


Fair Tree Connection Games with Topology-Dependent Edge Cost

Davide Bilò 

Department of Humanities and Social Sciences, University of Sassari,
Via Roma 151, 07100 Sassari (SS), Italy
davide.bilo@uniss.it

Tobias Friedrich 

Hasso Plattner Institute, University of Potsdam,
Prof.-Dr.-Helmert-Straße 2-3, 14482 Potsdam, Germany
tobias.friedrich@hpi.de

Pascal Lenzner 

Hasso Plattner Institute, University of Potsdam,
Prof.-Dr.-Helmert-Straße 2-3, 14482 Potsdam, Germany
pascal.lenzner@hpi.de

Anna Melnichenko 

Hasso Plattner Institute, University of Potsdam,
Prof.-Dr.-Helmert-Straße 2-3, 14482 Potsdam, Germany
anna.melnichenko@hpi.de

Louise Molitor 

Hasso Plattner Institute, University of Potsdam,
Prof.-Dr.-Helmert-Straße 2-3, 14482 Potsdam, Germany
louise.molitor@hpi.de

Abstract

How do rational agents self-organize when trying to connect to a common target? We study this question with a simple tree formation game which is related to the well-known fair single-source connection game by Anshelevich et al. (FOCS'04) and selfish spanning tree games by Gourvès and Monnot (WINE'08). In our game agents correspond to nodes in a network that activate a single outgoing edge to connect to the common target node (possibly via other nodes). Agents pay for their path to the common target, and edge costs are shared fairly among all agents using an edge. The main novelty of our model is dynamic edge costs that depend on the in-degree of the respective endpoint. This reflects that connecting to popular nodes that have increased internal coordination costs is more expensive since they can charge higher prices for their routing service.

In contrast to related models, we show that equilibria are not guaranteed to exist, but we prove the existence for infinitely many numbers of agents. Moreover, we analyze the structure of equilibrium trees and employ these insights to prove a constant upper bound on the Price of Anarchy as well as non-trivial lower bounds on both the Price of Anarchy and the Price of Stability. We also show that in comparison with the social optimum tree the overall cost of an equilibrium tree is more fairly shared among the agents. Thus, we prove that self-organization of rational agents yields on average only slightly higher cost per agent compared to the centralized optimum, and at the same time, it induces a more fair cost distribution. Moreover, equilibrium trees achieve a beneficial trade-off between a low height and low maximum degree, and hence these trees might be of independent interest from a combinatorics point-of-view. We conclude with a discussion of promising extensions of our model.

2012 ACM Subject Classification Theory of computation → Algorithmic game theory; Theory of computation → Quality of equilibria; Theory of computation → Convergence and learning in games; Theory of computation → Network formation

Keywords and phrases Network Design Games, Spanning Tree Games, Fair Cost Sharing, Price of Anarchy, Nash Equilibrium, Algorithmic Game Theory, Combinatorics



© Davide Bilò, Tobias Friedrich, Pascal Lenzner, Anna Melnichenko, and Louise Molitor;
licensed under Creative Commons License CC-BY

40th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2020).

Editors: Nitin Saxena and Sunil Simon; Article No. 15; pp. 15:1–15:15



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2020.15

Related Version A full version of the paper is available at [9], <http://arxiv.org/abs/2009.10988>.

Acknowledgements We thank Warut Suksompong for many interesting discussions. Moreover, we are grateful to our anonymous reviewers for their valuable suggestions. This work has been partly supported by COST Action CA16228 European Network for Game Theory (GAMENET).

1 Introduction

Network Design is an important optimization problem where for a given weighted host graph and a given set of terminal pairs the cheapest subgraph which connects all terminal pairs has to be found. Besides an abundance of research works with an optimization point-of-view, e.g. see the survey by Magnanti and Wong [28], a strategic version of the Network Design problem [5, 4] has kindled significant interest in recent years. In the *connection game*, a weighted host graph H is given and n agents with given terminal node pairs (s_i, t_i) , for $1 \leq i \leq n$, strategically select s_i - t_i -paths in H to connect their respective terminal nodes. The union of the selected paths forms a subgraph G of H which constitutes the actually designed network. The usage cost of each edge of H corresponds to its weight, and agents using some edge e in H have to pay this cost. If an edge e is used by more than one agent, then a cost-sharing protocol determines how the usage cost of e is split among its users. One of the most common cost-sharing protocols is Shapley cost-sharing where each agent pays a fair share of the edge cost, i.e., the cost-share is the edge cost divided by the number of users. This game-theoretic setting, called *fair connection game*, was investigated by Anshelevich et al. [4] and has since become an influential paper in Algorithmic Game Theory. An important special case is the setting in which all the strategic agents want to connect to a common source node. This variant, where $t_1 = \dots = t_n$ and where every other node is a terminal node of some agent, is usually denoted as the (*fair*) *single-source connection game*, with the interpretation that all the agents want to connect to a common source node to receive broadcast messages and that the edge cost for connecting to the common source is paid by the downstream users.

A similar related game-theoretic setting are *selfish spanning tree games* [21]. There a weighted complete host graph with $n + 1$ nodes, consisting of a common target node r and n nodes which correspond to selfish agents, is given and every selfish agents now selects an incident edge to connect to the common target node r either directly or indirectly via selected edges of other agents. The cost of an agent is then determined by its unique path to r . Thus, in any equilibrium the subgraph of all selected edges forms a spanning tree rooted at r .

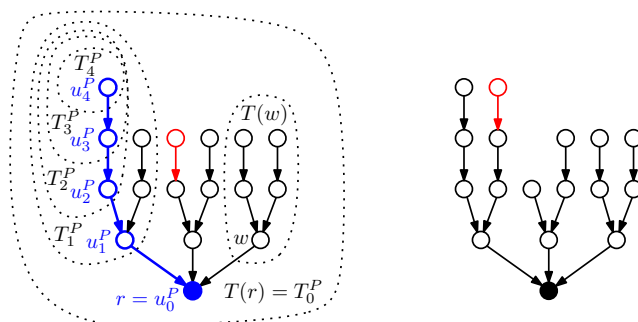
This paper sets out to investigate a game-theoretic Network Design model that is closely related to the fair single-source connection game and to selfish spanning tree games. The main novel feature of our model is the twist that the cost of the edges in the formed spanning tree depend on its topology. In particular, we consider dynamic edge costs which are proportional to the in-degree of the node they connect to. Network nodes with high in-degree can be considered as popular, and we assume that connecting to popular nodes is more expensive than connecting to unpopular nodes. These dynamic edge costs can also be understood as the internal cost of a node for coordinating data traffic coming from different connections. A node with many incoming edges and thus higher internal coordination cost can charge higher prices for serving each of the incoming edges.

To the best of our knowledge, we define and analyze the first (game-theoretic) Network Design model where the edge costs depend on the topology of the formed network. We believe that this model sheds light on settings where the actual charges for establishing links are determined by supply and demand and the agents act strategically to optimize their cost for receiving their desired service.

