

Trouble Comes in Threes: Core stability in Minimum Cost Connection Networks*

Jens Leth Hougaard^{†‡}

Mich Tvede[§]

Abstract

We consider a generalization of the Minimum Cost Spanning Tree (MCST) model, called the Minimum Cost Connection Network (MCCN) model, where network users have connection demands in the form of a pair of nodes they want connected directly or indirectly. For a fixed network, which satisfies all connection demands, the problem consists of allocating the total cost of the network among its users. Thereby every MCCN problem induces a cooperative cost game where the cost of every coalition of users is the cost of an efficient network satisfying the demand of the users in the coalition. Unlike the MCST-model, we show that the core of the induced cost game in the MCCN-model can be empty even when all locations are demanded. We therefore consider sufficient conditions for non-empty core. It is shown that: when the efficient network and the demand graph (i.e. the graph consisting of the direct connections between the pairs of demanded nodes) consist of the same components, the induced cost game has non-empty core (Theorem 1); and, when the demand graph consists of at most two components, the induced cost game has non-empty core (Theorem 2).

Keywords: Game Theory; Minimum Cost Connection Network; Spanning Tree; Cost Sharing; Fair allocation.

JEL Classification: C70, C72, D71, D85.

*We wish to thank three anonymous referees as well as Peter Sudhölter and audience at the GEM Workshop at University of Southern Denmark for comments on earlier versions of the paper. Hougaard is grateful for financial support from the Center for Blockchains and Electronic Markets funded by the Carlsberg Foundation (CF18-1112).

[†]University of Copenhagen, Rolighedsvej 25, DK-1958 Frederiksberg C, Denmark, email: jlh@ifro.ku.dk

[‡]NYU Shanghai, China, email: jlh21@nyu.edu

[§]University of East Anglia, Norwich Research Park, Norwich, NR4 7TJ, United Kingdom, email: m.tvede@uea.ac.uk

1 Introduction

The minimum cost spanning tree (MCST) model is a staple in combinatorial optimization with numerous papers dedicated to theoretical analysis as well as practical applications (see e.g. Korte and Vygen, 2018, Sharkey, 1995). A set of agents, with each agent being identified by a node (location), all want to be connected to the same target node (the source). Connections are costly to build and congestion free. The first problem consists of finding a cost minimizing network (a minimum cost spanning tree) that connects all agents to the source, directly or indirectly. The second problem consists of how to allocate the total cost of the efficient network between the agents (Claus and Kleitman, 1973, Bird, 1976).

Bird (1976) uses a conventional approach to fair allocation, mapping the MCST-model to a cooperative game where the cost of every coalition of agents S is the minimum cost of connecting all members of S to the source node. With the model represented by a cooperative game we can now apply solution concepts from cooperative game theory as justifiable means of allocating the common cost. Specifically, Bird (1976) proves that all MCST-games are balanced (i.e. have non-empty core) so it is possible to allocate the total cost such that no coalition of agents will be charged more than the minimum cost of satisfying their connection demands. Therefore, cost allocations in the core are stable in that no coalition of agents would like to break out and build their own network. In contrast, cost allocations not in the core are unstable because some coalitions of agents would like to break out and build their own networks. For the MCST-model it is well-known that the core is non-empty and that the Shapley value does not need to be in the core.

In the present paper, we study a recent generalization of the MCST-model, introduced in Moulin (2009, 2014) and named the Minimum Cost Connection Network (MCCN) model in Houggaard and Tvede (2015), where every agent wants to have an arbitrary pair of target nodes connected directly or indirectly. This seemingly minor modification of the model has radical consequences both for finding a cost minimizing network (which now becomes an NP-hard problem)¹ as well as for the ability to sustain the efficient network by suitable pricing of network usage. We focus on the latter issue related to balancedness of the induced cooperative game: by the Bondareva–Shapley theorem, a cooperative game has non-empty core if and only if it is balanced (see e.g. Peleg and Sudhölter, 2007).

Since the set of agents can no longer be identified by the set of nodes, the MCCN-model has to specify a set of agents as well as a set of nodes (or locations). In some cases each node will be a target node of some agent, while in other cases the model may include nodes that are not among the target nodes of any agent. These nodes, that are not target nodes of any agent, are called undemanded (or Steiner nodes). The so-called Minimum Cost Steiner Tree

¹See e.g. Karp (1972) showing that the special case of a minimum cost Steiner tree problem is NP-hard.

model is another generalization of the MCST-model where nodes that are neither locations of agents nor target nodes are added to the MCST model. These additional nodes can be seen as potential “hubs” that agents can use to obtain connection to the source node. Since the examples in Meggido (1978) and Tamir (1991), it has been well-known that the cost game induced by the Minimum Cost Steiner Tree model can have empty core with as few as three Steiner nodes.² Since the Minimum Cost Steiner Tree model is a special case of the MCCN-model, we can immediately infer that if there are undemanded nodes in the MCCN-model, then the MCCN-induced cost game can have empty core. In other words, the presence of undemanded nodes can be a source of instability in efficient network design.

