On total traffic domination in non-complete graphs

Pablo Pavón Mariño[a], Michał Pióro[b,c,*]

[a] Technical University of Cartagena, Plaza Cronista Isidoro Valverde s/n 30202 Cartagena, Spain
[b] Inst. of Telecom, Warsaw University of Technology, Nowowiejska 15/19, 00-665 Warszawa, Poland
[c] Dept. of Electrical and Information Technology, Lund University, Box 118, S-221 00 Lund, Sweden

Abstract

Given a graph $\mathcal{G}(V, E)$, a set of traffic matrices $\mathcal{H}$ and one additional traffic matrix h, we say that $\mathcal{H}$ totally dominates h if for each capacity reservation u supporting $\mathcal{H}$, u also supports h using the same routing pattern. It has been shown that if $|\mathcal{H}| = 1$ and $\mathcal{G}$ is a complete graph, $\mathcal{H}$ totally dominates h if, and only if, $\hat{h} \geq h$ component-wise. In this paper we give a generalized condition for $|\mathcal{H}| \geq 1$ and any connected undirected graph.

Keywords: network optimization, traffic matrices domination, multi-hour optimization

1. Introduction

In the paper we generalize the following necessary and sufficient condition for traffic domination presented in [5]. Given a graph $\mathcal{G}(V, E)$, and two traffic matrices $\hat{h}$ and h, $\hat{h}$ *totally dominates* h ($\hat{h} \succeq h$ in short) if for each capacity reservation $u : E \mapsto \mathcal{R}_+$ and for each flow pattern $f : \mathcal{P} \mapsto \mathcal{R}_+$ (where $\mathcal{P}$ is the set of routing paths) such that (u, f) supports matrix $\hat{h}$, the solution (u, f) does also support matrix h. Paper [5] shows that for a fully connected graph $\mathcal{G}$, $\hat{h} \succeq h$ if, and only if, $\hat{h} \geq h$ component-wise.

We consider a more general definition involving a set of (non-simultaneous) traffic matrices $\mathcal{H}$ and one additional traffic matrix h. We say that $\mathcal{H}$ *totally dominates* h ($\mathcal{H} \succeq h$) if for each pair (u, f) supporting every matrix in $\mathcal{H}$, the pair (u, f) does also support matrix h. After introducing basic definitions in Section 2, in Section 3 we generalize the sufficient

*Corresponding author: Michał Pióro, Warsaw University of Technology, Nowowiejska 15/19, 00-665 Warszawa, Poland, e-mail: mpp@tele.pw.edu.pl

condition of [5] by showing that if there exists a convex combination $\hat{h}$ of the matrices in $\mathcal{H}$ such that $\hat{h} \geq h$, then $\mathcal{H} \succeq h$. In the same section we demonstrate that if $\mathcal{G}$ possesses a certain technical property, the condition is also necessary. Then, in Section 4, we show that the property is fulfilled by 2-connected undirected graphs which leads to the main result of the paper: in 2-connected graphs $\mathcal{H} \succeq h$ if, and only if, there exists a convex combination $\hat{h}$ of the matrices in $\mathcal{H}$ such that $\hat{h} \geq h$. This characterization does extend to connected graphs (not necessarily 2-connected), but requires some adjustment of the involved traffic matrices $\mathcal{H}$ and h, as we finally demonstrate in Section 5.

We wish to mention that the reference paper [5] deals also with other types of traffic domination, in particular with the case when different flow patterns can be used for different traffic matrices, and gives a necessary and sufficient condition for this type of domination in fully connected networks. This case is not considered in our paper.

2. Definition of total domination

We assume that we are given an undirected graph $\mathcal{G} = \mathcal{G}(V, E)$ with the set of nodes V and the set of undirected links E. A capacity *reservation* $u : E \rightarrow \mathcal{R}_+$ specifies for each link $e \in E$ the amount of capacity $u(e)$ installed on e. In the sequel, the value $u(e)$ will be denoted by u_e, and function u will be represented by the *capacity reservation vector* $u = (u_e,\ e \in E)$.