1.1 Model, Definition, Notation

We consider a strategic game called *fair tree connection game with topology-dependent edge cost*, or *tree connection game (TCG)* for short. In the TCG we will consider a given unweighted complete directed host graph $H = (V, E)$, where V is the set of nodes and E is the set of edges of H . The host graph H consists of $n + 1$ nodes $V = \{r, v_1, \dots, v_n\}$ where node r is the common target node, also called the root, and every node v_i , for $1 \leq i \leq n$, corresponds to a selfish agent i striving to be connected to the root r . For this, every agent i strategically activates a single incident edge (v_i, s_i) , where $s_i \in V \setminus \{v_i\}$. Hence, the strategy space of each agent is the set of other nodes to connect to. Given a strategy profile $\mathbf{s} = (s_1, \dots, s_n)$, i.e., an n -dimensional vector where the j -th entry corresponds to the node to which agent j wants to activate her edge, we consider the directed network $T(\mathbf{s}) = (V, E(\mathbf{s}))$ which is induced by all the activated edges, i.e., $E(\mathbf{s}) = \{(v_i, s_i) \mid 1 \leq i \leq n\}$. We will see later that $T(\mathbf{s})$ is a spanning tree rooted at r if \mathbf{s} is an equilibrium state of the TCG, hence the name.

The cost of agent i in the network $T(\mathbf{s})$ depends on its unique path P_i in $T(\mathbf{s})$ to the root r (if such a path exists). In case of existence, the path P_i must be unique, since the out-degree of every node in $T(\mathbf{s})$ is at most 1. More precisely, let P_i be the directed path from v_i to r in $T(\mathbf{s})$, let $\text{indeg}_{T(\mathbf{s})}(v)$ denote the number of edges with endpoint v in $T(\mathbf{s})$, let $T(u)$ denote the subgraph of $T(\mathbf{s})$ rooted at node u , i.e., the subgraph of $T(\mathbf{s})$ induced by the nodes u and every node which has a directed path to u and let $|T(u)|$ denote the number of nodes in $T(u)$. See Figure 1.



■ **Figure 1** Left: $T(\mathbf{s})$ for $n = 16$ agents. The path P is colored blue and we have $d_0^P = 3$, $d_1^P = 2$, $d_2^P = d_3^P = 1$, $d_4^P = 0$ and $|T_1^P| = 6$, $|T_2^P| = 3$, $|T_3^P| = 2$, $|T_4^P| = 1$, $|T(w)| = 5$. Nodes w and u_1^P and also their corresponding agents are siblings. The shown network $T(\mathbf{s})$ is not stable since the agent colored red with cost $1 + \frac{2}{2} + \frac{3}{5} = \frac{13}{5}$ has an improving move. Right: the network after the agent colored red improved its cost to $1 + \frac{1}{2} + \frac{2}{3} + \frac{3}{7} = \frac{109}{42} < \frac{13}{5}$.

The cost of agent i in $T(\mathbf{s})$ then is $\text{cost}_{T(\mathbf{s})}(i) := \sum_{(u,v) \in P_i} \frac{\text{indeg}_{T(\mathbf{s})}(v)}{|T(u)|}$, if P_i exists and ∞ otherwise. This cost function has the following very natural interpretation: the cost of activating edge (u, v) from node u to node v is equal to node v 's in-degree, and this cost is fairly shared by all agents who use edge (u, v) on their path towards the root r , i.e., by all agents in $T(u)$. We assume that each agent activates a single edge strategically to minimize its cost in the induced network $T(\mathbf{s})$. Clearly, since every agent i can activate the edge (v_i, r) , i can enforce finite cost by enforcing that the path P_i exists.

Consider a strategy profile $\mathbf{s} = (s_1, \dots, s_{i-1}, s_i, s_{i+1}, \dots, s_n)$. We say that agent i has an *improving move* in \mathbf{s} if i has some alternative strategy $s'_i \neq s_i$ such that for the induced strategy profile $\mathbf{s}' = (s_1, \dots, s_{i-1}, s'_i, s_{i+1}, \dots, s_n)$ we have $cost_{T(\mathbf{s}')} (i) < cost_{T(\mathbf{s})} (i)$, i.e., agent i can strictly decrease its cost by activating a different outgoing edge. With this, we define the strategy profile \mathbf{s} to be in *pure Nash equilibrium (NE)* or to be *stable* if no agent has an improving move in \mathbf{s} . If the context is clear, we use strategy profiles and their induced network interchangeably, i.e., we say that the network $T(\mathbf{s})$ is in NE or stable, if \mathbf{s} is in NE. Moreover, when we refer to some network $T(\mathbf{s})$ we will from now on omit the reference to the strategy profile \mathbf{s} and call the network simply T . Every stable network T must be a spanning tree rooted at r , since every agent i can activate the edge (v_i, r) to achieve finite cost.

The *social cost* $SC(T)$ of a network T is simply the sum over all agents' costs, i.e.,

$$SC(T) = \sum_{i=1}^n cost_T(i) = \sum_{v_i \in V} \sum_{(u,v) \in P_i} \frac{indeg_T(v)}{|T(u)|} = \sum_{(u,v) \in E} \frac{indeg(v)}{|T(u)|} \cdot |T(u)| = \sum_{v \in V} (indeg(v))^2.$$

Note that $SC(T)$ nicely reflects the overall cost impact of the nodes' popularity or coordination costs which scales quadratically with the in-degree of a node. For a given number of agents n , let OPT_n denote the network which minimizes the social cost. Moreover, if stable networks exist for n agents, we let $worstNE_n$ denote the stable network with the highest social cost and $bestNE_n$ the stable network with the lowest social cost. We define the *Price of Anarchy (PoA)* [27] as $PoA = \sup_n \frac{SC(worstNE_n)}{SC(OPT_n)}$ and the *Price of Stability (PoS)* [4] as $PoS = \sup_n \frac{SC(bestNE_n)}{SC(OPT_n)}$, where the supremum is taken over all n that admit a stable network. Besides the PoA and the PoS, that both focus on the overall cost and compare with the cost of a centrally designed social optimum network, we use a measure of the quality of networks which focuses on the cost distribution among the agents, called the *Fairness Ratio (FR)*, analogously to the *utility uniformity* introduced in [19]. For a given network T , the $FR(T)$ is the ratio between the maximum and the minimum cost incurred by any agent, i.e., $FR(T) := \frac{\max_{v_i \in V} cost_T(i)}{\min_{v_i \in V} cost_T(i)}$.

Finally, we introduce some additional notation for arguing about the designed networks $T(\mathbf{s})$. (See Fig. 1 for an illustration). For our analysis we use directed paths in $T(\mathbf{s})$ which start at some non-root node z and end at the root r . Let P be such a path of length $\ell \in \mathbb{N}$. We denote by u_j^P the node on P which is at distance j to r , hence the root r is denoted by u_0^P and node z by u_ℓ^P . Moreover, let $T_j^P := T(u_j^P)$ and we use d_j^P for the in-degree of a node with distance j from the root r on path P , hence, $d_j^P := indeg_{T(\mathbf{s})}(u_j^P)$. We omit the reference to path P whenever it is clear from the context.

