Pricing Traffic in a Spanning Network

Hervé Moulin*, Department of Economics, Rice University
6100 Main Street, Houston TX, 77005; moulin@rice.edu

ABSTRACT

Each user of the network needs to connect a pair of target nodes. There are no variable congestion costs, only a direct connection cost for each pair of nodes. A centralized mechanism elicits target pairs from users, and builds the cheapest forest meeting all demands. We look for cost sharing rules satisfying

- *Routing-proofness:* no user can lower its cost by reporting as several users along an alternative path connecting his target nodes;
- *Stand Alone core stability*: no group of users pay more than the cost of a subnetwork meeting all connection needs of the group.

We construct first two core stable and routing-proof rules when connecting costs are all 0 or 1. One is derived from the random spanning tree weighted by the volume of traffic on each edge; the other is the weighted Shapley value of the Stand Alone cooperative game.

For arbitrary connecting costs, we prove that the core is non empty if the graph of target pairs connects all pairs of nodes. Then we extend both rules above by the piecewise-linear technique. The former rule is computable in polynomial time, the latter is not.

Category and Subject Descriptors:
C.2.4 [**Computer-Communication Networks**]: Distributed Systems; F.2.0 [**Analysis of Algorithms and Problem Complexity**]: General; J.4 [**Social and Behavioral Sciences**]: Economics; K.4.4 [**Electronic Commerce**]: Payment Schemes

General Terms: Algorithms, Economics, Theory

Keywords: minimal cost spanning forest, core stability, routing-proofness, cost sharing rules, network design

*The critical comments of seminar participants at Lund University, CORE, Boston College, Université de Laval and Université de Montreal have been very helpful.

EC'09, July 6–10, 2009, Stanford, California, USA.

1. INTRODUCTION

We consider a network between a set of nodes $i, j, \cdots$, serving the traffic demands of a set of users. Each user requests an (undirected) connection between a pair of *target nodes* i and j; several users may request the same connection. The network must connect, directly or indirectly, all target pairs. There are no variable congestion costs: building a given edge imposes a fixed cost (to cover the construction of some infrastructure, or some access fee) that may depend on the end points of the edge, but not on the amount of traffic on this edge. Users are satisfied with an indirect connection between their target nodes, therefore the efficient network is a minimal cost forest (disjoint union of trees) *spanning* all the traffic demands (connecting all target pairs).

A familiar special case is the minimal cost spanning tree problem (thereafter mcst: [4]) where users are attached to a node and each user needs to connect to a special node (the *source*). With n nodes, the $n-1$ types of users in the mcst problem become $\frac{n(n-1)}{2}$ types in our model. To the familiar applications of the mcst model to multicast transmission ([18],[14]), our model adds the special case of the network synthesis problem where links are congestion-free (as in the connection game of [1], and the capacity synthesis problem of [6]). Like in [1],[6], and unlike in [18],[23], users' demands are inelastic, ie., all demands must be served.

Given the complex externalities in this network design problem, the challenge is to propose a compelling division of the efficient cost among the participants. For over three decades, the most fruiteful approach in the mcst and other combinatorial optimization problems has been cooperative game theory, and the key concept Stand Alone core stability: no group of users should be charged more than the cost of a subnetwork meeting the needs of the group members ([4],,[20],[16]). Adding to core stability natural requirements of continuity and monotonicity with respect to costs ([12]) led, more recently, to the definition of a certain canonical cost sharing rule, invented independently by several authors ([15],[25],[2]), and dubbed the "folk" solution in [5]. Here we follow this tradition, to the extent that the two cost sharing rules we propose reduce to the folk solution in the mcst case.

We introduce a novel type of strategic maneuver of a non cooperative nature, quite different than the (real or virtual) coordinated secessions prevented by Stand Alone core stability. To implement a personalized price to users, as required by core stability, the system manager must be able to identify the target nodes of every user, for instance by "tagging" individual messages. However in many large networks the users can easily create multiple aliases without being detected, and in this sense they remain anonymous. Gaming a mechanisms by using multiple aliases has been discussed in a variety of contexts, including the ranking of web pages ([11], [8]), scheduling ([21]), and voting ([9]). In our model it means that a user with target nodes i,j may be able to lower his charge by reporting as *two* users, one requesting ik, the other kj, or more generally by reporting as several users, each requesting one edge of a path $ik_1, k_1k_2, \cdots, k_tj$ from i to j (we give simple examples at the beginning of section 2). We call a mechanism *routing-proof* if it is not vulnerable to such maneuvers.

Strategic routing may have an efficiency cost, because the optimal graph based on the reported traffic may include some edges that it is inefficient to build. An example is given in section 3. Even if routing maneuvers do not change the optimal graph, they distort the report of actual traffic and may create strategic unstability, including the possibility of conflicting reporting equilibria (see section 2). Because it eliminates gaming opportunities in the traffic reports, Routing-proofness is as desirable a property as the familiar strategy-proofness in other resource allocation mechanisms.

An inspiration for the paper is the recent work on *selfish routing* in networks, in particular the *global connection games* in [1] (see also [18],[14]), where the network technology is the same as here, each user chooses non cooperatively a route (a sequence of edges connecting his two target nodes), and the cost of each edge is divided equally between all users who choose a path containing this edge (while in our model all messages are routed along the efficient spanning forest). A Nash equlibrium always exists ([1]), strong equilibrium may ot may not exist ([13]), and in equilibrium the optimal spanning forest may or may not be built. Remarkably, the relative excess cost of the equilibrium forest can often be bounded a priori ([1], [27], [24] Chap 19).

Finally, profitable routing maneuvers when the network is a fixed tree are the subject of [17] and [10]. In the former, users are located on nodes, and comunicate with every other user of the network; in the later, the tree represents a linear highway and each user travels on a subinterval of the highway.

2. SUMMARY OF RESULTS, AND AN EXAMPLE

Recall that Stand Alone core stability is out of reach for some patterns of traffic and connection costs. This is already the case in the mcst model ([20]). The difficulty comes from *inactive nodes and edges* (a.k.a. *Steiner nodes and edges*), namely edges $e = ij$ such that no user has i, j for target pair, and nodes i such that no user has i for one of her target nodes, while they can be used in the optimal spanning forest[1].

We identify two families of connection games where the Stand Alone core is non empty. In each family we construct two cost sharing rules that are core stable and routing-proof (Theorems 1 and 2). The first family (section 4) places no restriction on the traffic pattern, but restricts all connecting costs c_e to be either 0 or 1. Although coalitions can use inactive nodes or edges, non emptiness of the Stand Alone core is a simple fact. The second family (section 5) allows for arbitrary connection costs but requires that the traffic be *spanning*, namely the graph of active edges is connected and reaches all nodes. Thus there are no inactive nodes, but coalitions may use inactive edges. An example is the standard mcst model, where all nodes other than the source are occupied by at least one agent. In the random graph model with no source, it takes only $n \log n$ edges to generate a spanning graph with very high probability ([7]), so we expect traffic to be spanning even when most edges are inactive. Showing that the core is non empty when traffic is spanning requires more work (Theorem 2).

We propose two cost sharing rules that are both routing proof and core stable in the two environments just discussed. For routing maneuvers to be relevant, agents must be able to create aliases without being detected, implying that the mechanism cannot distinguish identical demands. Thus our rules charge the same price to all users with the same target pair. A consequence of routing-proofness is that cost shares are sensitive to usage (although total costs are not). E.g., both rules satisfy the following property: if m users have target nodes i, j, and one more user with the same demand shows up, the *total* charge to the users with this demand weakly increases, whereas the charge to *any initial* user, whether he had the target nodes i, j or not, weakly decreases.

Both rules coincide with the folk solution in the special case of the mcst model ; they share with that solution two powerful monotonicity properties: one was mentioned in the previous paragraph; the other is that *all* cost shares increase weakly when *any* connection cost increases ([2]); see Section 6. Both rules are easy to define (though only one is easy to compute) when connecting costs are all 0 or 1 (Section 4), and are extended to arbitrary costs spanning traffic by means of the *piecewise-linear* technique pioneered by [15],[25] and developed in [5] and [6]; see Section 5.

Two simple examples show that core stability and routing-proofness are not "spontaneously" compatible. Consider the familiar *proportional* cost sharing rule, dividing the optimal cost between the demanded egdes (target pairs) in proportion to their respective Stand Alone costs (and dividing equally the charge to a demanded edge between all its users). Assume three nodes $1, 2, 3$, symmetric costs $c_{12} = c_{23} = c_{13} = \$10$, and all edges active. Write θ_{ij} for the (positive) number of users with target pair i, j. The rule charges $\$\frac{20}{3}$ to each edge, so that a user of ij pays $\frac{20}{3\theta_{ij}}$. If a user of 12 passes as one user of 13 *plus* one user of 23, her total charge is $\frac{20}{3(1+\theta_{13})} + \frac{20}{3(1+\theta_{23})}$. This is profitable if

$$\frac{1}{\theta_{12}} > \frac{1}{1+\theta_{13}} + \frac{1}{1+\theta_{23}}$$

i.e., if θ_{12} is sufficiently small w.r.t. θ_{13} and θ_{23}. This move benefits the users on 13 and 23, and hurts the remaining users on 12.