We therefore ask whether the MCCN induced cost game has non-empty core on the restricted domain of problems where all nodes are demanded so every node is a target node of some agent. By a simple three-agent example we demonstrate that even if all nodes are demanded we can still end up in a situation where the MCCN-induced cost game has empty core in contrast to the MCST model where the core is non-empty provided there are no undemanded nodes. In particular, the example points to an additional source of instability in efficient network design: if the demand graph, consisting of the direct connections between all pairs of demanded nodes, has at least three components, then the core of the induced cost game can be empty.

We aim to investigate under what conditions on the demand graph we can guarantee non-empty core of the induced cost game on the restricted domain of MCCN-problems where all locations are demanded. The first result (Theorem 1) shows that if the efficient network and the demand graph consists of the same components, then the induced cost game has non-empty core. In this sense disjoint demands are not a problem in itself provided the structure of the efficient network is disjoint as well, stability of the efficient network can still be ensured. Intuitively, if the demand graph is disjoint, then the efficient network is disjoint as well provided it is costly to connect different components of agents.

Our second result (Theorem 2) shows that if the demand graph consists of at most two components, then the core is non-empty enabling stability of efficient network design. The result allows us to infer the core is non-empty in many economically relevant models. For instance, it generalizes the seminal result of Bird (1976) for the MCST-model where the demand graph is a star with the source as center, and allows us to conclude that in multi-source spanning tree problems, core stability is ensured when every source is demanded by at least one agent and the demand graph has at most two components, see e.g. Rosenthal (1987), Bergantiños and Navarro-Ramos (2019). Moreover, communication games where all pairs of nodes can be interpreted as an agent, so the demand graph coincides with the

²A full characterization of the Minimum Cost Steiner Tree models for which the core of the induced cost game is non-empty does not exist.

complete graph, will also be stable, see e.g. Moretti (2018), Skorin-Kapov (2018).

While finding an MCCN may be computationally complex, we emphasize that determining the number of components of the demand graph is simple and can be done using a simple fast algorithm. Thus, the conditions in our theorems are easy to check even for large problems.

Our results relate to several strands of literature. Issues of fair division among users sharing a common network resource have attracted much attention over the past couple of decades: see e.g. Moulin (2014, 2019) and Hougaard (2018) for recent surveys. The standard approach to fair division has been to formulate an associated cooperative game (see e.g., Peleg and Sudhölter, 2007) and use solution concepts from game theory such as the core and the Shapley value to guide allocation of costs and revenues.

Since the seminal papers by Claus and Kleitman (1973) and Bird (1976), the minimum cost spanning tree model and its many variations have been particularly popular topics in cost and revenue sharing in networks (see e.g. Granot and Huberman, 1981, Tijs et al., 2006, Bergantiños and Vidal-Puga, 2007, Bergantiños and Martínez, 2014, Bogomolnaia and Moulin, 2010, Bogomolnaia et al., 2010, Hougaard et al. 2010, Trudeau, 2012, 2013). Implementation of minimum cost spanning trees has been studied in Bergantiños and Lorenzo (2004, 2005), Bergantiños and Vidal-Puga (2010), and Hougaard and Tvede (2012).

The more general MCCN-model is originally introduced in Moulin (2009, 2014), inspired by non-cooperative cost sharing network games analyzed in the computer science literature, e.g., Anshelevich et al. (2008) and Chen et al. (2010). In particular, Moulin (2009, 2014) analyze two types of cost sharing rules satisfying core stability and routing-proofness (a user cannot lower its cost share by reporting as multiple aliases along an alternative path connection her target nodes) where the induced MCCN-games are balanced. Juarez and Kumar (2013) consider Nash implementation in a game where users choose paths connecting their target nodes. Using a particular game form, Hougaard and Tvede (2015) show that the options for implementing MCCNs are much more limited than in the MCST-model. Ensuring a cost minimizing network by truthful reporting now implies compromising with individual rationality. Hougaard and Tvede (2019) introduce users with limited willingness to pay for connectivity and show that welfare maximizing networks with individually rational cost allocation are both Nash and strong Nash implementable.