1.2 Related Work

Our model is closely related to several models that have been intensively studied.

We start with the (fair) single-source connection game [5, 4] which we already briefly discussed in the introduction. The key feature of this game is that agents strategically select a set of edges to connect their respective terminals. The cost of each edge is shared among all the agents who selected the respected edge. While in [5] and later also in [24] arbitrary cost sharing is considered, the paper [4] focuses on fair cost sharing which can be derived from the Shapley value [33]. For this Anshelevich et al. [4] show that stable networks always exist since the game is a potential game [32], additionally they prove that the PoA is n and the PoS is upper bounded by H_n , where H_n is the n -th harmonic number. For a given directed host graph this bound on the PoS is tight but the case for undirected host networks is still a major open problem. More is known for single-source connection games on undirected

networks. Chekuri et al. [12] show that the PoA is in $\mathcal{O}(\sqrt{n} \log^2 n)$ if the agents join the game sequentially and play their respective best response. A PoS in $\mathcal{O}(\log \log n)$ was proven by Fiat et al. [20] for the special case where all nodes of the given network correspond to a terminal of some agent. Finally, Bilò et al. [8] prove a constant PoS for the fair single-source connection game on undirected networks. Moreover, Albers and Lenzner [1] show that the optimum is a H_n -approximate Nash equilibrium for the fair single-source connection game. In contrast to our model, the cost of an edge in the (fair single-source) connection game is given via a positively weighted host network. Hoefer and Krysta [25] investigate a variant with edge weights derived from a geometry.

Also selfish spanning tree games [21] are close to our model and we already briefly discussed them in the introduction. The key difference to our model is that a weighted complete network is given and that the cost of an agent is defined differently. Gourvès and Monnot [21] define three variants of the agents' cost function: either it is the weight of the first edge on the path to the common root r , or the minimum or maximum weight edge on the entire path towards r . Cost sharing is not considered. The authors prove bounds on the PoA which vary from unbounded to 1 depending on the exact setting. The games in [21] are inspired by the classical problem of allocating the cost of a spanning tree among its nodes by Claus and Kleitman [13] and its variant from cooperative game theory considered by Bird [10]. Later, Granot and Huberman [22, 23] considered minimum cost spanning tree games and different cost allocation protocols for this have been considered by Escoffier et al. [17]. The key difference of all these models to our model is that a cooperative game is considered which is a stark contrast to our non-cooperative setting. Also game-theoretic topology control problems are related to spanning tree games and our model. Eidenbenz et al. [16] consider a setting where a set of agents which correspond to wireless devices want to connect terminal nodes, whereas Mittal et al. [31] consider wireless access point selection by selfish agents.

Also classical network formation games [26, 6, 18] are related to our model. There the agents correspond to nodes in a network and every agent buys a set of incident edges to connect to other agents. The goal of each agent is to create a connected network and to occupy a central position in this network. For the influential network creation game of Fabrikant et al. [18], that has a parameter α for the trade-off between edge cost and distance costs, the PoA was shown to be constant for almost all values of α [14, 2]. For high values of α all equilibrium networks of these games are known to be trees [30, 29, 7]. A variant of the network creation game where agents can only buy a single edge was considered by Ehsani et al. [15]. Most notably, the topology dependent edge costs that we employ in our model were proposed by Chauhan et al. [11] for the network creation game [18]. To the best of our knowledge, this is the only setting where topology dependent edge costs have been considered.

1.3 Our Contribution

We study a novel game-theoretic model for the formation of a tree network which is related to the well-known fair single-source connection game by Anshelevich et al. [4, 5] and to selfish spanning tree games by Gourvès and Monnot [21]. The key difference of our model is that we consider dynamic edge costs which depend on the topology of the created spanning tree. In particular, the cost of an edge is equal to the in-degree of its endpoint. This specific choice was proposed in [11] for the classical network creation game [18] and we transfer this idea to the Network Design domain. Our analysis holds for any edge cost function of the form α times the in-degree of the target node, for any constant α . However, our general approach is valid also for edge cost functions that depend non-linearly on the degrees of the involved nodes.

Regarding the existence of stable trees we show that our model is in stark contrast to the models in [4, 18, 21] since in our model stable trees may not exist. In particular, we show that our game has no NE for $n = 16$ and $n = 18$ which implies that the TCG cannot admit a potential function. This is contrasted with a proof that for infinitely many n stable trees do exist, and we conjecture that we have found all examples for NE non-existence. Towards investigating the quality of the equilibrium networks of our model, we first provide a rigorous study of the structural properties of stable trees. We show that every stable tree consists of stable subtrees and that the height of any stable tree is in $\mathcal{O}\left(\frac{\log n}{\log \log n}\right)$. For the root r , which turns out to be the node with the highest in-degree in any stable network, we show that its in-degree is between $\Omega\left(\frac{\log n}{\log \log n}\right)$ and $2^{\mathcal{O}(\sqrt{\log n})}$. This shows that the maximum internal coordination overhead of a single node in any stable tree is rather small.

Our main results are on the quality of equilibrium trees. By using the established structural properties and a connection to the Riemann zeta function we obtain an upper bound on the PoA of 8.62 which is contrasted with a lower bound of 2.4317. For the PoS we derive a lower bound of $\frac{7}{5} - \varepsilon$. Moreover, we give for an infinite number of values for n an upper bound of 2.83 on the PoS. Regarding the Fairness Ratio, we first show that the socially optimal tree is rather unfair, i.e., having a Fairness Ratio of $n \cdot H_n$. In contrast, we prove that any equilibrium tree has a Fairness Ratio in $o(n)$.

This shows that stable trees have only slightly higher social cost compared to the social optimum. In particular, on average every agent pays only a constant factor more than the trivial lower bound for any spanning tree. At the same time stable trees are more fair, have low height and low in-degrees.

We conclude with a brief discussion of the path version extension of our model, where agents select paths as strategies as in [5, 4]. This extension seems promising for future work since we show that allowing a richer strategy space yields a larger set of equilibria and we give equilibria for $n = 16$ and $n = 18$. Hence, in the path-version equilibria may always exist, but the PoA could be higher.

We refer to [9] for all details which were omitted due to space constraints.

2 Structure and Properties of Equilibrium Trees

It is clear that each agent can compute her best response in polynomial time as the number of possible strategies for an agent is n , and the agent can easily compute her cost in linear time. In the following we show that any stable tree consists of stable subtrees, we prove an upper bound of $\mathcal{O}\left(\frac{\log n}{\log \log n}\right)$ to the number of edges of any leaf-to-root path of any stable network, and in the end, we provide bounds on the degree of the root. We start with the statement that any stable tree consists of stable subtrees.

► **Lemma 1.** *If T is stable, then any subtree $T(x)$ is stable in the corresponding subgame.*

Next, we will consider the height of a stable network and need the following technical lemmas.

