On Dynamics in Selfish Network Creation

Bernd Kawald
Department of Computer Science
Humboldt-Universität zu Berlin
Berlin, Germany
kawald@informatik.hu-berlin.de

Pascal Lenzner
Department of Computer Science
Humboldt-Universität zu Berlin
Berlin, Germany
lenzner@informatik.hu-berlin.de

ABSTRACT

We consider the dynamic behavior of several variants of the Network Creation Game, introduced by Fabrikant et al. [PODC'03]. Equilibrium networks in these models have desirable properties like low social cost and small diameter, which makes them attractive for the decentralized creation of overlay-networks. Unfortunately, due to the non-constructiveness of the Nash equilibrium, no distributed algorithm for *finding* such networks is known. We treat these games as sequential-move games and analyze if (uncoordinated) selfish play eventually converges to an equilibrium. Thus, we shed light on one of the most natural algorithms for this problem: distributed local search, where in each step some agent performs a myopic selfish improving move.

We show that fast convergence is guaranteed for all versions of Swap Games, introduced by Alon et al. [SPAA'10], if the initial network is a tree. Furthermore, we prove that this process can be sped up to an almost optimal number of moves by employing a very natural move policy. Unfortunately, these positive results are no longer true if the initial network has cycles and we show the surprising result that even one non-tree edge suffices to destroy the convergence guarantee. This answers an open problem from Ehsani et al. [SPAA'11] in the negative. Moreover, we show that on non-tree networks no move policy can enforce convergence. We extend our negative results to the well-studied original version, where agents are allowed to buy and delete edges as well. For this model we prove that there is no convergence guarantee – even if all agents play optimally. Even worse, if played on a non-complete host-graph, then there are instances where no sequence of improving moves leads to a stable network. Furthermore, we analyze whether cost-sharing has positive impact on the convergence behavior. For this we consider a version by Corbo and Parkes [PODC'05] where bilateral consent is needed for the creation of an edge and where edge-costs are shared among the involved agents. We show that employing such a cost-sharing rule yields even worse dynamic behavior.

Finally, we contrast our mostly negative theoretical results by a careful empirical study. Our simulations indicate two positive facts: (1) The non-convergent behavior seems to be confined to a small set of pathological instances and is unlikely to show up in practice. (2) In all our simulations we observed a remarkably fast convergence towards a stable network in $\mathcal{O}(n)$ steps, where n is the number of agents.

SPAA'13, July 23–25 2013, Montréal, Québec, Canada.

Categories and Subject Descriptors

F.2.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity—*Nonnumerical Algorithms and Problems*; G.2.2 [**Mathematics of Computing**]: Discrete Mathematics—*Graph Theory, Network Problems*

General Terms

Algorithms, Design, Theory, Experimentation

Keywords

Network Creation Games, Game Dynamics, Convergence, Stabilization, Distributed Local Search

1. INTRODUCTION

Understanding Internet-like networks and their implications on our life is a recent endeavor undertaken by researchers from different research communities. Such networks are difficult to analyze since they are created by a multitude of selfish entities (e.g. Internet Service Providers) which modify the infrastructure of parts of the network (e.g. their Autonomous Systems) to improve their service quality. The classical field of Game Theory provides the tools for analyzing such decentralized processes and from this perspective the Internet can be seen as an equilibrium state of an underlying game played by selfish agents.

Within the last decade several such games have been proposed and analyzed. We will focus on the line of works which consider Network Creation Games, as introduced by Fabrikant et al. [11]. These games are very simple but they contain an interesting trade-off between an agent's investment in infrastructure and her obtained usage quality. Agents aim to invest as little as possible but at the same time they want to achieve a good connection to all other agents in the network. Network Creation Games and several variants have been studied intensively, but, to the best of our knowledge, almost all these works exclusively focus on properties of the equilibrium states of the game. With this focus, the game is usually considered to be a one-shot simultaneous-move game. However, the Internet was not created in "one shot". It has rather evolved from an initial network, the ARPANET, into its current shape by repeated infrastructural changes performed by selfish agents who entered or left the stage at some time in the process. For this reason, we focus on a more dynamic point of view: We analyze the properties of the network creation *processes* induced by the sequential-move version of the known models of selfish network creation.

It is well-known that Network Creation Games have low price of anarchy, which implies that the social cost of the worst stable states arising from selfish behavior is close to the cost of the social optimum. Therefore these games are appealing for the decentralized and selfish creation of networks which optimize the service quality for all agents at low infrastructural cost, e.g. overlay networks created by selfish peers. But, to the best of our knowledge, it is not known how a group of agents can collectively *find* such a desirable stable state. Analyzing the game dynamics of Network Creation Games is equivalent to analyzing a very natural search strategy: (uncoordinated) distributed local search, where in every step some agent myopically modifies the network infrastructure to better suit her needs. Clearly, if at some step in the process no agent wants to modify her part of the network, then a stable network has emerged.

1.1 Models and Definitions

We consider several versions of a network creation process performed by n selfish agents. In all versions we consider networks, where every node corresponds to an agent and undirected links connect nodes in the network. The creation process is based on an underlying Network Creation Game (NCG) and can be understood as a dynamic process where agents sequentially perform strategy-changes in the NCG. In such games, the strategies of the agents determine which links are present in the network and any strategy-profile, which is a vector of the strategies of all n agents, determines the induced network. But this also works the other way round: Given some network $G = (V, E, o)$, where V is the set of n vertices, E is the set of edges and $o : E \to V$ is the *ownership-function*, which assigns the ownership of an edge to one of its endpoints, then G completely determines the current strategies of all n agents of the NCG. Starting from a network G_0, any sequence of strategy-changes by agents can thus be seen as a sequence of networks $G_0, G_1, G_2, \ldots$, where the network G_{i+1} arises from the network G_i by the strategy-change of exactly one agent. In the following, we will write xy or yx for the undirected edge $\{x, y\} \in E$. In figures we will indicate edge-ownership by directing edges away from their owner.

The creation process starts in an initial state G_0, which we call the *initial network*. A step from state G_i to state G_{i+1} consists of a *move* by one agent. A move of agent u in state G_i is the replacement of agent u's pure strategy in G_i by another *admissible* pure strategy of agent u. The induced network after this strategy-change by agent u then corresponds to the state G_{i+1}. We consider only improving moves, that is, strategy-changes which strictly decrease the moving agent's cost. The cost of an agent in G_i depends on the structure of G_i and it will be defined formally below. If agent u in state G_i has an admissible new strategy which yields a strict cost decrease for her, then we call agent u *unhappy in network* G_i and we let U_i denote the set of all unhappy agents in state G_i. Only one agent can actually move in a state of the process and this agent $u \in U_i$, whose move transforms G_i into G_{i+1}, is called *the moving agent in network* G_i. In any state of the process the *move policy* determines which agent is the moving agent. The process stops in some state G_j if no agent wants to perform a move, that is, if $U_j = \emptyset$, and we call the resulting networks *stable*. Clearly, stable networks correspond to pure Nash equilibria of the underlying NCG.

Depending on what strategies are admissible for an agent in the current state, there are several variants of this process, which we call *game types*:

- In the *Swap Game* (SG), introduced as "Basic Network Creation Game" by Alon et al. [2], the strategy S_u of an agent u in the network G_i is the set of neighbors of vertex u in G_i. The new strategy S_u^* is admissible for agent u in state G_i, if $|S_u| = |S_u^*|$ and $|S_u \cap S_u^*| = |S_u| - 1$. Intuitively, admissible strategies in the SG are strategies which replace one neighbor x of u by another vertex y. Note, that this corresponds to "swapping" the edge ux from x towards y, which is the replacement of edge ux by edge uy. Furthermore, observe, that in any state both endpoints of an edge are allowed to swap this edge. Technically, this means that the ownership of an edge has no influence on the agents' strategies or costs.

