On Dynamics in Basic Network Creation Games*

(full version)

Pascal Lenzner†

Abstract

We initiate the study of game dynamics in the Sum Basic Network Creation Game, which was recently introduced by Alon et al.[SPAA'10]. In this game players are associated to vertices in a graph and are allowed to "swap" edges, that is to remove an incident edge and insert a new incident edge. By performing such moves, every player tries to minimize her connection cost, which is the sum of distances to all other vertices. When played on a tree, we prove that this game admits an ordinal potential function, which implies guaranteed convergence to a pure Nash Equilibrium. We show a cubic upper bound on the number of steps needed for any improving response dynamic to converge to a stable tree and propose and analyse a best response dynamic, where the players having the highest cost are allowed to move. For this dynamic we show an almost tight linear upper bound for the convergence speed. Furthermore, we contrast these positive results by showing that, when played on general graphs, this game allows best response cycles. This implies that there cannot exist an ordinal potential function and that fundamentally different techniques are required for analysing this case. For computing a best response we show a similar contrast: On the one hand we give a linear-time algorithm for computing a best response on trees even if players are allowed to swap multiple edges at a time. On the other hand we prove that this task is NP-hard even on simple general graphs, if more than one edge can be swapped at a time. The latter addresses a proposal by Alon et al..

1 Introduction

The importance of the Internet as well as other networks has inspired a huge body of scientific work to provide models and analyses of the networks we interact with every day. These models incorporate game theoretic notions to be able to express and analyse selfish behavior within these networks. Such behavior by players can be the creation or removal of links to influence the network structure to better suit their needs. However, most of this work focused on *static* properties of such networks, like structural properties of solution concepts. Prominent examples are bounds on the Price of Anarchy or on the Price of Stability of (pure) Nash Equilibria in games that model network creation. The problem is, that such results do not explain how selfish and myopic players can actually *find* such desired states.

*See [9] for the original publication.

†Department of Computer Science, Humboldt-Universität zu Berlin, Unter den Linden 6, 10099 Berlin, Germany. Email: `lenzner@informatik.hu-berlin.de`

In this paper we focus on the process itself. That is, on the dynamic behavior of players which eventually leads to a state of the game having interesting properties like stability against unilateral deviations and low social cost. We initiate the study of myopic game dynamics in the SUM BASIC NETWORK CREATION GAME, which was introduced very recently by Alon et al. [2]. This elegant model incorporates important aspects of network design as well as network routing but is at the same time simple enough to provide insights into the induced dynamic process. The idea is to let players "swap" edges to resemble the natural process of weighing two decisions (possible edges) against each other. We investigate the convergence process of dynamics which allow players to myopically swap edges until a stable state of the game emerges. Furthermore, we take the mechanism design perspective and propose a specific dynamic, which yields near optimal convergence speed.

1.1 Related Work

The line of research which is closest to our work was initiated by Fabrikant et al. [5], who considered network creation with a fixed edge-cost of α. For some ranges of α they proved first bounds on the Price of Anarchy [8], which is the ratio of the social cost of the worst (pure) Nash Equilibrium and the minimum possible social cost achieved by central design. Subsequent work [1, 4, 10, 11] has shown, that this ratio is constant for almost all values of α. Only for $\alpha \in \Theta(n)$ there remains a gap. However, there is a downside of this model: As already observed in [5], computing a best response is NP-hard, which implies, that players cannot efficiently decide if the game has reached a stable state. This computational hardness prevents myopic dynamics from being applied to finding a pure Nash Equilibrium.

Very recently, Alon et al. [2] proposed a slightly different model, which no longer depends on the parameter α but still captures important aspects of network creation. The authors consider two different cost-measures, namely the sum of distances to all other players and the maximum distance to all other players and give bounds on the price of anarchy. Here, we adopt the former measure. Alon et al. proved that in this case the star is the only equilibrium tree. Interestingly, as observed in [11], it is not true that the class of equilibria in the model without parameter is a super-class of the equilibria in the original model. Nevertheless, we believe that the model of Alon et al. is still interesting, because it models the natural process of locally weighing alternatives against each other. Furthermore, it has another striking feature: Best responses can be computed efficiently. Thus, applying myopic dynamics seems a natural choice for the task of finding stable states in the game. The authors of [2] also propose to analyse the case where players are allowed to swap more than one edge at a time.

The work of Baumann and Stiller [3] is very similar in spirit to our work. They provide deep insights into the dynamics of a related network creation game and show various structural properties, e.g. sufficient and necessary conditions for stability.

Due to space constraints we refer for further work on selfish network creation to Jackson's survey [6] and to the references in Nisan et al. [13, Chapter 19].

1.2 Model and Definitions

The SUM BASIC NETWORK CREATION GAME is defined as follows: Given an undirected, connected graph $G = (V, E)$, where each vertex corresponds to a player. Every player $v \in V$

selfishly aims to minimize her connection cost by performing moves in the game. A player's *connection cost* $c(v)$ is the sum of all shortest-path distances to all other players. If the graph is disconnected, then we define $c(v)$ to be infinite. At any time, a player can "swap" an incident existing edge with an incident non-edge at no cost. More formally, let u be a neighbor of v and w be a non-neighbor of v, then the edge swap (u, w) of player v removes the edge vu and creates the edge vw. Let $\Gamma_G(v)$ denote the closed neighborhood of v in G, which includes v and all neighbors of v. The set of pure strategies for player v in G is $S_G(v) = \big(\Gamma_G(v) \setminus \{v\} \times V \setminus \Gamma_G(v)\big) \cup \{\perp\}$, where $\perp$ denotes, that player v does not swap. Note, that this set depends on the current graph G and that moves of players in the game modify the graph. We allow only pure strategies and call a pure strategy $s \in S_G(v)$, which decreases player v's current connection cost most, a *best response*. Sometimes we say that a vertex x is a best response of a player v, which abbreviates, that v has a best response of the form (y, x), for some $y \neq x$. Player v is called *active* if v can swap an edge to strictly decrease her connection cost. If no such swap is possible, then we call player v *passive*.

