The Impact of Cooperation in Bilateral Network Creation

Tobias Friedrich
tobias.friedrich@hpi.de
Hasso Plattner Institute
University of Potsdam
Potsdam, Germany

Hans Gawendowicz
hans.gawendowicz@hpi.de
Hasso Plattner Institute
University of Potsdam
Potsdam, Germany

Pascal Lenzner
pascal.lenzner@hpi.de
Hasso Plattner Institute
University of Potsdam
Potsdam, Germany

Arthur Zahn
arthur.zahn@student.hpi.de
Hasso Plattner Institute
University of Potsdam
Potsdam, Germany

ABSTRACT

Many real-world networks, like the Internet or social networks, are not the result of central design but instead the outcome of the interaction of local agents that selfishly optimize their individual utility. The well-known Network Creation Game by Fabrikant, Luthra, Maneva, Papadimitriou, and Shenker [23] models this. There, agents corresponding to network nodes buy incident edges towards other agents for a price of $\alpha > 0$ and simultaneously try to minimize their buying cost and their total hop distance. Since in many real-world networks, e.g., social networks, consent from both sides is required to establish and maintain a connection, Corbo and Parkes [14] proposed a bilateral version of the Network Creation Game, in which mutual consent and payment are required in order to create edges. It is known that this cooperative version has a significantly higher Price of Anarchy compared to the unilateral version. On the first glance this is counter-intuitive, since cooperation should help to avoid socially bad states. However, in the bilateral version only a very restrictive form of cooperation is considered.

We investigate this trade-off between the amount of cooperation and the Price of Anarchy by analyzing the bilateral version with respect to various degrees of cooperation among the agents. With this, we provide insights into what kind of cooperation is needed to ensure that socially good networks are created. As a first step in this direction, we focus on tree networks and present a collection of asymptotically tight bounds on the Price of Anarchy that precisely map the impact of cooperation. Most strikingly, we find that weak forms of cooperation already yield a significantly improved Price of Anarchy. In particular, the cooperation of coalitions of size 3 is enough to achieve constant bounds. Moreover, for general networks we show that enhanced cooperation yields close to optimal networks for a wide range of edge prices. Along the way, we disprove an old conjecture by Corbo and Parkes [14].

PODC '23, June 19–23, 2023, Orlando, FL, USA

ACM ISBN 979-8-4007-0121-4/23/06.
https://doi.org/10.1145/3583668.3594588

CCS CONCEPTS

• **Theory of computation → Network formation**; **Quality of equilibria**; **Social networks**; **Algorithmic game theory**.

KEYWORDS

Network Creation Games, Cooperation, Price of Anarchy

ACM Reference Format:
Tobias Friedrich, Hans Gawendowicz, Pascal Lenzner, and Arthur Zahn. 2023. The Impact of Cooperation in Bilateral Network Creation. In *ACM Symposium on Principles of Distributed Computing (PODC '23), June 19–23, 2023, Orlando, FL, USA.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3583668.3594588

1 INTRODUCTION

Many real-world problems are related to networks or connections between entities. Given this, research on networks is concerned with the creation of efficient networks with regard to various objective functions. Traditionally, such networks are created by a centralized algorithm, like in the cases of Minimum Spanning Trees [27], Topology Control Problems [40], and Network Design Problems [33]. Such central orchestration is a good model if the whole network belongs to one real-world entity, i.e., a firm designing its internal communication network, who governs the network and centrally pays for its infrastructure. However, most large real-world networks instead emerged from the interaction of multiple entities, each with their own goals, who each have control over a local part of the network. For example, the Internet is a network of networks where each subnetwork is centrally controlled by an Internet service provider. Hence, to understand the formation of such networks, game-theoretic agent-based models are needed. For such models, one of the prime questions is to understand the impact of the agents' selfishness on the overall quality of the created networks. This impact is typically measured by the Price of Anarchy (PoA) [31].

The Network Creation Game (NCG) by Fabrikant, Luthra, Maneva, Papadimitriou, Shenker [23] is one of the most prominent game-theoretic models for the formation of networks. There, the agents correspond to nodes in a network and each agent can unilaterally build incident edges to other nodes for an edge price of $\alpha > 0$. All agents simultaneously try to minimize (1) their total hop distance to everyone else and (2) the cost they incur for building connections.

As both incentives are in direct conflict, the challenge is to analyze how this tension is resolved in the equilibria of the game. The NCG found widespread appeal and many variants have recently been studied. One of the earliest variants is the Bilateral Network Creation Game (BNCG) by Corbo and Parkes [14], which models that not all real-world settings allow unilateral edge formation. For instance, social networks require mutual consent and both sides to invest their time and effort in order to maintain a connection. Consequently, the BNCG demands that both incident agents pay for an edge to establish it. From this arises a need for coordination. Thus, Corbo and Parkes [14] did not analyze the standard Pure Nash Equilibrium (NE) as solution concept, but instead focused on the well-known concept of Pairwise Stability (PS) [28], in which two agents are allowed to cooperate to form a mutual connection.

Interestingly, despite an abundance of literature on the NCG, its bilateral variant remains widely unexplored. This is even more astonishing given the known results on the PoA for both models: While the PoA with respect to NE of the NCG is known to be constant for most ranges of α, the PoA with respect to PS of the BNCG was proven to be high [14, 18]. Hence, the required cooperation for establishing edges leads to socially worse equilibrium states. But cooperation among the agents should be beneficial and should reduce the social cost of equilibrium states. Thus, the problem with the analysis of the BNCG seems to be the drastically limited amount of cooperation allowed by Pairwise Stability. This directly gives rise to a very natural question: How much cooperation among the agents is actually needed to ensure a low Price of Anarchy?

We answer this question by investigating the impact of different amounts of cooperation on the quality of the resulting equilibrium states, i.e., the social cost of the created networks. We establish the positive result that allowing slightly more cooperation than allowed by PS already has a significant impact on the PoA.

1.1 Model and Notation

For $i, j \in \mathbb{N}$, with $i \leq j$, we denote the set $\{k \in \mathbb{N} \mid i \leq k \leq j\}$ as $[i..j]$ and the set $[1..i]$ as $[i]$. Whenever we use the logarithm, the base is always 2 unless specified otherwise.

Graphs: We model networks as *graphs*, consisting of nodes and undirected edges, defined as a pair $G = (V_G, E_G)$, where V_G is the set of nodes and E_G the set of edges. The number of nodes $|V_G|$ is $n_G > 0$. We omit the subscript if the graph is clear from the context. Each edge $e \in E_G$ is a subset of V_G and consists of the two distinct nodes it connects. For two nodes $u, v \in V_G$, we use uv for $\{u, v\}$ and say that u and v are *neighbors* in G if $uv \in E_G$.

The *distance* $\text{dist}_G(u, v)$ between two nodes $u, v \in V$ in a graph G is defined as the number of edges on a shortest path from u to v in G. Since our graphs are undirected, we have $\text{dist}(u, v) = \text{dist}(v, u)$. Furthermore, we assume $\text{dist}(u, u) = 0$. If no path exists, we consider the distance to be extremely large. For technical reasons, we define the constant $M \in \mathbb{N}$ and say $\text{dist}(u, v) = M$ in such cases. We elaborate on this constant M below. The *total distance* of a node $u \in V$ to multiple other nodes $V' \subseteq V$ is defined by $\text{dist}(u, V') = \sum_{v \in V'} \text{dist}(u, v)$. In the special case $V = V'$, we use the shorthand $\text{dist}(u) = \text{dist}(u, V)$ and call this the total distance cost of node u. We define the *extended neighborhood* of a node $u \in V$ as follows: for $i \in \mathbb{N}$, we refer to the set of all nodes of distance at most i from u as the set $\text{Neigh}_G^{\leq i}(u) = \{v \in V \mid \text{dist}(u, v) \leq i\}$. Analogously, we define $\text{Neigh}^{=i}(u)$ and note that $\text{Neigh}^{=1}(u)$ denotes the *neighborhood* of u.

For a graph G and $u, v \in V$, the graph $G - uv$ refers to G without edge uv, i.e., $(V, E \setminus \{uv\})$. By contrast, if we consider adding uv to G, we denote the resulting graph as $G + uv$.

(Bilateral) Network Creation Game: The Network Creation Game consists of n selfish *agents* who are building a network among each other. These agents correspond to the node set V of the resulting graph, so we will use the terms agent and node interchangeably. Each agent $u \in V$ has a *strategy* $S_u \subseteq V \setminus \{u\}$, which specifies towards which other nodes the agent u wants to create edges. The individual strategies of all agents are combined into a *strategy vector* S, which encapsulates a state of the game. The *created graph* G is derived from this strategy vector. We consider two versions, the *Unilateral Network Creation Game (NCG)* [23] and the *Bilateral Network Creation Game (BNCG)* [14]. In the NCG, for two nodes $u, v \in V$, the edge uv is part of G exactly if $u \in S_v$ or $v \in S_u$. In contrast, in the BNCG, we have $uv \in E$ if and only if $u \in S_v$ and $v \in S_u$, i.e., both agents need to agree on the creation of the edge. In this work, we will mostly focus on the BNCG.