Finally, we note that the MCCN-model is relevant for many economic and engineering applications including various cost sharing issues related to the Internet and e-commerce (e.g. Jain and Mahdian, 2007), such as Multicasting and client-server networks, where clients may want multiple connections to different servers (e.g. Feigenbaum et al. 2001, Archer et al. 2004); and pricing network traffic, where users request traffic flows between

different pairs of network destinations, e.g. data centers (e.g. Anshelevich et al. 2008, Moulin, 2009, Shi et al. 2018).

The rest of the paper is organized as follows: In section 2, we set up the model and provide an example demonstrating a case with empty core on the domain where all locations are demanded. In section 3, we present our main results, Theorems 1 and 2, on the connection between the structure of the demand graph and balancedness of the induced cost game. Section 4 closes with final remarks and a couple of conjectures regarding the domain with undemanded (Steiner) locations.

2 Model

We first recall the MCCN-model (see e.g., Moulin, 2014, or Hougaard and Tvede, 2015). Let $M = \{1, \dots, m\}$ be a finite set of agents and \mathcal{N} a finite set of locations (nodes). The set of connections (edges) between pairs of locations is $\mathcal{N}^2 = \mathcal{N} \times \mathcal{N}$. A *cost structure* C describes costs of connecting locations and is defined by a map $c : \mathcal{N}^2 \rightarrow \mathbb{R}_+$ with: $c_{jj} = 0$ for every location j ; and, $c_{jk} > 0$ for every pair of locations (j, k) with $j \neq k$. Connections are undirected so $c_{kj} = c_{jk}$ for every pair of locations (j, k) . Connection costs are constant so the network is congestion free.

Every agent $i \in M$ has a connection *demand* $D_i = (a_i, b_i) \in \mathcal{N} \times \mathcal{N}$ with $a_i \neq b_i$, where (a_i, b_i) is a pair of locations that agent i wants to have connected directly or indirectly. A *demand structure* is a collection of demands $D = (D_i)_{i \in M}$. Note that in the classic Minimum Cost Spanning Tree (MCST) model all agents demand connection to the same location (the source). Agents can therefore be identified by the set of nodes with the source as an additional (non-involved) "agent". The MCST model is therefore a special case of the MCCN model. Given a demand structure D , define the *demand graph* $G^D = \cup_{i \in M} D_i$.

A *connection problem* (M, D, C) consists of a set of agents, a demand structure, and a cost structure.

Specifically, we focus on the domain of connection problems, Γ , where *all locations are demanded*, i.e., $\cup_{i \in M} D_i = \mathcal{N}$. Thus, for any problem in Γ , the number of locations is at most $2m$. If there are n locations and no two agents have the same demand, we can have at most $n(n-1)/2$ agents. Problems in Γ , for which the demand graph G^D has k components, involves at most $m+k$ nodes.

A *graph* g on \mathcal{N} is a set of connections $g \subset \mathcal{N}^2$. For a cost structure C , and a graph g , let $v(C, g) \geq 0$ be the total cost of the graph g

$$v(C, g) = \sum_{jk \in g} c_{jk}.$$

For a given connection problem (M, D, C) , a Connection Network (CN) is a graph g meeting the connection demand of every agent $i \in M$: for every agent $i \in M$ there is a path $p = \{n_1 n_2, n_2 n_3, \dots, n_{\ell-1} n_\ell\}$ with $n_1 = a_i$, $n_\ell = b_i$ and $n_j \neq n_k$ for every pair of locations (j, k) with $j \neq k$, such that $p \subseteq g$. Denote by \mathcal{CN} the set of CNs.

A Minimal Cost Connection Network (MCCN) is a CN that minimizes cost: that is, g is MCCN if

$$g \in \{\arg \min_{g \in \mathcal{CN}} v(C, g)\}.$$

The set of MCCNs is non-empty and finite because the set of CNs is non-empty and finite. Clearly, every MCCN is either a tree or a forest (a graph where every component is a tree).

A connection problem (M, D, C) induces a cooperative (cost) game (M, c) where, for every coalition of agents $S \subseteq M$, $c(S) = v(C|_S, g^S)$: with g^S being an MCCN of the S -projected connection problem $(S, D|_S, C|_S)$, i.e., the problem where only connections (and their cost) between locations demanded by agents in S are considered.

By construction, the game (M, c) is *subadditive* (i.e. for every $S, T \subseteq M$ such that $S \cap T = \emptyset$, $c(S) + c(T) \geq c(S \cup T)$).