We assume that players are lazy, in the sense that if for some player v the best possible edge-swap yields no decrease in connection cost, then player v prefers the strategy $\perp$, that is, not to swap. We say that G is *stable* or in *swap-equilibrium* if $\perp$ is a best response for every player.

Since the model does not include costs for edges, the utility of a player is simply the negative of her connection cost. Let $x \in G$ denote that G contains vertex x. The connection cost of player v in graph G is defined as $c_G(v) = \sum_{x \in G} d_G(v, x)$, where $d_G(v, x)$ is the number of edges on the shortest path from v to x in G. We omit the reference to G, if it is clear from the context. The *social cost* of a graph G is the sum of the connection costs of all players in G.

Furthermore, we use the convention, that for a graph G, we let $|G|$ denote the number of vertices in G and we define $G - x$ to be the graph G after the removal of vertex x.

1.3 Our Contribution

We provide a rigorous treatment of the induced game dynamics of the SUM BASIC NETWORK CREATION GAME on trees. For this case, Theorem 1 shows that the game dynamic has the desirable property that local improvements by players directly yield a global improvement in terms of the social cost. More formally, we show that the game on trees is an *ordinal potential game* [12], that is, there exists a function mapping states of the game to values with the property that pure Nash Equilibria of the game correspond to local minima of the function. A prominent feature of such games is, that a series of local improvements must eventually converge to a pure Nash Equilibrium – a stable state of the game in which no player wants to unilaterally change her strategy. Theorem 3 shows that this convergence is fast by providing a cubic upper bound on the number of steps any improving response dynamic needs to reach such a stable state. Furthermore we introduce and analyse a natural dynamic called *Max Cost Best Response Dynamic*. This dynamic is proven to be close to optimal in terms of convergence speed, since Theorem 4 shows that the number of steps needed by this dynamic almost matches the trivial lower bound. This implies, that the process of finding a pure Nash Equilibrium can be significantly sped up by introducing coordination and enforcing that best responses are played.

In contrast to these positive results on trees, Theorem 7 is a strong negative result for the SUM BASIC NETWORK CREATION GAME on general graphs. We show that in this case best response dynamics can cycle, which implies, that there cannot exist an ordinal potential function. Thus, any treatment of the game dynamics on general graphs requires fundamentally different techniques and is an interesting open problem for ongoing research.

Last, but not least, we use structural insights to obtain a linear-time algorithm for computing a best response on trees even for the case where players are allowed to swap multiple edges at a time. For the game on general graphs, we provide another sharp contrast by showing that computing a best response in the general case is NP-hard, if more than one edge can be swapped at a time. This is particularly interesting, since this addresses the proposal of Alon et al. [2] to analyse this case. Our results imply, that in this case best responses can be efficiently computed only if the game is played on trees or on very simple graphs.

2 Playing on a Tree

In this section we consider the special case where the given graph G is a tree. We show, that the SUM BASIC NETWORK CREATION GAME on trees belongs to the well-studied class of *ordinal potential games* [12]. This guarantees the desirable property that pure Nash Equilibria always exist and that such solutions can be found by myopic play.

Theorem 1. *The* SUM BASIC NETWORK CREATION GAME *on trees is an ordinal potential game.*

Before we prove the Theorem, we analyse the impact of an edge-swap on the connection cost of the swapping player and on the social cost.

Let $T = (V, E)$ be a tree having n vertices. Assume that player v performs the edge-swap vu to vw in the tree T. (Note, that this implies, that $vw \notin E$). Let T' be the tree obtained after this edge-swap. Let Φ and Φ' be the social cost of T and T', respectively. Let T_v and T_u be the tree T rooted at v and u, respectively. Let A be the subtree rooted at v in T_u and let B be the subtree rooted at u in T_v. See Fig. 1 for an illustration. Let $c_K(z) = \sum_{k \in K} d_K(z, k)$

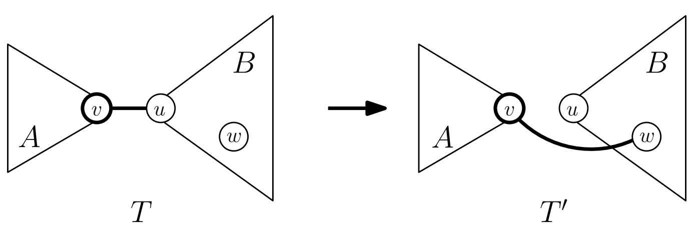

Figure 1: Player v swaps edge vu to edge vw.

denote the connection cost of player z within tree K.

Lemma 1. *The change in player v's connection cost induced by the edge-swap vu to vw is $\Delta(v) = c_B(u) - c_B(w)$.*

Proof. Let $c(v)$ and $c'(v)$ denote the connection cost of v in T and T', respectively. In T, we have

$$c(v) = \sum_{x \in A} d(v,x) + \sum_{y \in B} d(v,y) = \sum_{x \in A} d(v,x) + \sum_{y \in B} (1 + d(u,y)) = \sum_{x \in A} d(v,x) + |B| + c_B(u).$$

Analogously, we obtain $c'(v) = \sum_{x \in A} d(v,x) + |B| + c_B(w)$. Thus, we have

$$\Delta(v) = c(v) - c'(v) = c_B(u) - c_B(w).$$

□

The following Lemma implies the desired property, that local improvement of a player yields a global improvement in terms of social cost.