Here is an example where under the same proportional sharing rule, routing maneuvers may have an efficiency cost. We have four nodes $1, 2, 3, 4$ and three agents demanding respectively 12, 23 and 13. Connection costs are

$$c_{13} = c_{23} = c_{14} = c_{24} = \$10; \; c_{12} = c_{34} = \$80$$

[1] A familiar trick to recover core stability is to restrict the power of (proper) coalitions by forbidding them to use inactive nodes or edges when formulating an objection. Only the grand coalition is allowed such opportunity. This assumption makes little sense in our model, where users can fake a traffic demand involving any node or edge.

The optimal tree is $1 \longleftrightarrow 3 \longleftrightarrow 2$ with cost \$20, divided as $z_{13} = z_{23} = \$2, z_{12} = \$16$. Here agent 12 may pose as two users, one on 14 and one on 24, which will increase the optimal (reported) cost to \$30 (by connecting any three among the \$10 edges), but decrease this agent's charge to $\frac{10+10}{10+10+10+10}30 = \$15$.

It is easy to find a routing-proof rule for a problem of any size: simply split total cost *uniformly* among all users, irrespective of their target nodes. When a user of edge $\widetilde{e}$ reports as two users of edges connecting the end-nodes of $\widetilde{e}$, her charge goes from $\frac{v}{\sum_e \theta_e}$ to $\frac{2v'}{1+\sum_e \theta_e}$, where v and v' are the optimal costs respectively for the true and reported target pairs, so that $v \leq v'$.

But this rule is not core stable. In the three nodes example above the users of 12 pay more than their Stand Alone cost if

$$\frac{\theta_{12}}{\theta_{12} + \theta_{23} + \theta_{13}} 20 > 10 \Leftrightarrow \theta_{12} > \theta_{23} + \theta_{13}$$

3. THE MODEL

The finite set of nodes is Q, with generic element i. The set of undirected edges is $Q(2)$, with generic element e. Thus $Q(2)$ is also the complete (undirected) graph on Q, and a subset $E \subseteq Q(2)$ is called a **graph** or a **network**.

An agent needs to send traffic between the end-nodes of a certain edge e. We only consider mechanisms treating identical traffic demands equally, therefore a profile of such demands can be written as a vector $\theta \in \mathbb{N}^{Q(2)}$, where θ_e is the number of agents requesting traffic on e. We call θ the **traffic profile**, and θ_e the **traffic volume***,* or simply the **traffic***,* on e. We write $[\theta]$ for the *support* of θ, namely $e \in [\theta] \Leftrightarrow \theta_e > 0$, and we speak of the **traffic graph** $[\theta]$.

Given two graphs $E, F \subseteq Q(2)$, we say that F **spans** E if for all $e \in E$, the end-nodes of e are connected in F. Note that we do not require F to be contained in E. Given the demand profile θ, a graph E is **feasible** for θ if E spans $[\theta]$.

The last component of our model is the "*matrix*" of connecting costs $c \in \mathbb{R}_+^{Q(2)}$, where c_e is the cost of building the edge e. The cost of a graph E is $c(E) = \sum_{e \in E} c_e$. The minimal cost of sustaining the traffic graph E is

$$v(c, E) \stackrel{def}{=} \min_{F: F \text{ spans } E} c(F) \quad (1)$$

An **efficient (optimal) network** for the problem (Q, c, θ) is a feasible graph F for θ of minimal cost $v(c, [\theta])$. If $[\theta]$ is connected F is a tree, otherwise it is a *forest* (a union of trees with no common nodes). In both cases F may use edges in $Q(2) \diagdown E$. An **efficient allocation** is a vector of charges $z \in \mathbb{R}_+^{[\theta]}$ such that

$$\sum_{[\theta]} z_e = v(c, [\theta])$$

Here z_e is the *total* charge on edge e, i.e., each agent requesting e pays $\frac{z_e}{\theta_e}$.

Given the set Q of nodes, a **cost-sharing rule,** or simply a **rule**, assigns to every problem (Q, c, θ) an efficient allocation $z(c, \theta)$.

Given the traffic profile θ, a coalition of agents is some $\theta' \in \mathbb{N}^{Q(2)}$ such that $\theta' \leq \theta$. Its Stand Alone cost is $v(c, [\theta'])$.

Definition 1 *The allocation z is in the Stand Alone core (or simply the core) of problem (Q, c, θ) if it is efficient and*

$$\text{for all } \theta' \leq \theta : \sum_{e \in [\theta']} \frac{\theta'_e}{\theta_e} z_e \leq v(c, [\theta'])$$

Clearly the relevant coalitions θ' to test core stability are such that for all e, $\theta'_e = \theta_e$ or 0. So the core stability of z takes the following simpler form

$$\sum_{[\theta]} z_e = v(c, [\theta]) \text{ and for all } F \subset [\theta] : \sum_{e \in F} z_e \leq v(c, F) \quad (2)$$

This property only depends upon the traffic graph $[\theta]$, not upon the number of agents active on each edge e of $[\theta]$. Indeed the system (2) means that the allocation z is in the core of the cooperative game $([\theta], v(c, \cdot))$, where players are the active edges. That core may be empty(see for instance [28]).

We now define routing-proofness. Given an edge e, an e-path is a path $e^1, e^2, \cdots, e^K$ in $Q(2)$ connecting the end-nodes of e. Given two traffic profiles θ, θ' and an edge $e \in [\theta]$, we say that θ' is a *routing maneuver of the e-traffic at* θ, if there exists an e-path and a *positive* integer $x \leq \theta_e$, such that:

$$\theta'_e = \theta_e - x;\ \theta'_{e^k} = \theta_{e^k} + x,\ 1 \leq k \leq K;$$
$$\theta'_g = \theta_g \text{ if } g \in Q(2) \diagdown \{e, e^1, e^2, \cdots, e^K\}$$

Note that a routing maneuver may involve any subset of agents with the same target nodes, but not agents with different target nodes.

Definition 2 *A cost sharing rule on Q is routing-proof at c if for all θ,e, and all routing maneuvers θ' at θ of the e-traffic, we have*

$$\frac{1}{\theta_e} z_e(c, \theta) \leq \sum_1^K \frac{1}{\theta'_{e^k}} z_{e^k}(c, \theta') \quad (3)$$

We must multiply both sides of the above inequality by x to obtain the actual charge of the rerouting coalition. The size of x still plays a role in the computation of $z(c, \theta')$.

4. ELEMENTARY COST MATRICES

In this section we assume that cost matrices are "elementary" and design two cost sharing rules that are both core stable and routing-proof, for any traffic graph. A cost matrix c is *elementary* if $c_e = 0$ or 1 for all e. Their set is denoted $\mathcal{E}(Q)$.

Our first solution is based on the Stand Alone cooperative game of which the players are the active edges, and coalitions are allowed to use all edges (in particular free edges), whether or not they connect the target pair of a coalition member. For any $c \in \mathbb{R}_+^{Q(2)}$ and any $E \subseteq Q(2)$, this game $(E, v(c, \cdot))$ ((1)) is clearly monotone and subadditive. If c is elementary, we can say more.

Lemma 1 *If $c \in \mathcal{E}(Q)$, the game $(E, v(c, \cdot))$ is concave (submodular) for any $E \subseteq Q(2)$.*
We omit the straightforward proof (or see [22]).