- The *Asymmetric Swap Game* (ASG), recently introduced by Mihalák and Schlegel [16], is similar to the SG, but here the ownership of an edge plays a crucial role. Only the owner of an edge is allowed to swap the edge in any state of the process. The strategy S_u of agent u in state G_i is the set of neighbors in G_i to which u owns an edge and the strategy S_u^* is admissible for agent u in state G_i, if $|S_u| = |S_u^*|$ and $|S_u \cap S_u^*| = |S_u| - 1$. Hence, in the ASG the moving agents are allowed to swap one own edge.

- In the *Greedy Buy Game* (GBG), recently introduced by us [14], agents have more freedom to act. In any state, an agent is allowed to buy or to delete or to swap one own edge. Hence, the GBG can be seen as an extension of the ASG. The strategy S_u of agent u in state G_i is defined as in the ASG, but the set of admissible strategies is larger: S_u^* is admissible for agent u in state G_i if (1) $|S_u^*| = |S_u| + 1$ and $S_u \subset S_u^*$ or (2) if $|S_u^*| = |S_u| - 1$ and $S_u^* \subset S_u$ or (3) if $|S_u| = |S_u^*|$ and $|S_u \cap S_u^*| = |S_u| - 1$.

- The *Buy Game* (BG), which is the original version of an NCG and which was introduced by Fabrikant et al. [11], is the most general version. Here agents can perform arbitrary strategy-changes, that is, agents are allowed to perform any combination of buying, deleting and swapping of own edges. The strategy S_u of agent u in G_i is defined as in the ASG, but an admissible strategy for agent u is any set $S_u^* \subseteq V \setminus \{u\}$.

The *cost* of an agent u in network G_i has the form $c_{G_i}(u) = e_{G_i}(u) + \delta_{G_i}(u)$, where $e_{G_i}(u)$ denotes the *edge-cost* and $\delta_{G_i}(u)$ denotes the *distance-cost* of agent u in the network G_i. Each edge has cost $\alpha > 0$, which is a fixed constant, and this cost has to be paid fully by the owner, if not stated otherwise. Hence, if agent u owns k edges in the network

G_i, then $e_{G_i}(u) = \alpha k$. In the (A)SG we simply omit the edge-cost term in the cost function.

There are two variants of distance-cost functions capturing the focus on average or worst-case connection quality. In the SUM-version, we have $\delta_{G_i}(u) = \sum_{v \in V(G_i)} d_{G_i}(u,v)$, if the network G_i is connected and $\delta_{G_i}(u) = \infty$, otherwise. In the MAX-version, we have $\delta_{G_i}(u) = \max_{v \in V(G_i)} d_{G_i}(u,v)$, if G_i is connected and $\delta_{G_i}(u) = \infty$, otherwise. In both cases $d_{G_i}(u,v)$ denotes the shortest path distance between vertex u and v in the undirected graph G_i.

The move policy specifies for any state of the process, which of the unhappy agents is allowed to perform a move. From a mechanism design perspective, the move policy is a way to enforce coordination and to guide the process towards a stable state. We will focus on the *max cost policy*, where the agent having the highest cost is allowed to move and ties among such agents are broken arbitrarily. Sometimes we will assume that an adversary chooses the worst possible moving agent. Note, that the move policy only specifies who is allowed to move, not which specific move has to be performed. We do not consider such strong policies since we do not want to restrict the agents' freedom to act.

Any combination of the four game types, the two distance functions and some move policy together with an initial network completely specifies a network creation process. We will abbreviate names, e.g. by calling the Buy Game with the SUM-version of the distance-cost the SUM-BG. If not stated otherwise, edge-costs cannot be shared.

A cyclic sequence of networks $C_1, \dots, C_j$, where network $C_{i+1 \bmod j}$ arises from network $C_{i \bmod j}$ by an improving move of one agent is called a *better response cycle*. If every move in such a cycle is a *best response move*, which is a strategy-change towards an admissible strategy which yields the largest cost decrease for the moving agent, then we call such a cycle a *best response cycle*. Clearly, a best response cycle is a better response cycle, but the existence of a better response cycle does not imply the existence of a best response cycle.

1.2 Classifying Games According to their Dynamics

Analyzing the convergence processes of games is a very rich and diverse research area. We will briefly introduce two well-known classes of finite strategic games: *games having the finite improvement property* (FIPG) [17] and *weakly acyclic games* (WAG) [20].

FIPG have the most desirable form of dynamic behavior: Starting from any initial state, every sequence of improving moves must eventually converge to an equilibrium state of the game, that is, such a sequence must have finite length. Thus, in such games distributed local search is guaranteed to succeed. It was shown by Monderer and Shapley [17] that a finite game is a FIPG if and only if there exists a *generalized ordinal potential function* Φ, which maps strategy-profiles to real numbers and has the property that if the moving agent's cost decreases, then the potential function value decreases as well. Stated in our terminology, this means that $\Phi : \mathcal{G}_n \to \mathbb{R}$, where $\mathcal{G}_n$ is the set of all networks on n nodes, and we have

$$c_{G_i}(u) - c_{G_{i+1}}(u) > 0 \Rightarrow \Phi(G_i) - \Phi(G_{i+1}) > 0,$$

if agent u is the moving agent in the network G_i. Clearly, no FIPG can admit a better response cycle. An especially nice subclass of FIPG are games that are guaranteed to converge to a stable state in a number of steps which is polynomial in the size of the game. We call this subclass poly-FIPG.

Weakly acyclic games are a super-class of FIPG. Here it is not necessarily true that *any* sequence of improving moves must converge to an equilibrium but we have that from any initial state there exists *some* sequence of improving moves which enforces convergence. Thus, with some additional coordination distributed local search may indeed lead to stable states for such games. A subclass of WAG are games where from any initial state there exists a sequence of best response moves, which leads to an equilibrium. We call those games *weakly acyclic under best response*, BR-WAG for short. Observe, that if a game is not weakly acyclic, then there is *no* way of enforcing convergence if agents stick to playing improving moves.

The above mentioned classes of finite strategic games are related as follows:

$$\text{poly-FIPG} \subset \text{FIPG} \subset \text{BR-WAG} \subset \text{WAG}.$$

The story does not end here. Very recently, Apt and Simon [3] have classified WAG in much more detail by introducing a "scheduler", which is a moderating super-player who guides the agents towards an equilibrium.

1.3 Related Work

The original model of Network Creation Games, which we call the SUM-BG, was introduced a decade ago by Fabrikant et al. [11]. Their motivation was to understand the creation of Internet-like networks by selfish agents without central coordination. In the following years, several variants were proposed: The MAX-BG [9], the SUM-SG and the MAX-SG [2], the SUM-ASG and the MAX-ASG [16], the SUM-GBG and the MAX-GBG [14], a bounded budget version [10], an edge-restricted version [8, 4], a version with bilateral equal-split cost-sharing [6] and a version considering points in a metric space using a different distance measure [18]. All these works focus on properties of stable networks or on the complexity of computing an agent's best response. To the best of our knowledge, the dynamic behavior of most of these variants, including best response dynamics in the well-studied original model, has not yet been analyzed.