Lemma 2. *The change in social cost induced by the edge-swap vu to vw is*

$$\Delta(\Phi) = 2|A|\Delta(v) \ .$$

Proof. First, we analyse the social cost Φ in terms of the subtrees A and B:

$$\begin{aligned}
\Phi &= \sum_{x \in T} c(x) = \sum_{x \in A} c(x) + \sum_{x \in B} c(x) \\
&= \sum_{x \in A} \Big(\sum_{y \in A} d(x,y) + \sum_{y \in B} d(x,y) \Big) + \sum_{x \in B} \Big(\sum_{y \in B} d(x,y) + \sum_{y \in A} d(x,y) \Big).
\end{aligned}$$

For all $a \in A$ and $b \in B$, we have that the neighbors v and u lie on the shortest path between a and b. Hence, we have

$$\begin{aligned}
\Phi &= \sum_{x \in A} \sum_{y \in A} d(x,y) + \sum_{x \in A} \sum_{y \in B} (d(x,v) + d(v,y)) + \sum_{x \in B} \sum_{y \in B} d(x,y) + \sum_{x \in B} \sum_{y \in A} (d(x,u) + d(u,y)) \\
&= \sum_{x \in A} \sum_{y \in A} d(x,y) + \sum_{x \in A} \Big(|B| d(x,v) + \sum_{y \in B} (1 + d(u,y)) \Big) \\
&\qquad + \sum_{x \in B} \sum_{y \in B} d(x,y) + \sum_{x \in B} \Big(|A| d(x,u) + \sum_{y \in A} (1 + d(v,y)) \Big)
\end{aligned}$$

Let $\Phi_A = \sum_{a \in A} c_A(a)$ and $\Phi_B = \sum_{b \in B} c_B(b)$ denote the social cost of tree A and B, respectively. Observe, that $\Phi_A = \sum_{x \in A} \sum_{y \in A} d(x,y)$ and $\Phi_B = \sum_{x \in B} \sum_{y \in B} d(x,y)$. Putting all pieces together, this yields

$$\begin{aligned}
\Phi &= \Phi_A + |B| c_A(v) + |A||B| + |A| c_B(u) + \Phi_B + |A| c_B(u) + |A||B| + |B| c_A(v) \\
&= \Phi_A + \Phi_B + 2(|A||B|) + 2|B| c_A(v) + 2|A| c_B(u).
\end{aligned}$$

In an analogous way, we get $\Phi' = \Phi_A + \Phi_B + 2(|A||B|) + 2|B| c_A(v) + 2|A| c_B(w)$. Thus, the amount of the change in social cost of any edge-swap is $\Delta(\Phi) = \Phi - \Phi' = 2|A|(c_B(u) - c_B(w))$. By Lemma 1, we have that $\Delta(v) = (c_B(u) - c_B(w))$. Thus, $\Delta(\Phi) = 2|A|\Delta(v)$. □

Now we are ready to prove Theorem 1.

Proof of Theorem 1. By Lemma 2, we have that the social cost strictly decreases if and only if the connection cost of the swapping player strictly decreases. This implies, that the social cost Φ is an ordinal potential function for the SUM BASIC NETWORK CREATION GAME on trees. $\square$

Theorem 1 guarantees that a pure Nash Equilibrium of this game can be reached by myopic play, even if the players do not play in an optimal way. We only need one very natural ingredient for convergence: Whenever a player moves, this move must decrease this player's connection cost. We call every dynamic where a player strictly improves by making a move (or passing if no improving move is possible) an *improving response dynamic*(IRD). Such a dynamic stops if no player can strictly improve, which implies that any IRD stops if a stable graph is obtained.

2.1 Improving Response Dynamics on Trees

For trees it was shown by Alon et al. [2] that the star is the only stable tree. Using this observation and Theorem 1, we arrive at the following Corollary.

Corollary 1. *For every tree T, every IRD converges to a star.*

Having guaranteed convergence, the natural question to ask is how many steps are needed to reach the unique pure Nash Equilibrium by myopic play. The following Theorems provide a lower and an upper bound on that number.

Theorem 2. *Let P_n be a path having n vertices. Any IRD on P_n needs at least* $\max\{0, n-3\}$ *steps to converge.*

Proof. Let $n \geq 4$, since otherwise P is already a star. Since the leaf-players can perform an improving move (every swap to an inner vertex strictly decreases their connection cost), we have that P cannot be stable. By Corollary 1, any IRD converges to a star on n vertices. Clearly, such a star contains a vertex having degree $n-1$. In any stepwise transformation of P into a star, some vertex will become the center of the star. Assume that an inner vertex v of P is the designated center. Since v has degree 2 there are $n-3$ non-neighbors of v. Since in every step of the dynamic only one edge can be swapped, it follows that at least $n-3$ steps are needed, to connect all of these non-neighbors to v. If a leaf of P is the designated center, then one additional step is needed. $\square$

Lemma 3. *P_n is the tree on n vertices which has maximum social cost.*

Proof. Let T have at least four vertices and assume towards a contradiction, that T has more than two leaves and has maximum social cost. Consider a leaf l of T, which has minimum connection cost $c(l)$ among all leaves of T. Let k be the neighbor of l in T and observe, that $c(l) = c(k) + (n-2)$, since k is the first vertex on l's shortest paths to all other $n-2$ vertices of T. Let u be a leaf of T, which has maximum connection cost $c(u)$ among all leaves of T. Thus, we have $c(u) \geq c(l)$. Now consider the tree $T-l$ and let $c'(k)$ and $c'(u)$ denote the connection cost of player k and u within $T-l$. Since l is a neighbor of k in T, we have $c'(k) = c(k) - 1 = c(l) - (n-1)$. Furthermore, we have $c'(u) = c(u) - d_T(u, l)$. The tree T has at least three leaves, which implies that the longest path of T can have length at most $n-2$. Thus, $d_T(u,l) < n-1$, together with $c(u) \geq c(l)$, this implies $c'(u) > c'(k)$. Consider the edge-swap lk to lu by player l and let T'' be the obtained tree. Let $c''(l)$ and $c''(u)$ be the connection cost of l and u in T''. We have $c''(l) = c''(u)+(n-2) = c'(u)+(n-1) > c'(k)+(n-1) = c(l)$. Thus, the edge-swap lk to lu strictly increases player l's connection cost. By Lemma 2, it follows that the social cost of T'' is strictly larger than the social cost of T, which is a contradiction. Hence, every tree with maximum social cost must have exactly two leaves. □

Theorem 3. *Any IRD on trees having n vertices converges in $\mathcal{O}(n^3)$ steps.*

Proof. The idea is to start with the tree having the highest potential and to bound the number of steps any IRD needs by analysing the number of steps needed if this potential is decreased by the smallest possible amount per step. By Lemma 3, we have that P_n has the maximum social cost Φ_{P_n}. Observe, that $\Phi_{P_n} = \sum_{i=1}^{n-1} 2i(n-i) = \frac{n^3-n}{3}$. Let X_n be a star having n vertices. We have $\Phi_{X_n} = 2n^2 - 4n + 2$. To transform P_n into X_n any IRD has to decrease the social cost by $\Phi_{P_n} - \Phi_{X_n} = \frac{n^3}{3} - 2n^2 + \frac{11n}{3} - 2$. Since we have an IRD, every moving player decreases her connection cost by at least 1. By Lemma 2, we have that the minimum decrease in social cost by any move is 2. Hence, at most $\frac{n^3}{6} - n^2 + \frac{11n}{6} - 1 \in \mathcal{O}(n^3)$ steps are needed to transform P_n into X_n. □

2.2 Best Response Dynamics on Trees

It is reasonable to assume, that players greedily try to decrease their connection cost most, whenever swapping an edge. In this section we analyse dynamics, where every move of a player is a best response move.