► **Lemma 2.** *Let $k \in \mathbb{N}$ be the length of a fixed leaf-to-root-path P in a stable network T . Then, for every $1 < i < k$, $d_{i-1} \geq \frac{|T_i|}{|T_i| - |T_{i+1}|}(d_i - 1)$.*

Since $|T_{i+1}| > 0$, Lemma 2 yields that the sequence d_0, d_1, \dots, d_k , is monotonically decreasing.

► **Corollary 3.** *Let $k \in \mathbb{N}$ be the length of a leaf-to-root-path P in a stable network T . Then, for every $1 < i < k$, $d_i \geq d_{i+1}$.*

The next lemma shows that the in-degree of nodes strictly decreases after a constant number of hops.

► **Lemma 4.** *Let $k \in \mathbb{N}$ be the length of a fixed leaf-to-root-path P in a stable network T . Then, for every subtree $T(v)$ with $|T(v)| > 4$ and for every $1 < i < k - 2$ we have $d_{i-1} > d_{i+1}$.*

In the following we investigate upper and lower bounds on the in-degree of the root in stable trees. More precisely, we show an upper bound of $2^{O(\sqrt{\log n})}$ and a lower bound of $\Omega(\log n / \log \log n)$.

► **Theorem 5.** *The in-degree of the root in a stable network T is at least $\frac{\ln(4\sqrt{n/5})}{\ln \ln(4\sqrt{n/5})}$.*

To give an upper bound on the in-degree of the root, we first have to provide the following technical lemmas. The first technical lemma bounds the in-degree of the parent of any leaf.

► **Lemma 6.** *In a stable network T the in-degree of the parent of any leaf is 1.*

The second technical lemma shows how the in-degrees of two sibling nodes are related.

► **Lemma 7.** *Consider a subtree $T(x)$ of a stable network T . Then $\text{indeg}(x) \leq \text{indeg}(v) \cdot \left(1 + \frac{|T(u)|}{|T(v)|}\right) + 1$, where v and u are different children of x .*

From Lemma 7, we derive the following remark and corollary.

► **Remark 8.** Consider a subtree $T(x)$ in a stable network T . Then $\text{indeg}(x) \leq 2 \cdot \text{indeg}(v) + 1$, where v is a root of the second smallest subtree of $T(x)$.

► **Corollary 9.** *If T is a stable network, then every node u in T has at least $\text{indeg}(u) - 1$ children of in-degree at least $(\text{indeg}(u) - 1)/2$.*

Now we can prove an upper bound to the in-degree of the root of any stable tree.

► **Theorem 10.** *The in-degree of the root in a stable network T is $2^{O(\sqrt{\log n})}$.*

Proof. Let T be a stable tree of height h . Let v_h, \dots, v_0 be a leaf-to-root path. Note that the in-degree of the root v_0 is maximal if the in-degree of each node in the v_h - v_0 -path is maximal, i.e., by Lemma 7 and 6, it corresponds to the in-degree sequence $D := (0, 1, d_{h-2}, \dots, d_0)$, where $d_{i-1} = 2d_i + 1$.

Next, we show that nodes at distance $h - 2$ from the root can have an in-degree of at most 2. Assume to the contrary that there is a node u having an edge to a node x such that $\text{indeg}(x) = 3$ and x is at distance $h - 2$ from the root v_0 . As we have proved above, the in-degree of all children of x is at most 1. Thus, u can swap to any leaf node of the subtree $T(x)$. Let T' be the tree obtained after u swapped. If u swaps to a child of x , it decreases its cost by $\text{cost}_T(u) - \text{cost}_{T'}(u) = \frac{3}{2} - \frac{1}{2} - \frac{2}{3} > 0$, i.e., it is an improving move. The swap to a leaf node at distance 2 from x implies an improvement by $\text{cost}_T(u) - \text{cost}_{T'}(u) = \frac{3}{2} - \frac{1}{2} - \frac{1}{3} - \frac{2}{4} > 0$, i.e., it is an improvement. Since T is stable, we get a contradiction. Thus, $D = (0, 1, 2, 5, 11, \dots, d_0)$, i.e.,

$$d_i = 3 \cdot 2^{h-i-2} - 1 \text{ for } i \leq h - 3, \text{ where } d_h = 0, d_{h-1} = 1, d_{h-2} = 2. \quad (1)$$

We now estimate the minimum possible number of nodes in the tree T . By Corollary 9 if the in-degree of a node is equal to k , then it has at least $k - 1$ children with an in-degree of at least $(k - 1)/2$. Thus, starting from the root, the in-degrees of the nodes on each level decrease no more than twice. Hence, the total size of the tree is at least

$$\sum_{i=1}^{h-1} \left(d_i \prod_{j=0}^{i-1} (d_j - 1) \right) > \sum_{i=1}^{h-1} \prod_{j=0}^{i-1} 2^{h-j-2} > 2^{\sum_{j=0}^{h-3} (h-j-2)} = 2^{\frac{(h-1)(h-2)}{2}},$$

where h is the height of T . Thus, $h < \frac{3 + \sqrt{1 + 8 \log n}}{2}$. With equation (1), this implies $d_0 \in 2^{O(\sqrt{\log n})}$. \blacktriangleleft

Now we are able to show that the length of any node-to-root path is $O\left(\frac{\log n}{\log \log n}\right)$.

► **Theorem 11.** *If T is a stable network, then its height $h \in O\left(\frac{\log n}{\log \log n}\right)$.*

Proof. Consider a leaf-to-root path P in T . We show that there are $O\left(\frac{\log n}{\log \log n}\right)$ indices such that $|T_i^P| - |T_{i+1}^P| \geq \sqrt[3]{\log n} \cdot |T_{i+1}^P|$ and $O\left(\frac{\log n}{\log \log n}\right)$ indices such that $|T_i^P| - |T_{i+1}^P| < \sqrt[3]{\log n} \cdot |T_{i+1}^P|$.

Let k be the number of indices i that satisfy $|T_i^P| - |T_{i+1}^P| \geq \sqrt[3]{\log n} \cdot |T_{i+1}^P|$. Then we have that $|T_i^P| = |T_i^P| - |T_{i+1}^P| + |T_{i+1}^P| > \sqrt[3]{\log n} \cdot |T_{i+1}^P|$. Since $|T_i^P| > |T_{i+1}^P|$ for every i , and because $|T_i^P| \leq n$, we have that $|T_0^P| > (\log n)^{k/3}$ and $|T_0^P| = n + 1$, from which we derive $(\log n)^{k/3} \leq n$, i.e., $k = O\left(\frac{\log n}{\log \log n}\right)$.

