithms for provisioning virtual private networks in the hose model*, IEEE/ACM Transactions on Networking **10** (2002), no. 4, 565–578.