Further, we suppose that we are given an ordered set D of demands. Each demand $d \in D$ is identified by its end nodes $s(d)$ and $t(d)$ (without loss of generality, $s(d) \neq t(d)$). Note that there can be more than one demand between the same pair of nodes (or no demand at all), so D is indeed a *multiset*. A function $h : D \rightarrow \mathcal{R}_+$ specifies for each demand $d \in D$ its *traffic volume* $h(d)$. Denoting $h(d)$ by h_d we form the *traffic vector* $h = (h_d,\ d \in D)$ that will represent function h. If h and h' are both traffic vectors, when we write $h \geq h'$, we mean that $h_d \geq h'_d$, for each $d \in D$. (We prefer to use the notion of the traffic vector instead of the traffic matrix used in the introduction.)

We also assume that we are given, for each demand $d \in D$, a set $\mathcal{P}_d$ of elementary paths between $s(d)$ and $t(d)$ that we call *admissible* (a path is *elementary* if all its nodes are different; observe that an elementary path p can be identified with the set of its links, i.e., $p \subseteq E$). Note that different sets of admissible paths can be given for a demand d and a demand d' that share the same end of nodes. It is convenient to think of $\mathcal{P}_d$ as a *given* subset of elementary paths between $s(d)$ and $t(d)$, that have been selected for carrying traffic of demand d. Finally, we let $\mathcal{Q}_{ed} \subseteq \mathcal{P}_d$ be the set of all admissible paths for demand d that contain link e: $\mathcal{Q}_{ed} = \{p \in \mathcal{P}_d : \ e \in p\}$.

Given the set of demands D and, for each $d \in D$, the set $\mathcal{P}_d$ of admissible paths, a *flow pattern* f specifies the (fractional) flow $f_{dp} \in \mathcal{R}_+$ assigned for each $d \in D$ to each path $p \in \mathcal{P}_d$, and will be identified with the vector $f = (f_{dp}, \ d \in D, p \in \mathcal{P}_d)$. For any flow pattern we require that $\sum_{p \in \mathcal{P}_d} f_{dp} = 1, \ d \in D$, so that for a given demand d the values $f_{dp}, p \in \mathcal{P}_d$, specify the fractions of the demand volume h_d assigned to the paths in $\mathcal{P}_d$ for each $d \in D$.

Consider a finite set of traffic vectors $\mathcal{H} = \{h^t : \ t \in T\}$ to be supported by the network, and one additional traffic vector $h = (h_d, d \in D)$. Set $\mathcal{H}$ can represent $|T|$ different traffic vectors observed at different hours, or a set of observed realizations of a random traffic vector. For every $h^t, t \in T$, its traffic volume for demand $d \in D$ is denoted by h^t_d so that $h^t = (h^t_d, \ d \in D)$.

We say that a capacity reservation u and a flow pattern f, defined on the given graph G, *support* a set of traffic vectors $\mathcal{H}$ with respect to the given family of admissible path sets $\mathcal{P} = \{\mathcal{P}_d : \ d \in D\}$, if the link loads induced by f do not exceed link capacity reservations for any traffic matrix in $\mathcal{H}$, i.e., when (u, f) satisfy the following linear constraints:

$$\sum_{p \in \mathcal{P}_d} f_{dp} = 1, \qquad d \in D \tag{1a}$$

$$\sum_{d \in D} \sum_{p \in \mathcal{Q}_{ed}} h^t_d f_{dp} \leq u_e, \qquad e \in E, \ t \in T \tag{1b}$$

$$f_{dp} \geq 0, \qquad d \in D, \ p \in \mathcal{P}_d. \tag{1c}$$

For a given family of admissible path sets $\mathcal{P}$ and a set of traffic vectors $\mathcal{H}$, the feasible set of (u, f) defined by (1) (a polyhedron in $\mathcal{R}_+^{|E| + \sum_{d \in D} |\mathcal{P}_d|}$) will be abbreviated by $\mathbb{P}_{\mathcal{P}}(\mathcal{H})$.

Now, we are ready to introduce a formal definition of total domination, extending the one from [5] to a set of traffic vectors $\mathcal{H}$ (instead of a single traffic vector).