Since a best response is always an improving response, we have that every dynamic where every move is a best response must converge to a star for every tree T. We are left with the question of how fast best response dynamics converge. In the following, we analyse a specific best response dynamic, called *Max Cost Best Response Dynamic*(mcBRD), whose convergence speed almost matches the lower bound provided by Theorem 2. Hence, for best response dynamics we can significantly improve the upper bound of Theorem 3.

Definition 1. *The* Max Cost Best Response Dynamic *on a graph G is a dynamic, where in every step the active player having the highest connection cost is allowed to play a best response. If two or more players are active and have maximum connection cost, then one of them is chosen uniformly at random.*

In this section we show the following upper bound on the speed of convergence for the Max Cost Best Response Dynamic. Surprisingly, mcBRD behaves differently depending on whether the number of vertices in the tree is odd or even.

Theorem 4. *Let T be a tree having n vertices. The following holds:*

- *If n is even, then mcBRD(T) converges after at most* $\max\{0, n-3\}$ *steps and every player moves at most once.*
- *If n is odd, then at most* $\max\{0, n+\lfloor n/2\rfloor-5\}$ *steps are needed and every player moves at most twice.*

In order to prove Theorem 4, we first show some useful properties of the convergence process induced by the mcBRD-rule.

We begin with characterizing a player's best response on a tree. Here, the notion of a *center-vertex* is crucial.

Definition 2. *A center-vertex of a graph G is a vertex x, which satisfies*

$$x \in \arg\min_{v \in G} c(v) \ .$$

Lemma 4. *Let v be an arbitrary vertex of a tree T and let $F = T - v = \bigcup_{j=1}^{l} T_l$, where the trees T_j are connected components in the forest F. Let $u_1, \ldots, u_l$ be the neighbors of v in T, where u_j is a vertex of T_j for all $1 \le j \le l$. Let w_j be a center-vertex of the tree T_j. The best response of v in T is the edge-swap vu_j to vw_j, where $j \in \arg\max_j \{c_{T_j}(u_j) - c_{T_j}(w_j)\}$.*

Proof. Let T' be the tree obtained after player v's swap. Observe, that if player v removes the edge vu_i for some $i \in \{1, \ldots, l\}$, then it must by replaced with an edge vx_i, where $x_i \in T_i$, since otherwise T' would be disconnected. Thus, if the edge vu_i is removed, then player v has to choose which of the vertices of T_i to connect to. By Lemma 1, player v's change in connection cost is $\Delta(v) = c_{T_i}(u_i) - c_{T_i}(x_i)$, if the edge vu_i is removed and vx_i is build. Since v's best response yields the largest decrease in connection cost, it follows that x_i must be chosen such that $c_{T_i}(x_i) \le c_{T_i}(y_i)$ holds for all vertices $y_i \in T_i$. Thus, x_i must be a center-vertex of T_i. Player v can swap only one edge. Hence, v's best response is to connect to a center-vertex x_j of a tree T_j, which maximizes $c_{T_j}(u_j) - c_{T_j}(x_j)$. □

The next Lemma provides a very useful property of neighbors in a tree.

Lemma 5. *Let u and w be neighbors in a tree T. Let T_u and T_w denote the tree T rooted at vertex u and w, respectively. Let U be the set of vertices in the subtree rooted at u in T_w. Analogously, let W be the set of vertices in the subtree rooted at w in T_u. Then we have $c(u) \le c(w) \iff |U| \ge |W|$ and $c(u) < c(w) \iff |U| > |W|$.*

Proof. Let u,w,U,W,T_u and T_w be defined as in the Lemma. Let $d(a,b)$ be the length of the shortest path from a to b in T. Then, since T is a tree, we have

$$\begin{aligned}
&c(u) \le c(w)\\
&\qquad \iff \sum_{x \in U} d(u,x) + \sum_{x \in W} (1 + d(w,x)) \le \sum_{x \in W} d(w,x) + \sum_{x \in U} (1 + d(u,x))\\
&\qquad \iff \sum_{x \in W} (1 + d(w,x)) - \sum_{x \in W} d(w,x) \le \sum_{x \in U} (1 + d(u,x)) - \sum_{x \in U} d(u,x)\\
&\qquad \iff |W| \le |U|.
\end{aligned}$$

If the inequality of the connection costs is strict, then the proof is similar. □

We can use Lemma 5, to show an important property of the mcBRD-process.

Lemma 6. *Let T be a tree. Every player who moves in a step of mcBRD(T) must be a leaf.*

Proof of Lemma 6. In every step of mcBRD the player with the largest connection cost is allowed to move. Assume towards a contradiction, that an inner vertex u has the largest connection cost c^* in T. Let $x_1, \dots, x_l$ be the neighbors of u. By Lemma 5, we have that at most one of the neighbors of u can have the same connection cost c^*.

If u has no neighbor having connection cost c^*, then all neighbors must have strictly smaller connection cost than u. But Lemma 5 yields, that at most one neighbor of any vertex can have smaller connection cost. Since u has at least two neighbors, there must be a neighbor of u having larger connection cost and we have a contradiction.