By Lemma 4 and Corollary 3, there are $O(\sqrt[3]{\log n}) = O\left(\frac{\log n}{\log \log n}\right)$ indices i such that $d_i^P \leq 4\sqrt[3]{\log n}$. We now prove that there are $O\left(\frac{\log n}{\log \log n}\right)$ indices such that $|T_i^P| - |T_{i+1}^P| < \sqrt[3]{\log n} \cdot |T_{i+1}^P|$ and such that $d_i^P \geq 4\sqrt[3]{\log n}$. By Lemma 2 and using the fact that $d_i^P \geq 4\sqrt[3]{\log n}$ and $n \geq 2$, we have that

$$d_{i-1}^P \geq \frac{|T_i^P|}{|T_i^P| - |T_{i+1}^P|} (d_i^P - 1) \geq \frac{1 + \sqrt[3]{\log n}}{\sqrt[3]{\log n}} (d_i^P - 1) \geq \sqrt{\frac{1 + \sqrt[3]{\log n}}{\sqrt[3]{\log n}}} d_i^P.$$

By Corollary 3 we have that $d_{i-1}^P \geq d_i^P$ for every i . Hence $d_0 \geq \left(\frac{1 + \sqrt[3]{\log n}}{\sqrt[3]{\log n}}\right)^{k/2}$. Moreover, by Theorem 10, $d_0^P \leq 2^\alpha \sqrt{\log n}$ for some constant $\alpha > 0$. As a consequence, we have that $\left(\frac{1 + \sqrt[3]{\log n}}{\sqrt[3]{\log n}}\right)^{k/2} \leq 2^\alpha \sqrt{\log n}$, i.e., $2^{\frac{k}{2} \log \frac{1 + \sqrt[3]{\log n}}{\sqrt[3]{\log n}}} \leq 2^\alpha \sqrt{\log n}$, which implies, $k \leq 2\alpha \frac{\sqrt{\log n}}{\log(1 + \sqrt[3]{\log n})}$.

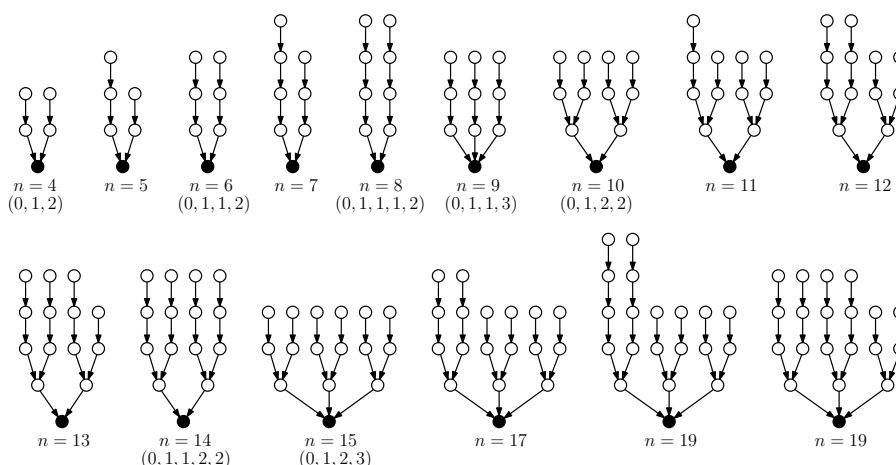
We complete the proof by showing that $\frac{\sqrt{\log n}}{\log \frac{1 + \sqrt[3]{\log n}}{\sqrt[3]{\log n}}} \leq 3 \frac{\log n}{\log \log n}$, i.e., we have to show that $\log \log n \leq 3\sqrt{\log n} \cdot \log \frac{1 + \sqrt[3]{\log n}}{\sqrt[3]{\log n}}$. Let $M = \sqrt[3]{\log n}$. We have to prove that

$$\log M \leq M^{3/2} \log \frac{1 + M}{M} = \log \left(\frac{1 + M}{M} \right)^{M^{3/2}}. \quad (2)$$

By Bernoulli's inequality, $(1 + \frac{1}{M})^{M^{3/2}} \geq 2^{M^{1/2}} \geq M$ for $M \geq 16$. Thus, inequality (2) is satisfied. \blacktriangleleft

3 Existence of Equilibrium Trees

In this section we analyze whether the TCG admits equilibrium trees for all agent numbers n . We first show that in general equilibrium existence is not guaranteed since for $n = 16$ and $n = 18$ no stable tree exists. We contrast this negative result with a NE existence proof for infinitely many agent numbers n . This positive result is achieved for so-called balanced trees, i.e., trees where all nodes with the same distance to the root have the same in-degree. We believe that our positive results can be strengthened to proving that stable trees exist for all n except $n = 16$ and $n = 18$, and we leave this as an intriguing open problem. Figure 2 shows sample equilibrium trees for small n .



■ **Figure 2** Sample equilibrium trees for $n = 4$ to $n = 19$. All depicted trees for $n < 19$ are the unique equilibria for the respective n . For $n = 19$ two equilibrium trees exist. No stable tree exists for $n = 16$ and $n = 18$. The stable trees for $n = 4, 6, 8, 9, 10, 14, 15$ are balanced trees and are annotated with their identifying in-degree sequence of all leaf-to-root paths. (See Section 3.1 for definitions.)

► **Theorem 12.** *For $n = 16$ there exists no stable network.*

The non-existence of a stable tree for $n = 16$ directly implies that the TCG cannot have the finite improvement property, which states that every sequence of improving moves must be finite, i.e., reaches a Nash equilibrium. Thus, since the finite improvement property is equivalent to the game admitting a potential function [32] this implies the following statement.

► **Corollary 13.** *The TCG is not a potential game.*

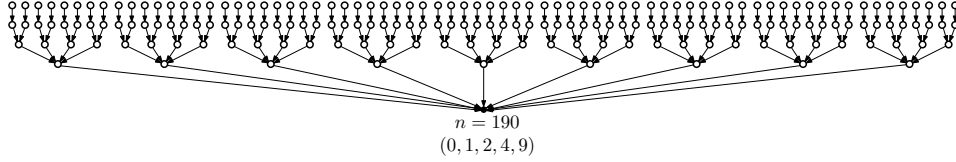
► **Remark 14.** By computational experiments we have obtained equilibrium trees for the TCG for $1 \leq n \leq 100$, except for $n = 16$ and $n = 18$. For $n = 18$ we have verified via a brute-force search over all possible trees that no stable tree exists. Interestingly, for $n \geq 19$ equilibrium trees are no longer unique and in general the number of non-isomorphic equilibrium trees grows as n grows.

3.1 Balanced Trees

Despite the negative result of the non-existence of a stable tree for $n = 16$, in this section we prove the existence of NE's for infinitely many values of n . We prove this result by showing an interesting set of conditions for ruling out potential edge swaps; the proved conditions

15:10 Fair Tree Connection Games with Topology-Dependent Edge Cost

altogether allow us to show that there are infinitely many (balanced) trees that are stable. More precisely, we say that T is *balanced* if any two nodes at the same distance from the root r have equal in-degrees. Note, that any balanced tree T of height h can be uniquely encoded by a sequence of node degrees $(0, d_{h-1}, \dots, d_0)$, where d_i is an in-degree of nodes at level i , i.e., at distance i from the root. In this section we show that all the balanced trees of the form $(0, 1, 2, 4, d_{h-4}, \dots, d_0)$ such that $d_i < d_{i-1} \leq 2d_i + 1$, for every $1 \leq i \leq h - 4$ are stable. (See Figure 3 for an example.)