Definition 1. *Let* $\mathcal{P} = \{\mathcal{P}_d : \ d \in D\}$ *be a given family of admissible path sets in an arbitrary graph* $\mathcal{G}(V, E)$*. We say that* $\mathcal{H}$ totally dominates h *with respect to* $\mathcal{P}$ *(*$\mathcal{H} \succeq_{\mathcal{P}} h$ *in short) if, and only if, each feasible solution in* $\mathbb{P}_{\mathcal{P}}(\mathcal{H})$ *is also a feasible solution in* $\mathbb{P}_{\mathcal{P}}(\{h\})$*, i.e., if, and only if,* $\mathbb{P}_{\mathcal{P}}(\mathcal{H}) \subseteq \mathbb{P}_{\mathcal{P}}(\{h\})$*.*

(In the sequel we will skip subscript $\mathcal{P}$ in $\mathbb{P}_{\mathcal{P}}$ and in $\succeq_{\mathcal{P}}$ when the family $\mathcal{P}$ of admissible path sets is fixed. We will also write $\mathbb{P}_{\mathcal{P}}(h)$ instead of $\mathbb{P}_{\mathcal{P}}(\{h\})$.)

Observe that our setting for the flow patterns corresponds to what is called oblivious [1], or stable [2], or static [3] routing, that is, it assumes splitting the demand volumes to admissible paths in the same proportion for any traffic vector. As already mentioned, traffic vector dependent (i.e., dynamic) flow patterns could as well be considered in the context of traffic domination (see [5]).

3. A sufficient and a necessary condition for total domination

The following sufficient condition for total domination holds for an arbitrary (undirected) graph $\mathcal{G}(V, E)$ with an arbitrary demand set D.

Proposition 1. *Consider a set of traffic vectors* $\mathcal{H}$ *and an additional traffic vector* h*. Assume that there exists* $\hat{h} \in conv(\mathcal{H})$ *such that* $\hat{h} \geq h$*. Then* $\mathcal{H}$ *totally dominates* h *with respect to any family of admissible path sets* $\mathcal{P}$*.*

Above, $\text{conv}(\mathcal{H})$ denotes the set of all convex combinations of the vectors in $\mathcal{H}$ (recall that each traffic vector is indeed a function $h : D \to \mathcal{R}_+$). The proof of Proposition 1 consists in showing that if $(u, f) \in \mathbb{P}_{\mathcal{P}}(\mathcal{H})$, then also $(u, f) \in \mathbb{P}_{\mathcal{P}}(\hat{h})$ and $(u, f) \in \mathbb{P}_{\mathcal{P}}(h)$, and is based on standard convexity arguments.

The sufficient condition for total domination given in Proposition 1 is in general not necessary, as illustrated in Fig.1 depicting a 3-node, 2-link graph with three demands $\{1, 3\}, \{1, 2\}, \{2, 3\}$, and with the family of admissible path sets $\mathcal{P} = \{\{123\}, \{12\}, \{23\}\}$. Assume $\mathcal{H} = \{\hat{h}\}, \hat{h} = (1, 0, 0)$ ($\hat{h}_{13} = 1, \hat{h}_{12} = \hat{h}_{23} = 0$), and $h = (0, 1, 1)$ ($h_{13} = 0, h_{12} = h_{23} = 1$). Since the network is a tree, only one flow pattern f exists and trivially every pair

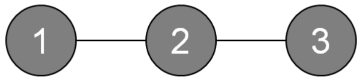

Figure 1: A two-link network.

(u, f) supporting $\hat{h}$ also supports h. Hence, $\hat{h}$ totally dominates h with respect to $\mathcal{P}$, and, in fact, vice versa. Still, neither $\hat{h} \geq h$ nor $h \geq \hat{h}$.

Below we will derive a (technical) property of the family $\mathcal{P}$ under which the condition $\exists\, \hat{h} \in \mathrm{conv}(\mathcal{H}),\ \hat{h} \geq h$ of Proposition 1 is also necessary for the total domination $\mathcal{H} \succeq_{\mathcal{P}} h$.