Building edges is not free in the (B)NCG. Instead, there is a parameter $\alpha \in \mathbb{R}_{>0}$ which determines the buying cost of each edge. An agent $u \in V$ incurs the buying cost $\text{buy}_S(u) = \alpha\,|S_u|$, so she has to pay for each target node in her strategy. Note that in the BNCG an agent u has to pay for $v \in S_u$ even if $u \notin S_v$ and thus $uv \notin E$. Similarly, edges might be paid for twice in the NCG if $u \in S_v$ and $v \in S_u$. However, such inefficiencies will not arise in equilibria and hence we will ignore them. Given this, we have a bijection between strategy vectors and created graphs in the BNCG. Thus, we can abstract away from the underlying strategy vector S and directly consider the created graph G.

Each agent u aims to minimize her *total cost* in graph G, denoted by $\text{cost}_G(u)$ and defined as the sum of her buying cost and her total hop distance to all other agents, i.e.,

$$\text{cost}_G(u) = \text{buy}_G(u) + \text{dist}_G(u) = \alpha|S_u| + \sum_{v \in V_G} \text{dist}_G(u, v).$$

Remember that we define the distance between two disconnected nodes to be M. We set M to some value larger than αn^3 to enforce that, for $i \in [n]$, an agent should prefer any graph where she can reach i agents over any graphs where she can reach at most $i - 1$ agents, but given the same reachability she should prefer to minimize her buying and distance cost.

Solution Concepts: As all agents are selfish, the created graphs are the result of decentralized decisions instead of central coordination. When analyzing the properties of such games, we are especially interested in *equilibria*, which describe strategy vectors that are stable against specific types of strategy changes. Non-equilibrium states might not persist, because the agents would defect from such states in order to decrease their cost. We consider different solution concepts, where each of them is characterized by the types of improving strategy changes it is stable against. We say that a strategy change is *improving* for some agent if the change yields strictly lower cost for this agent.

The most prominent solution concept for the NCG is the *Pure Nash Equilibrium (NE)*. A strategy vector S is in NE if and only if no agent $u \in V$ can strictly decrease her cost by changing her strategy S_u. However, the NE is not suited for the BNCG, since a node $u \in V$ may delete edges unilaterally by removing them from her strategy, but she cannot add new edges by changing only her own strategy. Hence, a meaningful solution concept for the BNCG requires that multiple agents are able to coordinate to change their strategies in an atomic step. We consider the following solution concepts, presented in order of increasing amount of cooperation. Newly introduced concepts are marked with a "⋆", concepts adapted from the unilateral version are marked with "†".

- **Remove Equilibrium⋆(RE):** A strategy vector S is a *Remove Equilibrium (RE)* if no agent $u \in V$ can improve by removing a single node from her strategy S_u.
- **Bilateral Add Equilibrium†(BAE):** A strategy vector S is a *Bilateral Add Equilibrium (BAE)*[1] if there are no $u, v \in V$ such that both improve by adding u to S_v and v to S_u.
- **Pairwise Stability (PS) [28]:** A strategy vector S is *pairwise stable (PS)*, if it is in Remove Equilibrium and in Bilateral Add Equilibrium.
- **Bilateral Swap Equilibrium†(BSwE):** A strategy vector S is a *Bilateral Swap Equilibrium (BSwE)*[2] if there are no nodes $u, v, w \in V$ with $u \in S_v$, $v \in S_u$ and $u \notin S_w$ or $w \notin S_u$ such that u and w can improve by replacing v by w in S_u and adding u to S_w.
- **Bilateral Greedy Equilibrium†(BGE):** A strategy vector S is in *Bilateral Greedy Equilibrium (BGE)*[3] if it is pairwise stable and in Bilateral Swap Equilibrium.
- **Bilateral Neighborhood Equilibrium⋆(BNE):** A strategy vector S is in *Bilateral Neighborhood Equilibrium (BNE)* if there is no agent $u \in V$ with the following type of improving move. Let $R \subseteq S_u$ and let $A \subseteq V \setminus S_u$. Removing the edges between u and R and adding the edges between u and A is an improving move if and only if u and all nodes in A strictly benefit from the whole change[4].
- **Bilateral (k-)Strong Equilibrium†((k-)BSE):** A strategy vector S is in *Bilateral k-Strong Equilibrium (k-BSE)*[5] if there is no coalition $\Gamma \subseteq V$ of size at most k such that there is the following type of improving move. The move can delete a subset of edges $R \subseteq E$, as long as for each edge $uv \in R$ it holds that $uv \cap \Gamma \neq \emptyset$. At the same time, it can add a set of new edges $A \subseteq 2^V \setminus E$, if for all $uv \in A$ it holds that u and v are both included in Γ. The move is improving if all nodes in Γ strictly benefit from it. A graph G is in *Bilateral Strong Equilibrium (BSE)* if it is in n-BSE.

Throughout the different solutions concepts above, we can see the trend of increasingly enhanced coordination, with the BSE admitting the strongest form of agent collaboration.

Quality of Equilibria: The *social cost* of G is defined as the total cost incurred by all agents, denoted by $\text{cost}(G) = \sum_{u \in V} \text{cost}(u)$. We define the *total distance cost* of G as $\text{dist}(G) = \sum_{u \in V} \text{dist}(u)$ and the *total buying cost* of G as $\text{buy}(G) = \sum_{u \in V} \text{buy}(u)$. Graphs with minimum social cost for given n and α are called *social optima*.

For a graph G with n nodes and parameter α, we define the *social cost ratio* $\rho(G)$ as follows. Let OPT be a social optimum for n and α, then $\rho(G) = \frac{\text{cost}(G)}{\text{cost}(\text{OPT})}$, i.e., the ratio of the social costs of G and OPT. In particular, $\rho(G)$ is 1 exactly if G is a social optimum.

The *Price of Anarchy (PoA)* [31] for a specific solution concept is a function of $n \in \mathbb{N}$ and $\alpha \in \mathbb{R}_{>0}$. For given n and α, let $\mathcal{G}$ denote the set of all graphs with n nodes which meet the considered equilibrium definition. Then, the PoA for that n and α is defined as $\max\{\rho(G) \mid G \in \mathcal{G}\}$. We will mainly consider the PoA of tree networks, which is defined as the general PoA but there $\mathcal{G}$ is further refined to only contain equilibrium networks that are trees. The PoA is a measure for the worst-case inefficiency and we are particularly interested in the asymptotics for the different solution concepts, i.e., in bounds on the PoA that depend on the number of agents n and the edge price α.

1.2 Related Work

The NCG was introduced by Fabrikant, Luthra, Maneva, Papadimitriou, and Shenker [23]. They showed that trees in NE have a PoA of at most 5 and conjectured that there is a constant $A \in \mathbb{N}$ such that all graphs in NE are trees if $\alpha \geq A$. While the original tree conjecture was disproven [1], it was reformulated to hold for $\alpha \geq n$. This is best possible, since non-tree networks in NE exist for $\alpha < n$ [34]. A recent line of research [1–3, 10, 20, 34, 38] proved the adapted conjecture to hold for $\alpha > 3n-3$, further refined the constant upper bounds on the PoA for tree networks, and showed that the PoA is constant if $\alpha > n(1+\varepsilon)$, for any $\varepsilon > 0$. The currently best general bound on the PoA was established by Demaine, Hajiaghayi, Mahini, and Zadimoghaddam [19]. They derived a PoA bound of $o(n^{\varepsilon})$, for any $\varepsilon > 0$, and gave a constant upper bound for $\alpha \in O(n^{1-\varepsilon})$. Besides bounds on the PoA, the complexity of computing best response strategies and the convergence to equilibria via sequences of improving moves have been studied in [23] and [30], respectively.

The NCG was also analyzed with regard to Strong Equilibria (SE), for which the PoA was shown to be at most 2 in general [4] and at most $\frac{3}{2}$ for $\alpha \geq 2$ [17]. Greedy Equilibria (GE) for the NCG have been introduced by Lenzner [32] who showed that any graph in GE is in 3-approximate NE and that for trees GE and NE coincide. Also, it is shown that if an agent can improve by buying multiple edges, then she can improve by buying a single edge.

Many variants of the NCG have been studied: a version where agents want to minimize their maximum distance [19], variants with weighted traffic [1], or a version on a host graph [18]. Also geometric variants have been considered, for example a variant where the agents want to minize their average stretch [39] or a version on a host graph with weighted edges [6, 24]. Moreover, also NCG variants with non-uniform edge prices [6, 7, 11, 12, 16, 24, 35], locality [8, 9, 15], and robustness [13, 21, 36].

[1] Chauhan, Lenzner, Melnichenko, and Molitor [12] considered Add-Only Equilibria for a unilateral variant of the NCG. Also there, only edge additions are allowed. A bilateral version was also considered by Bullinger, Lenzner, and Melnichenko [11].
[2] Mihalák and Schlegel [37] study Asymmetric Swap Equilibria for the NCG. The BSwE extends this.
[3] For the NCG, Lenzner [32] defined the Greedy Equilibrium (GE), where no agent can improve by unilaterally adding, removing, or swapping a single incident edge. The BGE is the natural extension.
[4] Note that the allowed strategy changes in the BNE mirror the types of improving moves considered for NE in the NCG. Hence, the BNE is the natural extension of the NE to bilateral edge formation.
[5] For the NCG, Andelman, Feldman, and Mansour [4] and de Keijzer and Janus [17] investigated the Strong Equilibrium (SE), which is stable against unilateral strategy changes by any coalition of agents. The BSE is the natural extension.