Concavity of $([\theta], v(c, \cdot))$ implies that for any ordering σ of the set $[\theta]$, the *marginal contribution* allocation z^σ

$$z_e^\sigma = \partial_e v(c, F)$$

where F is the set of edges in $[\theta]$ preceding e in σ, is core stable. This cost sharing rule ignores the volume of traffic on each edge, therefore it is vulnerable to routing. But any weighted average of such rules remains core stable, and a judicious choice of the weights will ensure routing-proofness. We choose the *weighted Shapley value* of the game $([\theta], v(c,\cdot))$, where the weight of edge e is the traffic θ_e. Recall ([19]) that the corresponding random ordering $\sigma(\theta)$ of the edges $e \in [\theta]$ is defined recursively for $t = 1, \cdots, |[\theta]|$ as follows:

if $e^1, \cdots, e^{t-1}$ have been chosen and not yet e, e is the t-th edge with probability $\frac{\theta_e}{\sum_{[\theta] \setminus \{e^1, \cdots, e^{t-1}\}} \theta_{e'}}$

Definition 3 *In a problem (Q, c, θ) where $c \in \mathcal{E}(Q)$, the weighted Shapley rule $z^{wsh}(c,\theta)$ charges to edge e the expectation of its marginal contributions in the random ordering $\sigma(\theta)$:*

$$z_e^{wsh}(c,\theta) = Exp\{\partial_e v(c, E) | E \text{ precedes } e \text{ in } \sigma(\theta)\} \quad (4)$$

To define our second rule we introduce the *null graph* of c : $\mathcal{N}(c) = \{e \in Q(2) | c_e = 0\}$, and $\mathcal{L}(c)$ the set of edges in $Q(2)$ linking two connected components of $\mathcal{N}(c)$. We can choose a minimal cost forest F spanning $[\theta]$ such that every edge of $F \cap \mathcal{L}(c)$ is in $[\theta]$, and every other edge is free (it may or may not be in $[\theta]$). In other words we can choose an efficient forest in which all inactive edges are free. For such a forest define the allocation z^F

$$z_e^F = 1 \text{ if } e \in F \cap \mathcal{L}(c); \;\; z_e^F = 0 \text{ if } e \in [\theta] \setminus \{F \cap \mathcal{L}(c)\}$$

We check that z^F is core stable. For any $E \subseteq F \cap \mathcal{L}(c)$ we have $v(c, E) = |E| = \sum_{e \in E} z_e^F$. For any $E \subseteq [\theta]$, monotonicity of $v(c,\cdot)$ implies

$$v(c, E) \geq v(c, E \cap F \cap \mathcal{L}(c)) = \sum_{e \in E \cap F \cap \mathcal{L}(c)} z_e^F = \sum_{e \in E} z_e^F$$

as desired. Averaging any number of allocations z^F respects core stability, and as for our first rule, an appropriate choice of their relative weights guarantees routing-proofness. Write $\mathcal{SP}(c, [\theta])$ for the set of efficient (minimal cost) forests F for the problem (Q, c, θ), such that $F \cap \mathcal{L}(c) \subseteq [\theta]$ (hence $c_e = 0$ for any $e \in F \setminus \mathcal{L}(c)$). We attach to every $F \in \mathcal{SP}(c, [\theta])$ the weight $\mu(F) = \prod_{e \in F \cap \mathcal{L}(c)} \theta_e$, and take the μ-average of the rules z^F.

Definition 4 *In a problem (Q, c, θ) where $c \in \mathcal{E}(Q)$, the* **weighted spanning rule** $z^{wsp}(c,\cdot)$ *is:*

$$z^{wsp}(c,\theta) = \frac{1}{\sum_{F \in \mathcal{SP}(c,[\theta])} \mu(F)} \sum_{F \in \mathcal{SP}(c,[\theta])} \mu(F) \cdot z^F \Leftrightarrow$$

$$z_e^{wsp}(c,\theta) = \begin{cases} \frac{\sum_{F \in \mathcal{SP}(c,[\theta]), e \in F} \mu(F)}{\sum_{F \in \mathcal{SP}(c,[\theta])} \mu(F)} & \text{for } e \in [\theta] \cap \mathcal{L}(c) \\ 0 & \text{for } e \in [\theta] \setminus \mathcal{L}(c) \end{cases} \quad (5)$$

By core stability the users of an edge e whose target nodes are in the same connected component of $\mathcal{N}(c)$, the null graph of c, pay nothing. Moreover the volume of traffic θ_e has no impact on any other cost share. Ignoring those users, a more intuitive interpretation of both rules (also helpful in many proofs) uses the *contracted multigraph* $\Delta(c,\theta)$, where we contract each component A_k of $\mathcal{N}(c)$ to a single node k, and keep between the nodes k, l as many edges as users with one target node in A_k and one in A_l. Formally the set of nodes of $\Delta(c,\theta)$ is $\overline{Q} = \{1, \cdots, K\}$, and we set $\mathcal{J}_{kl} = \{e \in Q(2) | e \text{ links } A_k \text{ to } A_l\}$. Then for any k, l the number of copies of edge $\overline{e} = kl$ is $\overline{\theta}_{\overline{e}} = \sum_{e \in \mathcal{J}_{kl}} \theta_e$.

It is enough to define cost shares $z_{\overline{e}}^{\varepsilon}(\overline{c}, \overline{\theta})$ in the contracted problem $(\overline{Q}, \overline{c}, \overline{\theta})$ where $\overline{c}$ is the constant cost matrix $\overline{c}_{\overline{e}} \equiv 1$. Then we recover the shares in the initial problem as follows

$$\{e \in \mathcal{J}_{kl} \text{ and } \overline{e} = kl\} \Rightarrow z_e^{\varepsilon}(c,\theta) = \frac{\theta_e}{\overline{\theta}_{\overline{e}}} z_{\overline{e}}^{\varepsilon}(\overline{c}, \overline{\theta}), \quad (6)$$

where $\varepsilon = wsh, wsp$. For the weighted spanning rule, $z_{\overline{e}}^{wsp}(\overline{c}, \overline{\theta})$ is now the probability that $\overline{e}$ belongs to the *uniform* random spanning forest of the multigraph $\Delta(c,\theta)$. Indeed the denominator in (5) is the number of spanning forests in the multigraph, and the numerator is the number of spanning forests containing a (necessarily unique) copy of edge $\overline{e}$.

Theorem 1 *For any problem (Q, c, θ) with $c \in \mathcal{E}(Q)$, the weighted Shapley rule ((4)) and the weighted spanning rule ((5)) are core stable and routing-proof. The latter is of polynomial complexity w.r.t. the dimension $|Q|$.*

The proof of routing-proofness is in Section 8. We show now that computing the weighted spanning rule is of polynomial complexity. By (6) we only need to consider the contracted multigraph $\Delta(c,\theta)$, with all connecting costs $\overline{c}_e \equiv 1$. For simplicity we write θ instead of $\overline{\theta}$ for the contracted traffic pattern, and Δ instead of $\Delta(c,\theta)$. We can restrict attention to a connected component of the graph $[\theta]$. Indeed if $[\theta]^1, \cdots, [\theta]^T$ are these components with corresponding traffic profiles $\theta^1, \cdots, \theta^T$, one checks easily

for all $e : z_e^{\varepsilon}(c,\theta) = z_e^{\varepsilon}(c,\theta^t)$ if $e \in [\theta]^t$ (where $\varepsilon = wsh, wsp$)

To compute the weighted spanning allocation for a *connected* graph $[\theta]$, we see from (5) that $1 - z_e^{wsp}(\overline{c}, \theta)$ is the ratio of the number of spanning trees of Δ not including e, to the number of spanning trees of Δ. Classic results ([3], see also [6]) state that the latter (resp. the former) number is any cofactor of the Lagrange matrix of Δ (resp. of $\Delta(-e)$ obtained by deleting all copies of e in Δ). Computing a determinant is of polynomial complexity in the matrix size.

By contrast, no algorithm to compute the weighted Shapley allocation is known to be polynomial.

We illustrate our two rules for a problem with three nodes $1, 2, 3$, and traffic $\theta_{12}, \theta_{23}, \theta_{13}$ such that $\theta_{23} > 0$. The weighted Shapley rule charges

$$z_{23}^{wsh} = \frac{\theta_{23}}{\theta_{12} + \theta_{23} + \theta_{13}} \left(1 + \frac{\theta_{12}}{\theta_{23} + \theta_{13}} + \frac{\theta_{13}}{\theta_{23} + \theta_{12}}\right) \quad (7)$$

The weighted spanning rule charges

$$z_{23}^{wsp} = \begin{cases} \frac{\theta_{23}(\theta_{12} + \theta_{13})}{\theta_{23}\theta_{12} + \theta_{23}\theta_{13} + \theta_{13}\theta_{12}} & \text{if } \theta_{12} + \theta_{13} > 0 \\ 1 & \text{if } \theta_{12} = \theta_{13} = 0 \end{cases} \quad (8)$$

Note that for both rules, if $\theta_e \to \infty$, while θ_{12}, θ_{13} remain fixed, then $z_e \to 1$. This is clearly a general property of the two rules for any size of Q and any traffic profile θ.

Remark. The case of a mcst problem. Assume that there is a distinguished node 1, the source, such that $[\theta] = \{1i | i = 2, \cdots, n\}$, and denote by A_1 the connected component of the source in the graph $\mathcal{N}(c)$ of null costs. It is easy to check that both rules charge nothing to an agent $i \in A_1$, and charge $\frac{1}{|A_k|}$ to an agent in A_k. This is precisely the folk solution discussed for instance in [5].