Proposition 2. *Let $\mathcal{H}$ and h be as in Proposition 1 and suppose that a family $\mathcal{P} = \{\mathcal{P}_d : d \in \mathcal{P}_d\}$ of admissible path sets is given. Suppose also that there exists a link $\hat{e} \in E$ (the so called* enabling link*) with the following property: for every demand $d \in D$ there exists a path $p' \in \mathcal{P}_d$ such that $\hat{e} \in p'$ and a path $p'' \in \mathcal{P}_d$ such that $\hat{e} \notin p''$.*

Under these assumptions, if $\mathcal{H}$ totally dominates h with respect to $\mathcal{P}$, then there exists $\hat{h} \in conv(\mathcal{H})$ such that $\hat{h} \geq h$.

Proof. Consider an enabling link $\hat{e} \in E$ and the following linear programming (LP) problem specified on polyhedron $\mathbb{P}_{\mathcal{P}}(\mathcal{H})$ (see (1)):

$$\text{minimize} \quad u_{\hat{e}} - \sum_{d \in D} \sum_{p \in \mathcal{Q}_{\hat{e}d}} h_d f_{dp} \tag{2a}$$

$$[\alpha_{et}] \quad \sum_{d \in D} \sum_{p \in \mathcal{Q}_{ed}} h_d^t f_{dp} - u_e \leq 0, \qquad e \in E,\ t \in T \tag{2b}$$

$$[\gamma_d] \quad \sum_{p \in \mathcal{P}_d} f_{dp} = 1, \qquad d \in D \tag{2c}$$

$$f_{dp} \geq 0, \qquad d \in D,\ p \in \mathcal{P}_d \tag{2d}$$

Using the dual variables specified in the square parentheses, the problem dual to (2) can be written down as:

$$\text{maximize} \quad \sum_{d\in D} \gamma_d \tag{3a}$$

$$\sum_{t\in T} \alpha_{et} = 0, \qquad e \in E \setminus \{\hat{e}\} \tag{3b}$$

$$\sum_{t\in T} \alpha_{\hat{e}t} = 1 \tag{3c}$$

$$\alpha_{et} \geq 0, \qquad e \in E,\ t \in T \tag{3d}$$

$$\gamma_d + h_d \leq \sum_{t\in T} \alpha_{\hat{e}t} h_d^t, \qquad d \in D:\ \mathcal{Q}_{\hat{e}d} \neq \emptyset \tag{3e}$$

$$\gamma_d \leq 0, \qquad d \in D:\ \mathcal{P}_d \setminus \mathcal{Q}_{\hat{e}d} \neq \emptyset. \tag{3f}$$

Note that we can write (3e) and (3f) in this form since it follows from (3b) and (3d) that $\alpha_{et} = 0$ for every $e \neq \hat{e}$ and $t \in T$. Let (α^*, γ^*) be an optimal solution of (3). Because link $\hat{e}$ is enabling, $\forall\ d \in D,\ \mathcal{Q}_{\hat{e}d} \neq \emptyset$ and $\forall\ d \in D,\ \mathcal{P}_d \setminus \mathcal{Q}_{\hat{e}d} \neq \emptyset$. The latter condition and inequality (3f) imply that $\gamma_d^* \leq 0$ for all $d \in D$. Further, we have that $\sum_{d\in D} \gamma_d^* \geq 0$, since optimal objective (3a) of the dual is equal to optimal objective (2) of the primal, and the latter is non-negative by the assumption $\mathcal{H} \succeq_{\mathcal{P}} h$. Hence, $\gamma_d^* = 0$ for all $d \in D$. This, the condition $\forall\ d \in D,\ \mathcal{P}_d \setminus \mathcal{Q}_{\hat{e}d} \neq \emptyset$, and (3e) imply that for all $d \in D$, $h_d \leq \sum_{t\in T} \alpha_{\hat{e}t}^* h_d^t$.

Hence, thanks to (3c) and (3d), total domination implies existence of a convex combination $\hat{h} = \sum_{t\in T} \alpha_{\hat{e}t}^* h_d^t$ of the vectors in $\mathcal{H}$ such that $\hat{h} \geq h$. ∎

4. A necessary and sufficient condition for two-connected networks

An undirected graph $\mathcal{G}(V, E)$ is called *2-connected* if it has at least three nodes and does not contain any *cut vertex*, i.e., any vertex $v \in V$ such that $\mathcal{G} \setminus v$ has more connected components than $\mathcal{G}$. (Below, we follow definitions and results given in Chapter 11 of [4].)