We focus on the BNCG, proposed by Corbo and Parkes [14]. Pairwise Stability dates back to Jackson and Wolinsky [28]. So far, the PoA in the BNCG has only been analyzed for PS. In particular, an upper bound of $O(\min\{\sqrt{\alpha}, n/\sqrt{\alpha}\})$ was shown [14] and proven to be tight [19]. Corbo and Parkes [14] conjecture that all graphs in NE are also pairwise stable. Recently, Bilò, Friedrich, Lenzner, Lowski, and Melnichenko [5] study a variant of the BNCG with non-uniform edge price. Also, a version with inverted cost function modeling social distancing was proposed [25].

The BNCG has also been generalized by introducing cost-sharing of the edge prices [1]. In this variant, every agent specifies for every possible edge a cost-share she is willing to pay. Then edges with total cost-shares of at least α are formed. Moreover, Demaine, Hajiaghayi, Mahini, and Zadimoghaddam [18] introduced the Collaborative Equilibrium (CE), which is in-between PS and SE. A CE is stable against strategy changes by any coalition of agents that concern the joint cost-shares of any single edge. In contrast to the strategy changes that we focus on, this implies that (coalitions of) agents can also create non-incident edges.

1.3 Our Contribution

We analyze the Bilateral Network Creation Game with regard to different solution concepts in order to evaluate the impact of different levels of cooperation on the Price of Anarchy.

For this, we introduce various natural solution concepts with a wide range of allowed cooperation among the agents. We compare the subset relationships and we discuss the relationship between solution concepts for the NCG and the BNCG. Thereby, we disprove the old conjecture by Corbo and Parkes [14] that all graphs in NE are also pairwise stable.

Our main contribution is a thorough investigation of the PoA for various degrees of allowed cooperation among the agents. See Table 1 for an overview. As a first step in this direction, we mainly focus on tree networks in equilibrium. Such networks are of prime importance in the (B)NCG research, since the existence of equilibria that are trees is guaranteed for most of the parameter space[6], and because the PoA of tree networks in the unilateral NCG is constant [23]. Moreover, real-world networks are typically sparse and tree-like. Thus, the analysis of tree equilibria serves as a natural first step towards understanding more complicated equilibria.

Equilibrium	PoA on Trees	Source
PS	$\Theta(\min\{\sqrt{\alpha}, n/\sqrt{\alpha}\})$	O: [14], Ω: [19]
BSwE	$\Theta(\log\alpha)$	Section 3.2.1
BGE	$\Theta(\log\alpha)$	Section 3.2.2
BNE	$\Theta(\log\alpha)$, for $\alpha \geq n^{1/2+\varepsilon}$,	Section 3.2.3
	$\Theta(1)$, for $\alpha \leq \sqrt{n}$	Section 3.2.3
3-BSE	$\Theta(1)$	Section 3.2.4
Equilibrium	**PoA on General Graphs**	**Source**
BSE	$\Theta(1)$, for $\alpha \leq n^{1-\varepsilon}$	Section 3.3
	$\Theta(1)$, for $\alpha \geq n\log n$	Section 3.3
	$O(\frac{\log n}{\log\log\log n})$, otherwise	Section 3.3

Table 1: Asymptotic PoA bounds for $\alpha \geq 1$ and $\alpha < n^{2-\varepsilon}$, for any $\varepsilon > 0$.

Most importantly, we show that studying tree networks yields valuable insights on the impact of cooperation for the BNCG. In particular, we observe multiple improvements to the asymptotic PoA as we increase the allowed amount of cooperation. While the PoA is known to be in $\Theta(\min\{\sqrt{\alpha}, n/\sqrt{\alpha}\})$ for PS [14, 19], which only allows cooperation for creating a single edge, we show that allowing cooperative edge swaps, as in Bilateral Swap Equilibria or in Bilateral Greedy Equilibria, already improves the PoA to $\Theta(\log\alpha)$. For $\alpha \geq n^{1/2+\varepsilon}$, with $\varepsilon > 0$, we get the same PoA for trees in Bilateral Neighborhood Equilibrium (BNE), which is the bilateral version of the NE in the unilateral NCG. However, on the positive side, we show that the PoA for trees in BNE surprisingly turns constant if edges are cheap, i.e., for $\alpha \leq \sqrt{n}$.

Our most significant result is a constant PoA for tree equilibria using the Bilateral 3-Strong Equilibrium. At least on tree networks, this implies that very little agent cooperation is needed to ensure socially good stable states. Thus, for a system designer it is quite easy to enable the agents to escape from socially bad stable states: simply allow cooperation of coalitions of size 3. Contrasting this, we also show that the PoA for Bilateral 2-Strong Equilibria is in $\Omega(\log\alpha)$, i.e., no constant PoA can be achieved by considering smaller coalitions. Thus, we exactly pin-point the minimum amount of cooperation needed to ensure a constant PoA on tree networks.

For general networks, we investigate the impact of enhanced cooperation on the PoA by studying Bilateral Strong Equilibria. Our results are similar to the best known PoA bounds for NE in the NCG. We prove a constant PoA for Bilateral Strong Equilibria for $\alpha \leq n^{1-\varepsilon}$, with $\varepsilon > 0$, and for $\alpha \geq n\log n$. For the α range in-between, we show that the PoA is in $o(\log n)$. It remains open whether the PoA is constant for all $\alpha > 0$ and we show that our techniques cannot be applied to obtain a constant PoA bound for $\alpha = n$.

All omitted details can be found in the full version [26].

2 RELATIONSHIPS OF EQUILIBRIA

We explore the subset relationships among the introduced solution concepts. See Figure 1a for our results. The relationships follow directly from the definitions, but showing that subsets are proper or showing non-comparability requires a rich set of examples.

We find that the well-known Pairwise Stability [28] is a superset of many of the solution concepts that we study, i.e., the BGE and BNE can be understood as stronger refinements of PS.

We compare the NCG and the BNCG and with regard to Remove Equilibria and Add Equilibria. In the NCG, the corresponding definition of an Add Equilibrium considers that a single agent might add a single edge without any strategy changes of other agents. For simplicity, we assume for the NCG that each edge of the graph G is owned by exactly one incident agent. This allows us to model the *edge assignment* as a function $f : E \to V$, where each edge is mapped to one of its incident nodes. Under these assumptions, a graph G and edge assignment f completely capture the strategy vector of the NCG.

Proposition 2.1. *Let graph G with edge assignment f be in Add Equilibrium for the unilateral NCG, then G is also in BAE in the BNCG. However, the reverse direction does not hold.*

[6]For $\alpha \geq 1$ a star on n nodes is an equilibrium for all considered solution concepts.

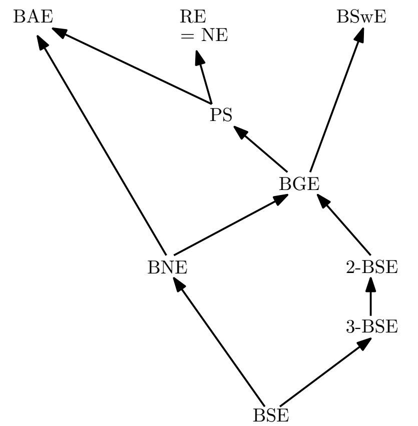

(a) The subset relationships between the considered solution concepts. Arrows point from subset to superset, all subset relations are proper.

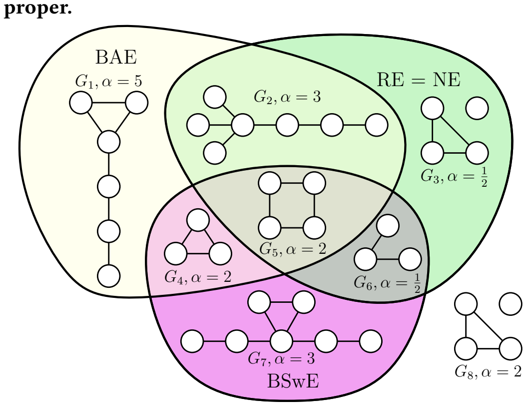

(b) Venn diagram showing the relation of RE, BAE, and BSwE. They are pairwise incomparable.

Figure 1: Overview over the subset relationships between the solution concepts.

As expected, there is no difference regarding Remove Equilibria.

Proposition 2.2. *A graph G is in Remove Equilibrium in the BNCG exactly if it is in Remove Equilibrium in the unilateral NCG for every edge assignment.*

This brings us to refuting the conjecture by Corbo and Parkes [14], which states that every graph G in NE in the NCG is also pairwise stable in the BNCG. For this, we consider a graph in NE which is not in unilateral Remove Equilibrium for a different edge assignment.

Proposition 2.3. *There exists a graph G and edge assignment f such that the graph with this assignment is in NE in the unilateral NCG but G is not pairwise stable in the BNCG.*

Proof. Figure 2 shows such a graph G. While the graph is in NE in the unilateral NCG, it is not pairwise stable in the BNCG because agent b profits from removing the edge ba. □

Another contrast between unilateral and bilateral equilibria is the existence of non-tree equilibrium networks for high α. Let C_n denote a cycle of n nodes. [14] already showed that for any $n \geq 3$,

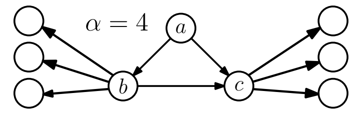

Figure 2: Example of a graph that is in NE in the NCG but not PS in the BNCG. The edge owner assignment for the NCG is depicted by edges pointing away from their owner.

there is a range of $\alpha \in \Theta(n^2)$ for which C_n is pairwise stable. We show that these α ranges can be refined further, so that they even apply for BSE. This implies that, contrary to the unilateral NCG, no tree conjecture is possible in the BNCG, as there can be non-tree equilibria for $\alpha \in \Theta(n^2)$.