The *core* of the game (M, c) is given by the set of allocations,

$$\text{core}(M, c) = \{x \in \mathbb{R}^M \mid \sum_{i \in M} x_i = c(M), \sum_{i \in S} x_i \leq c(S), \text{ for all } S \subset M\}. \quad (1)$$

Given the set of agents M , a collection $\mathcal{B} = \{S_1, \dots, S_k\}$ of non-empty subsets of M is called *balanced* if there exists positive numbers $\delta_1, \dots, \delta_k$ such that $\sum_{j: i \in S_j} \delta_j = 1$, for all $i \in M$. By the Bondareva–Shapley theorem, $\text{core}(M, c) \neq \emptyset$ if and only if for each balanced collection and each system of weights δ , that

$$\sum_{S \subseteq \mathcal{B}} \delta_S c(S) \geq c(M). \quad (2)$$

Games satisfying (2) are called *balanced*.


A game (M, c) is said to be *concave* if, for every $S, T \subseteq M$,

$$c(S \cup T) + c(S \cap T) \leq c(S) + c(T). \quad (3)$$

A game is concave if and only if, for each $i \in M$, i 's marginal cost $m_i(S) = c(S \cup \{i\}) - c(S)$ is non-increasing in S . Concave games are balanced.

By examples in Meggido (1978) and Tamir (1991) it is known that the core may be empty for MCCM problems with *undemanded* locations (Steiner nodes in these examples). The following example shows that the core may even be empty for games induced by connection problems where all locations are demanded (i.e. for problems in our domain Γ).

Example 1: Consider six locations $\mathcal{N} = \{a, b, c, d, e, f\}$ and three agents $M = \{A, B, C\}$ with connection demands (a, b) , (c, d) , and (e, f) , respectively: so all locations are demanded with the demand graph G^D consisting of three components. Connection costs are given as follows: $c_{af} = c_{bf} = c_{ae} = c_{de} = c_{df} = c_{ce} = c_{bc} = c_{ac} = c_{bd} = 1$, and $c_{ij} = 10$ otherwise. In the graph below only the (relevant) edges with cost equal to 1 are illustrated.



The induced cost game (M, c) has empty core since $c(AB) = c(AC) = c(BC) = 3$ and $c(ABC) = 5$ (so all agents must pay at least 2, but the total cost is only 5). As Theorem 2 will show, this example is minimal (on Γ) in the sense that we need at least three components of the demand graph (and thereby at least three agents) in order to produce an example of an induced cost game with empty core. \square

3 Core Stability

In this section we identify classes of connection problems for which the induced cost games are balanced (always have non-empty core). As indicated by Example 1, the shape of the demand graph G^D plays a key role.

Indeed, in the special case of the MCST-model, Bird (1976) demonstrated that the induced cost games are balanced. For MCST-problems the demand graph is star-shaped with

all nodes demanding connection to the same source node. Our first result, Theorem 1 below, generalizes Bird's result to hold for arbitrarily connected demand graphs, and not just star-shaped demand graphs. In fact, Theorem 1 goes much further than that, and states that whenever there exists an MCCN g with the same number of components as the demand graph, the induced cost game is balanced. So even though disjoint demand graphs may result in empty cores (as shown by Example 1 above), then as long as the efficient network is disjoint too (separated into the same number of components), then no matter how many components, core stability is still ensured.

The following Lemma constitutes the first step towards a proof.

Lemma 1 *For $(M, C, D) \in \Gamma$ suppose G^D consists of one component and (M, c) is balanced. For every $(M', C, D') \in \Gamma$, where $M \subset M'$ and $D_i = D'_i$ for every $i \in M$, (M', c) is balanced.*

Proof: Since G^D consists of one component, if g is a MCCN for (M, C, D) , then g is a tree and g is a MCCN for all (M', C, D') where $M \subset M'$ and $D_i = D'_i$ for every $i \in M$. Moreover, for every $S' \subset M'$, $c(S') \geq c(S' \cap M)$. Hence, balancedness of (M', C, D') follows from balancedness of (M, C, D) . \square

With Lemma 1 we are ready to present and prove our first main result.

Theorem 1 *Let $(M, C, D) \in \Gamma$. Suppose there is a MCCN g for which the number of components is equal to the number of components in G^D then (M, c) is balanced.*

Proof: Consider first $(M, D, C) \in \Gamma$ with MCCN g having as many components as G^D and where every component of g is a spanning tree. The components of G^D and g must be the same since no MCCN g can “cut” a component of G^D . Indeed, otherwise there would be some agent i for which a_i belongs to one component of g and b_i belongs to another component of g , so the demand of agent i would not be satisfied. Denote by $\{K_1, \dots, K_l\}$ the partition of M given by the components of g and denote by g^j the cost minimizing spanning tree of the j 'th component. Thus, $|M| = m = \sum_{j=1}^l |K_j|$. We first consider the case where, for every component j , there are $|K_j|$ links in g^j , involving $|K_j| + 1$ nodes.