Previous work, e.g. [11, 1, 9, 15], has shown that the price of anarchy for the SUM-BG and the MAX-BG is constant for a wide range of α and in $2^{\mathcal{O}(\sqrt{\log n})}$ in general. For the SUM-(A)SG the best upper bound is in $2^{\mathcal{O}(\sqrt{\log n})}$ as well [2, 16], whereas the MAX-SG has a lower bound of $\Omega(\sqrt{n})$ [2]. Interestingly, if played on trees, then the SUM-SG and the MAX-SG have constant price of anarchy [2], whereas the SUM-ASG and the bounded budget version on trees has price of anarchy in $\Theta(\log n)$ [10, 16]. Moreover, it is easy to show that the MAX-ASG on trees has price of anarchy in $\Theta(n)$. Thus, we have the desirable property that selfish behavior leads to a relatively small deterioration in social welfare for most of the proposed versions.

In earlier work [13] we studied the game dynamics of the SUM-SG and showed that if the initial network G_0 is a tree on n nodes, then the network creation process is guaranteed to converge in $\mathcal{O}(n^3)$ steps. By employing the max cost policy, this process can be sped up significantly to $\mathcal{O}(n)$ steps, which is asymptotically optimal. For the SUM-SG on general networks we showed that there exists a best response cycle, which implies that the SUM-SG on arbitrary initial networks is not a FIPG.

Very recently, Cord-Landwehr et al. [7] studied a variant of the MAX-SG, where agents have communication interests, and showed that this variant admits a best response cycle on a tree network as initial network. Hence the restricted-interest variant of the MAX-SG is not a FIPG – even on trees.

Brandes et al. [5] were the first to observe that the SUM-BG is not a FIPG and they prove this by providing a better response cycle. Very recently, Bilò et al. [4] gave a better response cycle for the MAX-BG which implies the same statement for this version. Note, that both proofs contain agents who perform a sub-optimal move at some step in the better response cycle. Hence, these two results do not address the convergence behavior if agents play optimally.

1.4 Our Contribution

In this work, we study Network Creation Games, as proposed by Fabrikant et al. [11], and several natural variants of this model from a new perspective. Instead of analyzing properties of equilibrium states, we apply a more constructive point of view by asking if and how fast such desirable states can be found by selfish agents. For this, we turn the original model and its variants, which are originally formulated as one-shot simultaneous-move games, into more algorithmic models, where moves are performed sequentially.

For the MAX Swap Game on trees, we show that the process must converge in $\mathcal{O}(n^3)$ steps, where n is the number of agents. Furthermore, by introducing a natural way of coordination we obtain a significant speed-up to $\Theta(n \log n)$ steps, which is almost optimal. We show that these results, combined with results from our earlier work [13], give the same bounds for the Asymmetric Swap Game on trees in both the SUM- and the MAX-version.

These positive results for initial networks which are trees are contrasted by several strong negative results on general networks. We show that the MAX-SG, the SUM-ASG and the MAX-ASG on general networks are *not* guaranteed to converge if agents repeatedly perform best possible improving moves and, even worse, that *no* move policy can enforce convergence. We show that these games are not in FIPG, which implies that there cannot exist a generalized ordinal potential function which "guides" the way towards an equilibrium state. For the SUM-ASG we show the even stronger negative result that it can happen that *no* sequence of best response moves may enforce convergence, that is, the SUM-ASG is not even weakly acyclic under best response. If not all possible edges can be created, that is if we have a non-complete *host graph* [8, 4], then we show that the SUM-ASG and the MAX-ASG on non-tree networks is not weakly acyclic. Moreover, we map the boundary between convergence and non-convergence in ASGs and show the surprising result that cyclic behavior can already occur in n-vertex networks which have n edges. That is, even one non-tree edge suffices to completely change the dynamic behavior of these games. In our constructions we have that every agent owns exactly one edge, which is equivalent to the uniform-budget case introduced by Ehsani et al. [10]. In their paper [10] the authors raise the open problem of determining the convergence speed for the bounded-budget version. Thus, our results answer this open problem – even for the simplest version of these games – in the negative, since we show that no convergence guarantee exists.

We provide best response cycles for all versions of the Buy Game, which implies that these games have no convergence guarantee – even if agents have the computational resources to repeatedly compute best response strategies. To the best of our knowledge, the existence of *best* response cycles for all these versions was not known before. Furthermore, we investigate the version where bilateral consent is needed for edge-creation and where the edge-cost is shared equally among its endpoints. We show that this version exhibits a similar undesirable dynamic behavior as the unilateral version. Quite surprisingly, we can show an even stronger negative result in the SUM-version which implies the counter-intuitive statement that cost-sharing may lead to worse dynamic behavior. Our findings nicely contrast a result of Corbo and Parkes [6] who show guaranteed convergence if agents repeatedly play best response strategies against *perturbations* of the other agents' strategies. We show, that these perturbations are necessary for achieving convergence.

Finally, we present a careful empirical study of the convergence time in the ASG and in the GBG. Interestingly, our simulations show that our negative theoretical results seem to be confined to a small set of pathological instances. Even more interesting may be that our simulations show a remarkably fast convergence towards stable networks in $\mathcal{O}(n)$ steps, where n is the number of agents. This indicates that despite our negative results distributed local search may be a suitable method for selfish agents for collectively finding equilibrium networks.

2. MAX SWAP GAMES

In this section we focus on the game dynamics of the MAX-SG. Interestingly, we obtain results which are very similar to the results shown in our earlier work [13] but we need entirely different techniques to derive them. Omitted proofs can be found in the full version [12].

2.1 Dynamics on Trees

We will analyze the network creation process in the MAX-SG when the initial network is a tree. We prove that this process has the following desirable property:

THEOREM 1. *The* MAX*-SG on trees is guaranteed to converge in $\mathcal{O}(n^3)$ steps to a stable network. That is, the* MAX*-SG on trees is a poly-FIPG.*

Before proving Theorem 1, we analyze the impact of a single edge-swap. Let $T = (V, E)$ be a tree on n vertices and let agent v be unhappy in network T. Assume that agent v can decrease her cost by performing the edge-swap vu to vw, for some $u, w \in V$. This swap transforms T into the new network $T' = (V, (E \setminus \{vu\}) \cup \{vw\})$. Let $c_T(v) = \max_{x \in V(T)} d_T(v, x)$ denote agent v's cost in the network T. Let $c_{T'}(u)$ denote her respective cost in T'. Let A denote the tree of $T'' = (V, E \setminus \{vu\})$ which contains v and let B be the tree of T'' which contains u and w. It is easy to see, that we have $d_T(x, y) = d_{T'}(x, y)$, if $x, y \in V(A)$ or if $x, y \in V(B)$.

LEMMA 1. *For all $x \in V(A)$ there is no $y \in V(A)$ such that $c_T(x) = d_T(x, y)$.*

Lemma 1 directly implies the following statement:

COROLLARY 1. *For all $x \in V(A)$: $c_T(x) > c_{T'}(x)$.*

Hence, we have that agent v's improving move decreases the cost for all agents in $V(A)$. For agents in $V(B)$ this may not be true: The cost of an agent $y \in V(B)$ can increase by agent v's move. Interestingly, the next result guarantees that such an increase cannot be arbitrarily high.