■ **Figure 3** Sample of an extremal balanced tree with degree sequence $(0, 1, 2, 4, 9)$.

► **Theorem 15.** *The balanced tree T with degree sequence $(0, 1, 2, 4, d_{h-4}, \dots, d_0)$, where $d_{j+1} < d_j \leq 2d_{j+1} + 1$ for every $j \leq h - 4$, is stable.*

From Theorem 15 we derive the following corollary.

► **Corollary 16.** *The TCG with n agents admits a NE for infinitely many values of $n \in \mathbb{N}$.*

By Corollary 16, we observe that NE exists for all n that admit an existence of a balanced tree. Intuitively, a minor modification of a balanced tree, e.g., removing a subset of leaf nodes, keep the tree stable. Moreover, for $n \geq 19$ we have found several non-isomorphic equilibrium trees in each case. The number of non-isomorphic equilibria grows with n , which indicates that for growing n also the number of possibilities how to combine suitable equilibrium trees into larger equilibrium trees grows. Therefore, we conjecture the existence of stable trees for all values of n except for $n = 16$ and $n = 18$. We believe that this conjecture can be proven by a dynamic programming approach that exploits the different possibilities of how equilibrium sub-trees can be combined into larger equilibrium trees.

► **Conjecture 1.** *For any $n \in \mathbb{N}$, with $n \neq 16$ and $n \neq 18$ a pure NE exists in the TCG.*

4 Quality of Equilibrium Trees

In this section we provide results on the quality of stable networks. In particular, we prove a constant upper bound on the PoA and give lower bounds on the PoA and PoS. Furthermore, we prove an upper bound on the PoS for certain balanced trees. We first observe that any network in which at least one node has in-degree 2 is not a social optimum. Hence, a Hamilton path is the social optimum.

► **Theorem 17.** *Any Hamiltonian path having the root r as one endnode is a social optimum.*

4.1 Price of Anarchy

In every network T for all $v \in V$, $\text{indeg}_T(v) \leq n$, since there are exactly n edges. Hence, the cost of an agent is upper bounded by n and the star graph yields a trivial upper bound of n for the PoA. However, we prove next a constant upper bound on the PoA.

► **Theorem 18.** *The PoA is at most 8.62.*

Proof. Consider a stable network $T = (V, E)$. By Theorem 17, the social optimum is a path of cost n . Hence, it is enough to show that in T the maximum cost of an agent is upper bounded by a constant. We clearly have that the cost incurred by a non-leaf agent is strictly smaller than the cost incurred by any of its descendants. Therefore, the maximum costs is achieved by a leaf agent.

Consider two leafs u and v in T such that u pays the maximum cost. Let P_v be the node-to-root path starting from the parent of v . Since T is stable, $cost_T(u) < 1 + 1/2 + \sum_{(i,j) \in P_v} \frac{indeg(j)}{|T(i)|} = 1/2 + cost_T(v)$. Therefore, we only have to show that there exists a leaf agent v with a constant cost value.

We now prove that such a leaf agent always exists. By Corollary 9, each node of in-degree d has at least $d - 1$ children of in-degree at least $\lceil (d - 1)/2 \rceil$. Consider a root-to-leaf path $P = (r = v_0, \dots, v_h = v)$ where each next hop goes always towards the smallest appended subtree where the root has an in-degree of at least half of the node's in-degree minus one, i.e., for any $v_i \in P$, $v_{i+1} = \operatorname{argmin}\{|T(w)| : (w, v_i) \in E \text{ and } indeg(w) \geq (indeg(v_i) - 1)/2\}$. Then for every $0 \leq i \leq h - 1$, $|T(v_i)| \geq (indeg(v_i) - 1)|T(v_{i+1})| + 2$.

Denote by $|t_k|$ the size of the minimum stable tree with a root of in-degree k . Then by Corollary 9 and Lemma 6 it holds that

$$|t_0| \geq 1, |t_1| \geq 2, t_k \geq (k - 1) \cdot |t_{\lceil (k-1)/2 \rceil}| + 2. \quad (3)$$

We show via induction that for any $k \geq 11$, $|t_k| \geq (2k + 1)k^2$. Indeed, it holds that $|t_{k+1}| \geq k \cdot |t_{\lceil k/2 \rceil}| + 2 > (k + 1)k^3/2^2 \geq (2(k + 1) + 1)(k + 1)^2$, where the last inequality holds for all $k \geq 11$.

The overall cost incurred by the leaf v is at most the costs incurred by v for all edges $(v_i, v_{i-1}) \in P$ where in-degree of v_{i-1} is less than the cost incurred by v for all other edges in P plus 2.

By Lemma 4, each leaf-to-root path has at most three nodes of in-degree 1, which implies that v pays at most $p_1 := \frac{11}{6}$ for all edges ending in a node with in-degree equals 1.

By Lemma 4, the in-degrees of the nodes in the leaf-to-root path P strictly increase with at least every second hop. This implies that for $i \leq h - 4 - (11 - 1) \cdot 2 = h - 24$ it is guaranteed that $indeg(v_i) \geq 11$. Hence, starting from the first node having in-degree at least 11 in P , agent v pays

$$\begin{aligned} p_2 &:= \sum_{i=1}^{h-24} \frac{indeg(v_{i-1})}{|T(v_i)|} \leq \sum_{i=1}^{h-24} \frac{2indeg(v_i) + 1}{|t_{indeg(v_i)}|} \leq \sum_{i=1}^{h-24} \frac{1}{(indeg(v_i))^2} \\ &\leq 2 \cdot \sum_{i=11}^{\infty} \frac{1}{i^2} < 2 \left(\zeta(2) - \sum_{i=1}^{10} \frac{1}{i^2} \right), \end{aligned}$$

where $\zeta(s)$ is the Riemann zeta function. Hence, $p_2 < 0.2$.

Finally, we need to evaluate the cost of the path P for all nodes v_i with the in-degree $2 \leq indeg(v_i) \leq 10$. Since for every $0 \leq i \leq h - 1$, $|T(v_i)| \geq (indeg(v_i) - 1)|T(v_{i+1})| + 2$, each edge (v_i, v_{i+1}) in the path P costs at most

$$\frac{2indeg(v_i) + 1}{|T(v_i)|} \leq \frac{2indeg(v_i) + 1}{(indeg(v_i) - 1)|T(v_{i+1})|} = \frac{2}{|T(v_{i+1})|} + \frac{3}{(indeg(v_i) - 1)|T(v_{i+1})|}.$$

By applying the inequality (3) and since the in-degrees of the nodes in P increase at most with every second level, it holds that the total cost of the subpath is at most $p_3 :=$

15:12 Fair Tree Connection Games with Topology-Dependent Edge Cost

$2 \sum_{i=2}^9 \frac{2}{t_i} + \frac{2}{t_1} + \frac{2}{t_{10}} + \sum_{i=2}^{10} \left(\frac{3}{(i-1)t_{i-1}} + \frac{3}{(i-1)t_i} \right) < 3.12 + 2.975 < 6.01$. Therefore, the total cost of the path P paid by an agent v is strictly less than $p_1 + p_2 + p_3 < 8.12$. This implies that the PoA is at most 8.62. ◀

We now prove a lower bound to the PoA using the extremal stable balanced trees of Theorem 15. (See Figure 3.) For the rest of this section, let \mathcal{T}_h denote the extremal balanced tree of height $h \geq 1$ and degree sequence $d_h = 0$, $d_{h-1} = 1$, $d_{h-2} = 2$ (if $h \geq 2$), $d_{h-3} = 4$ (if $h \geq 3$), and $d_i = 2d_{i+1} + 1$ for every $i \leq h-4$. We will denote by sc_h and n_h the social cost and the number of nodes (root included) of \mathcal{T}_h .