Lemma 2.4. *Let $n \in \mathbb{N}_{\geq 3}$, then the cycle C_n is in BSE for some range of $\alpha \in \Theta(n^2)$.*

3 COOPERATION AND THE PRICE OF ANARCHY

We present our main results on the impact of the degree of cooperation on the PoA. We start with preliminaries, then analyze the PoA of trees before we investigate general networks.

3.1 Preliminaries

The social optima of the BNCG have been identified by Corbo and Parkes [14]. For $\alpha < 1$, the clique is the only social optimum and has a cost of $\text{cost}(OPT) = n(n-1)(1+\alpha)$. However, the case with $\alpha < 1$ is of little interest, since it is always beneficial to buy an edge. In the following, we assume $\alpha \geq 1$, unless explicitly noted otherwise.

For $\alpha \geq 1$, the star is a social optimum, and for $\alpha > 1$, it is the only social optimum. There are $n-1$ edges, so the total buying cost is $2(n-1)\alpha$. The total distance cost among the outer nodes is $2(n-1)(n-2)$ while the total distance from and to the center node is $2(n-1)$. Hence,

$$\text{cost}(OPT) = 2(n-1)\alpha + 2(n-1)(n-2) + 2(n-1) = 2(n-1)(\alpha+n-1).$$

Based on a similar proof by Albers, Eilts, Even-Dar, Mansour, and Roditty [1], we can upper bound the PoA by analyzing the distance cost of a single node. This is very powerful as this implies that it is sufficient to show a small distance cost for a single node, while the buying cost automatically follows suit.

Proposition 3.1. *Let G be in RE and connected, then for any node $u \in V$ the PoA can be upper bounded by $\rho(G) \leq \frac{\alpha + dist(u)}{\alpha+n-1}$.*

Since $\text{dist}(u) < n^2$ trivially holds in all connected graphs, we get the following corollary which implies a constant PoA for $\alpha \in \Omega(n^2)$.

Corollary 3.2. *Let G be a in RE and connected, then $\rho(G) \leq 1 + \frac{n^2}{\alpha}$.*

3.2 The Price of Anarchy for Tree Networks

In the following, we always root the tree G at a node $r \in V$. Based on this root r, the *layer* of a node $u \in V$ is $\ell(u) = \text{dist}(r,u)$. Each edge $uv \in E$ connects two nodes of adjacent layers, i.e., it holds that $|\ell(u) - \ell(v)| = 1$, and if $\ell(u) = \ell(v) - 1$, we say that u is the *parent* of v and v is the *child* of u. Each node has exactly one parent, the root r has none. The *depth* of a tree is $\text{depth}(G) = \max_{v \in V} \ell(v)$, i.e., the maximum distance between the root r and any other node. For $u \in V$, let T_u denote the subtree rooted at u containing u and all of its descendants, i.e., exactly the set of nodes $v \in V$ for which the unique v-r-path in the tree contains node u. In particular, the subtree T_r is the original graph G. We abuse notation by treating T_u like a node set, so we will write $v \in T_u$ when referring to a node in the subtree. We write $\text{depth}(T_u)$ when referring to $\max \{\text{dist}(u, v) \mid v \in T_u\}$.

Another important concept for our proofs is the *1-median* [29] (referred to as center node in [23]). A 1-median in a tree is a node $u \in V$ with the lowest distance cost. Equivalently, a 1-median can also be defined as a node whose removal from a tree with n nodes will create connected components of size at most $\frac{n}{2}$. Each tree has exactly one or two 1-medians. When referring to the 1-median of a subtree T, we also call it a *T-1-median*. In the following proofs, we always consider G to be rooted at a 1-median $r \in V$. So, for any $u \in V$ with $u \neq r$, it holds that $|T_u| \leq \frac{n}{2}$. In particular, this implies for each non-root $u \in V$ that at least $\frac{n}{2}$ shortest paths contain r. Thus, getting closer to r can lead to large cost reductions for the nodes. Many of our proofs are based on this.

3.2.1 Bilateral Swap Equilibrium on Trees. We upper bound the PoA for trees in BSwE by $\Theta(\log \alpha)$. The proof for the asymptotically tight lower bound is presented in Section 3.2.2. This result shares some parallels with the results for unilateral Asymmetric Swap Equilibria [22, 37], where it is shown that they have a diameter in $O(\log n)$ and that there is an edge assignment under which a complete binary tree is stable.

A key insight for our upper bound is that subtrees quickly fan out into small subtrees.

Lemma 3.3. *If the tree G with root and 1-median r is in BSwE, then for $u \in V$ there is a T_u-1-median $v \in T_u$ with $\ell(v) \leq \ell(u) + \frac{2\alpha}{n}$.*

Proof. If there are two T_u-1-medians $v_1, v_2 \in T_u$, then we choose the one closer to u, i.e., we choose v with $\ell(v) = \min\{\ell(v_1), \ell(v_2)\}$. If $u = r$, we have $u = v$ and the claim holds.

Otherwise, there is a parent $p \in V$ of u. If u is not a 1-median of T_u, then p prefers swapping pu for pv since that would reduce $\text{dist}(p, T_u)$. For agent v, accepting the proposed swap decreases her distance to r by $\text{dist}(v, p) - 1 = \ell(v) - \ell(u)$. By definition of the 1-median, the path from v to at least $\frac{n}{2}$ nodes must contain r, hence we get $(\ell(v) - \ell(u))\frac{n}{2} \leq \alpha$ as G is in BSwE. Rearranging this inequality concludes the proof. □

This allows us to obtain a bound on the depth of subtrees.

Lemma 3.4. *If the tree G with root and 1-median r is in BSwE, then for $u \in V$ we have that $\text{depth}(T_u) \leq \left(1 + \frac{2\alpha}{n}\right) \log |T_u|$.*

This already allows us to upper bound the depth and diameter of G by $O(\frac{\alpha}{n} \log n)$, which implies a PoA in $O(\log n)$. However, we can improve this bound further for cases where α is significantly smaller than n. The next lemma states that the tree fans out very quickly in the beginning, specifically, all subtrees rooted in layer 2 contain at most α nodes.

Lemma 3.5. *If the tree G with root and 1-median r is in BSwE, then it holds for $u \in V$ with $\ell(u) \geq 2$ that $|T_u| \leq \frac{\alpha}{\ell(u)-1}$.*

Finally, we can now combine the above lemmas to get a $O(\log \alpha)$ upper bound on the PoA. This allows us to obtain a $O(\log \alpha)$ upper bound on the PoA.

Theorem 3.6. *If the tree G is in BSwE, then $\rho(G) \leq 2 + 2\log \alpha$.*

Proof. We root G at a 1-median $r \in V$. We start by bounding $\text{depth}(G)$. If $\text{Neigh}^{=2}(r) = \emptyset$, we have $\text{depth}(G) \leq 1$. Otherwise, we choose $u \in \text{Neigh}^{=2}(r)$ such that $\text{depth}(T_u)$ is maximal, as this implies that $\text{depth}(G) = 2 + \text{depth}(T_u)$. By Lemma 3.5, we can upper bound $|T_u|$ by α and with Lemma 3.4 it follows that $\text{depth}(T_u) \leq \left(1 + \frac{2\alpha}{n}\right) \log \alpha$. Thus, we get

$$\text{depth}(G) \leq 2 + \left(1 + \frac{2\alpha}{n}\right) \log \alpha.$$

Now, we upper bound $\text{dist}(r)$ by $(n-1)\text{depth}(G)$. This gives

$$\text{dist}(r) \leq (n-1)\left(2 + \left(1 + \frac{2\alpha}{n}\right) \log \alpha\right) \leq (n-1)(2+\log \alpha) + 2\alpha \log \alpha.$$

To conclude the proof, we apply Proposition 3.1 to upper bound $\rho(G)$ based on $\text{dist}(r)$ with the inequality

$$\rho(G) \leq \frac{\alpha + \text{dist}(r)}{\alpha + n - 1} \leq \frac{\alpha + (n-1)(2 + \log \alpha) + 2\alpha \log \alpha}{\alpha + n - 1} = \frac{\alpha(1 + 2\log \alpha) + (n-1)(2 + \log \alpha)}{\alpha + n - 1} \leq 2 + 2\log \alpha. \quad \square$$

This result shows that on tree networks swapping an edge is more powerful than only adding or removing an edge. This is good news, as organizing a swap only requires little coordination. However, combining all three operations as considered in the study of (Bilateral) Greedy Equilibria does not grant additional asymptotic improvements, as we shall see next.

3.2.2 Bilateral Greedy Equilibrium on Trees. In Section 3.2.1, we have shown that the PoA for BSwE on trees is in $O(\log \alpha)$. This bound also carries over to BGE as they are a subset of BSwE. Now, we show that this bound is asymptotically tight. Note that the lower bound only applies for $\alpha \leq n^{2-\varepsilon}$, with $\varepsilon > 0$, since by Corollary 3.2, for $\alpha \in \Omega(n^2)$ the PoA of any connected graph is constant.

To motivate BGE further, we start with showing that on trees they are equivalent to 2-BSE.