LEMMA 2. *Let $x \in V(A)$, $y \in V(B)$ such that $d_{T'}(x, y) = c_{T'}(y)$. It holds that $c_T(x) > c_{T'}(y)$.*

PROOF. In tree T we have $c_T(x) = d_T(x, v) + d_T(u, z) + 1$. Furthermore, in tree T' we have $c_{T'}(y) = d_{T'}(x, v) + d_{T'}(w, y) + 1$. Since $c_T(v) > c_{T'}(v)$, we have $d_T(w, y) < d_T(u, z)$, where $z \in V(B)$ is a vertex having maximum distance to v in T. Hence, this implies $c_T(x) - c_{T'}(y) = d_T(u, z) - d_T(w, y) > 0$. $\square$

Towards a generalized ordinal potential function we will need the following:

Definition 1. (Sorted Cost Vector and Center-Vertex) Let G be any network on n vertices. The *sorted cost vector* of G is $\overrightarrow{c_G} = (\gamma_G^1, \ldots, \gamma_G^n)$, where γ_G^i is the cost of the agent, who has the i-th highest cost in the network G. An agent having cost γ_G^n is called *center-vertex* of G.

LEMMA 3. *Let T be any tree on n vertices. The sorted cost vector of T induces a generalized ordinal potential function for the* MAX-*SG on T.*

PROOF. Let v be any agent in T, who performs an edge-swap which strictly decreases her cost and let T' denote the network after agent v's swap. We show that $c_T(v) - c_{T'}(v) > 0$ implies $\overrightarrow{c_T} >_{\text{lex}} \overrightarrow{c_{T'}}$, where $>_{\text{lex}}$ is the lexicographic order on $\mathbb{N}^n$. The existence of a generalized ordinal potential function then follows by mapping the lexicographic order on $\mathbb{N}^n$ to an isomorphic order on $\mathbb{R}$.

Let the subtrees A and B be defined as above and let $c_T(v) - c_{T'}(v) > 0$. By Lemma 1 and Lemma 2, we know that there is an agent $x \in V(A)$ such that $c_T(x) > c_{T'}(y)$, for all $y \in V(B)$. By Lemma 1 and Corollary 1, we have that $c_T(x) > c_{T'}(x)$, which implies that $\overrightarrow{c_T} >_{\text{lex}} \overrightarrow{c_{T'}}$. $\square$

In the following, a special type of paths in the network will be important.

Definition 2. (Longest Path) Let G be any connected network. Let v be any agent in G having cost $c_G(v) = k$. Any simple path in G, which starts at v and has length k is called a *longest path of agent v*.

As we will see, center-vertices and longest paths are closely related.

LEMMA 4. *Let T be any connected tree and let v^* be a center-vertex of T. Vertex v^* must lie on all longest paths of all agents in $V(T)$.*

PROOF. Let P_{xy} denote the path from vertex x to vertex y in T. We assume towards a contradiction that there are two vertices $v, w \in V(T)$, where $c_T(v) = d_T(v, w)$, and that $v^* \notin V(P_{vw})$. Let $z \in V(T)$ be the only shared vertex of the three paths $P_{vv^*}, P_{wv^*}, P_{vw}$. We have $d_T(v, z) < d_T(v, v^*) \leq c_T(v^*)$ and $d_T(w, z) < d_T(w, v^*) \leq c_T(v^*)$. We show that $c_T(z) < c_T(v^*)$, which is a contradiction to v^* being a center-vertex in T.

Assume that there is a vertex $u \in V(T)$ with $d_T(u, z) \geq c_T(v^*)$. It follows that $V(P_{vz}) \cap V(P_{zu}) = \{z\}$, since otherwise $d_T(v^*, u) = d_T(v^*, z) + d_T(z, u) > c_T(v^*)$. But now, since $d_T(z, w) < c_T(v^*) \leq d_T(z, u)$, we have $d_T(v, u) > c_T(v)$, which clearly is a contradiction. Hence, we have $d_T(z, u) < c_T(v^*)$, for all $u \in V(T)$, which implies that $c_T(z) < c_T(v^*)$. $\square$

Lemma 4, leads to the following observation.

OBSERVATION 1. *Let G be any connected network on n nodes and let $\overrightarrow{c_G} = (\gamma_G^1, \ldots, \gamma_G^n)$ be its sorted cost vector. We have $\gamma_G^1 = \gamma_G^2$ and $\gamma_G^n = \left\lceil \frac{\gamma_G^1}{2} \right\rceil$.*

Now we are ready to provide the key property which will help us upper bound the convergence time.

LEMMA 5. *Let $T = (V, E)$ be a connected tree on n vertices having diameter $D \geq 4$. After at most $\frac{nD - D^2}{2}$ moves of the* MAX-*SG on T one agent must perform a move which decreases the diameter.*

PROOF. Let $v, w \in V$ such that $d_T(v, w) = D \geq 4$ and let P_{vw} be the path from v to w in T. Clearly, if no agent in $V(P_{vw})$ makes an improving move, then the diameter of the network does not change. On the other hand, if the path P_{vw} is the unique path in T having length D, then any improving move of an agent in $V(P_{vw})$ must decrease the diameter by at least 1. The network creation process starts from a connected tree having diameter $D \geq 4$ and, by Lemma 3, must converge to a stable tree in a finite number of steps. Moreover, Lemma 3 guarantees that the diameter of the network cannot increase in any step of the process. It was shown by Alon et al. [2] that any stable tree has diameter at most 3. Thus, after a finite number of steps the diameter of the network must strictly decrease, that is, on all paths of length D some agent must have performed an improving move which reduced the length of the respective path. We fix the path P_{vw} to be the path of length D in the network which survives longest in this process.

It follows, that there are $|V \setminus V(P_{vw})| = n - (D+1)$ agents which can perform improving moves without decreasing the diameter. We know from Observation 1 and Lemma 4 that each one of those $n - (D+1)$ agents can decrease her cost to at most $\left\lceil \frac{D}{2} \right\rceil + 1$ and has to decrease her cost by at least 1 for each edge-swap. We show that an edge-swap of such an agent does not increase the cost of any other agent and use the minimum possible cost decrease per step to conclude the desired bound.

Let $u \in V(T) \setminus V(P_{vw})$ be an agent who decreases her cost by swapping the edge ux to uy and let T' be the tree after this edge-swap. Let $a, b \in V(T)$ be arbitrary agents. Clearly, if $\{u, y\} \not\subseteq V(P_{ab})$ in T', then $d_T(a, b) = d_{T'}(a, b)$. Let A be the tree of $T'' = (V, E \setminus \{uy\})$ which contains u and let B be the tree of T'' which contains y. W.l.o.g. let $a \in V(A)$ and $b \in V(B)$. By Corollary 1, we have $c_T(z) > c_{T'}(z)$ for all $z \in V(A)$ and it follows that $V(A) \cap V(P_{vw}) = \emptyset$. Hence, it remains to analyze the change in cost of all agents in $V(B)$.

If no vertex on the path P_{ab} is a center-vertex in T', then, by Lemma 4, we have that $d_{T'}(a, b) < c_{T'}(b)$. It follows that every longest path of agent b in T' lies entirely in subtree B which implies that $c_{T'}(b) \leq c_T(b)$.

If there is a center-vertex of T' on the path P_{ab} in T', then let v^* be the last such vertex on this path. We have assumed that the diameters of T' and T are equal, which implies that P_{vw} is a longest path of agent v in T'. Since, by Lemma 4, any center-vertex of T' must lie on all longest paths, it follows that v^* is on the path P_{vw} and we have $v^* \in V(B)$. W.l.o.g. let $d_{T'}(v, b) \geq d_{T'}(w, b)$. We have $d_{T'}(a, b) = d_{T'}(a, v^*) + d_{T'}(v^*, b) \leq d_{T'}(v, v^*) + d_{T'}(v^*, b)$. Hence, we have $d_{T'}(a, b) \leq c_{T'}(b)$. Since the path P_{bv} is in subtree B, we have $c_{T'}(b) \leq c_T(b)$.