► **Theorem 19.** *The PoA is at least 2.4317.*

Next, we prove an upper bound to the average agent's cost in \mathcal{T}_h and provide an interesting conjecture. We define $a_h := sc_h / (n_h - 1)$ as the average agent's cost in \mathcal{T}_h .

► **Lemma 20.** *For every $h \geq 1$, $a_h \leq 2.4318$.*

► **Conjecture 2.** *The PoA is equal to $\lim_{h \rightarrow \infty} a_h$.*

4.2 Price of Stability

We now turn our focus to the PoS and prove a lower bound.

► **Theorem 21.** *The PoS is at least $\frac{7}{5} - \varepsilon$, for $\varepsilon \in \Theta(1/n)$.*

Next, we investigate the PoS in certain balanced trees and prove an upper bound which is strictly better than the upper bound on the PoA.

► **Theorem 22.** *For all $n \in \mathbb{N}$ such that there is a balanced tree T of size n with the in-degree sequence $(0, 1, 2, 4, d_{h-4}, \dots, d_0)$, where $d_i \leq 2d_{i+1} + 1$ for $i \leq h-4$, the PoS is at most 2.83.*

4.3 Fairness measure

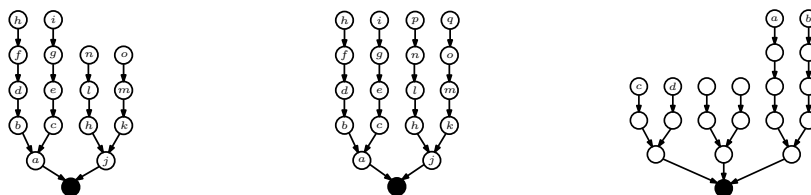
We investigate the Fairness Ratio which considers the cost distribution among the agents. We show that stable trees admit a more fair cost-sharing compared with the social optimum.

► **Theorem 23.** *The Fairness Ratio for OPT_n is nH_n , where $H_n = \sum_{i=1}^n \frac{1}{i}$ is the n -th harmonic number.*

We now turn our focus to the analysis of the class of all stable trees and prove that the FR is in $o(n)$.

► **Theorem 24.** *The Fairness Ratio for any NE is at most $\frac{8.62(n-2) \cdot \ln \ln(4\sqrt{n/5})}{\ln(4\sqrt{n/5})} - 2^{13}$.*

$\left(1 - \frac{2 \ln \ln(4\sqrt{n/5})}{\ln(4\sqrt{n/5})}\right) \cdot \left(\frac{\ln(4\sqrt{n/5})}{\ln \ln(4\sqrt{n/5})}\right)^{\log \left(\sqrt{\frac{\ln(4\sqrt{n/5})}{\ln \ln(4\sqrt{n/5})}}\right) - 5.5}$, which is at most $8.62 \cdot \frac{(n-2) \cdot \ln \ln(4\sqrt{n/5})}{\ln(4\sqrt{n/5})}$.



■ **Figure 4** Left: A path-TCG NE that is not a TCG NE for $n = 16$. Middle: A path-TCG NE that is not a TCG NE for $n = 18$. Right: A TCG NE that is not a strong path-TCG NE.

► **Theorem 25.** *The Fairness Ratio for a stable tree is at least $n \cdot 2^{-2\sqrt{2\log(n)}}$.*

Finally, we investigate the class of stable balanced trees and prove a more precise upper bound.

► **Theorem 26.** *The Fairness Ratio for a stable balanced tree with the in-degree sequence $(0, 1, 2, 4, d_{h-4}, \dots, d_0)$, where $d_i \leq 2d_{i+1} + 1$ for $i \leq h - 4$, is at most $\frac{2.4318n \cdot (\ln \ln(4\sqrt{n/5}))^2}{(\ln(4\sqrt{n/5}))^2}$.*

5 Extensions for Future Work: The Path Version and Coalitions

A natural extension of our model is to allow for a richer strategy space. Instead of selecting a single outgoing edge, agents could strategically select a complete path towards the root r . This version, called the *path-TCG*, is closer to the fair single-source connection game by Anshelevich et al. [5, 4]. See the full version [9] for a formal definition of the path-TCG.

We give some preliminary results relating the equilibria of the TCG to the equilibria of the path-TCG. Our results indicate that studying the path-TCG, in particular its PoA and PoS, is a promising next step. We start with showing that also in the path-TCG all equilibria must be trees.

► **Lemma 27.** *Any equilibrium network in the path-TCG is a tree.*

Now we show that the TCG can be considered as a special case of the path-TCG since all equilibrium trees of the TCG are equilibria in the path-TCG but not vice versa.

► **Theorem 28.** *The set of NE in the path-TCG is a superset of the set of NE in the TCG.*

We showed for the TCG that for $n = 16$ and $n = 18$ there exists no stable network. We contrast this negative result with a NE existence proof for the path-TCG for the corresponding values. Figure 4 (left and middle) show equilibrium trees for the path-TCG for $n = 16$ and $n = 18$, respectively.

► **Theorem 29.** *For $n = 16$ and $n = 18$ there exists a stable network for the path-TCG.*

Together with Theorem 29 and since any NE in the TCG is a NE in the path-TCG, we go along with Conjecture 1 and believe that for all values of n stable trees exist for the path-TCG.

► **Conjecture 3.** *For any $n \in \mathbb{N}$ a pure NE exists in the path-TCG.*

An agent a in the TCG benefits from the fact that if a changes her strategy and switches her edge towards another node the costs of the new edge is also shared among all of a 's ancestors. It seems natural to consider a strategy change in the TCG as a coalitional strategy change in the path-TCG by the coalition consisting of agent a and all her ancestors. So NE in the TCG could be in strong NE [3] for the pathTCG. However, we show that this is not true, see Figure 4 (right).

► **Theorem 30.** *There is a NE in the TCG which is not in strong NE for the path-TCG.*

6 Conclusion

We have studied a tree formation game to investigate how selfish agents self-organize to connect to a common target in the presence of dynamic edge costs that are sensitive to node degrees. This mimics settings in which nodes can charge prices for offering their routing service and where these prices are guided by supply and demand, i.e., more popular nodes with higher in-degree can charge higher prices to make up for their increased internal coordination cost.