Proposition 3.7. *Let G be a tree. Then G is in BGE if and only if it is in 2-BSE.*

Now we introduce the *stretched binary tree* that will be used for the PoA bound of $\Omega(\log \alpha)$. A *stretched binary tree* T with parameters $d \in \mathbb{N}$ and $k \in \mathbb{N}_{\geq 1}$ is defined as follows. Let B be a complete binary tree of depth d with root r. For $u \in V_B \setminus \{r\}$, we define $P_u = \{u^i\}_{i \in [k-1]} \cup \{u\}$ and define the node set of T as $V_T = \{r\} \cup \bigcup_{u \in V_B \setminus \{r\}} P_u$. For $uv \in E_B$, where u is the parent, the tree T contains the edges $uv^1, v^1v^2, \ldots, v^{k-1}v$. See Figure 3 for an example of a 3-stretched binary tree. The resulting graph has $(|V_B| - 1)k + 1 = \left(2^{d+1} - 2\right)k + 1$ nodes and is a tree. For $u, v \in V_B$, it holds that $\text{dist}_T(u, v) = k \cdot \text{dist}_B(u, v)$ and thus $\text{depth}(T) = k \cdot \text{depth}(B)$. We also call T a *k-stretched binary tree*. For $i \in [k-1]$ and $u \in B$, we refer to the subtree rooted at u^i as T_u^i, as this is easier to read as T_{u^i}. Moreover, we use u^k as an alias for u.

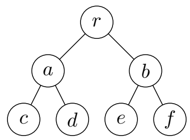

(a) Complete binary tree B with $d = 2$.

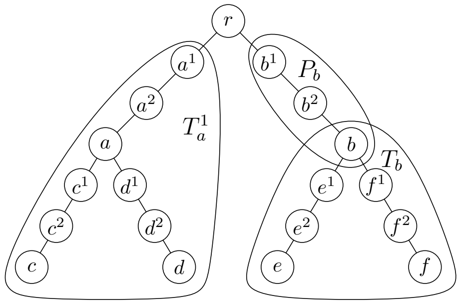

(b) Stretched binary tree T with $d = 2, k = 3$.

Figure 3: A complete binary tree B (top) and the associated 3-stretched binary tree T (bottom).

The intuition for using k-stretched binary trees is as follows. A binary tree has logarithmic depth in the number of nodes. But since the PoA formula divides the cost by $\alpha + n - 1$, a depth of $\log n$ will get dominated by the edge cost for $\alpha \geq n \log n$. Hence, we stretch the tree to preserve the distance cost. In particular, we will later choose $k \in \Theta(\frac{\alpha}{n})$.

Proposition 3.8. *Let T be a k-stretched binary tree, then T is in BGE for $\alpha \geq 7kn$.*

Finally, we can lower bound the PoA for large α.

Proposition 3.9. *For sufficiently large $\eta \in \mathbb{N}$ and $\eta \leq \alpha \leq \eta^{2-\gamma}$ with $0 < \gamma \leq 1$, there exists a stretched tree G with $\frac{\eta}{42} \leq n \leq \frac{\eta}{14}$ nodes such that G is in BGE and $\rho(G) \geq \frac{25}{32} + \frac{1}{96}\gamma \log \eta$.*

Since we construct a lower bound for a large range of α, we need to be able to scale our graph, so we can provide many different values of n for a given α. Thus, the next definition combines many stretched trees into a single graph.

We define a *stretched tree star G* with stretch factor $k \in \mathbb{N}_{\geq 1}$, target subtree size $t \in \mathbb{R}_{\geq 2k+1}$ and target size $\eta \in \mathbb{N}_{\geq 2t+1}$ as follows. Let T be a stretched tree with parameter k and d maximal subject to $|T| \leq t$. Then G consists of a root r with $\left\lceil \frac{\eta - 1}{|T|} \right\rceil$ copies of T as child subtrees.

Finally, we can use stretched tree stars to lower bound the PoA with respect to BGE.

Theorem 3.10. *For sufficiently large α and $\eta \in \mathbb{N}_{\geq \alpha}$, a stretched tree star G with $\eta \leq n \leq \frac{3\eta}{2}$ nodes exists, such that G is in BGE and $\rho(G) \geq \frac{1}{4}\log \alpha - \frac{17}{8}$.*

Proof Sketch. We construct G as a stretched tree star with the parameters $k = 1$, $t = \frac{\alpha}{15}$, and η as provided. The graph G is evidently in RE since it is a tree.

In order to show that G is in BAE, consider $u, v \in V$ with $\ell(u) \leq \ell(v)$ and $uv \notin E$. Moreover, let $a, b \in V$ be children of r such that $u, v \in T_a \cup T_b \cup \{r\}$. Let T' denote the subgraph induced by $T_a \cup T_b \cup \{r\}$. As u does not get any closer to the root r by adding the edge uv, it suffices to only consider T' in our analysis, which is a binary tree with at most $\frac{\alpha}{7}$ nodes (for sufficiently large α). With additional observations we conclude that G is in BAE, as $7k|T'| \leq \alpha$.

Analogously, for a swap $u, v, w \in V$ with $uv \in E, uw \notin E$ and $\ell(u) \leq \ell(w)$, we can again define T' and conclude that the change is restricted to this complete binary tree. With additional observations we conclude that G is in BSwE.

Hence, the graph G is in BGE. It remains to show the logarithmic lower bound on the PoA. In the full version [26] we provide the lower bound

$$\rho(G) \geq \frac{nk\left(\log\left(\frac{t}{k}\right) - \frac{9}{2}\right)}{2(\alpha + n - 1)} = \frac{n(\log t - \frac{9}{2})}{2(\alpha + n - 1)}.$$

We conclude the proof by upper bounding $2(\alpha + n - 1)$ in the denominator by $4n$ and simplifying as follows:

$$\rho(G) \geq \frac{n\left(\log t - \frac{9}{2}\right)}{4n} = \frac{\log \frac{\alpha}{15} - \frac{9}{2}}{4} = \frac{1}{4}\log \alpha - \frac{17}{8}. \qquad \square$$

Remember that networks in BGE are a subset of the networks in BSwE and hence the PoA upper bound of $O(\log \alpha)$ from Theorem 3.6 also holds for networks in BGE. Thus, Theorem 3.10 establishes a tight bound on the PoA for tree networks in BGE.

3.2.3 Bilateral Neighborhood Equilibrium on Trees. Here we prove that the PoA of BNE is in $\Theta(\log \alpha)$ for $\alpha \geq n^{1/2+\varepsilon}$, with $\varepsilon > 0$, so it remains asymptotically unchanged in comparison to BSwE and BGE. However, the PoA surprisingly changes to being constant for $\alpha \leq \sqrt{n}$. Thus, the additional coordination improves the asymptotic PoA for small values of α.

Since networks in BNE are also in BSwE, the PoA upper bound of $O(\log \alpha)$ from Theorem 3.6 carries over. Now we show that this bound is tight for $\alpha \geq n^{1/2+\varepsilon}$, with $\varepsilon > 0$. For this, we again use stretched tree stars, but this time we have to check for stability with respect to BNE.

Lemma 3.11. *Let G be a stretched tree star with parameter $k \in \mathbb{N}$ based on a stretched tree T. Let $k = 1$ or $\alpha \geq 6kn$. Then the graph G is in BNE if $\frac{3n \cdot depth(G)}{\alpha} + 1 \leq \frac{\alpha}{3|T| \cdot depth(G)}$.*

Now, we can use Lemma 3.11 to show a lower bound on the PoA for BNE on trees.

Theorem 3.12. *The following statements hold:*

(i) *For $0 < \varepsilon \leq 1$, sufficiently large $\eta \in \mathbb{N}$ and $9\eta \leq \alpha \leq \eta^{2-\varepsilon}$, there exists a BNE G, with $\eta \leq n \leq \frac{3}{2}\eta$, such that the inequality $\rho(G) \geq \frac{\varepsilon}{168}\log \alpha - \frac{3}{28}$ holds.*
(ii) *For $0 < \varepsilon \leq \frac{1}{2}$, sufficiently large $\eta \in \mathbb{N}$ and $\eta^{1/2+\varepsilon} \leq \alpha \leq \eta$, there exists a BNE G, with $\eta \leq n \leq \frac{3}{2}\eta$, such that the inequality $\rho(G) \geq \frac{1}{4}\varepsilon \log \alpha - \frac{9}{8}$ holds.*

Looking at the range for α in Theorem 3.12, we see that no lower bound for $\alpha \leq \sqrt{n}$ is derived. In fact, we cannot apply Lemma 3.11 for $\alpha \leq \sqrt{n}$, as inserting $\alpha = \sqrt{n}$ gives the inequality

$$3\sqrt{n} \cdot \mathrm{depth}(G) + 1 \leq \frac{\sqrt{n}}{3\,|T| \cdot \mathrm{depth}(G)},$$

which does not hold for any legal parameters. This is not a fault of our technique, but instead, we show that the PoA is actually constant for this range of α.

Theorem 3.13. *Let G be a tree with $n > 15$. If G is in BNE for $\alpha \leq \sqrt{n}$, then $\rho(G) \leq 4$.*

Proof. Let r denote the 1-median and root of G. Let $u \in V$ be a node of maximum layer. If $\ell(u) \leq 2$, the claim holds. So, we assume that $\ell(u) \geq 3$. Moreover, it holds that $\alpha < \frac{n}{2}$. We assume that $n > 4$ and thus $\alpha < \frac{n}{2}$, as otherwise the diameter is at most 3 and the claim holds.