Now we can easily conclude the upper bound on the number of moves which do not decrease the diameter of T. Each of the $n-(D+1)$ agents with cost at most D may decrease their cost to $\lceil \frac{D}{2} \rceil + 1$. If we assume a decrease of 1 per step, then this yields the following bound:

$$(n-(D+1))\left(D - \left(\left\lceil \frac{D}{2} \right\rceil + 1\right)\right) < \frac{nD - D^2}{2}. \quad \square$$

Proof of Theorem 1. By Lemma 3, we know there exists a generalized ordinal potential function for the Max-SG on trees. Hence, we know that this game is a FIPG and we are left to bound the maximum number of improving moves needed for convergence. It was already shown by Alon et al. [2], that the only stable trees of the Max-SG on trees are stars or double-stars. Hence, the process must stop at the latest when diameter 2 is reached.

Let $N_n(T)$ denote the maximum number of moves needed for convergence in the Max-SG on the n-vertex tree T. Let $D(T)$ be the diameter of T. Let $D_{i,n}$ denote the maximum number of steps needed to decrease the diameter of any n-vertex tree having diameter i by at least 1. Hence, we have

$$N_n(T) \leq \sum_{i=3}^{D(T)} D_{i,n} \leq \sum_{i=3}^{n-1} D_{i,n},$$

since the maximum diameter of a n-vertex tree is $n-1$. By applying Lemma 5 and adding the steps which actually decrease the diameter, this yields

$$N_n(T) \leq \sum_{i=3}^{n-1} D_{i,n} < \sum_{i=3}^{n-1} \left(\frac{ni - i^2}{2} + 1\right)$$
$$< n + \frac{n}{2}\left(\sum_{i=1}^{n} i\right) - \frac{1}{2}\left(\sum_{i=1}^{n} i^2\right) \in \mathcal{O}(n^3). \quad \square$$

The following result shows that we can speed up the convergence time by employing a very natural move policy. The speed-up is close to optimal, since it is easy to see that there are instances in which $\Omega(n)$ steps are necessary.

Theorem 2. *The Max-SG on trees with the max cost policy converges in $\Theta(n \log n)$ moves.*

We prove Theorem 2, by proving the lower and the upper bound separately, starting with the former. Since we analyze the max cost policy, we need two additional observations.

Observation 2. *An agent having maximum cost in a tree T must be a leaf of T.*

Observation 3. *Let u be an unhappy agent in the tree $T = (V, E)$ and let u be a leaf of T and let v be u's unique neighbor. Let B be the tree of $T' = (V, E \setminus \{uv\})$ which contains v. The edge-swap uv to uw, for some $w \in V(B)$ is a best possible move for agent u if w is a center-vertex of B.*

Lemma 6. *There is a tree T on n vertices where the Max-SG on T with the max cost policy needs $\Omega(n \log n)$ moves for convergence.*

Proof. We will consider the path on n-vertices $P_n = v_1 v_2 \dots v_n$ of length $n-1$. We apply the max cost policy and for breaking ties we will always choose the vertex having the smallest index among all vertices having maximum cost. If a maximum cost vertex has more than one best response move, then we choose the edge-swap towards the new neighbor having the smaller index. With these assumptions and with Observation 2 and Observation 3, we have that the center-vertex having the smallest index will "shift" towards a higher index, from $v_{\lceil n/2 \rceil}$ to v_{n-2}. Finally, agent v_n is the unique agent having maximum cost and her move transforms the tree to a star.

We start by analyzing the change in costs of agent v_1. Clearly, $c_0 = c_{P_n}(v_1) = n-1$. By Observation 3, we know that v_1's best swap connects to the minimum index center-vertex of the tree without vertex v_1. Hence after the best move of v_1 this agent has cost $c_1 = \lceil \frac{c_0 - 1}{2} \rceil + 1 > \frac{c_0}{2}$. When v_1 is chosen to move again, her cost can possibly decrease to $\lceil \frac{c_1 - 1}{2} \rceil + 1 > \frac{c_0}{4}$. After the i-th move of v_1 her cost is at least $\left\lceil \frac{c_{i-1} - 1}{2} \right\rceil + 1 > \frac{c_0}{2^i}$. Thus, the max cost policy allows agent v_1 to move at least $\log \frac{c_0}{3}$ times until she is connected to vertex v_{n-2}, the center of the final star, where she has cost 3.

The above implies, that the number of moves of every agent allowed by the max cost policy only depends on the cost of that agent when she first becomes a maximum cost agent. Moreover, since all moving agents are leaves, no move of an agent increases the cost of any other agent. By construction, the cost of every moving agent is determined by her distance towards vertex v_n. Since agent v_n does not move until in the last step of the process, we have that a move of agent v_i does not change the cost of any other agent $v_j \neq v_n$ who moves after v_i. It follows, that we can simply add up the respective lower bounds on the number of moves of all players, depending on the cost when they first become maximum cost agents. It is easy to see, that agent v_i becomes a maximum cost agent, when the maximum cost is $n-i$. Let $M(P_n)$ denote the number of moves of the Max-SG on P_n with the max cost policy and the above tie-breaking rules. This yields

$$M(P_n) > \sum_{c_0 = n-1}^{4} \log \frac{c_0}{3} \in \Omega(n \log n). \quad \square$$

Lemma 7. *The Max-SG on a n-vertex tree T with the max cost policy needs $\mathcal{O}(n \log n)$ moves to converge to a stable tree.*

Proof. Consider any tree T on n vertices. By Observation 2, we know that only leaf-agents are allowed to move by the max cost policy, which implies that no move of any agent increases the cost of any other agent. Observation 3 guarantees that the best possible move of a leaf-agent u having maximum cost c decreases agent u's cost to at most $\lceil \frac{c}{2} \rceil + 1$. Hence, after $\mathcal{O}(\log n)$ moves of agent u her cost must be at most 3. If the tree converges to a star, then agent u may move one more time. If we sum up over all n agents, then we have that after $\mathcal{O}(n \log n)$ moves the tree must be stable. $\square$

2.2 Dynamics on General Networks

In this section we show that allowing cycles in the initial network completely changes the dynamic behavior of the Max-SG.

Theorem 3. *The Max-SG on general networks admits best response cycles. Moreover, no move policy can enforce convergence. The first result holds even if agents are allowed to perform multi-swaps.*

3. ASYMMETRIC SWAP GAMES

In this section we consider the SUM-ASG and the MAX-ASG. Note, that now we assume that each edge has an owner and only this owner is allowed to swap the edge. We show that we can directly transfer the results from above and from [13] to the asymmetric version if the initial network is a tree. On general networks we show even stronger negative results. Omitted proofs can be found in the full version [12].

Observe, that the instance used in the proof of Theorem 3 and the corresponding instance in [13] show that best response cycles in the Swap Game are not necessarily best response cycles in the Asymmetric Swap Game. We will show the rather counter-intuitive result that this holds true for the other direction as well.

3.1 Dynamics in ASGs on Trees

The results in this section follow from the respective theorems in [13] and from the results in Section 2.1 and are therefore stated as corollaries.