For every component and arbitrary coalition $T \subseteq K_j$, let κ^T be the minimum cost of satisfying the demands of agents $K_j \setminus T$ using links from the efficient graph g^j added to the demand subgraph of T , $G_T^D = \cup_{i \in T} D_i$ (this is well defined since g^j is a spanning tree). In other words, κ^T equals $c(K_j)$ minus the sum of the $|T|$ most expensive links in g^j that can be replaced by G_T^D subject to the new graph remaining a spanning tree. In particular, $\kappa^\emptyset = c(K_j)$ and $\kappa^{K_j} = 0$.

For arbitrary coalitions $S \subseteq M$ consider the game defined by

$$\bar{c}(S) = c(M) - \sum_{K_j} \kappa^{S \cap K_j}. \quad (4)$$

Clearly, $c(M) = \bar{c}(M)$. We claim that $c(S) \geq \bar{c}(S)$ for every $S \subset M$. Indeed, suppose $c(S) < \bar{c}(S)$. Thus $c(S) < c(M) - \sum_{K_j} \kappa^{S \cap K_j} \Leftrightarrow c(S) + \sum_{K_j} \kappa^{S \cap K_j} < c(M)$, which by definition of κ contradicts that g is MCCN.

We claim that the game (M, \bar{c}) is concave. Indeed, the marginal cost is given by $\bar{m}_i(S) = \bar{c}(S \cup \{i\}) - \bar{c}(S) = \sum_{K_j} \kappa^{S \cap K_j} - \sum_{K_j} \kappa^{(S \cup \{i\}) \cap K_j}$.

Suppose $i \in K_j$ and $K_j \cap S = \emptyset$. Then

$$\bar{m}_i(S) = c(K_j) - \kappa^i = \max\{c_{l_z} \mid l_z \in g^j \text{ and } g^j - l_z + a_i b_i \text{ is a spanning tree}\}$$

by definition of κ . First, for every T with $S \subset T$ and $T \cap K_j = \emptyset$, $\bar{m}_i(T) = \bar{m}_i(S)$. Second for every T with $S \subset T$ and $T \cap K_j \neq \emptyset$,

$$\begin{aligned} \bar{m}_i(S) &= \max\{c_{l_z} \mid l_z \in g^j \text{ and } g^j - l_z + a_i b_i \text{ is a spanning tree}\} \\ &\geq \max\{c_{l_z} \mid l_z \in g^{K_j \cap T} \text{ and } g^{K_j \cap T} - l_z + a_i b_i \text{ is a spanning tree}\} = \bar{m}_i(T). \end{aligned}$$

Suppose $i \in K_j$ and $K_j \cap S \neq \emptyset$. Then $\bar{m}_i(S) = \kappa^{S \cap K_j} - \kappa^{(S \cup \{i\}) \cap K_j}$ which by definition of κ is weakly decreasing in the size of S . Indeed, let $g^{K_j \cap S}$ denote a spanning tree obtained from g^j by replacing $|S \cap K_j|$ links in g^j with the demand graph of $S \cap K_j$ such that the total cost of the links removed is maximized. Then

$$\bar{m}_i(S) = \max\{c_{l_z} \mid l_z \in g^{K_j \cap S} \text{ and } g^{K_j \cap S} - l_z + a_i b_i \text{ is a spanning tree}\},$$

which is weakly decreasing in the size of S .

To conclude, (M, \bar{c}) is concave and thus satisfies (2). Consequently, since $c(M) = \bar{c}(M) \leq \sum_{S \subseteq \emptyset} \delta_S \bar{c}(S) \leq \sum_{S \subseteq \emptyset} \delta_S c(S)$, (M, c) is balanced. By Lemma 1, this extends to problems where components of G^D span less than $|K_i| + 1$ nodes. \square

In our model agents accept any MCCN, but in some practical cases there may be limitations on the shape of the efficient network, for instance, due to various types of privacy concerns. In the extreme case where no agent has confidence in other nodes than the two target nodes they demand, the demand graph coincides with the only eligible MCCN inducing an additive cost game and additive games are balanced (see e.g. Peleg and Sudhölter, 2007). Theorem 1, can be viewed along these lines: if, for instance, privacy concerns result in a clustering of eligible nodes demanded by subsets of agents, this can actually help creating core stability by making sure that eligible MCCNs are partitioned in the same components as the agents with mutually admissible target nodes.