With $i = \left|\text{Neigh}^{=3}(r)\right|$, let $\left\{c_j\right\}_{j \in [i]}$ denote the nodes in layer 3 sorted descendingly by their subtree size. Let $k = \min\left\{\left\lfloor \frac{n}{\alpha} - 1 \right\rfloor, i\right\}$. We consider the following change around u. Agent u buys an edge towards r and towards the nodes in $\left\{c_j\right\}_{j \in [k]} \setminus \{u\}$.

Agent u pays for at most $\frac{n}{\alpha}$ additional edges, so her buying cost increases by at most n. Connecting to r decreases her distance cost by at least $2\frac{n}{2}$. Further, agent u profits from the new direct connections to the nodes in $\left\{c_j\right\}_{j \in [k]} \setminus \{u\}$, so the overall change is beneficial for u. Each agent $c \in \left\{c_j\right\}_{j \in [k]}$ profits because her distance to r decreases by 1, so her distance cost decreases by at least $\frac{n}{2}$ while she only has to pay for one additional edge. Then, agent r must not benefit from the proposed change, as G is in BNE. As agent r decreases her distance to each node in $c \in \left\{c_j\right\}_{j \in [k]} \setminus \{u\}$ by 1 and to u by 2, it must hold that $\sum_{j=1}^{k} \left|T_{c_j}\right| < \alpha$. Then, if $i > k$, this allows us to upper bound $\left|T_{c_{k+1}}\right|$ as follows

$$\left|T_{c_{k+1}}\right| < \frac{\alpha}{k} < \frac{\alpha}{\frac{n}{\alpha} - 2} \leq \frac{\alpha}{\frac{n}{2\alpha}} = 2.$$

So, if the node c_{k+1} exists, then it has no children. Thus, only the agents $c \in \left\{c_j\right\}_{j \in [k]}$ have descendants beyond layer 3. Carrying over from BGE, we get for $c \in \left\{c_j\right\}_{j \in [k]}$ that depth$(T_c) \leq \log \alpha$. We conclude that dist$(r) \leq 3(n-1) + \alpha \log \alpha$ and apply Proposition 3.1 to get

$$\rho(G) \leq \frac{\alpha + 3(n-1) + \alpha \log \alpha}{\alpha + (n-1)} \leq 3 + \frac{\sqrt{n}\left(1 + \log \sqrt{n}\right)}{n} \leq 4. \quad \square$$

Note that Theorem 3.13 at least partially recovers a well known positive result from the unilateral NCG using the NE as solution concept: that the PoA of tree networks is constant. While for the unilateral NCG with NE this holds for all edge prices α, we get the contrasting result that for the BNCG using the BNE, this is only true for $\alpha \leq \sqrt{n}$. However, in the following we will see that we actually can guarantee a constant PoA on trees for all α if we allow coalitions of size 3 to cooperate.

3.2.4 Bilateral 3-Strong Equilibrium on Trees. We show that the PoA for 3-BSE on trees is constant. Thus, allowing coalitions of three agents provides us with the same asymptotics as NE in the unilateral NCG, at least on trees.

The intuition behind this result is derived from Section 3.2.3, in particular from the PoA lower bound for BNE. In the proof of Lemma 3.11, we encountered a situation where an agent $u \in V$ from a deep layer attempted to connect to a node $v \in V$ from a layer closer to the root. Adding the edge uv would reduce the distance cost of agent u significantly, so agent u was willing to buy many extra edges to incentivize agent v to accept the connection, but it ultimately failed to provide enough value to offset the increased buying cost of agent v.

For 3-BSE we consider the following: What if agent u does not need to convince agent v to buy an extra edge, but instead to swap an existing one? The swap is possible since agents u and v are part of a coalition, and agent u can collaborate with a third member of the coalition to provide the incentive. This idea leads us to our key lemma, which states that all but one child-subtrees of a node must be shallow.

Lemma 3.14. *Let G be a tree in 3-BSE with root and 1-median r, then every node $u \in V$ has at most one child $c \in V$ such that depth$(T_c) > 2\left\lceil \frac{4\alpha}{n} \right\rceil + 1$.*

Proof. Assume towards a contradiction that there are two children c, c' of node u whose subtrees have at least depth $2\left\lceil \frac{4\alpha}{n} \right\rceil + 2$. Then, there is a node $z \in T_c$ such that $\ell(z) = \ell(u) + 2\left\lceil \frac{4\alpha}{n} \right\rceil + 3$. Moreover, on the path from c to z, there are nodes $x, y \in T_c$ such that y is the child of x and $\ell(x) = \ell(u) + \left\lceil \frac{4\alpha}{n} \right\rceil + 2$. Analogously, we define the nodes $x', y', z' \in T_{c'}$. These paths are visualized in Figure 4 along with an annotation of their layers relative to u.

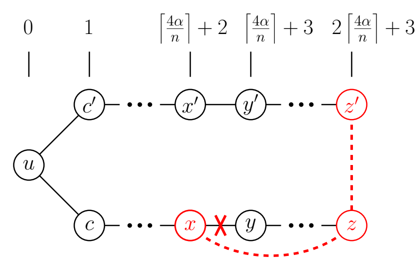

Figure 4: The paths considered in the proof of Lemma 3.14. The proposed coalitional change is indicated in red (dashed red lines are to be built, the edge xy is to be removed). Nodes outside of the two paths are omitted for simplicity. The labels on top show the layer of the nodes relative to the layer of u.

We now consider a coalitional move by agents x, z and z', in which the edges xz and zz' get added while edge xy is removed. This move decreases the distance from z to x by $\left\lceil \frac{4\alpha}{n} \right\rceil$ and by extension the distance to u. Since the path from z to at least $\frac{n}{2}$ nodes must contain u and since the path to x also gets shorter, the distance cost for z decreases by at least

$$\left\lceil \frac{4\alpha}{n} \right\rceil \cdot \left(\frac{n}{2} + 1\right) > 2\alpha.$$

Thus, even when considering the increased buying cost, agent z profits from this coalitional move. Likewise, the distance cost for z' decreases by at least

$$\left(\left\lceil \frac{4\alpha}{n} \right\rceil - 1\right) \cdot \frac{n}{2} > \alpha,$$

so z' also profits by the proposed change.

As G is in 3-BSE, agent x must not benefit from the coalition. The buying cost of agent x remains unchanged, so her distance cost must not improve. The only nodes towards which the distance from x may increase are the nodes in T_y. This increase per node can be upper bounded by the distance increase towards y, which is $\lceil \frac{4\alpha}{n} \rceil$ as the new shortest path goes through z. On the other hand, the distance from x to x' decreases from $2\lceil \frac{4\alpha}{n} \rceil + 4$ down to $2 + \text{dist}(z', x') = \lceil \frac{4\alpha}{n} \rceil + 3$, this is an improvement by $\lceil \frac{4\alpha}{n} \rceil + 1$. And by extension, the distance to each node in $T_{x'}$ also improves by at least this amount.

As we assume that the distance cost of agent x is not reduced by the coalitional move, we deduce that

$$\left(\left\lceil \frac{4\alpha}{n} \right\rceil + 1\right) |T_{x'}| \leq \left\lceil \frac{4\alpha}{n} \right\rceil \left|T_y\right|,$$

which implies that $|T_{x'}| < \left|T_y\right|$. But, by symmetry, we can also consider the coalition x', z', z, which gives us the bound $|T_x| < \left|T_{y'}\right|$. In combination with the inequalities $\left|T_y\right| < |T_x|$ and $\left|T_{y'}\right| < |T_{x'}|$, we conclude that this is a contradiction, so there cannot be two different child-subtrees with a depth greater than $2\lceil \frac{4\alpha}{n} \rceil + 1$. □

With this insight, we now proceed to show a constant PoA.

Theorem 3.15. *Let G be a tree in 3-BSE, then $\rho(G) \leq 25$.*

Proof. Let $d = \text{depth}(G)$ and $r \in V$ be the 1-median and root of G. For a node u with $\ell(u) = d$, consider the path from r to u denoted by $r = u^0, u^1, \ldots, u^d = u$. For $i \in [d]$, we denote the subtree rooted in u^i by T_u^i for better readability.

We derive an upper bound on $\text{dist}(r)$ based on the path to u. To do so, for any $i, j \in [d]$ with $i + j \leq d$, we split $\text{dist}(r, T_u^i)$ into its two components $\text{dist}(r, T_u^{i+j})$ and $\text{dist}(r, T_u^i \setminus T_u^{i+j})$.

For any node $v \in T_u^i \setminus T_u^{i+j}$, there is a $k \in [i, i+j]$ such that u^k is the deepest common ancestor of u and v. By Lemma 3.14, we have that the layer of v is at most $k + 2\lceil \frac{4\alpha}{n} \rceil + 2$, as otherwise u^k would have two child-subtrees with a layer exceeding the threshold. To keep our formulas simple, we define $x = 2\lceil \frac{4\alpha}{n} \rceil + 2$ and follow that $\ell(v) \leq i + j + x$.