COROLLARY 2. *The* SUM*-ASG and the* MAX*-ASG on n-vertex trees are both a poly-FIPG and both must converge to a stable tree in* $\mathcal{O}(n^3)$ *steps.*

COROLLARY 3. *Using the max cost policy and assuming a n-vertex tree as initial network, we have that*

- *the* SUM*-ASG converges in* $\max\{0, n-3\}$ *steps, if n is even and in* $\max\{0, n + \lceil n/2 \rceil - 5\}$ *steps, if n is odd. Moreover, both bounds are tight and asymptotically optimal.*
- *the* MAX*-ASG converges in* $\Theta(n \log n)$ *steps.*

3.2 Dynamics in ASGs on General Graphs

If we move from trees to general initial networks, we get a very strong negative result for the SUM-ASG: There is no hope to enforce convergence if agents stick to playing best responses even if multi-swaps are allowed.

THEOREM 4. *The* SUM*-ASG on general networks is not weakly acyclic under best response. Moreover, this result holds true even if agents can swap multiple edges in one step.*

PROOF. We give a network which induces a best response cycle. Additionally, we show that in each step of this cycle exactly one agent can decrease her cost by swapping an edge and that the best possible swap for this agent is unique in every step. Furthermore, we show that the moving agent cannot outperform the best possible single-swap by a multi-swap. This implies that if agents stick to best response moves then *no* best response dynamic can enforce convergence to a stable network and allowing multi-swaps does not alter this result.

Fig. 1 shows the best response cycle consisting of the networks G_1, G_2, G_3 and G_4. We begin with showing that in $G_1, \ldots, G_4$ all agents, except agent b and agent f, cannot perform an improving strategy change even if they are allowed to swap multiple edges in one step.

In $G_1, \ldots, G_4$ all leaf-agents do not own any edges and the agents c and e cannot swap an edge since otherwise the network becomes disconnected. For the same reason, agent d cannot move the edge towards d_1. Agent d owns three other edges, but they are optimally placed since they are connected to the vertices having the most leaf-neighbors.

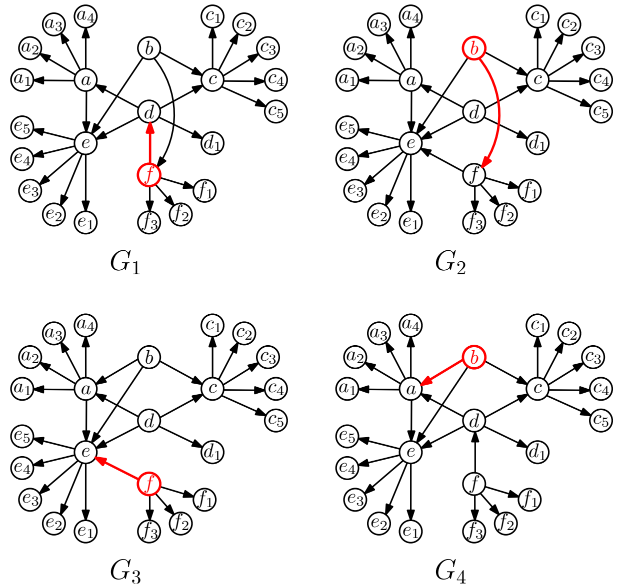

Figure 1: The steps of a best response cycle for the Sum-ASG on general networks. Edge directions indicate edge-ownership. All edges are two-way.

It follows, that agent d cannot decrease her cost by swapping one edge or by performing a multi-swap. Note, that this holds true for all networks $G_1, \ldots, G_4$, although the networks change slightly. Agent a cannot move her edges towards a_i, for $1 \leq i \leq 4$. On the other hand, it is easy to see that agent a's edge towards vertex e cannot be swapped to obtain a strict cost decrease since the most promising choice, which is vertex c, yields the same cost in G_1 and G_4 and even higher cost in G_2 and G_3. Trivially, no multi-swap is possible for agent a.

Now, we consider agent b and agent f. First of all, observe that in $G_1, \ldots, G_4$ agent f owns exactly one edge which is not a bridge. Thus, agent f cannot perform a multi-swap in any step of the best response cycle. Agent b, although owning three edges, is in a similar situation: Her edges to vertex c and e can be considered as fixed, since swapping one or both of them does not yield a cost decrease in $G_1, \ldots, G_4$. Hence, agent b and agent f each have one "free" edge to operate with. In G_1 agent b's edge towards f is placed optimally, since swapping towards a or d does not yield a cost decrease. In G_3, agents b's edge towards a is optimal, since swapping towards d or f does not decrease agent b's cost. Analogously, agent f's edge towards e in G_2 and her edge towards d in G_4 are optimally placed.

Last, but not least, we describe the best response cycle: In G_1 agent f can improve and her unique best possible edge-swap in G_1 is the swap from d to e, yielding a cost decrease of 4. In G_2 agent b has the swap from f to a as unique best improvement which yields a cost decrease of 1. In G_3 have agent f being unhappy with her strategy and the unique best swap is the one from e to d yielding an improvement of 1. In G_4 it is agent b's turn again and her unique best swap is from a to f which decreases her cost by 3. After agent b's swap in G_4 we arrive again at network G_1, hence $G_1, \ldots, G_4$ is a best response cycle where in each step exactly one agent has a single-swap as unique best possible improvement. □

Note, that the best response cycle presented in the proof of Theorem 4 is not a best response cycle in the SUM-SG. The swap fb to fe of agent f in G_1 yields a strictly larger cost decrease than her swap fd to fe.

Compared to Theorem 4, we show a slightly weaker negative result for the max-version.

THEOREM 5. *The* MAX-*ASG on general networks admits best response cycles. Moreover, no move policy can enforce convergence.*

If played on a non-complete host-graph, then we get the worst possible dynamic behavior.

COROLLARY 4. *The* SUM-*ASG and the* MAX-*ASG on a non-complete host graph are not weakly acyclic.*

3.3 The Boundary between Convergence and Non-Convergence

In this section we explore the boundary between guaranteed convergence and cyclic behavior. Quite surprisingly, we can draw a sharp boundary by showing that the undesired cyclic behavior can already occur in n-vertex networks having exactly n edges. Thus, one non-tree edge suffices to radically change the dynamic behavior of Asymmetric Swap Games. Our constructions are such that each agent owns exactly one edge, which corresponds to the uniform unit budget case, recently introduced by Ehsani et al. [10]. Hence, even if the networks are build by identical agents having a budget the cyclic behavior may arise. This answers the open problem raised by Ehsani et al. [10] in the negative.

THEOREM 6. *The* SUM-*ASG and the* MAX-*ASG admit best response cycles on a network where every agent owns exactly one edge.*

PROOF OF THEOREM 6, SUM-VERSION. The network inducing a best response cycle and the steps of the cycle are shown in Fig. 2. Let n_k denote the number of vertices having the form k_j, for some index j.

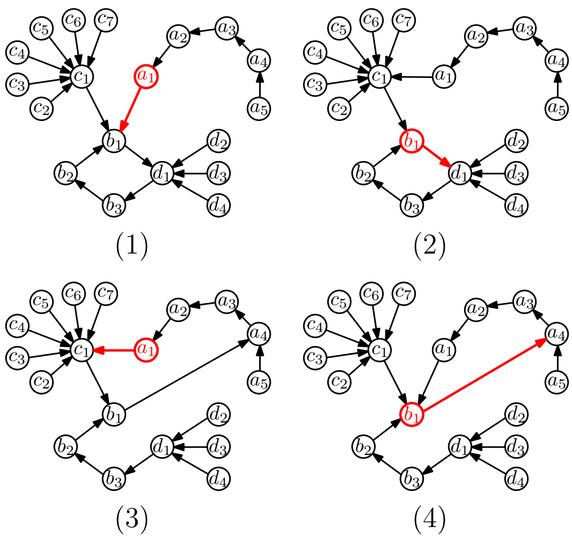

Figure 2: The steps of a best response cycle for the Sum-ASG where each agent owns exactly one edge.