Example 2: Recall the situation in Example 1 above. Here $G^D = \{ab\} \cup \{cd\} \cup \{ef\}$ consists of three components while MCCNs g are spanning trees (e.g. $g = \{af, bf, bc, ce, de\}$ consisting of one component). Now, to illustrate the content of Theorem 1 we can modify connection costs such that the MCCN coincides with the demand graph. For instance, this will be the case if $c_{ab} = c_{cd} = c_{ef} = 1.5$ (notice that the sum of any two connections ab , cd , and ef cannot exceed 3, which is the minimum cost of connecting any two-agent coalition). With the new costs for direct connections ab , cd , and ef , it is easy to see that $g = G^D$ is MCCN and the corresponding cost game (M, c) is additive and thereby balanced. \square

Our second result shows that if the agents' demands are relatively clustered, in the sense that the demand graph has at most two components, then core stability is also ensured. This generalizes earlier findings on balancedness of induced games for specific instances of the MCCN-model where the demand graph consists of one component: for instance, MCST-games (Bird, 1976) as mentioned above, connection games (e.g., Moretti, 2018, Skorin-Kapov, 2018) when the demand graph is connected as when it coincides with the complete graph, multi-source spanning tree games (e.g. Rosenthal, 1987) when all sources are demanded by some agent, etc.

Theorem 2 *Let $(M, C, D) \in \Gamma$. If G^D has at most two components, then (M, c) is balanced.*

Proof: By Theorem 1 we know that if G^D has one component, then (M, c) is balanced. Suppose G^D has two components and $m+2$ nodes so G^D is a forest and every agent demand contains a node that is not demanded by any other agent. According to Theorem 1, if a MCCN g has two components, then (M, c) is balanced. Therefore, we can restrict attention to the case where g has one component, so g is a spanning tree connecting all $m+2$ nodes.

Consider an arbitrary MCCN g spanning all $m+2$ nodes. For every coalition $S \subseteq M$ suppose g^S is a minimum cost graph among the graphs satisfying all demands and consisting of G_S^D and a selection of links from $g - G_S^D$. Let κ^S be the cost of the selected links in $g - G_S^D$, $\kappa^S = \sum_{lz \in \{g^S - G_S^D\}} c_{lz}$, then $\kappa^\emptyset = c(M)$ and $\kappa^M = 0$. In other words, κ^S equals $c(M)$ minus the sum of the $|S|$ most expensive links in g that can be replaced by G_S^D while ensuring that the resulting graph remains a connection network.

Define the game (M, \bar{c}) by

$$\bar{c}(S) = c(M) - \kappa^S$$

for every $S \subseteq M$. To see that $c(S) \geq \bar{c}(S)$ for every $S \subseteq M$ with $c(M) = \bar{c}(M)$ suppose $c(S) < \bar{c}(S) = c(M) - \kappa^S$ for some $S \subseteq M$. Then $c(S) + \kappa^S < c(M)$ contradicting that g is a MCCN.

To see that the game (M, \bar{c}) is concave let $\bar{m}_i(S) = \bar{c}(S \cup \{i\}) - \bar{c}(S) = \kappa^S - \kappa^{S \cup \{i\}}$. Since there are two components in the demand graph and the MCCN g is a spanning tree,

$$\bar{m}_i(\emptyset) = \max\{c_{lz} \mid lz \in g \text{ and } g - lz + a_i b_i \text{ is a spanning tree}\}$$

so by construction of κ , the link $a_i b_i$ can replace exactly one link, $lz \in g$ while making sure all demands remain satisfied; $m_i(\emptyset)$ is equal to the most costly link with this property. Note that if $a_i b_i \in g$ then $m_i(\emptyset) = c_{a_i b_i}$ as the link $a_i b_i$ “replaces” itself. Moreover,

$$\bar{m}_i(S) = \max\{c_{lz} \mid lz \in g^S \text{ and } g^S - lz + a_i b_i \text{ is a spanning tree}\}.$$

Therefore, the marginal cost of agent i is non-increasing for T such that $S \subset T$. To conclude, (M, \bar{c}) is concave so (M, c) is balanced. By Lemma 1, this extends to problems where components of G^D span less than $|K_i| + 1$ nodes. \square

The proof technique used in Theorem 2 is basically similar to that used in the proof of Theorem 1. Yet, there is one important difference: the proof of Theorem 1 uses explicitly the context where the MCCN g has as many components as the demand graph G^D . In Theorem 2, the demand graph has only two components while the MCCN g spans the entire set of locations. This is an important difference since if there are three or more components in G^D , the marginal cost \bar{m}_i is no longer non-increasing: adding the demanded link $a_i b_i$ of agent i may enable deletion of more than one link whilst keeping the resulting graph a spanning tree. Therefore, the game may not be concave as illustrated by Example 1.