5. SPANNING TRAFFIC

We extend the two solutions, weighted Shapley and weighted spanning rules, to arbitrary cost matrices. We use the piecewise-linear extension technique developed in [5] for the mcst problem, and in [6] for the related capacity synthesis problem. This step requires that the traffic graph $[\theta]$ spans the entire network $Q(2)$. The construction is not straightforward, so we illustrate it first in the case of a three-node problem.

Here $[\theta]$ is spanning if and only if at most one of $\theta_{23}, \theta_{13}, \theta_{12}$ is zero, which we assume. Connecting costs are c_{23}, c_{13}, c_{12} and we assume $c_{23} \le c_{13} \le c_{12}$. Consider the canonical non negative linear decomposition of c:

$$\begin{pmatrix} c_{23} \\ c_{13} \\ c_{12} \end{pmatrix} = c_{23} \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} + (c_{13} - c_{23}) \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix} + (c_{12} - c_{13}) \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \quad (9)$$

For the problem with costs $(1,1,1)$ the cost shares z^{wsh} and z^{wsp} are given respectively by (7) and (8). If the costs are $(0,0,1)$ the three nodes can be connected for free so by core stability nobody is charged. If the costs are $(0,1,1)$, traffic on 23 pays nothing (its stand alone cost is 0), and we *contract* nodes $2,3$ into a single node $\overline{23}$. In the two-node problem between node 1 and node $\overline{23}$, the efficient cost \$1 is shared equally between all users. Note that if we charge more to those with target nodes $1,2$, the routing maneuver by these agents

$$\theta_{23}, \theta_{13}, \theta_{12} \to \theta_{23} + \theta_{12}, \theta_{13} + \theta_{12}, 0$$

lowers their cost. Hence for both rules the cost shares

$$\begin{pmatrix} z_{23} \\ z_{13} \\ z_{12} \end{pmatrix} ((0,1,1),\theta) = \begin{pmatrix} 0 \\ \frac{\theta_{13}}{\theta_{13}+\theta_{12}} \\ \frac{\theta_{12}}{\theta_{13}+\theta_{12}} \end{pmatrix}$$

Finally the cost shares in the initial problem c are the linear combination (9) of the cost shares in the three "elementary" problems:

$$\begin{aligned} z(c,\theta) = c_{23} z((1,1,1),\theta) + (c_{13} - c_{23}) z((0,1,1),\theta) \\ + (c_{12} - c_{13}) z((0,0,1),\theta) \quad (10) \end{aligned}$$

These cost shares are non negative and they cover exactly the minimal cost because

$$\begin{aligned} c_{23} + c_{13} &= v(c,\theta) \\ &= c_{23} v((1,1,1),\theta) + (c_{13} - c_{23}) v((0,1,1),\theta) \\ &\quad + (c_{12} - c_{13}) v((0,0,1),\theta) \quad (11) \end{aligned}$$

Therefore

$$z^{wsp} = c_{23} \begin{pmatrix} \frac{\theta_{23}(\theta_{12}+\theta_{13})}{\theta_{23}\theta_{12}+\theta_{23}\theta_{13}+\theta_{13}\theta_{12}} \\ \frac{\theta_{13}(\theta_{23}+\theta_{12})}{\theta_{23}\theta_{12}+\theta_{23}\theta_{13}+\theta_{13}\theta_{12}} \\ \frac{\theta_{12}(\theta_{13}+\theta_{23})}{\theta_{23}\theta_{12}+\theta_{23}\theta_{13}+\theta_{13}\theta_{12}} \end{pmatrix} + (c_{13} - c_{23}) \begin{pmatrix} 0 \\ \frac{\theta_{13}}{\theta_{13}+\theta_{12}} \\ \frac{\theta_{12}}{\theta_{13}+\theta_{12}} \end{pmatrix}$$

$$\Leftrightarrow z^{wsp} = \begin{pmatrix} \frac{\theta_{23}(\theta_{12}+\theta_{13})}{\sigma} c_{23} \\ \frac{\theta_{13}}{\theta_{13}+\theta_{12}} (c_{13} + \frac{(\theta_{12})^2}{\sigma} c_{23}) \\ \frac{\theta_{12}}{\theta_{13}+\theta_{12}} (c_{13} + \frac{(\theta_{13})^2}{\sigma} c_{23}), \end{pmatrix}$$

where $\sigma = \theta_{23}\theta_{12} + \theta_{23}\theta_{13} + \theta_{13}\theta_{12}$ and

$$\begin{aligned} z^{wsh} =& c_{23} \begin{pmatrix} \frac{\theta_{23}}{\theta_{12}+\theta_{23}+\theta_{13}} (1 + \frac{\theta_{12}}{\theta_{23}+\theta_{13}} + \frac{\theta_{13}}{\theta_{23}+\theta_{12}}) \\ \frac{\theta_{13}}{\theta_{12}+\theta_{23}+\theta_{13}} (1 + \frac{\theta_{23}}{\theta_{13}+\theta_{12}} + \frac{\theta_{12}}{\theta_{13}+\theta_{23}}) \\ \frac{\theta_{12}}{\theta_{12}+\theta_{23}+\theta_{13}} (1 + \frac{\theta_{13}}{\theta_{12}+\theta_{23}} + \frac{\theta_{23}}{\theta_{12}+\theta_{13}}) \end{pmatrix} \\ &+ (c_{13} - c_{23}) \begin{pmatrix} 0 \\ \frac{\theta_{13}}{\theta_{13}+\theta_{12}} \\ \frac{\theta_{12}}{\theta_{13}+\theta_{12}} \end{pmatrix} \end{aligned}$$

$$\Leftrightarrow z^{wsh} = \begin{pmatrix} \frac{\theta_{23}}{S} (1 + \frac{\theta_{12}}{\theta_{23}+\theta_{13}} + \frac{\theta_{13}}{\theta_{23}+\theta_{12}}) c_{23} \\ \frac{\theta_{13}}{\theta_{13}+\theta_{12}} c_{13} + \frac{\theta_{13}\theta_{12}}{S(\theta_{23}+\theta_{13})} c_{23} \\ \frac{\theta_{12}}{\theta_{13}+\theta_{12}} c_{13} + \frac{\theta_{13}\theta_{12}}{S(\theta_{23}+\theta_{12})} c_{23}, \end{pmatrix}$$

where $S = \theta_{23} + \theta_{13} + \theta_{12}$.

It is important to note that this construction fails if $[\theta]$ is not spanning, i.e., all traffic is concentrated on one edge. One way to see this is to check that if two of $\theta_{23}, \theta_{13}, \theta_{12}$ are zero, the formulas above involve $\frac{0}{0}$. A better explanation is the fact that if $\theta_{23} = \theta_{13} = 0$, the optimal cost $v(c,\theta) = \min\{c_{23} + c_{13}, c_{12}\}$ does not satisfy equation (11) as $v((1,1,1),\theta) = v((0,1,1),\theta) = v((0,0,1) = 1$. We say below that $c \to v(c,\theta)$ is not piecewise-linear.

The general construction of our two rules starts with the definition of piecewise-linear functions on $\mathbb{R}_+^{Q(2)}$. Label arbitrarily the $p = \frac{|Q|(|Q|-1)}{2}$ edges of $Q(2)$, so that a cost matrix $c \in \mathbb{R}_+^{Q(2)}$ takes the form $c = (c_{e^1}, ..., c_{e^p})$. For any permutation σ of $\{1,...,p\}$, define $K_\sigma = \{c \in \mathbb{R}_+^{Q(2)} : c_{e^{\sigma(1)}} \le c_{e^{\sigma(2)}} \le ... \le c_{e^{\sigma(p)}}\}$, the closed convex cone of cost matrices with the (weak) ordering σ of its coordinates. Note that $\bigcup_\sigma K_\sigma = \mathbb{R}_+^{Q(2)}$ and the only points that belong to more than one cone are frontier points, namely those for which at least one inequality in $c_{e^{\sigma(1)}} \le c_{e^{\sigma(2)}} \le ... \le c_{e^{\sigma(p)}}$ is an equality.