In the first step, depicted in Fig. 2 (1), agent a_1 has only one improving move, which is the swap from b_1 to c_1. This swap reduces agent a_1's cost by 1, since $n_c = n_b + n_d + 1$. After this move, shown in Fig. 2 (2), agent b_1 is no longer happy with her edge towards d_1, since by swapping towards a_4 she can decrease her cost by 2. This is a best possible move for agent b_1 (note, that a swap towards a_3 yields the same cost decrease). But now, in the network shown in Fig. 2 (3), by swapping back towards vertex b_1, agent a_1 can additionally decrease her distances to vertices a_4 and a_5 by 1. This yields that agent a_1's swap from c_1 to b_1 decreases her cost by 1. This is true, since all distances to c_j vertices increase by 1 but all distances to b_i and d_l vertices and to a_4 and a_5 decrease by 1 and since we have $n_c = n_b + n_d + 1$. Note, that this swap is agent a_1's unique improving move. By construction, we have that after agent a_1 has swapped back towards b_1, depicted in Fig. 2 (4), agent b_1's edge towards a_4 only yields a distance decrease of 7. Hence, by swapping back towards d_1, agent b_1 decreases her cost by 1, since her sum of distances to the d_j vertices decreases by 8. This swap is the unique improving move of agent b_1 in this stage. Now the best response cycle starts over again, with agent a_1 moving from b_1 to c_1. □

3.4 Empirical Study of the Bounded-Budget Version

We have conducted extensive simulations of the convergence behavior and the obtained results provide a sharp contrast to our mostly negative theoretical results for both versions of the ASG. Our experiments show for the bounded-budget version a surprisingly fast convergence in at most $5n$ steps under the max cost policy or by choosing the moving agents uniformly at random. Despite millions of trials we have not found any best response cycle in our experiments. This indicates that our negative results may be only very rare pathological examples. We refer to the full version [12] for a detailed description of our simulations and results.

4. (GREEDY) BUY GAMES

We focus on the dynamic behavior of the Buy Game and the Greedy Buy Game. Remember, that we assume, that each edge can be created for the cost of $\alpha > 0$.

4.1 Convergence Results

We show that best response cycles exist, even if arbitrary strategy-changes are allowed. However, on the positive side, we were not able to construct best response cycles where only one agent is unhappy in every step. Hence, the right move policy may have a substantial impact in (Greedy) Buy Games. In contrast to this, we rule out this glimmer of hope if played on a non-complete host-graph.

THEOREM 7. *The* SUM-*(G)BG and the* MAX-*(G)BG admit best response cycles.*

We sketch the proof by providing the best response cycle for the SUM-(G)BG in Fig. 3 and for the MAX-(G)BG in Fig. 4. See the full version [12] for the complete proof.

If we restrict the set of edges which can be build, then we get the worst possible dynamic behavior. In this case there is no hope for convergence if agents are only willing to perform improving moves.

COROLLARY 5. *The* SUM-*(G)BG and the* MAX-*(G)BG on general host graphs is not weakly acyclic.*

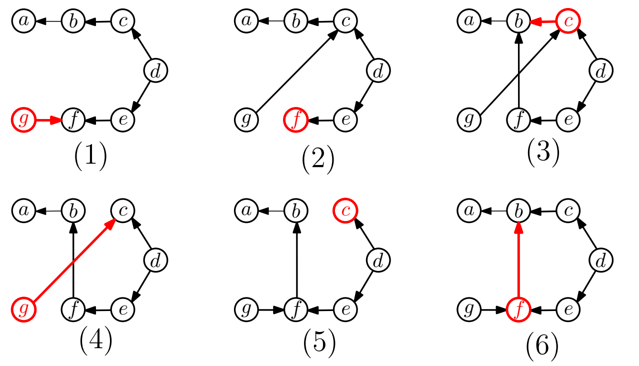

Figure 3: The steps of a best response cycle for the Sum-(G)BG for $7 < \alpha < 8$.

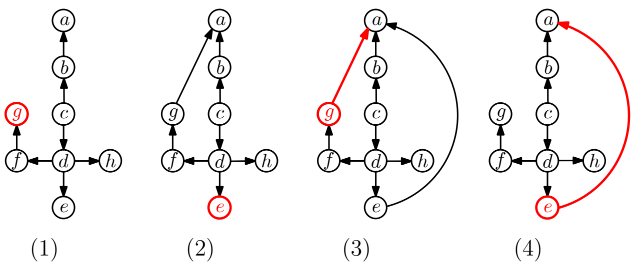

Figure 4: The steps of a best response cycle for the Max-(G)BG for $1 < \alpha < 2$.

4.2 Empirical Study of Greedy Buy Games

We give empirical results for the convergence time for both versions of the GBG. Our focus is on the GBG, since a best response for both versions of the GBG can be computed in polynomial time [14], whereas this problem is well-known [11, 15] to be NP-hard for the BG. Due to space constraints we can only sketch the setup and give a glimpse of the results. See the full version [12] for a detailed discussion and additional plots.

4.2.1 Experimental Setup

One run of our simulations consists of the generation of a random initial network and then the max cost or the random policy is employed in the GBG until the process converges to a stable network. We measure the number of steps needed for this convergence to happen and take the maximum over 5000 such runs for each configuration.

The initial networks are generated as follows: Starting from an empty graph on n vertices we first generate a random spanning tree to enforce connectedness of our networks. Then we randomly insert edges until the desired number of edges is present. Note, that we do not allow multi-edges. The ownership of every edge is chosen uniformly at random among the endpoints.

We have considered networks having n agents, where n ranges between 10 and 100. In order to investigate the impact of the density of the initial network on the convergence time, we fix the number of edges in the initial network to be n, $2n$ and $4n$, respectively. The impact of the edge-cost parameter α is investigated by setting α to $n/10$, $n/4$, $n/2$ and n, respectively. Demaine et al. [9] argue that this is the most interesting range for α, since implies that the average distance is roughly on par with the creation cost of an edge.

4.2.2 Experimental Results

We have observed a remarkably small number of steps needed for convergence in these games, which indicates that distributed local search is a practical method for selfishly creating stable networks. For the SUM-GBG no run took longer than $7n$ steps to converge, whereas for the MAX-version we always observed less than $8n$ steps until convergence, see Fig. 5. It can be seen that the convergence time grows roughly linear in n for all configurations, which implies that these processes scale very well.

The number of edges in the initial network has an impact on the convergence time: All curves for $m = 4n$ are well above the respective curves for $m = n$. The reason for this may be the relatively high value for α compared to the diameter of the resulting networks. We have not found any stable network having a diameter larger than 4.

The choice of α influences the convergence time in the SUM-version but not in the MAX-version. In the former we see that a smaller α generally yields a higher convergence time. In the latter, there may be no influence since the relatively high values of α yield that the edge-cost dominates the distance-cost for most of the agents.

Moreover, as in the simulations for the ASG, despite several millions of trials we did not encounter a cyclic instance. This indicates that such instances are rather pathological and may never show up in practice.