Inserting this into our upper bound yields the following recursive inequality

$$\text{dist}(r, T_u^i) \leq \text{dist}(r, T_u^{i+j}) + \left|T_u^i \setminus T_u^{i+j}\right| (i + j + x).$$

We extend our notation for $i > d$, such that T_u^i denotes an empty subgraph. This allows us to transform the recursion into an infinite sum. Moreover, we fix $i = 0$ since we are interested in $\text{dist}(r) = \text{dist}(r, T_u^0)$. Thus, we get

$$\begin{aligned}\text{dist}(r) &\leq \sum_{k=0}^{\infty} \left|T_u^{kj} \setminus T_u^{(k+1)j}\right| ((k+1)j + x) \\ &\leq xn + j \sum_{k=0}^{\infty} \left|T_u^{kj} \setminus T_u^{(k+1)j}\right| (k+1).\end{aligned}$$

Now, we choose $j = \lceil \frac{2\alpha}{n} \rceil + 1$, as this ensures, by Lemma 3.3, for any $i \in \mathbb{N}$ that T_u^{i+j} contains at most half as many nodes as T_u^i. While this already gives us the bound $\left|T_u^{kj}\right| \leq 2^{-k} n$, we insert $2^{-(k+1)} n$ for $\left|T_u^{kj} \setminus T_u^{(k+1)j}\right|$ in the sum since this represents the worst case where the deep trees contain as many nodes as possible. This yields

$$\text{dist}(r) \leq xn + j \sum_{k=0}^{\infty} 2^{-(k+1)} n(k+1) = xn + jn \sum_{k=1}^{\infty} 2^{-k} k = xn + 2jn.$$

Substituting x and j by their underlying values finally yields the result

$$\begin{aligned}\text{dist}(r) &\leq \left(2\left\lceil \frac{4\alpha}{n} \right\rceil + 2\right) n + 2\left(\left\lceil \frac{2\alpha}{n} \right\rceil + 1\right) n \\ &\leq 16\alpha + 2n + 8\alpha + 2n = 24\alpha + 4n.\end{aligned}$$

For $n < 5$, the proposed bound holds for any tree, and for $n \geq 5$, we can upper bound $4n$ by $5(n-1)$. Finally, applying Proposition 3.1 yields

$$\rho(G) \leq \frac{\alpha + 24\alpha + 5(n-1)}{\alpha + n - 1} \leq 25. \qquad \square$$

Thus, we have proven that already a very limited form of cooperation, i.e., joint coalitional moves of coalitions of size at most 3, guarantees a constant PoA for tree networks. This is a strongly positive result since from an agents' point of view, negotiating changes within such small coalitions seems feasible. Such coalitional moves might be much easier to coordinate than an improving move of some agent in the BNE setting, where potentially many agents must simultaneously evaluate and agree to the change.

3.3 The Price of Anarchy for General Networks

We now consider general graphs. This means our arguments cannot rely anymore on reducing the distance to a 1-median or on the path between two nodes being unique. This makes reasoning about the equilibria significantly more difficult.

As a first step in this direction, we investigate if cooperation of the agents can guarantee good equilibria at all, i.e., we focus our analysis on the most powerful solution concept, the BSE, for which strategy changes by coalitions of arbitrary size are permitted, as long as all members of the coalition benefit from the change. We show that the PoA is constant for $\alpha \geq n \log n$ and for $\alpha \leq n^{1-\varepsilon}$, with $\varepsilon > 0$, but demonstrate that our technique cannot yield constant bounds for $\alpha = n$.

We start by showing that the set of social optima and BSE coincide for $\alpha \leq 1$. Then, for $\alpha > 1$, we show a constant PoA for most ranges of α.

It is already known that for $\alpha \leq 1$ any pairwise stable graph is socially optimal [14], so any BSE also must be socially optimal. Hence, the following proposition shows that BSE exist for such α. Moreover, it shows that the situation for $\alpha > 1$ is more complicated.

Proposition 3.16. *For $\alpha < 1$, the clique is the only BSE. For $\alpha = 1$, only graphs with diameter at most 2 are in BSE. For $\alpha > 1$, the star is in BSE, but so are other graphs.*

Next, we consider the PoA for $\alpha > 1$. For any given n and α, we show that the agent with the highest cost in an arbitrary graph implies an upper bound on the PoA. The considered graph does not even need to be in BSE. Consequently, it suffices to identify graphs where all agents have a low cost in order to bound the PoA.

Lemma 3.17. *Let G be a graph with $n > 1$ and let $u \in V_G$ be the agent with the highest cost. Then, for every graph H in BSE with n agents and the same α it holds that $\rho(H) \leq \frac{cost_G(u)}{\alpha + n - 1}$.*

Equipped with Lemma 3.17, we can now construct graphs where the worst-off agent has a low cost. We do so by building trees in which the cost is distributed evenly over the agents.

Lemma 3.18. *For $d \in \mathbb{N}_{\geq 2}$ and $n \in \mathbb{N}$, let G be an almost complete d-ary tree with n agents, then for all $u \in V$ holds that $cost(u) \leq (d+1)\alpha + 2(n-1)\log_d n$.*

Now we can set d depending on α. For $\alpha \geq n \log n$, it suffices to consider binary trees to obtain a good PoA, as the logarithmic distances get dominated by the buying cost.

Theorem 3.19. *Let G be in BSE and $\alpha \geq n \log n$, then $\rho(G) \leq 5$.*

Proof. We apply Lemma 3.18 with $d = 2$ and combine it with Lemma 3.17 to get

$$\rho(G) \leq \frac{3\alpha + 2(n-1)\log n}{\alpha + n - 1} \leq 3 + \frac{2(n-1)\log n}{n \log n} \leq 3 + 2 = 5. \quad \square$$

For α significantly smaller than n, it is necessary to keep the distances small. This can only be achieved by scaling the node degrees. This is possible since the buying cost is low.

Theorem 3.20. *For $\varepsilon \in \mathbb{R}_{>0}$ and $\alpha \leq n^{1-\varepsilon}$, let G be in BSE, then $\rho(G) \leq 3 + \frac{2}{\varepsilon}$.*

Proof. We apply Lemma 3.18 with $d = \lceil n^{\varepsilon} \rceil$ and combine it with Lemma 3.17 to get

$$\rho(G) \leq \frac{(\lceil n^{\varepsilon} \rceil + 1)\,\alpha + 2(n-1)\log_{\lceil n^{\varepsilon} \rceil} n}{\alpha + n - 1}.$$

We upper bound the first summand by using $\lceil n^{\varepsilon} \rceil + 1 \leq n^{\varepsilon} + 2$ and inserting the bound for α to get that $(n^{\varepsilon} + 2)\,\alpha \leq n + 2n^{1-\varepsilon}$. For the second summand, we do a change of base to convert $\log_{n^{\varepsilon}} n = \frac{\log n}{\log n^{\varepsilon}} = \frac{1}{\varepsilon}$ and upper bound with $\frac{2(n-1)}{\varepsilon}$. This gives us the inequality

$$\rho(G) \leq \frac{n + 2n^{1-\varepsilon} + \frac{2(n-1)}{\varepsilon}}{\alpha + n - 1}.$$

Since $\alpha \geq 1$, we lower bound the denominator by n and conclude the proof with

$$\rho(G) \leq 1 + 2n^{-\varepsilon} + \frac{2}{\varepsilon} \leq 3 + \frac{2}{\varepsilon}. \quad \square$$

This bound, however, is only constant with regard to a specific ε, so $\frac{2}{\varepsilon}$ can become arbitrarily large. Moreover, the preceding two theorems still leave a gap for $n^{1-\varepsilon} < \alpha < n \log n$, which is similar to the gap for NE in the NCG, for which no constant PoA has yet been established [3]. By using binary trees, we can show a PoA in $o(\log n)$ for general α.

Theorem 3.21. *Let G be in BSE, then*

$$\rho(G) \leq 2 + \log\log n + 2\frac{\log n}{\log\log\log n}.$$

Proof. We apply Lemma 3.18 with $d = \lceil \log\log n \rceil$ and combine it with Lemma 3.17 to get

$$\begin{aligned}\rho(G) &\leq \frac{(\lceil \log\log n \rceil + 1)\alpha + 2(n-1)\log_{\lceil \log\log n \rceil} n}{\alpha + n - 1} \\ &\leq \log\log n + 2 + \frac{2(n-1)\log_{\log\log n} n}{n-1}.\end{aligned}$$

By a change of base, we convert $\log_{\log\log n} n$ to $\frac{\log n}{\log\log\log n}$. The proof concludes with

$$\begin{aligned}\rho(G) &\leq 2 + \log\log n + \frac{2(n-1)\frac{\log n}{\log\log\log n}}{n-1} \\ &= 2 + \log\log n + 2\frac{\log n}{\log\log\log n}. \quad \square\end{aligned}$$

Theorem 3.21 gives a $O(\frac{\log n}{\log\log\log n})$ bound which is the same as $O(\frac{\log \alpha}{\log\log\log \alpha})$ for the remaining gap $n^{1-\varepsilon} < \alpha < n \log n$. We suspect that additional improvements can be achieved by refining the parameter d. Nonetheless, it remains open whether there is a constant PoA for α close to n. What we do know, however, is that such a bound cannot be obtained by our technique. The reason is that for a constant PoA for $\alpha = n$, the cost cannot remain evenly distributed across the agents as we scale-up n, as we see now.