Each cone K_σ has a canonical basis made of elementary matrices. For any $k \in \{1,...,p\}$, define $\mathbf{b}^k \in \mathcal{E}(Q)$ by $\mathbf{b}^k_{e^{\sigma(1)}} = \mathbf{b}^k_{e^{\sigma(2)}} = ... = \mathbf{b}^k_{e^{\sigma(k-1)}} = 0$, while $\mathbf{b}^k_{e^{\sigma(k)}} = \mathbf{b}^k_{e^{\sigma(k+1)}} = ... = \mathbf{b}^k_{e^{\sigma(p)}} = 1$. In words, $\mathbf{b}^k$ is the zero-one cost matrix with cost 1 on all edges except the first $(k-1)$ in the permutation σ. We have for all $c \in \mathbb{R}_+^{Q(2)}$

$$c = \sum_{k=1}^{p} (c_{e^{\sigma(k)}} - c_{e^{\sigma(k-1)}})\, \mathbf{b}^k, \text{ with the convention } c_{e^{\sigma(o)}} = 0 \quad (12)$$

If $c \in K_\sigma$, and only then, this decomposition has non negative coordinates (an example is (9)). Then equation (12) admits a more compact integral formulation. Given $c \in \mathbb{R}_+^{Q(2)}$, and $t > 0$, define $c^{[t]} \in \mathcal{E}(Q)$ as follows

$$c^{[t]}_e = 0 \text{ if } c_e < t; c^{[t]}_e = 1 \text{ if } t \le c_e$$

The identity

$$c = \int_0^{\max_e c_e} c^{[t]} dt$$

is precisely (12) for the ordering(s) σ of the coordinates of c.

We call a real valued function $c \to f(c)$ on $\mathbb{R}_+^{Q(2)}$ **piecewise linear**, if it is linear in K_σ for all σ and continuous on $\mathbb{R}_+^{Q(2)}$. Any real valued function f^0 defined on $\mathcal{E}(Q)$ has a unique piece-wise linear extension to $\mathbb{R}_+^{Q(2)}$ given by

$$f(c) = \int_0^{\max_e c_e} f^0(c^{[t]})dt \text{ for all } c \quad (13)$$

Our two cost sharing rules z^{wsh}, z^{wsp} on $\mathcal{E}(Q)$ extend as follows:

$$z^{\varepsilon}(c,\theta) = \int_0^{\max_e c_e} z^{\varepsilon}(c^{[t]},\theta)dt \text{ for all } c, \text{ where } \varepsilon = wsh, wsp \quad (14)$$

(an example is (10)). As illustrated in the three-node example, this does not define a bona fide cost sharing rule for all traffic patterns.

Recall that the graph E spans $Q(2)$ if any two nodes in Q are connected by E, or equivalently if E contains a tree spanning Q. In this case we simply say that E is **spanning.**

Lemma 2 *If $[\theta]$ is spanning, equation (14) defines two cost sharing rules.*

Proof. We fix c and check $z^{\varepsilon} \in \mathbb{R}_+^{[\theta]}$ and $\sum_{[\theta]} z_e = v(c,[\theta])$. Each $z^{\varepsilon}(c^{[t]},\theta)$ is non negative, and integration respects this property. Budget balance follows from the fact that the function $c \to v(c,\theta)$ is piecewise-linear. Suppose $c \in K_{\sigma}$, then a minimal cost spanning tree of $Q(2)$ obtains by Prim's algorithm ([26]): build $e^{\sigma(1)}$ and $e^{\sigma(2)}$, then $e^{\sigma(3)}$ unless it forms a cycle with the previous two, and so on... Therefore in K_{σ}, $v(c,\theta) = v(c,Q(2))$ is the total cost of a fixed set of $|Q|-1$ edges $e^{\sigma(k)}$, a linear function of c. Continuity for all c is clear. Piecewise-linearity means that (13) is true with the same function $c \to v(c,\theta)$ on both sides. Therefore budget balance of $z^{\varepsilon}(c,\theta)$ on $\mathcal{E}(Q)$ and (14) give

$$\begin{aligned}\sum_{e\in[\theta]} z_e^{\varepsilon}(c,\theta) &= \int_0^{\max_e c_e} \{\sum_{e\in[\theta]} z_e^{\varepsilon}(c^{[t]},\theta)\}dt \\ &= \int_0^{\max_e c_e} v(c^{[t]},\theta)dt \\ &= v(c,\theta).\end{aligned}$$

∎

For the extension of the weighted Shapley value, we speak of the *pseudo-Shapley* solution, because it is no longer the Shapley value of a canonical cooperative game. Indeed $i)$ the Shapley value of the game $([\theta], v(c,\cdot))$ is not piecewise-linear in c because for a non spanning coalition E, $v(c,E)$ may not be piecewise linear; and $ii)$ the Shapley value of the auxiliary game $([\theta], \widetilde{v}(c,\cdot))$ (see step 1 in the proof of Theorem 2) *is* piecewise-linear in c, but it may charge a negative cost share to a player (which is ruled out in our model).

Theorem 2 *If the traffic graph $[\theta]$ is a spanning graph, the pseudo-Shapley and weighted spanning rules defined by (14) are core stable and routing-proof. The weighted spanning rule is of polynomial complexity in $|Q|$.*

Proof in section 8.

Remark. The case of a mcst problem. The traffic graph is clearly spanning in this case (see Remark 1). Both rules coincide with the folk solution for elementary matrices, and the latter is piecewise-linear ([5]). Therefore both rules coincide with the folk solution for any cost matrix.

6. TWO MONOTONICITY PROPERTIES

On both domains where we define them, our two cost sharing rules satisfy additional desirable properties. The first one is **continuity** of the mappings $c \to z(c,\theta)$ for all c. This is important because arbitrarily small changes in the connection costs can change the minimal cost spanning forest. In fact some of the mcst cost sharing rules in the literature are not continuous (see e.g. [12]).

Next we have **Population Monotonicity:** if one or several new users are added to the network, the charge of each incumbent user weakly decreases. Formally

$$\text{for all } \theta,\theta'\colon\ \theta \le \theta' \Rightarrow \frac{1}{\theta_e} z_e(c,\theta) \ge \frac{1}{\theta'_e} z_e(c,\theta') \text{ for all } e \in [\theta] \quad (15)$$

This property has important incentives consequences. Suppose users can walk away from the network if the charge exceeds their benefit from connecting their two target nodes. Then the following direct revelation mechanism is strategyproof, even group-strategyproof: each user reports his benefit and the PM rule serves the largest subset of users whose benefit exceeds their cost share. See [23].

The last property is normatively less compelling than PM, but it has incentives implications of its own. **Cost Solidarity** states that when the cost of any edge increases, the charge to *all* users on *all* edges weakly increases:

$$\text{for all } c, c'\colon\ c \le c' \Rightarrow z(c,\theta) \le z(c',\theta)$$

Thus all participants are motivated to seek cost savings anywhere in the network. See also the role of Cost Solidarity to characterize the folk solution in [2].

Proposition 1 *The pseudo-Shapley and weighted spanning rules meet Population Monotonicity and Cost Solidarity for all problems where $c \in \mathcal{E}(C)$. For problems (Q,c,θ) where θ is spanning, these two rules are continuous in c; they meet Cost Solidarity, and Population Monotonicity.*

The proof is available in [22].

Another consequence of adding new users to the network was mentioned in the introduction: if more users demanding e show up, the *total* charge z_e^{ε} to the users of e *in*creases weakly (whereas the charge to an individual user weakly decreases by PM). Direct inspection of Definitions 1 and 2 proves the claim for elementary matrices, and the inequality is preserved by piecewise-linear extension.

7. AN IMPOSSIBILITY RESULT

Suppose routing maneuvers may involve coalitions of users with different target nodes, who can in addition use side-payments. The uniform cost sharing rule (dividing the efficient cost equally among all users) is not vulnerable to such moves, but any core stable rule is. This holds true even if all connecting costs are identical.

Proposition 2 *All core stable cost sharing rules are vulnerable to cooperative routing maneuvers involving users with different traffic demands and side payments. This is true even for elementary cost matrices.*

Proof. Fix three nodes with identical connecting costs $c_e \equiv 1$. Pick an arbitrary cost sharing rule $\theta \to z(\theta)$, where we write $\theta = (\theta_{23}, \theta_{13}, \theta_{12})$. Consider the traffic profile $\theta = (5,1,1)$. A coalition of four users: the ones with target pairs $1,3$ and $1,2$, and two of the 5 users with target nodes $2,3$, contemplates posing as two agents demanding $1,3$, and two demanding $1,2$. This move is profitable, given appropriate side-payments, if we have

$$(z_{13}+z_{12})(3;3,3) < (\frac{2}{5}z_{23}+z_{13}+z_{12})(5;1,1)$$

$$\Leftrightarrow 2-z_{23}(3;3,3) < 2-\frac{3}{5}z_{23}(5;1,1)$$

To rule out such move we need

$$\frac{1}{5}z_{23}(5,1,1) \geq \frac{1}{3}z_{23}(3,3,3)$$

Core stability implies $z_{23}(5,1,1) \leq 1$, so that $z_{23}(3,3,3) \leq \frac{3}{5}$. This and two similar inequalities for z_{13} and z_{12} contradicts budget-balance $z_{23}+z_{13}+z_{12}=2$. ∎

8. PROOFS

8.1 Theorem 1: routing-proofness

Step 1. Weighted Shapley rule.
By (6) and the discussion preceding Theorem 1, it is enough to prove the result in the contracted problem, i.e., we can assume that all connecting costs are $\overline{c}_e \equiv 1$. Indeed a user with a target pair inside a connected component of the null graph of c pays nothing, and so he has no incentive to reroute. Moreover the number of such users has no impact on anyone's cost share.