Example 3: In order to illustrate the content of Theorem 2, we can modify the situation from Example 1 as follows. Consider only five locations $\mathcal{N}' = \{a, b, c, d, e\}$ with demands $d_A = (a, b)$, $d_B = (c, d)$ and $d_C = (d, e)$ (so all locations are demanded). Thus, the demand graph consists of two components $G^D = \{a, b\} \cup \{cd, de\}$ and not three components as in Example 1. As in the original problem, let demanded direct connections cost 10 ($c_{ab} = c_{cd} = c_{de} = 10$) and all other connections cost 1. One MCCN is $g = \{ae, ad, bd, bc\}$ with a total cost of 4. The induced cost game is therefore given by stand-alone costs $c(\{i\}) = 10$ for all i , and total cost $c(ABC) = 4$. For two-agent coalitions $c(AB) = c(AC) = 3$ and $c(BC) = 11$, so clearly the core is non-empty (containing e.g. the equal split). Now, to reduce the core we can lower direct costs. However, making sure that $g = \{ae, ad, bd, bc\}$ remains MCCN we must have $c(AB) \geq 3$, $c(AC) \geq 3$ and $c(BC) \geq 2$. Thus, even in the “worst” case the core will be non-empty since $c(AB) + c(AC) + c(BC) \geq 2c(ABC)$ (the other conditions (2) are satisfied by c being sub-additive by construction). \square

4 Final Remarks

So far we have focused on the domain Γ of MCCN-problems for which all locations are demanded. We close with a few remarks on the presence of undemanded nodes (i.e. Steiner nodes) in the MCCN-model.

As mentioned in the introduction, examples in Megiddo (1978) and Tamir (1991) imply that games with empty core can occur in minimum cost Steiner tree problems (where the demand graph is star-shaped) with at least three agents by including three (or more) Steiner nodes. We can therefore directly infer that this holds by including at least three undemanded nodes in the MCCN-model as well. In Tamir's example, it is straightforward to observe that omitting one Steiner node is enough to retain balancedness of the resulting cost game. Tamir's example is minimal in that sense. More generally, we here conjecture that for any connected demand graph, the induced MCCN game is balanced if at most two undemanded nodes are introduced to the problem.

In particular, if we consider connected demand graphs with the shape of a connected chain (i.e., agents' demands form a path with $b_i = a_{i+1}$ for $i = 1, \dots, m - 1$) we conjecture that the induced MCCN game is balanced even when allowing for the presence of any number of undemanded nodes. Loosely speaking, the chain structure limits how much individual coalitions can save by using undemanded locations compared to the savings of the grand coalition, which thereby seems to ensure compliance with the balancedness conditions.

References

- Anshelevich, E., A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler and T. Roughgarden (2008), The price of stability for network design with fair cost allocation, *SIAM Journal of Computing*, 38 (4), 1602–1623.
- Archer, A, J. Feigenbaum, A. Krishnamurthy, R. Sami, and S. Shenker (2004), Approximation and collusion in multicast cost sharing, *Games and Economic Behavior*, 47 (1), 36–71.
- Bergantiños, G., and L. Lorenzo (2004), A non-cooperative approach to the cost spanning tree problem, *Mathematical Methods of Operations Research*, 59, 393–403.
- Bergantiños, G., and L. Lorenzo (2005), Optimal equilibria in the non-cooperative game associated with cost spanning tree problems, *Annals of Operations Research*, 137, 101–115.
- Bergantiños, G., R. Martinez (2014), Cost allocation in asymmetric trees, *European Journal of Operational Research*, 237 (3), 975-987.