Proposition 3.22. *There is no constant $p \in \mathbb{R}$ such that for all $n \in \mathbb{N}$ and $\alpha = n$ there exists a graph G such that for all agents $u \in V$ it holds that $\frac{cost_G(u)}{\alpha + n - 1} \leq p$.*

As an extension of this proposition, a constant PoA for $\alpha = n$ would imply that some graphs can only reach an equilibrium state through a series of multiple improving coalitional moves. In particular, we have already seen that for $\alpha = n$, every agent in an almost complete binary tree has costs in $\Theta(n \log n)$. But if the PoA is constant, any graph in BSE must have a node with degree $\Omega(\sqrt[4]{n})$ and thus costs in $\Omega(n \sqrt[4]{n})$, which is asymptotically worse. Evidently, this transition cannot be achieved in a single move, as only agents within the coalition can increase their degree, but they would only do so if they benefit from the change.

4 CONCLUSION AND OUTLOOK

We analyzed the BNCG under different amounts of agent cooperation. Previously, only Pairwise Stability, where cooperation is restricted to single edge additions, has been studied.

On tree networks our results convey the general picture that the PoA improves asymptotically as we progress towards more cooperation among the agents. In particular, the significant improvements achieved by BSwE, BGE, BNE and 3-BSE give valuable insights for system designers. When defining the protocol or contracts by which agents establish and remove connections, a system designer should try to allow for edge swaps or even joint changes by coalitions of size at least 3, instead of only permitting single edges to be removed or added. This is a positive result, as coordinating a swap or a 3-ways contract is presumably significantly less demanding than coordinating joint changes by larger coalitions.

Our results for tree networks raise the most pressing open question of whether these bounds on the PoA also carry over to general networks. For the latter, we showed that BSE have a constant PoA for most ranges of α and we conjecture that this extends to the full range of α. However, coordinating simultaneous infrastructural changes by large coalitions might prove an inconceivable effort in many practical applications. Hence, settling the PoA for solution concepts with more restricted agent coordination, like the 3-BSE, seems even more relevant.

REFERENCES

[1] Susanne Albers, Stefan Eilts, Eyal Even-Dar, Yishay Mansour, and Liam Roditty. 2014. On Nash equilibria for a network creation game. *ACM TEAC* 2, 1 (2014), 1–27.

[2] Carme Àlvarez and Arnau Messegué. 2017. Network creation games: Structure vs anarchy. *arXiv preprint arXiv:1706.09132* (2017).

[3] Carme Àlvarez and Arnau Messegué. 2019. On the Price of Anarchy for High-Price Links. In *WINE 2019*. 316–329.

[4] Nir Andelman, Michal Feldman, and Yishay Mansour. 2009. Strong Price of Anarchy. *Games and Economic Behavior* 65, 2 (2009), 289–317.

[5] Davide Bilò, Tobias Friedrich, Pascal Lenzner, Stefanie Lowski, and Anna Melnichenko. 2021. Selfish Creation of Social Networks. In *AAAI 2021*. 5185–5193.

[6] Davide Bilò, Tobias Friedrich, Pascal Lenzner, and Anna Melnichenko. 2019. Geometric network creation games. In *SPAA 2019*. 323–332.

[7] Davide Bilò, Tobias Friedrich, Pascal Lenzner, Anna Melnichenko, and Louise Molitor. 2020. Fair Tree Connection Games with Topology-Dependent Edge Cost. In *FSTTCS 2020*. 15:1–15:15.

[8] Davide Bilò, Luciano Gualà, Stefano Leucci, and Guido Proietti. 2016. Locality-Based Network Creation Games. *ACM Trans. Parallel Comput.* 3, 1 (2016), 6:1–6:26.

[9] Davide Bilò, Luciano Gualà, Stefano Leucci, and Guido Proietti. 2021. Network Creation Games with Traceroute-Based Strategies. *Algorithms* 14, 2 (2021), 35.

[10] Davide Bilò and Pascal Lenzner. 2020. On the tree conjecture for the network creation game. *Theory of Computing Systems* 64, 3 (2020), 422–443.

[11] Martin Bullinger, Pascal Lenzner, and Anna Melnichenko. 2022. Network Creation with Homophilic Agents. In *IJCAI 2022*. 151–157.

[12] Ankit Chauhan, Pascal Lenzner, Anna Melnichenko, and Louise Molitor. 2017. Selfish network creation with non-uniform edge cost. In *SAGT 2017*. 160–172.

[13] Ankit Chauhan, Pascal Lenzner, Anna Melnichenko, and Martin Münn. 2016. On Selfish Creation of Robust Networks. In *SAGT 2016*. 141–152.

[14] Jacomo Corbo and David Parkes. 2005. The price of selfish behavior in bilateral network formation. In *PODC 2005*. 99–107.

[15] Andreas Cord-Landwehr and Pascal Lenzner. 2015. Network Creation Games: Think Global - Act Local. In *MFCS 2015*. 248–260.

[16] Andreas Cord-Landwehr, Alexander Mäcker, and Friedhelm Meyer auf der Heide. 2014. Quality of Service in Network Creation Games. In *WINE 2014*. 423–428.

[17] Bart de Keijzer and Tomasz Janus. 2019. On strong equilibria and improvement dynamics in network creation games. *Internet Mathematics* 1, 1 (2019), 6946.

[18] Erik D Demaine, MohammadTaghi Hajiaghayi, Hamid Mahini, and Morteza Zadimoghaddam. 2009. The Price of Anarchy in cooperative network creation games. *ACM SIGecom Exchanges* 8, 2 (2009), 1–20.

[19] Erik D Demaine, MohammadTaghi Hajiaghayi, Hamid Mahini, and Morteza Zadimoghaddam. 2012. The Price of Anarchy in network creation games. *ACM TALG* 8, 2 (2012), 1–13.

[20] Jack Dippel and Adrian Vetta. 2021. An Improved Bound for the Tree Conjecture in Network Creation Games. *arXiv preprint arXiv:2106.05175* (2021).

[21] Hagen Echzell, Tobias Friedrich, Pascal Lenzner, and Anna Melnichenko. 2020. Flow-Based Network Creation Games. In *IJCAI 2020*. 139–145.

[22] Shayan Ehsani, Saber Shokat Fadaee, MohammadAmin Fazli, Abbas Mehrabian, Sina Sadeghian Sadeghabad, Mohammadali Safari, and Morteza Saghafian. 2015. A bounded budget network creation game. *ACM TALG* 11, 4 (2015), 1–25.

[23] Alex Fabrikant, Ankur Luthra, Elitza N. Maneva, Christos H. Papadimitriou, and Scott Shenker. 2003. On a network creation game. In *PODC 2003*. 347–351.

[24] Wilhelm Friedemann, Tobias Friedrich, Hans Gawendowicz, Pascal Lenzner, Anna Melnichenko, Jannik Peters, Daniel Stephan, and Michael Vaichenker. 2021. Efficiency and Stability in Euclidean Network Design. In *SPAA 2021*. 232–242.

[25] Tobias Friedrich, Hans Gawendowicz, Pascal Lenzner, and Anna Melnichenko. 2022. Social Distancing Network Creation. In *ICALP 2022*. 62:1–62:21.

[26] Tobias Friedrich, Hans Gawendowicz, Pascal Lenzner, and Arthur Zahn. 2022. *The Impact of Cooperation in Bilateral Network Creation.* Technical Report 2207.03798. arXiv. https://doi.org/10.48550/arXiv.2207.03798 Full version of this paper.

[27] Ronald L Graham and Pavol Hell. 1985. On the history of the minimum spanning tree problem. *Annals of the History of Computing* 7, 1 (1985), 43–57.

[28] Matthew O Jackson and Asher Wolinsky. 1996. A strategic model of social and economic networks. *Journal of economic theory* 71, 1 (1996), 44–74.

[29] Oded Kariv and S Louis Hakimi. 1979. An algorithmic approach to network location problems. I: The p-centers. *SIAM J. Appl. Math.* 37, 3 (1979), 513–538.

[30] Bernd Kawald and Pascal Lenzner. 2013. On dynamics in selfish network creation. In *SPAA 2013*. 83–92.

[31] Elias Koutsoupias and Christos H. Papadimitriou. 2009. Worst-case equilibria. *Comput. Sci. Rev.* 3, 2 (2009), 65–69.

[32] Pascal Lenzner. 2012. Greedy Selfish Network Creation. In *WINE 2012*. 142–155.

[33] Thomas L Magnanti and Richard T Wong. 1984. Network design and transportation planning: Models and algorithms. *Transportation science* 18, 1 (1984), 1–55.

[34] Akaki Mamageishvili, Matús Mihalák, and Dominik Müller. 2015. Tree Nash Equilibria in the Network Creation Game. *Internet Mathematics* 11, 4-5 (2015), 472–486.

[35] Eli A. Meirom, Shie Mannor, and Ariel Orda. 2014. Network formation games with heterogeneous players and the internet structure. In *EC 2014*. 735–752.

[36] Eli A. Meirom, Shie Mannor, and Ariel Orda. 2015. Formation games of reliable networks. In *INFOCOM 2015*. 1760–1768.

[37] Matúš Mihalák and Jan Christoph Schlegel. 2012. Asymmetric swap-equilibrium: A unifying equilibrium concept for network creation games. In *MFCS 2012*. 693–704.

[38] Matúš Mihalák and Jan Christoph Schlegel. 2013. The Price of Anarchy in Network Creation Games Is (Mostly) Constant. *TCS* 53, 1 (2013), 53–72.

[39] Thomas Moscibroda, Stefan Schmid, and Roger Wattenhofer. 2006. On the topologies formed by selfish peers. In *PODC 2006*. 133–142.

[40] Rajmohan Rajaraman. 2002. Topology control and routing in ad hoc networks: A survey. *ACM SIGACT News* 33, 2 (2002), 60–73.