Following a remark due to [19], in this proof we find it useful to recover the weighted Shapley value $z_{\overline{e}}^{wsh}(\overline{c},\overline{\theta})$ from the *ordinary* Shapley value of the game $(N,\widehat{v})$ in which N is the set of edges in $\Delta(c,\theta)$, with cardinality $\sum_{Q(2)}\theta_e$, and for all $S \subseteq N$:

$$\widehat{v}(S)=v(\overline{c},\overline{S}) \text{ where } \overline{S}=\{\overline{e}\in\overline{Q}(2)|S \text{ contains a copy of } \overline{e}\}$$

If y is the Shapley value of $(N,\widehat{v})$, we have

$$\{e\in\mathcal{J}_{kl} \text{ and } a\in N \text{ connects } A_k \text{ and } A_l\}$$
$$\Rightarrow y_a=\frac{1}{\theta_e}z_e^{wsh}(c,\theta)=\frac{1}{\theta_{\overline{e}}}z_e^{wsh}(\overline{c},\overline{\theta}) \quad (16)$$

Fix $e,e^1,\cdots,e^K,\theta,\theta'$ as in the premises of (3), and assume first that only one user with demand e reroutes ($x=1$). We must prove that if $a,a^1,\cdots,a^K$ are corresponding edges in the multigraphs $\Delta(\overline{c},\theta)$ and $\Delta(\overline{c},\theta')$, we have

$$y_a(\theta) \leq \sum_{k=1}^{K} y_{a^k}(\theta') \quad (17)$$

We write θ^* for the multigraph with one less copy of e than $\theta: \theta^*_e=\theta_e-1, \theta^*_{e'}=\theta_{e'}$ else. We pick an arbitrary ordering σ of the set N^* of edges in $\Delta(\overline{c},\theta^*)$. It is enough to prove (17) conditional on the ordering σ of N^*. Write the corresponding expected cost shares as $y_a(\theta|\sigma), y_{a^k}(\theta'|\sigma)$, where the expectation bears on the ranking of a (or $a^1,\cdots,a^K$) w.r.t. the edges in N^*. Also let S_t be the set of the first t edges in N^* according to σ, and $t(e)$, resp. $t(e^k)$, be the first index such that the graph S_t spans (connects the endnodes of) e, resp. e^k. Then $y_a(\theta|\sigma)$ is simply the probability that a is ranked before the edge ranked $t(e)$ in σ. There are $n^*=|N^*| \overset{def}{=} \sum_{[\theta]}\theta_e-1$ edges in N^* therefore

$$y_a(\theta|\sigma)=\frac{t(e)}{n^*+1} \quad (18)$$

Now consider an arbitrary ranking of $N^*\cup\{a^1,\cdots,a^K\}$ compatible with σ. We set $t^*=\max_k t(e^k)$ and $R=\arg\max_k t(e^k) \subset (1,\cdots,K\}$. Let a^{k^*} be ranked first among the $a^k, k\in R$. Suppose that a^{k^*} is also ranked *before* the edge ranked t^* in N^*. Then the graph formed by the set of edges E preceding a^{k^*} in $N^*\cup\{a^1,\cdots,a^K\}$ cannot span more edges than S_{t^*-1}, therefore it does not span e^{k^*}. This implies $\partial_{a^{k^*}}v^*(E)=1$, i.e., the edge a^{k^*} is charged. Thus the set of rankings of $N^*\cup\{a^1,\cdots,a^K\}$ compatible with σ is such that: if $\sum_{k=1}^{K}y_{a^k}=0$, *all* the edges a^k, $k\in R$, are ranked *after* the edge ranked t^* in N^*; otherwise we have $\sum_{k=1}^{K}y_{a^k}\geq 1$. The probability of all $a^k,k\in R$ ranked after t^* is no more than $\frac{n^*+1-t^*}{n^*+1}$, therefore

$$\sum_{k=1}^{K}y_{a^k}(\theta'|\sigma) \geq 1-\frac{n^*+1-t^*}{n^*+1}=\frac{t^*}{n^*+1} \quad (19)$$

The path $e^1,\cdots,e^K$ connects the end nodes of e, therefore S_{t^*} spans e as well, implying $t(e)\leq t^*$. Now (18) and (19) imply (17) conditional on σ, as desired.

The proof when x players on e jointly reroute through $e^1,\cdots,e^K$ is entirely similar. We set θ^*: $\theta^*_e=\theta_e-x, \theta^*_{e'}=\theta_{e'}$ else, and pick an ordering σ of N^*. For any ordering of $N^*\cup\{x \text{ copies of } a\}$ compatible with σ, the total charge to the x deviant players is 1 if at least one of the x copies of a is ranked before the edge ranked $t(e)$ in σ, and 0 otherwise. Therefore:

$$x\cdot y_a(\theta|\sigma)=1-(\frac{n^*+1-t(e)}{n^*+1})^x$$

where $n^* \overset{def}{=} \sum_{[\theta]}\theta_e-x$. Similarly, in any ranking of $N^*\cup\{x$ copies of $a^1,\cdots,a^K\}$ compatible with σ, the total charge to the $x\cdot K$ deviant players is zero only if *all* copies of the edges a^k, $k\in R$, are ranked *after* the edge ranked t^* in N^*; otherwise this total charge is no less than 1. Therefore:

$$x\cdot\sum_{k=1}^{K}y_{a^k}(\theta'|\sigma)\geq 1-(\frac{n^*+1-t^*}{n^*+1})^x$$

Step 2. Weighted spanning rule
We work again, for the same reasons, in the contracted multigraph where all connecting costs are 1.
Step 2.1. Preliminary results. We start with some notation. We write $\delta(\theta)$ for the number of subforests of the multigraph θ spanning $[\theta]$. If $\theta^1,\cdots,\theta^T$ are the (multigraphs) connected components of θ, we have $\delta(\theta)=\prod_1^T\delta(\theta^t)$, where $\delta(\theta^t)$ is the number of subtrees of θ^t spanning $[\theta]^t$. If θ' is a submultigraph of θ (i.e., $\theta'\leq\theta$), we write $\delta(\theta;\theta')$ for the number of subforests of θ' spanning $[\theta]$. For any multigraph θ and edge e, $(\theta|^e y)$ is the multigraph with y copies of edge e, and otherwise identical to θ. Finally we call edge e *critical* in $[\theta]$ if $e\in[\theta]$ and every spanning forest of $[\theta]$ contains e. Thus e is critical *iff* $\delta(\theta;(\theta|^e 0))=0$.

the graph $[\theta]\diagdown\{e\}$ has one less connected component than $[\theta]$.

The proof of the following identity, for any multigraph θ and edge e (where θ_e may be zero)

$$\delta(\theta)=\delta(\theta;(\theta|^e 0))+\theta_e(\delta(\theta|^e 1)-\delta(\theta;(\theta|^e 0))) \quad (20)$$

is left to the reader.

As explained after Definition 4, in equation (5) defining $z_e^{wsp}(\theta)$ the denominator is $\delta(\theta)$ and the numerator is $\delta(\theta)-\delta(\theta;(\theta|^e 0))$, namely the number of spanning forests of θ containing a copy of e:

$$z_e^{wsp}(\theta)=\frac{\delta(\theta)-\delta(\theta;(\theta|^e 0))}{\delta(\theta)} \tag{21}$$

By the decomposition property (6) and $(\prod_{s\neq t}\delta(\theta^s))=\frac{\delta(\theta)}{\delta(\theta^t)}=\frac{\delta(\theta;(\theta|^e 0))}{\delta(\theta^t;(\theta^t|^e 0))}$, this equation remains true if we restrict attention to a connected component of $[\theta]$. Moreover (20) implies

$$z_e^{wsp}(\theta)=\theta_e\frac{\delta(\theta|^e 1)-\delta(\theta;(\theta|^e 0))}{\delta(\theta)} \tag{22}$$

Next we fix $\theta, e\in[\theta]$ and $x\leq\theta_e$, and prove the following fact

$$\frac{x}{\theta_e}z_e^{wsp}(\theta)=\frac{\delta(\theta)-\delta(\theta|^e\theta_e-x)}{\delta(\theta)}$$
$$\text{if } \{x<\theta_e\} \text{ or } \{x=\theta_e; e \text{ is not critical in } \theta\} \tag{23}$$

If $x<\theta_e$ we have $\delta((\theta|^e\theta_e-x);(\theta|^e 0))=\delta(\theta;(\theta|^e 0))$ so that (20) implies:

$$\delta(\theta|^e\theta_e-x)-\delta(\theta;(\theta|^e 0))=(\theta_e-x)\{\delta(\theta|^e 1)-\delta(\theta;(\theta|^e 0))\}$$

Combining this equality with (20) gives

$$\delta(\theta)-\delta(\theta|^e\theta_e-x)=x\{\delta(\theta|^e 1)-\delta(\theta;(\theta|^e 0))\} \tag{24}$$

and (23) follows. Assume next $x=\theta_e$ but e is not critical in θ. Then $\delta(\theta|^e 0)=\delta(\theta;(\theta|^e 0))$; together with (21) this gives (23).