If u has a neighbor w having connection cost c^*, then, by Lemma 5, all other neighbors of w must have smaller connection cost. If there is more than one such neighbor, then again, we have a contradiction. Thus, assume that there is exactly one such neighbor z. Let T_u,T_w and T_z denote the tree T rooted at vertex u, w and z, respectively. Let U and W_1 denote the subtree rooted at u and w, respectively, in tree T_z. Let W_2 denote the tree rooted at w in tree T_u and let Z denote the subtree rooted at z in tree T_w. By Lemma 5, we have that $|Z| > |W_1| \geq |U|$. Furthermore, we have $|W_2| > |Z|$. Hence, we have $|W_2| > |U|$ and thus, again by Lemma 5, it follows that $c(u) > c(w)$, which is a contradiction. $\square$

The following Lemma provides the key to analysing mcBRD. It shows, that at some point in the dynamic a certain behavior is "triggered", which forces the dynamic to converge quickly.

Lemma 7 (First Trigger Lemma)**.** *Let T be a tree. If the player who moves in step i of mcBRD(T) has a unique best response vertex w, then all players who move in a later step of mcBRD(T) will connect to vertex w.*

Proof. Let T be any tree. Let T^s denote the tree obtained after step s of mcBRD(T) and let v^s denote the player who moves in step s. Consider step i of mcBRD(T) and assume that player v^i has maximum connection cost in T^{i-1}. Let the edge-swap towards w be the unique best response of player v^i in this step. We show for any step $j \geq i+1$ of mcBRD(T) that if T^{j-1} is not a star, then player v^j will connect to vertex w, if player v^{j-1} did.

Consider the tree T^{j-2} in which player v^{j-1} has maximum connection cost. It follows by Lemma 6, that v^{j-1} must be a leaf of T^{j-2}. Since T^{j-1} is not a star, we have that T^{j-2} is not a star. Assume that the unique best response of player v^{j-1} is to connect to vertex w. By Lemma 4, it follows that w must be the unique center-vertex of the tree $T'' = T^{j-2} - \{v^{j-1}\}$. Let $x_1, \dots, x_k$ be the neighbors of w in T''. Let T_w be the tree T'' rooted at vertex w and let $X_1, \dots, X_k$ be the subtrees of T_w rooted at $x_1, \dots, x_k$, respectively. Using the fact that w is the unique center-vertex of T'' and Lemma 5, we obtain that $1 + \sum_{p \neq q} |X_p| > |X_q|$ for any $q \in \{1, \dots, k\}$. After his move, player v^{j-1} will end up as the $k+1$'th neighbor of w in the tree T^{j-1}. Since, by assumption, this tree is not stable, there is a leaf v^j of T^{j-1} who swaps an edge in step j. Clearly, we have $v^j \in X_r$ for some $r \in \{1, \dots, k\}$. Now consider $T''' = T^{j-1} - \{v^j\}$ and let $X'_1, \dots, X'_{k+1}$ be defined analogously as above for the tree T'''. We have that $|X_i| = |X'_i|$, for all $i \in \{1, \dots, k\} \setminus \{r\}$, and $|X_r| = |X'_r| + 1$. The new tree X_{k+1} contains only vertex v^{j-1} and thus compensates the loss of tree $|X'_r|$. Hence, we have $1 + \sum_{p \neq q} |X_p| > |X_q|$, for $q \in \{1, \dots, k+1\}$. By Lemma 5, this implies that w is the unique center-vertex of T''' and thus, player v^j will connect to w in step j. $\square$

Lemma 8. *In any tree T on n vertices, there are at most two center-vertices. If this is the case, then they are neighbors and n must be even.*

Proof. Assume that T contains exactly two vertices x_1 and x_2, which both have minimum connection cost c^* but there is no edge x_1x_2 in T. Since T is connected, there is a path P from x_1 to x_2 of length at least 2. Let z_1 and z_2 be the neighbors of x_1 and x_2, respectively, on path P. Let $T_{x_1}, T_{x_2}, T_{z_1}$ and T_{z_2} be the tree T rooted at vertex x_1, x_2, z_1 and z_2, respectively. Let X_1 be the set of vertices in the subtree of T_{z_1}, which is rooted at vertex x_1. Analogously, X_2 denotes the set of vertices in the subtree rooted at x_2 in T_{z_2}. We define Z_1 and Z_2 in the same way, for the trees T_{x_1} and T_{x_2}, respectively. We apply Lemma 5, which yields $c(x_1) \leq c(z_1) \iff |X_1| \geq |Z_1|$ and $c(x_2) \leq c(z_2) \iff |X_2| \geq |Z_2|$. Since T is a tree and x_1 and x_2 are non-neighbors, we have that $|Z_1| > |Y|$ and $|Z_2| > |X|$. Using Lemma 5, this implies $c(z_1) < c(y)$ and $c(z_2) < c(x)$, which is a contradiction. The only feasible solution is that $z_1 = x_2$ and $z_2 = x_1$ and thus, x_1 and x_2 have to be neighbors. Furthermore, we have that $c(x_1) = c_(x_2) = c^*$. By Lemma 5, it follows that $|X_1| = |X_2|$, which implies that n must be even.

If there are more than two vertices having minimum connection cost, then the above argumentation implies, that all of them must be pairwise neighbors. Since T is a tree, this is impossible. □

Now we are ready, to prove the first part of Theorem 4.

Proof of Theorem 4, Part 1. We show, that if the number of vertices in a tree T is even, then mcBRD needs at most $\max\{0, n-3\}$ to converge and every player moves at most once.

If T has two vertices, then it is already a star and no player will move in mcBRD(T). Thus, let T be a tree having at least $n \geq 4$ vertices, where n is even. By Lemma 6, we have that in every step of mcBRD(T) a leaf l of the current tree is allowed to move. By Lemma 4, we know that player l will connect to a center-vertex of $T' - l$, where T' is the tree before player l moves. Observe, that the tree $T' - l$ has an odd number of vertices. By Lemma 8, we have that any tree having an odd number of vertices must have a unique center vertex. It follows, that the leaf who moves in the first step of mcBRD(T) has a unique best response. Let this best response be the edge-swap towards vertex w. Lemma 7 implies, that all players who move in a later step of mcBRD(T), will connect to vertex w as well. Furthermore, again by Lemma 7, after the first step of mcBRD(T) it holds, that every vertex who is already connected to vertex w will never move again. Hence, every vertex moves at most once.