- Bergantiños, G., and A. Navarro-Ramos (2019), The folk rule through a painting procedure for minimum cost spanning tree problems with multiple sources, *Mathematical Social Sciences*, 99, 43-48.
- Bergantiños, G., and J.J. Vidal-Puga (2007), A fair rule in minimum cost spanning tree problems, *Journal of Economic Theory*, 137 (1), 326–352 .
- Bergantiños, G. and J.J. Vidal-Puga (2010), Realizing fair outcomes in minimum cost spanning tree problems through non-cooperative mechanisms, *European Journal of Operational Research*, 201 (3), 811–820.
- Bird, C.G. (1976), On cost allocation for a spanning tree: a game theoretic approach, *Networks*, 6, 335–350.
- Bogomolnaia, A., and H. Moulin (2010), Sharing a minimal cost spanning tree: Beyond the Folk solution, *Games and Economic Behavior*, 69 (2), 238–248.
- Bogomolnaia, A., R. Holzman and H. Moulin (2010), Sharing the cost of a capacity network, *Mathematics of Operations Research*, 35 (1), 173–192.
- Chen, H.-L., T. Roughgarden and G. Valiant (2010), Designing network protocols for good equilibria, *SIAM Journal of Computing*, 39 (5), 1799–1832.
- Claus, A. and D.J. Kleitman (1973), Cost allocation for a spanning tree, *Networks*, 3, 289–304.
- Feigenbaum, J., C. Papadimitriou, and S. Shenker (2001), Sharing the cost of multicast transmissions, *Journal of Computer and Systems Sciences*, 63 (1), 21–41.
- Granot, D. and G. Huberman (1981), Minimum cost spanning tree games, *Mathematical Programming*, 21, 1–18.
- Hougaard, J.L. (2018), *Allocation in Networks*, MIT Press.
- Hougaard, J.L. and H. Moulin (2014), Sharing the cost of redundant items, *Games and Economic Behavior*, 87, 339–352.
- Hougaard, J.L., H. Moulin and L.P. Østerdal (2010), Decentralized pricing in minimum cost spanning trees, *Economic Theory*, 44, 293–306.
- Hougaard, J.L. and M. Tvede (2012), Truth-telling and Nash equilibria in minimum cost spanning tree models, *European Journal of Operational Research*, 222 (3), 566–570.
- Hougaard, J.L. and M. Tvede (2015), Minimum cost connection networks: truth-telling and implementation, *Journal of Economic Theory*, 157, 76–99.

- Hougaard, J.L. and M. Tvede (2019), Implementation of optimal connection networks, Manus.
- Jain, K. and M. Mahdian (2007), Cost sharing, In Nisan et al. (eds), *Algorithmic Game Theory*, Ch. 15, Cambridge University Press.
- Juarez, R. and R. Kumar (2013), Implementing efficient graphs in connection networks, *Economic Theory*, 54, 359–403.
- Karp, R.M. (1972), Reducibility among combinatorial problems. In *Complexity of Computer Computations*, p. 85-103, Springer.
- Korte, B. and J. Vygen (2018), *Combinatorial Optimization: Theory and Algorithms*, Sixth Edition, Springer.
- Megiddo, N., (1978), Cost allocation for Steiner trees, *Networks*, 8, 1–6.
- Moretti S. (2018), On Cooperative Connection Situations Where the Players Are Located at the Edges. In: Belardinelli F., Argente E. (eds) *Multi-Agent Systems and Agreement Technologies*. EUMAS 2017, AT 2017. *Lecture Notes in Computer Science*, vol 10767, Springer.
- Moulin, H., (2009), Pricing traffic in a spanning network, *Proceedings of the 10th ACM conference on Electronic commerce*, 21–30.
- Moulin, H., (2014), Pricing traffic in a spanning network, *Games and Economic Behavior*, 86, 475–490.
- Moulin, H., (2019), Fair division in the Internet age, *Annual Review of Economics*, 11.
- Peleg, B. and P. Sudhölter (2007), *Introduction to the Theory of Cooperative Games*, Second Edition, Kluwer.
- Rosenthal, E.C. (1987), The minimum spanning forest game, *Economics Letters*, 23 (4), 355–357.
- Tijs, S., R. Branzei, S. Moretti and H. Norde (2006), Obligation rules for minimum cost spanning tree situations and their monotonicity properties, *European Journal of Operational Research*, 175 (1), 121–134.
- Shi, W., C. Wu, and Z. Li (2018), A Shapley-Value mechanism for bandwidth on demand between datacenters, *IEEE Transactions on Cloud Computing*, 6, 19–32.
- Skorin-Kapov, D. (2018), Social enterprise tree network games, *Annals of Operations Research*, 268, 5–20.

- Tamir, A., (1991), On the core of network synthesis games, *Mathematical Programming*, 50, 123–135.
- Trudeau, C., (2012), A new stable and more responsive cost sharing solution for minimum cost spanning tree problems, *Games and Economic Behavior*, 75 (1), 402–412.
- Trudeau, C. (2013), Characterizations of the Kar and folk solutions for minimum cost spanning tree problems, *International Game Theory Review*, 15 (2), 134–143.
- Young H.P. (1998), Cost allocation, demand revelation, and core implementation, *Mathematical Social Sciences*, 36 (3), 213–228.