Proposition 3. *Suppose that an undirected graph $\mathcal{G}$ is 2-connected and that for every $d \in D$ the admissible set $\mathcal{P}_d$ contains all elementary paths between nodes $s(d)$ and $t(d)$. Then*

$$\mathcal{H} \succeq h \ \Leftrightarrow \exists\ \hat{h} \in conv(\mathcal{H}),\ \hat{h} \geq h. \tag{4}$$

Proof. 2-connected graphs enjoy the two following properties: (i) $\mathcal{G}$ is 2-connected if, and only if, any two nodes $v, w \in V$ are connected by two node disjoint paths ([4], Theorem 11.1.1), and (ii) $\mathcal{G}$ is 2-connected if, and only if, for any two nodes $v, w \in V$ and link $e \in E$ there exists an elementary path between v and w containing link e ([4], Theorem 11.3.1).

These properties imply that in a 2-connected graph each link $\hat{e} \in E$ is enabling for any demand set D when, as assumed, each set $\mathcal{P}_d$ ($d \in D$) contains all elementary paths between $s(d)$ and $t(d)$. Consider a fixed link $\hat{e}$ and an arbitrary demand $d \in D$. There must be a path $p \in \mathcal{P}_d$ such that $\hat{e} \notin p$. Otherwise, all elementary paths for d would contain $\hat{e}$, contradicting property (i). Also, by property (ii), there must be a path $p' \in \mathcal{P}_d$ such that $\hat{e} \in p'$. ∎

We observe that when graph $\mathcal{G}$ is composed of a single link (i.e., isomorphic to $\mathcal{K}_2$) and thus not 2-connected, then (4) is not true when there are multiple demands between the two nodes of the graph. In this case the following obvious property holds instead.

Proposition 4. *Suppose that an undirected graph $\mathcal{G}$ is composed of a single link. Then*

$$\mathcal{H} \succeq h \;\Leftrightarrow \exists\, \hat{h} \in conv(\mathcal{H}),\; \sum_{d \in D} \hat{h}_d \geq \sum_{d \in D} h_d.$$

5. General characterization of total domination

In this section we will present the main result of the paper – a general necessary and sufficient condition for total domination in an arbitrary undirected connected graph.

A connected undirected graph $\mathcal{G}$ containing cut vertices is called *separable*. The maximal induced subgraphs of $\mathcal{G}$ that are not separable are called *blocks*. A block is either 2-connected or is isomorphic to $\mathcal{K}_2$. The blocks of a graph are unique. Any two blocks intersect in at most one node and this node must be a cut vertex. Two nodes are in the same block if, and only if, they are connected by a path not traversing a cut vertex. Finally, any cycle must be contained in a block (see [4], page 338).

Consider a connected separable undirected graph $\mathcal{G} = \mathcal{G}(V, E)$ with the demand set D. Suppose that graph $\mathcal{G}$ is composed of B blocks $\mathcal{G}^b = \mathcal{G}^b(V^b, E^b)$, $b \in B$. Clearly, each elementary path between a given pair of nodes traverses the same sequence of blocks and the same sequence of cut vertices (otherwise, there would be a cycle not contained in a single block). The set of the cut vertices traversed by the paths between the end nodes of demand $d \in D$ will be denoted by $C(d)$. We treat each block as a separate graph $\mathcal{G}^b$ with its own set of demands D^b and the family of admissible path sets $\mathcal{P}^b = \{\mathcal{P}^b_d : \ d \in D^b\}$, induced respectively by D and $\mathcal{P} = \{\mathcal{P}_d : \ d \in D\}$). The demand set for a block is induced by the overall demand set in a natural way: the demands of block b are all those demands in D that either have both end nodes in V^b, or only one end in V^b (that is not a cut vertex), or both ends outside V^b. More precisely, $D^b = D^b_0 \cup D^b_1 \cup D^b_2$ (with a slight abuse of notation since these sets are indeed multisets) where:

$$D^b_0 = \{d \in D : \ |\{s(d), t(d)\} \cap V^b| = 0, \ |C(d) \cap V^b| = 2\}$$
$$D^b_1 = \{d \in D : \ |\{s(d), t(d)\} \cap V^b| = 1, \ |\{s(d), t(d)\} \cap (V^b \setminus C(d))| = 1, \ |C(d) \cap V^b| = 1\}$$
$$D^b_2 = \{d \in D : \ |\{s(d), t(d)\} \cap V^b| = 2\}.$$