By Lemma 5, we have that w must be an inner vertex of T. Thus, w has at most $n-3$ non-neighbors, which implies that the dynamic mcBRD(T) will need at most $n-3$ steps to converge to a star having w as its center-vertex. □

The next Theorem shows a lower bound on the speed of convergence for mcBRD on trees having an odd number of vertices. Surprisingly, the behavior of the dynamic on such instances is much more complex. The lower bound for odd n is roughly 50% greater than the upper bound for even n. Furthermore, the following Theorem together with Theorem 2 implies, that the analysis of mcBRD is tight.

Theorem 5. *There is a family of trees having an odd number of vertices greater than* 5*, where mcBRD can take* $n + \lfloor n/2 \rfloor - 5$ *steps to converge. Furthermore, every player moves at most twice.*

Figure 2 shows an example of a tree which belongs to the above mentioned family of trees and it sketches the convergence process induced by mcBRD.

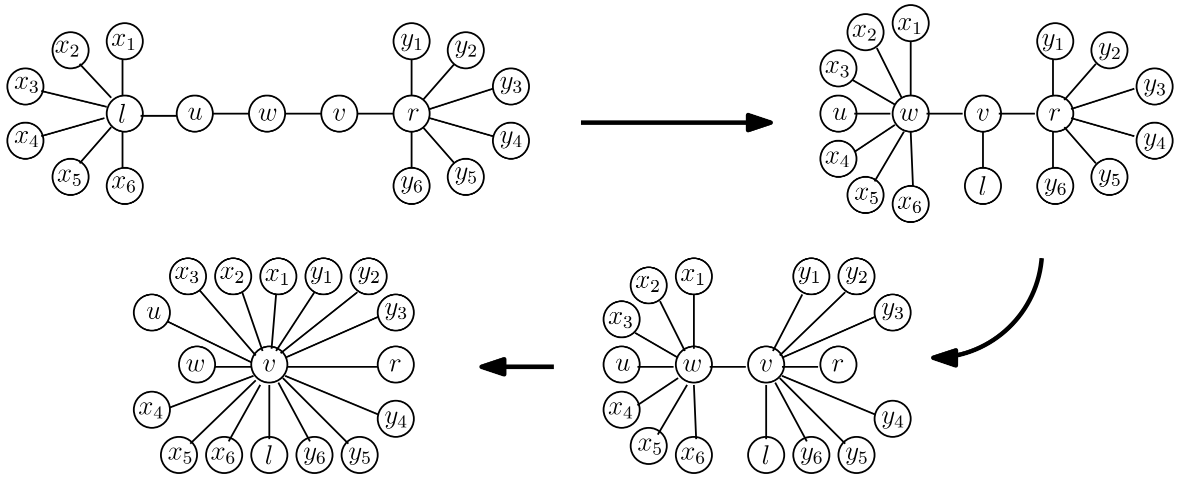

Figure 2: Example of a tree T having 17 vertices, where mcBRD(T) takes $n + \lfloor n/2 \rfloor - 5 = 20$ steps to converge. The vertices $x_1, \ldots, x_6, u$ move twice.

Proof of Theorem 5. A member of the family is constructed as follows: We start with a path having 5 vertices. Let the leaves of this path be l and r, let the center be w and let z be the vertex between w and r. Fix an even number $k \geq 2$ and connect $k/2$ vertices to l and r, respectively. Let $x_1, \ldots, x_{k/2}$ be the vertices having l as neighbor and $y_1, \ldots, y_{k/2}$ are the vertices connected to r. An example is shown top left in figure 2.

Let T be a tree constructed in the described way. During mcBRD(T) some players will have two best responses and the number of steps towards convergence depends on which best response is chosen. Note, that this implies, that a local decision has a global impact. We show that these choices can be made such that mcBRD(T) takes $n + \lfloor n/2 \rfloor - 5$ steps until a star emerges.

Since T is symmetric, all leaves are equal in the first step of mcBRD(T), that is, they all have the same connection cost. By Lemma 6, one of those leaves moves in the first step of mcBRD(T). Let T^i denote the tree, which is obtained after the i-th step of $mcBRD$(T).Let T^i_w and T^i_z denote the tree T^i rooted at w and z, respectively. Let W^i be the set of vertices in the tree rooted at w in T^i_z and let Z^i be the set of vertices in the tree rooted at z in T^i_w. Before the first step of mcBRD(T), we have that $|W^0| = |Z^0| + 1$.

The convergence proceeds in three stages:

Stage 1: Without loss of generality, assume that player x_1 moves first. Consider the tree $T^0 \setminus \{x_1\}$. This tree has two center-vertices, namely w and z. Hence, player x_1 has two best responses. Assume that x_1 chooses to connect to vertex w. Thus, we have that $W^0 = W^1$ and $Z^0 = Z^1$. In the next step, we have that vertices $x_2, \ldots, x_{k/2}$ have largest connection cost. Each of those players has two best responses, namely to connect to w or z. This is true by Lemma 5, since $|W^1| - 1 = |Z^1|$. Let x_2 move towards vertex w, which implies $W^1 = W^2$ and $Z^1 = Z^2$. This process iterates until all x_i-vertices are connected to w. Thus, we have that $W^0 = W^{k/2}$ and $Z^0 = Z^{k/2}$. In the following step, player l is allowed to move and, again, there are the two best responses w and z. Let l choose the connection towards z. This implies $W^{k/2+1} = W^0 \setminus \{l\}$ and $Z^{k/2+1} = Z^0 \cup \{l\}$. Observe, that $|W^{k/2+1}| < |Z^{k/2+1}|$ holds. The top right graph in figure 2 illustrates the result of the steps mentioned so far.

Stage 2: Now, all y_i-vertices have largest connection cost. Again, they have the two best responses w and z and we let them choose the vertex which is closer, that is z. After another $k/2$ steps a graph similar to the one bottom right in figure 2 is obtained. Furthermore, we have that $W^{k+1} = W^{k/2+1}$ and $Z^{k+1} = Z^{k/2+1}$. Let N_z^Z denote the number of neighbors of z in Z^{k+1}. Let N_w^W be defined analogously. Observe, that $N_z^Z > N_w^W$ holds.