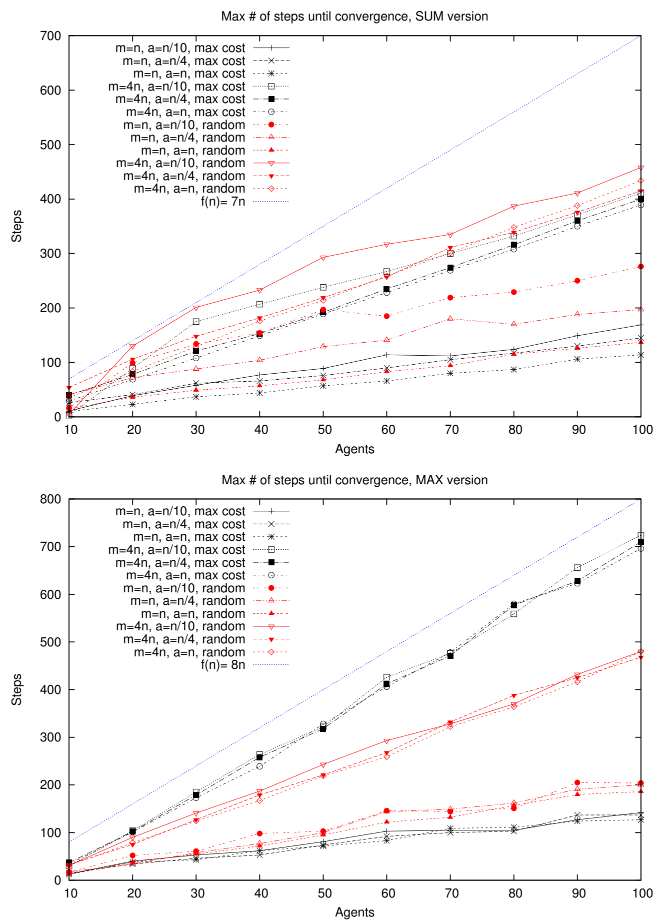

Figure 5: Experimental results for both versions of the GBG. Each point is the maximum over the steps needed for convergence of 5000 trials with random initial networks having m edges and $\alpha = a$.

5. BILATERAL BUY GAMES WITH COST-SHARING

We consider "bilateral network formation", as introduced by Corbo and Parkes [6], which we call the *bilateral equal-split BG*. This version explicitly models that bilateral consent is needed in order to create an edge, which is a realistic assumption in some settings. The cost of an edge is split equally among its endpoints and edges are build only if *both* incident agents are willing to pay half of the edge-price. This model implicitly assumes coordination among coalitions of size two and the corresponding solution concept is therefore the pairwise Nash equilibrium, which can be understood as the minimal coalitional refinement of the pure Nash equilibrium. The authors of [6] show that this solution concept is equivalent to Meyerson's proper equilibrium [19], which implies guaranteed convergence if the agents repeatedly play best response strategies against *perturbations* of the other players' strategies, where costly mistakes are made with less probability. We show in this section that these perturbations are necessary for achieving convergence by proving that the bilateral equal-split BG is not weakly acyclic in the SUM-version and that it admits best response cycles in the MAX-version. Interestingly, the first result is stronger than the result for the SUM-(G)BG, which yields the counter-intuitive observation, that sharing the cost of edges can lead to worse dynamic behavior.

THEOREM 8. *The* SUM *bilateral equal-split Buy Game is not weakly acyclic.*

For the MAX-version, we can show a slightly weaker result.

THEOREM 9. *The* MAX *bilateral equal-split Buy Game admits best response cycles.*

6. REFERENCES

[1] S. Albers, S. Eilts, E. Even-Dar, Y. Mansour, and L. Roditty. On nash equilibria for a network creation game. In *Proceedings of the seventeenth annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '06, pages 89–98, New York, NY, USA, 2006. ACM.

[2] N. Alon, E. D. Demaine, M. Hajiaghayi, and T. Leighton. Basic network creation games. In *SPAA '10: Proceedings of the 22nd ACM Symposium on Parallelism in Algorithms and Architectures*, pages 106–113, New York, NY, USA, 2010. ACM.

[3] K. Apt and S. Simon. A classification of weakly acyclic games. In M. Serna, editor, *Algorithmic Game Theory*, volume 7615 of *Lecture Notes in Computer Science*, pages 1–12. Springer Berlin / Heidelberg, 2012.

[4] D. Bilò, L. Gualà, S. Leucci, and G. Proietti. The max-distance network creation game on general host graphs. In P. Goldberg, editor, *Internet and Network Economics*, Lecture Notes in Computer Science, pages 392–405. Springer Berlin Heidelberg, 2012.

[5] U. Brandes, M. Hoefer, and B. Nick. Network creation games with disconnected equilibria. In C. Papadimitriou and S. Zhang, editors, *Internet and Network Economics*, volume 5385 of *Lecture Notes in Computer Science*, pages 394–401. Springer Berlin Heidelberg, 2008.

[6] J. Corbo and D. Parkes. The price of selfish behavior in bilateral network formation. In *Proceedings of the twenty-fourth annual ACM symposium on Principles of distributed computing*, PODC '05, pages 99–107, New York, NY, USA, 2005. ACM.

[7] A. Cord-Landwehr, M. Hüllmann, P. Kling, and A. Setzer. Basic network creation games with communication interests. In M. Serna, editor, *Algorithmic Game Theory*, volume 7615 of *Lecture Notes in Computer Science*, pages 72–83. Springer Berlin / Heidelberg, 2012.

[8] E. D. Demaine, M. T. Hajiaghayi, H. Mahini, and M. Zadimoghaddam. The price of anarchy in cooperative network creation games. *SIGecom Exch.*, 8(2):2:1–2:20, Dec. 2009.

[9] E. D. Demaine, M. T. Hajiaghayi, H. Mahini, and M. Zadimoghaddam. The price of anarchy in network creation games. *ACM Trans. on Algorithms*, 8(2):13, 2012.

[10] S. Ehsani, M. Fazli, A. Mehrabian, S. S. Sadeghabad, M. Safari, M. Saghafian, and S. ShokatFadaee. On a bounded budget network creation game. In *SPAA*, pages 207–214, 2011.

[11] A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. Shenker. On a network creation game. In *Proceedings of the twenty-second annual Symposium on Principles of Distributed Computing*, PODC '03, pages 347–351, New York, NY, USA, 2003. ACM.

[12] B. Kawald and P. Lenzner. On dynamics in selfish network creation. *CoRR*, abs/1212.4797, 2012.

[13] P. Lenzner. On dynamics in basic network creation games. In G. Persiano, editor, *Algorithmic Game Theory*, volume 6982 of *Lecture Notes in Computer Science*, pages 254–265. Springer Berlin / Heidelberg, 2011.

[14] P. Lenzner. Greedy selfish network creation. In P. Goldberg, editor, *Internet and Network Economics*, Lecture Notes in Computer Science, pages 142–155. Springer Berlin Heidelberg, 2012.

[15] M. Mihalák and J. C. Schlegel. The price of anarchy in network creation games is (mostly) constant. In *Proceedings of the Third international conference on Algorithmic Game Theory*, SAGT'10, pages 276–287, Berlin, Heidelberg, 2010. Springer-Verlag.

[16] M. Mihalák and J. C. Schlegel. Asymmetric swap-equilibrium: A unifying equilibrium concept for network creation games. In B. Rovan, V. Sassone, and P. Widmayer, editors, *Mathematical Foundations of Computer Science 2012*, volume 7464 of *LNCS*, pages 693–704. Springer Berlin / Heidelberg, 2012.

[17] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124 – 143, 1996.

[18] T. Moscibroda, S. Schmid, and R. Wattenhofer. On the topologies formed by selfish peers. In *Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*, PODC '06, pages 133–142, New York, NY, USA, 2006. ACM.

[19] R. Myerson. *Game theory: analysis of conflict.* Harvard University Press, 1997.

[20] H. P. Young. The evolution of conventions. *Econometrica*, 61(1):pp. 57–84, 1993.