Note that in block b, the end nodes of the demands in D^b_0 and D^b_1 are defined as follows:

$$\{s^b(d), t^b(d)\} = C(d) \cap V^b, \qquad d \in D^b_0$$
$$\{s^b(d), t^b(d)\} = (V^b \cap \{s(d), t(d)\}) \cup (V^b \cap C(d)), \qquad d \in D^b_1.$$

The induced admissible path set $\mathcal{P}^b_d$ for a $d \in D^b$ is composed of the sub-paths of the paths in $\mathcal{P}_d$ traversing $\mathcal{G}^b$. (Note that a path in $\mathcal{P}^b_d$ can be induced by more than one path in $\mathcal{P}_d$.)

A traffic vector h in $\mathcal{G}$ induces a traffic vector h^b in $\mathcal{G}^b$ by assuming $h^b_d = h_d, \ d \in D^b$. Similarly, a set $\mathcal{H}$ of traffic vectors in $\mathcal{G}$ induces a set of traffic vectors $\mathcal{H}^b = \{h^b : \ h \in \mathcal{H}\}$. Also, a capacity reservation vector u defined in $\mathcal{G}$ induces a capacity reservation vector u^b defined in $\mathcal{G}^b$ by assuming $u^b_e = u_e, \ e \in E^b$. Finally a flow pattern f in $\mathcal{G}$ induces a flow pattern f^b in $\mathcal{G}^b$ as follows:

$$f^b_{dp} = \sum_{q \in \mathcal{P}_d, \ p \subseteq q} f_{dq}, \qquad d \in D^b, \ p \in \mathcal{P}^b_d. \tag{5}$$

Note that the so defined mapping $(u, f) \mapsto ((u^b, f^b), \ b \in B)$ preserves capacity reservation (by definition) and, for an arbitrary traffic vector h in $\mathcal{G}$, the link loads, since (5) implies

$$\sum_{d \in D} \sum_{q \in \mathcal{Q}_{ed}} h_d f_{dq} = \sum_{d \in D^b} \sum_{p \in \mathcal{Q}^b_{ed}} h^b_d f^b_{dp}, \qquad e \in E^b, \ b \in B.$$

Hence, if (u, f) supports h in $\mathcal{G}$, then (u^b, f^b) supports h^b in every block $b \in B$, i.e., if $(u, f) \in \mathbb{P}_{\mathcal{P}}(h)$, then $(u^b, f^b) \in \mathbb{P}^b_{\mathcal{P}^b}(h^b)$, where polyhedron $\mathbb{P}_{\mathcal{P}}(h)$ is defined by (1) for $\mathcal{G}$, $\mathcal{P}$ and h, and polyhedra $\mathbb{P}^b_{\mathcal{P}^b}(h^b)$ are defined by (1) for $\mathcal{G}^b$, $\mathcal{P}^b$ and h^b ($b \in B$).

Lemma 5. *Consider a connected undirected graph $\mathcal{G}$ split into blocks $\mathcal{G}^b, b \in B$. Let $\mathcal{H}$ be a set of traffic matrices, and let h be an additional traffic vector in $\mathcal{G}$. Then,*

$$\mathcal{H} \succeq_{\mathcal{P}} h \Leftrightarrow \forall\, b \in B,\ \mathcal{H}^b \succeq_{\mathcal{P}^b} h^b. \tag{6}$$