Stage 3: In the next steps, all neighbors of w in W have highest connection cost, but this time, there is only one best response, which is connection to vertex z. The dynamic stops when all vertices are connected to z, which happens after N_w^W many steps.

Now we analyse the number steps of mcBRD(T). In Stage 1, there are $k/2$ steps, where a player x_i moves and one step where l swaps an edge. In Stage 2, all y_i-vertices move, which implies $k/2$ steps. In Stage 3, there are N_w^W steps. Since $|W^{k+1}| = |W^0| - 1$, we have that $N_w^W = |W_0| - 2 = \lceil n/2 \rceil - 2$. Observe, that in Stage 3 all x_i-vertices move a second time. There are no other players, which move twice.

Hence, in total there are $k/2 + 1 + k/2 + \lceil n/2 \rceil - 2 = k + \lfloor n/2 \rfloor$ steps. By construction, we have that $k = n - 5$ and thus mcBRD(T) takes $n + \lfloor n/2 \rfloor - 5$ steps to converge and every player moves at most twice. □

For proving the second part of Theorem 4 we need two additional properties of the mcBRD-process.

Lemma 9 (Second Trigger Lemma)**.** *Let T be an unstable tree having n vertices. If after any step i in mcBRD(T) a vertex w of T^i has degree $\lceil n/2 \rceil$, then this vertex will be the unique best response to connect to for all players moving in a later step of mcBRD(T).*

Proof. Let T be any unstable tree having n vertices. Observe, that whenever a player moves in mcBRD(T), the degrees of exactly two vertices change by an amount of 1. Let step i be the first step of mcBRD(T) after which a vertex w having degree $\lceil n/2 \rceil$ occurs. Using Lemma 7, it suffices to show, that if vertex w has degree $\lceil n/2 \rceil$ after step i of mcBRD(T), then the player v^{i+1} who moves in the $i+1$'th step will have w as its unique best response vertex. To prove this, we show, that v^{i+1} cannot be a neighbor of w in T^i. By Lemma 5, this implies that w is the unique best response vertex of v^{i+1}.

Assume towards a contradiction, that player x is a neighbor of w in T^i and that player x moves in the $i+1$'th step of mcBRD(T). Thus, player x must be a leaf an can swap an edge to decrease her connection cost. This implies that w is not a vertex having minimum connection cost in $T' = T^i - x$. Observe, that vertex w has degree $\lfloor n/2 \rfloor$ in T'. Let u be a vertex having minimum connection cost in T'. Let T'_w denote the tree T' rooted at w. We have, that u lies in a subtree rooted at some neighbor y of w in T'_w. Let T'_y denote the tree T' rooted at y. Let W be the set of vertices in the subtree rooted at w in T'_y and let Y denote the set of vertices in the subtree rooted at y in T'_w. Since, w has degree $\lfloor n/2 \rfloor$, we have that

$|W| \geq \lfloor n/2 \rfloor$. Thus, $|Y| \leq |T'| - |W| = \lfloor n/2 \rfloor$. By Lemma 5, it follows, that w cannot have higher connection cost than u and we have a contradiction. □

Lemma 10. *Let T be an unstable tree having an odd number of vertices. Only vertices which are best responses of the player who moves in the first step of mcBRD(T) will be best responses in any step of mcBRD(T).*

Proof. Let T be an unstable tree having n vertices, where n is odd. Let v be the player who moves in the first step of mcBRD(T). Let T^i denote the resulting tree after step i of mcBRD(T).

If player v has a unique best response w, then, by Lemma 7, vertex w will be the unique best response of any player who moves in a later step of mcBRD(T).

The only case left is the one where player v has two best responses u and w. We show that the set of best responses of any player y who moves in step $j \geq 2$ is a subset of the set of best responses of the player x who moved in step $j-1$. Settling this, implies the Lemma.

If in step $j-1$ player x has a unique best response, then, by Lemma 7, the claim is true. Thus, suppose that in step $j-1$ player x has two best responses p and q. Let T^{j-2} be the tree T after step $j-2$. Let T_p^{j-2} and T_q^{j-2} be the tree T^{j-2} rooted at p and q, respectively. Let P^{j-2} be the subtree rooted at p in T_q^{j-2} and let Q^{j-2} be the subtree rooted at q in T_p^{j-2}. Suppose that x is a leaf of P^{j-2}. Since both p and q are best responses for x we have, by Lemma 5, that $|P^{j-2} - x| = |Q^{j-2}|$. In step $j-1$ player x will connect either to p or to q. Suppose x connects to p. Now consider the moving player y in step j. Define T^{j-1}, P^{j-1} and Q^{j-1} analogous to the respective trees after step $j-2$. There are two cases:

If y is a leaf of P^{j-1}, then $|P^{j-1} - y| = |Q^{j-1}|$ and thus, by Lemma 4 and Lemma 5, player y has p and q as its best responses.

If y is a leaf of Q^{j-1}, then we claim that vertex p is the unique vertex having minimum connection cost in $T^{j-1} - y$ and thus vertex p must be player y's unique best response. To prove the claim, it suffices to show that every neighbor of p in $T^{j-1} - y$ has higher connection cost than p itself. Since $|P^{j-1}| > |Q^{j-1} - y|$, this holds, by Lemma 5, for vertex q. Furthermore it is trivially true for x, which is a leaf connected to p. Let z be any other neighbor of p in $T^{j-1} - y$. In tree $T^{j-2} - x$, by assumption, p and q are the vertices having minimum connection cost. This implies, that player z must have higher connection costs than player p in $T^{j-2} - x$. Since player x, who is missing in $T^{j-2} - x$ connects in step $j-1$ to vertex p, this difference increases further, which settles the claim.

The case where x connects in step $j-1$ to vertex q and both subcases where x is a leaf of Q^{j-2} are analogous. □

Finally, we have set the stage to prove the second part of Theorem 4.

Theorem 4, Part 2. We show that if a tree T has an odd number of vertices, then mcBRD(T) takes at most $\max\{0, n + \lfloor n/2 \rfloor - 5\}$ steps to converge and every player moves at most twice.