Step 2.2. We fix $e,e^1,\cdots,e^K,x,\theta,\theta'$ as in the premises of (3). We distinguish two cases.

Case 1: e is critical in θ and $x=\theta_e$. Then every spanning forest of θ contains a copy of e therefore $z_e^{wsp}(\theta)=\frac{x}{\theta_e}z_e^{wsp}(\theta)=1$. Moreover the set $I=\{k\in\{1,\cdots,K\}|\theta_{e^k}=0\}$ is non empty, else e is not critical in θ. Finally in the graph $[\theta']\diagdown\{e^k,k\in I\}$, the end nodes of e are not connected, therefore in every spanning forest of $[\theta']$, there is at least one copy of some $e^k, k\in I$. This implies $\sum_{k\in I}z_{e^k}^{wsp}(\theta')\geq 1$, hence the desired inequality (3) because $\theta'_{e^k}=x$ for all $k\in I$.

Case 2. $x<\theta_e$ and/or e is not critical in θ. Now (23) holds for $z_e^{wsp}(\theta)$ and for $z_{e^k}^{wsp}(\theta')$ we have

$$\frac{\delta(\theta')-\delta(\theta'|^{e^k}\theta_{e^k})}{\delta(\theta')}\leq\frac{x}{\theta'_{e^k}}z_{e^k}^{wsp}(\theta')$$

Indeed this is equation (23) applied to θ' if $\theta_{e^k}=\theta'_{e^k}-x>0$ and/or e^k is not critical in θ'; and if $\theta'_{e^k}=x$ and e^k is critical in θ', the RHS is 1.

Thus (3) is true if we show

$$\frac{\delta(\theta)-\delta(\theta|^e\theta_e-x)}{\delta(\theta)}\leq\sum_{k=1}^{K}\frac{\delta(\theta')-\delta(\theta'|^{e^k}\theta_{e^k})}{\delta(\theta')}$$

$$\Leftrightarrow\frac{1}{\delta(\theta')}\sum_1^K\delta(\theta'|^{e^k}\theta_{e^k})\leq K-1+\frac{\delta(\theta|^e\theta_e-x)}{\delta(\theta)} \tag{25}$$

Set $\theta^*=(\theta|^e\theta_e-x)$ and $\alpha=\delta(\theta|^e 1)-\delta(\theta;(\theta|^e 0))$. Note that α is the number of subforests F of $(\theta|^e 0)$ (equivalently, of θ^*) such that adding e to F makes it a spanning forest of θ. In view of (24) the RHS of (25) is

$$K-1+\frac{\delta(\theta^*)}{\delta(\theta^*)+x\alpha} \tag{26}$$

We evaluate now the LHS of (25). For any $\omega\subseteq\{1,\cdots,K\}$, including $\omega=\varnothing$, let $\gamma(\omega)$ be the number of subforests G of θ^* such that adding to G the edges $e^k,k\in\omega$, makes it a spanning forest of θ'. In particular $\gamma(\varnothing)=\delta(\theta';\theta^*)$. Distinghishing among spanning forests of θ' by the subsets of $e^1,\cdots,e^K$ they contain, we have:

$$\delta(\theta')=\gamma(\varnothing)+\sum_{t=1}^{K}x^t\sum_{\omega:|\omega|=t}\gamma(\omega)$$

and for all $k=1,\cdots,K$

$$\delta(\theta'|^{e^k}\theta_{e^k})=\gamma(\varnothing)+\sum_{t=1}^{K-1}x^t\sum_{\omega:|\omega|=t,k\notin\omega}\gamma(\omega)$$

$$\Rightarrow\sum_1^K\delta(\theta'|^{e^k}\theta_{e^k})=K\gamma(\varnothing)+\sum_{t=1}^{K}(K-t)x^t\sum_{\omega:|\omega|=t}\gamma(\omega)$$

$$\Rightarrow\sum_1^K\delta(\theta'|^{e^k}\theta_{e^k})\leq K\gamma(\varnothing)+(K-1)\sum_{t=1}^{K}x^t\sum_{\omega:|\omega|=t}\gamma(\omega)$$

$$\Rightarrow\frac{1}{\delta(\theta')}\sum_1^K\delta(\theta'|^{e^k}\theta_{e^k})\leq K-1+\frac{\gamma(\varnothing)}{\gamma(\varnothing)+\sum_{t=1}^{K}x^t\sum_{\omega:|\omega|=t}\gamma(\omega)}$$

Clearly $\gamma(\varnothing)=\delta(\theta';\theta^*)\leq\delta(\theta^*)$, therefore

$$\frac{1}{\delta(\theta')}\sum_1^K\delta(\theta'|^{e^k}\theta_{e^k})\leq K-1+\frac{\delta(\theta^*)}{\delta(\theta^*)+\sum_{t=1}^{K}x^t\sum_{\omega:|\omega|=t}\gamma(\omega)}$$

Comparing with the RHS (26), it remains to check

$$x\alpha\leq\sum_{t=1}^{K}x^t\sum_{\omega:|\omega|=t}\gamma(\omega) \tag{27}$$

where we only need to look at the case $x=1$. Recall that α is the number of subforests G of θ^* such that $G\cup\{e\}$ is a spanning forest of θ. Consider such a G: it is not spanning in θ', else it would span e as well, because the end nodes of e are connected in θ', and $G\cup\{e\}$ would not be a forest. On the other hand $G\cup\{e^1,\cdots,e^K\}$ spans θ', though it may not be a forest. Hence there is a non empty subset $\omega\subseteq\{1,\cdots,K\}$ such that $G\cup\{e^k,k\in\omega\}$ is a spanning forest of θ': therefore G is counted in the term $\gamma(\omega)$ and the inequality (27) follows.

8.2 Theorem 2

Step 1 Preliminary definition and property (31).

Given a graph $E\subseteq Q(2)$, we write $\widehat{E}$ for the set of edges e of which the end-nodes are connected in E. With the terminology of section 3, $\widehat{E}$ is the largest graph F such that E spans F. Thus $E\subseteq\widehat{E}$, and $\widehat{\widehat{E}}=\widehat{E}$; moreover $F\subseteq\widehat{E}\Leftrightarrow\widehat{F}\subseteq\widehat{E}$. Recall that E is spanning if $\widehat{E}=Q(2)$, or equivalently if E contains a tree spanning Q.

Write $Core(c,E)$ for the core of the cooperative game $(E,v(c,\cdot))$ where players are the edges of E. It is defined by

(2) where E replaces $[\theta]$. Recall from section 1 that it may be empty. Note that $Core(c,E)$ is embedded in $Core(c,\widehat{E})$ by adding zero coordinates on $\widehat{E}\diagdown E$:

$$z \in Core(c,E) \Rightarrow (z, 0_{[\widehat{E}\diagdown E]}) \in Core(c,\widehat{E}) \qquad (28)$$

This follows at once from the identity $v(c,E) = v(c,\widehat{E})$. The converse property is also true, for the same reason. An allocation in $Core(c,\widehat{E})$ that charges only the edges in E, projects into an allocation of $Core(c,E)$:

$$\{z \in Core(c,\widehat{E}) \text{ and } \{z_e = 0 \text{ for all } e \in \widehat{E}\diagdown E\}\} \Rightarrow z_{[E]} \in Core(c,E) \qquad (29)$$

where $z_{[E]}$ is the projection of $z \in \mathbb{R}_+^{\widehat{E}}$ into $\mathbb{R}_+^{E}$.