Our main findings are that our game admits equilibrium trees with intriguing properties like low height, low maximum degree, almost optimal cost, and a somewhat fair distribution of the total cost among the agents. The set of equilibrium trees seems to be combinatorially rich, and characterizing stable trees that are not balanced seems an exciting and challenging problem for future research. It would also be interesting to study the degree distribution in stable trees and to evaluate possible connections with power-law degree distributions which are ubiquitous in real-world networks.

We note in passing that our model can easily be generalized to settings with more than one target node as long as every possible incident edge may be activated. In this case, several disjoint trees, one for each target node, will be formed. Things change if target nodes and agent nodes may be co-located, and exploring this variant might be interesting.

References

- 1 Susanne Albers and Pascal Lenzner. On Approximate Nash Equilibria in Network Design. *Internet Mathematics*, 9(4):384–405, 2013.
- 2 Carme Àlvarez and Arnau Messegué. On the Price of Anarchy for High-Price Links. In *WINE'19*, pages 316–329. Springer, 2019.
- 3 Nir Andelman, Michal Feldman, and Yishay Mansour. Strong price of anarchy. *Games and Economic Behavior*, 65(2):289–317, 2009.
- 4 Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Eva Tardos, Tom Wexler, and Tim Roughgarden. The Price of Stability for Network Design with Fair Cost Allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.
- 5 Elliot Anshelevich, Anirban Dasgupta, Éva Tardos, and Tom Wexler. Near-Optimal Network Design with Selfish Agents. *Theory of Computing*, 4(1):77–109, 2008.
- 6 Venkatesh Bala and Sanjeev Goyal. A Noncooperative Model of Network Formation. *Econometrica*, 68(5):1181–1229, 2000.
- 7 Davide Bilò and Pascal Lenzner. On the Tree Conjecture for the Network Creation Game. *Theory of Computing Systems*, 64(3):422–443, 2020.
- 8 Vittorio Bilò, Michele Flammini, and Luca Moscardelli. The Price of Stability for Undirected Broadcast Network Design with Fair Cost Allocation is Constant. *Games and Economic Behavior*, 2014.
- 9 Davide Bilò, Tobias Friedrich, Pascal Lenzner, Anna Melnichenko, and Louise Molitor. Fair tree connection games with topology-dependent edge cost, 2020. [arXiv:2009.10988](https://arxiv.org/abs/2009.10988).

- 10 Charles G. Bird. On Cost Allocation for a Spanning Tree: A Game Theoretic Approach. *Networks*, 6(4):335–350, 1976.
- 11 Ankit Chauhan, Pascal Lenzner, Anna Melnichenko, and Louise Molitor. Selfish Network Creation with Non-uniform Edge Cost. In *SAGT'17*, pages 160–172. Springer, 2017.
- 12 Chandra Chekuri, Julia Chuzhoy, Liane Lewin-Eytan, Joseph Naor, and Ariel Orda. Non-cooperative Multicast and Facility Location Games. *IEEE Journal on Selected Areas in Communications*, 25(6):1193–1206, 2007.
- 13 Armin Claus and Daniel J. Kleitman. Cost Allocation for a Spanning Tree. *Networks*, 3(4):289–304, 1973.
- 14 Erik D. Demaine, Mohammad Taghi Hajiaghayi, Hamid Mahini, and Morteza Zadimoghaddam. The Price of Anarchy in Network Creation Games. *ACM Transactions on Algorithms*, 8(2):13, 2012.
- 15 Shayan Ehsani, Saber Shokat Fadaee, MohammadAmin Fazli, Abbas Mehrabian, Sina Sadeghian Sadeghabad, Mohammad Ali Safari, and Morteza Saghafian. A Bounded Budget Network Creation Game. *ACM Transactions on Algorithms*, 11(4):1–25, 2015.
- 16 Stephan Eidenbenz, Sritesh Kumar, and Sibylle Züst. Equilibria in Topology Control Games for Ad hoc Networks. *Mobile Networks and Applications*, 11(2):143–159, 2006.
- 17 Bruno Escoffier, Laurent Gourvès, Jérôme Monnot, and Stefano Moretti. Cost Allocation Protocols for Network Formation on Connection Situations. In *ICST'12*, pages 228–234, 2012.
- 18 Alex Fabrikant, Ankur Luthra, Elitza Maneva, Christos H. Papadimitriou, and Scott Shenker. On a network creation game. In *PODC'03*, pages 347–351. ACM, 2003.
- 19 Michal Feldman, Kevin Lai, and Li Zhang. The Proportional-share Allocation Market for Computational Resources. *IEEE Transactions on Parallel and Distributed Systems*, 20(8):1075–1088, 2008.
- 20 Amos Fiat, Haim Kaplan, Meital Levy, Svetlana Olonetsky, and Ronen Shabo. On the Price of Stability for Designing Undirected Networks with Fair Cost Allocations. In *ICALP'06*, pages 608–618. Springer, 2006.
- 21 Laurent Gourvès and Jérôme Monnot. Three Selfish Spanning Tree Games. In *WINE'08*, pages 465–476. Springer, 2008.
- 22 Daniel Granot and Gur Huberman. Minimum Cost Spanning Tree Games. *Mathematical programming*, 21(1):1–18, 1981.
- 23 Daniel Granot and Gur Huberman. On the Core and Nucleolus of Minimum Cost Spanning Tree Games. *Mathematical programming*, 29(3):323–347, 1984.
- 24 Martin Hoefer. Non-Cooperative Tree Creation. *Algorithmica*, 53(1):104–131, 2009.
- 25 Martin Hoefer and Piotr Krysta. Geometric Network Design with Selfish Agents. In *COCOON'05*, pages 167–178, 2005.
- 26 Matthew O. Jackson and Asher Wolinsky. A Strategic Model of Social and Economic Networks. *Journal of Economic Theory*, 71(1):44–74, 1996.
- 27 Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *STACS'99*, pages 404–413. Springer-Verlag, 1999.
- 28 Thomas L. Magnanti and Richard T. Wong. Network Design and Transportation Planning: Models and Algorithms. *Transportation Science*, 18(1):1–55, 1984.
- 29 Akaki Mamageishvili, Matúš Mihalák, and Dominik Müller. Tree Nash Equilibria in the Network Creation Game. In *WAW'13*, pages 118–129. Springer, 2013.
- 30 Matúš Mihalák and Jan Christoph Schlegel. The Price of Anarchy in Network Creation Games is (Mostly) Constant. In *SAGT'10*, pages 276–287. Springer, 2010.
- 31 Kimaya Mittal, Elizabeth M. Belding, and Subhash Suri. A Game-theoretic Analysis of Wireless Access Point Selection by Mobile Users. *Computer Communications*, 31(10):2049–2062, 2008.
- 32 Dov Monderer and Lloyd S. Shapley. Potential Games. *Games and Economic Behavior*, 14(1):124–143, 1996.
- 33 Hervé Moulin and Scott Shenker. Strategyproof Sharing of Submodular Costs: Budget Balance versus Efficiency. *Economic Theory*, 18(3):511–533, 2001.