Proof. ($\Rightarrow$) Assume that $\mathcal{H} \succeq_{\mathcal{P}} h$ and consider, for some $b \in B$, a solution $(u(b), f(b)) \in \mathbb{P}^b_{\mathcal{P}^b}(\mathcal{H}^b)$. Omitting the details, we observe that it is straightforward to construct a solution $(u, f) \in \mathbb{P}_{\mathcal{P}}(\mathcal{H})$ (provided $\mathbb{P}_{\mathcal{P}}(\mathcal{H})$ is not empty) such that $u^b = u(b)$ and $f^b = f(b)$. By assumption, (u, f) supports h and hence (u^b, f^b) supports h^b. Thus, $\mathcal{H}^b \succeq_{\mathcal{P}^b} h^b$.
($\Leftarrow$) Assume that for every $b \in B$, $\mathcal{H}^b \succeq_{\mathcal{P}} h^b$ and consider a solution $(u, f) \in \mathbb{P}_{\mathcal{P}}(\mathcal{H})$, i.e., $(u, f) \in \mathbb{P}_{\mathcal{P}}(\bar{h})$ for every $\bar{h} \in \mathcal{H}$. It follows that every (u^b, f^b) supports each $\bar{h}^b \in \mathcal{H}^b$ and therefore, by assumption, (u^b, f^b) supports h^b. Due to (6), this means that (u, f) supports h. ∎

Since the blocks $b \in B$ are 2-connected (denote the set of such blocks by B') or isomorphic to $\mathcal{K}_2$ (denote the set of such blocks by B''), Lemma 5, Proposition 3 and Proposition 4 lead to the main result of the paper:

Proposition 6. *If the admissible path sets $\mathcal{P}_d, d \in D$ are composed of all elementary paths in a connected undirected graph $\mathcal{G}$ then $\mathcal{H} \succeq h$ if, and only, if*

$$(\forall\, b \in B'\ \exists\, \hat{h}^b \in conv(\mathcal{H}^b),\ \hat{h}^b \geq h^b) \wedge (\forall\, b \in B''\ \exists\, \hat{h}^b \in conv(\mathcal{H}^b),\ \sum_{d \in D^b} \hat{h}^b_d \geq \sum_{d \in D^b} h^b_d).$$

In the example from Fig.1, the graph is split into two one-link blocks $\mathcal{G}^1$ and $\mathcal{G}^2$ with $V^1 = \{1, 2\}$ and $V^2 = \{2, 3\}$, and the adjusted traffic vectors are: $\hat{h}^1 = (1, 0), h^1 = (0, 1), \hat{h}^2 = (1, 0), h^2 = (0, 1)$. Hence, $\hat{h} \succeq h$ because (as required by Proposition 4) $\hat{h}^1_1 + \hat{h}^1_2 \geq h^1_1 + h^1_2$ for $\mathcal{G}^1$, and $\hat{h}^2_1 + \hat{h}^2_2 \geq h^2_1 + h^2_2$ for $\mathcal{G}^2$ (and vice versa, $h \succeq \hat{h}$ because $h^1_1 + h^1_2 \geq \hat{h}^1_1 + \hat{h}^1_2$ on G^1 and $h^2_1 + h^2_2 \geq \hat{h}^2_1 + \hat{h}^2_2$ on G^2).

Acknowledgement: The authors acknowledge a significant contribution of the anonymous reviewer to the final shape of the paper. In particular, he pointed out the relation between existence of an enabling link and 2-connectivity, leading to Proposition 3, and suggested the construction resulting in Lemma 5. Moreover, his comments helped us to improve the original text. Also, contribution of Walid Ben-Ameur in simplifying the proof of Proposition 2 is kindly acknowledged.

P. Pavón was supported by the FP7 BONE project, the MEC project TEC2010-21405-C02/TCM CALM, and "Programa de Ayudas a Grupos de Excelencia de la R. de Murcia, F. Séneca". M. Pióro was supported by the Polish Ministry of Science and Higher Education (grants no. 280/N-DFG/2008/0 and N517 397334), and by the Swedish Research Council (grant no. 621-2006-5509).

References

[1] D. Applegate and E. Cohen. Making intra-domain routing robust to changing and uncertain traffic demands: undertanding fundamental tradeoffs. *in Proc. of SIGCOMM'03 ACM, New York (USA)*, 210:313–324, 2003.

[2] W. Ben-Ameur and H. Kerivin. Routing of uncertain traffic demands. *Optimization and Engineering*, 6:283–313, 2005.

[3] C. Chekuri, F. B. Shepherd, G. Oriolo, and M. G. Scutellà. Hardness of robust network design. *Networks*, 50:50–54, 2007.

[4] D. Jungnickel. *Graphs, Networks and Algorithms.* Springer, 1999.

[5] G. Oriolo. Domination between traffic matrices. *Mathematics of Operations Research*, 33(1):91–96, 2008.