In the **auxiliary game** $(E, \widetilde{v}(c,\cdot))$, a coalition $F \subseteq E$ can only use edges in $\widehat{F}$ to cover its comunication needs:

$$\widetilde{v}(c,F) = \min\{c(\Gamma)|\widehat{\Gamma} = \widehat{F}\} \geq v(c,F) = \min\{c(\Gamma)|\widehat{\Gamma} \supseteq F\}$$

We have for any $F \subseteq Q(2)$

$$\widehat{F} = Q(2) \Rightarrow \widetilde{v}(c,F) = v(c,F) \qquad (30)$$

Indeed we saw earlier that $\widehat{\Gamma} \supseteq F$ implies $\widehat{\Gamma} \supseteq \widehat{F}$, therefore if $\widehat{F} = Q(2)$, $\widehat{\Gamma} \supseteq F$ implies $\widehat{\Gamma} = \widehat{F}$.

Write $\widetilde{Core}(c,E)$ for the core of the cooperative game $(E, \widetilde{v}(c,\cdot))$. Clearly $Core(c,E) \subseteq \widetilde{Core}(c,E)$. The opposite inclusion is not true in general, except in the case where E is a spanning graph:

$$\widehat{E} = Q(2) \Rightarrow \widetilde{Core}(c,E) = Core(c,E) \qquad (31)$$

We prove (31) first for $E = Q(2)$. We pick $z \in \widetilde{Core}(c,Q(2))$ such that $z \notin Core(c,Q(2))$ and derive a contradiction. As $\widetilde{v}(c,Q(2)) = v(c,Q(2))$ ((30)), there must exist $E \subset Q(2)$ such that $\sum_E z_e > v(c,E)$. Choose a forest F such that $v(c,E) = c(F)$ and $\widehat{F} \supseteq E$. From $\widehat{\widehat{F}} = \widehat{F}$ and the definition of $\widetilde{v}$, we have $\widetilde{v}(c,\widehat{F}) \leq c(F)$; and $z \in \widetilde{Core}(c,Q(2))$ gives $\sum_{\widehat{F}} z_e \leq \widetilde{v}(c,\widehat{F})$. Combining these inequalities and equality we get $\sum_E z_e > \sum_{\widehat{F}} z_e$, a contradiction of $E \subseteq \widehat{F}$ and $z \geq 0$.

Now choose E such that $\widehat{E} = Q(2)$. Properties (28) and (29) together mean that $Core(c,E)$ is the projection on $\mathbb{R}_+^E$ of $Core(c,Q(2)) \cap K$, where $K = \{z \in \mathbb{R}_+^{Q(2)} | z_e = 0$ for $e \in Q(2)\diagdown E\}$. One checks easily that $\widetilde{Core}(c,E)$ is similarly the projection of $\widetilde{Core}(c,Q(2)) \cap K$. This concludes the proof of (31).

Step 2 Core stability of z^{wsh}, z^{wsp}

To show $z^\varepsilon(c,\theta) \in Core(c,[\theta])$, we cannot extend the inequalities $\sum_{e\in E} z_e^\varepsilon(c^{[t]},\theta) \leq v(c^{[t]},E)$ by piece-wise linearity because for proper subsets E of $[\theta]$, $v(c,E)$ is typically *not* piece-wise linear (as in the three node example in section 5).

On the other hand, the function $c \to \widetilde{v}(c,E)$ *is* piecewise linear for any fixed $E \subset Q(2)$, because the cheapest spanning trees (or forests) of $\widehat{E}$ are known once we know the ordering of the costs $c_e, e \in \widehat{E}$. Therefore (13) and $z^\varepsilon(c^{[t]},\theta) \in \widetilde{Core}(c,[\theta])$ (Theorem 1) imply $z^\varepsilon(c,\theta) \in \widetilde{Core}(c,[\theta])$:

$$\{\text{for all } t: \sum_{e\in E} z_e^\varepsilon(c^{[t]},\theta) \leq v(c^{[t]},E) \leq \widetilde{v}(c^{[t]},E)\}$$
$$\Rightarrow \sum_{e\in E} z_e^\varepsilon(c,\theta) \leq \widetilde{v}(c,E)$$

The desired conclusion follows by (31).

Step 3 End of proof

For routing-proofness, it is enough to check that inequality (3) is preserved by piecewise-linear combinations.

For the polynomial complexity of the weighted spanning rule, observe that the canonical decomposition (12) is itself of polynomial complexity (see [6] for details).

9. REFERENCES

[1] Anshelevich, E., Dasgupta, J., Kleinberg, J., Tardos, E., Wexler, T., and T. Roughgarden, The price of anarchy for network design with fair cost allocation, *Proceedings IEEE Symposium on foundations of Computer Science, 2004, pp 295-304.*

[2] Bergantiños, G.and J.J. Vidal-Puga, A Fair rule in minimum cost spanning tree problems, forthcoming, *Journal of Economic Theory 2008.*

[3] Berge, C., *Graphs and Hypergraphs*, North-Holland, 1973.

[4] Bird, C. J., On cost allocation for a spanning tree: a game theoretic approach. *Networks 6 (1976), 335-350.*

[5] Bogomolnaia, A., and H. Moulin, Sharing a minimal cost spanning tree: beyond the Folk solution, *mimeo, Rice university, September 2008.*

[6] Bogomolnaia, A., Holzman, R. and H. Moulin, Sharing the cost of a capacity network, *mimeo, Rice University and Technion, 2008.*

[7] Bollobas, B., *Random Graphs,* Cambridge University press, 2001.

[8] Cheng, A. and E. Friedman, 2005. Sybilproof Reputation Mechanisms, *SIGCOMM'05 Workshops* August 22-26, Philadelphia, USA.

[9] Conitzer, V., Anonymity-proof voting rules, *Fourth Workshop on Internet and Network Economics (WINE-08), Shanghai, 2008.*

[10] Dong, B., Gou, G., and Y. Wang, Highway toll pricing, *mimeo, University of windsor, October 2008.*

[11] Douceur, J., 2002. The Sybil Attack, *Proceedings of the 1st International Workshop on Peer-to-Peer Systems* (IPTPS 2002), LNCS 2429, pp 251-260.

[12] Dutta, B. and A. Kar, Cost monotonicity, consistency, and minimum cost spanning tree games, *Games Econ. Behav., 48 (2004), 223-248.*

[13] Epstein, A., Feldman, M., and Mansour, Y., Strong Equilibrium in cost sharing connection games, *ACM conference on Electronic Commerce, San Diego, (2007).*

[14] Feigenbaum, J., Papadimitriou, C., and S. Shenker, Sharing the cost of multicast transmission, *Journal of Computing and System Science, 63 (2001), 21-41.*

[15] Feltkamp, T., Tijs, S. and S. Muto, On the irreducible core and the equal remaining obligations rule of minimum cost spanning extension problems, *mimeo University of Tilburg, CentER DP 94106, 1994.*

[16] Granot, D. and G. Huberman, 1984, On the core and nucleolus of minimum cost spanning tree games, *Mathematical Programming 29, 323-347.*

[17] Henriet, D. and H. Moulin, Traffic Based Cost Allocation in a Network, *Rand Journal of Economics, 27, 2 (1996), 332–345.*

[18] Herzog, S., Shenker, S., and D. Estrin, Sharing the

cost of multicast trees: an axiomatic analysis, *IEEE/ACM Transactions on Networking, December 1997.*

[19] Kalai, E., and D. Samet, On weighted Shapley values, *International Journal of Game Theory, 16 (1987), 205-222.*

[20] Megiddo, N., Cost allocations for Steiner trees, *Networks, 8 (1978), 1-6.*

[21] Moulin, H. ,On scheduling fees to prevent merging, splitting and transferring of jobs, *Mathematics of Operations Research", 32, 2 (2007), 266-283.*

[22] Moulin, H. , Pricing Traffic in a Spanning Network: long version, http://www.ruf.rice.edu/~econ/faculty/Moulin/routingpf2F.pdf

[23] Moulin, H. and S. Shenker, Strategyproof Sharing of Submodular Costs: Budget Balance versus Efficiency, *Economic Theory*, 18, 3, 511-533, 2001.

[24] Nisan, N., Roughgarden, T., Tardos, E., and V. Vazirani (Editors), *Algorithmic Game Theory,* Cambridge U. Press, 2007.

[25] Norde, H., Moretti, S. and S. Tijs, Minimum cost spanning tree games and population monotonic allocation schemes, *European Journal of Operations Research 154, 1 (2001), 84-97.*

[26] Prim, R. C., Shortest connection network and some generalization, *Bell System Technical Journal 36 (1957), 1389-1401.*

[27] Roughgarden, T., *Selfish Routing And The Price Of Anarchy*, MIT Press, 2005.

[28] Sharkey, W. W., Network models in economics, in: *Ball, M. O., Magnanti, T. L., Nonma, C. L., Nemhauser, G. L. (Eds.), Handbooks in Operation Research and Management Science, Elsevier, New York, pp. 713-765, 1995.*