If $n = 5$, then the worst case instance is a path and thus the convergence takes at most 2 steps. Hence, we assume for the following that $n \geq 7$. Observe, that there are two events that force the dynamic to converge: Let E_1 be the event, where for the first time in the convergence process a vertex w becomes the unique best response of a moving player. Let E_2 be the event, where for the first time a vertex w has degree $\lceil n/2 \rceil$.

If event E_1 occurs in step j, then, by Lemma 7, all non-neighbors of the vertex w will connect to w in the subsequent steps of mcBRD(T). Thus, mcBRD(T) will converge in at most $j + |V \setminus \Gamma(w)|$ steps, where $\Gamma(w)$ is the closed neighborhood of w. If event E_2 occurs in step j, then, by Lemma 9, all non-neighbors of w will connect to w in the subsequent steps. Thus, in this case $j + \lfloor n/2 \rfloor - 1$ steps are needed for mcBRD(T) to converge.

Let T be any tree and v be the first player to move and assume that v has two best responses p and q, since otherwise the dynamic will converge in at most $n - 3$ steps. By Lemma 10, we have that in any step of mcBRD(T) a player will connect either to p or to q. Let $t_1(T)$ denote the number of steps until event E_1 is the first event to occur in mcBRD(T). Analogously, let $t_2(T)$ denote the number of steps until E_2 is the first occurring event. Let $r_1(T)$ denote the number of steps needed for convergence after event E_1. Hence, the maximum number of steps needed until mcBRD(T) converges is

$$t(T) = \max\{t_1(T) + r_1(T), t_2(T) + \lfloor n/2 \rfloor - 1\} \ .$$

We claim, that $t_1(T) + r_1(T) \leq n + \lfloor n/2 \rfloor - 5$. Observe, that $r_1(T) \leq n - 3$, since the vertex that becomes the center of the star must be an inner vertex of T and, thus, can have at most $n - 3$ non-neighbors. Furthermore, if $t_1(T) \leq \lfloor n/2 \rfloor - 2$, then the claim is true. Now let $t_1(T) > \lfloor n/2 \rfloor - 2$. Note, that both p and q must be inner vertices of T. Thus, they have at least degree 2. Since event E_2 did not occur in the first $t_1(T)$ steps of mcBRD(T) we have that not all players who moved within the first $t_1(T)$ steps can be connected to p. Thus, at least $x = t_1(T) - (\lfloor n/2 \rfloor - 2)$ players have connected to q. This yields $t_1(T) + r_1(T) \leq t_1(T) + n - 3 - x \leq n + \lfloor n/2 \rfloor - 5$. On the other hand, since all players move either to p or q and both p and q have degree at least 2, it follows that $t_2(T) \leq 2(\lfloor n/2 \rfloor - 2)$. Hence, $t_2(T) + \lfloor n/2 \rfloor - 1 \leq n + \lfloor n/2 \rfloor - 5$.

Observe, that any player x who is a neighbor of either p or q will not move again until event E_1 or E_2 happens. This holds because every leaf, which is not a neighbor of p or q must have higher connection cost than x and will therefore move before x. Thus, every player moves at most twice. □

2.3 Computing a Best Response on Trees

Observe, that Lemma 4 directly yields an algorithm for computing a best response move of a player v: Compute the connection-costs of all other vertices in $T - v$ within their respective connected component to find a center-vertex for every component. Then choose the center-vertex, which yields the greatest cost decrease for player v. Clearly, the connection-cost of a player can be obtained using a BFS-computation. However the above naive approach of computing a center-vertex yields an algorithm with running time quadratic in n, since $\Omega(n)$ BFS-computations can occur. The following Lemma shows, that a center-vertex can be computed in linear time, which is clearly optimal. The algorithm crucially uses the structural property provided by Lemma 5.

Lemma 11. *Let T be a tree having n vertices. A center-vertex of T and its connection-cost can be computed in $\mathcal{O}(n)$ time.*

Proof. We give a linear time algorithm, which computes a center-vertex of T and its connection-cost. Let L be the set of leaves of T. Clearly, L can be computed in $\mathcal{O}(n)$ steps by inspecting every vertex.

Given T and L, the algorithm proceeds in two stages:

1. The algorithm computes for every vertex v of T two values n_v and c_v. This is done in reverse BFS-order: We define n_v to be the number of vertices in the already processed subtree T_v containing v and c_v to v's connection-cost to all vertices in T_v. For every leaf $l \in L$ we set $n_l := 1$ and $c_l := 0$. Let i be an inner vertex and assume that we have already processed all but one neighbor of i. Let $a_1, \ldots, a_s$ denote these neighbors. We set $n_i := 1 + n_{a_1} + \cdots + n_{a_s}$ and $c_i := n_i - 1 + c_{a_1} + \cdots + c_{a_s}$. By breaking ties arbitrarily, this computation terminates at a root-vertex r, for which all neighbors are already processed. Let $b_1, \ldots, b_q$ denote these neighbors. We set $n_r := n$ and $c_r := n - 1 + c_{b_1} + \cdots + c_{b_q}$.

2. Starting from vertex r, the algorithm performs a local search for the center-vertex with the help of Lemma 5. For all neighbors $b_i \in \{b_1, \ldots, b_q\}$ of r, the algorithm checks if $n_{b_i} > n_r - n_{b_i}$. Since T is a tree, this can hold for at most one neighbor x. In this case, x will be considered as new root-vertex. Let $c_1, \ldots, c_s, r$ be the neighbors of x. By setting $n_x := n$ and $c_x := n - 1 + c_1 + \cdots + c_s + c_r - c_x$ we arrive at the same situation as before and we now check for all neighbors $c_j \neq r$ if $n_{c_j} > n_x - n_{c_j}$ holds and proceed as above. Once no neighbor of the current root-vertex satisfies the above condition, the algorithm terminates and the current root-vertex is the desired center-vertex.

The correctness of the above algorithm follows by Lemma 5. Step 1 clearly takes time $\mathcal{O}(n)$. Step 2 takes linear time as well, since the condition is checked exactly once for every edge towards a neighbor and there are only $n-1$ edges in T. $\square$

Theorem 6. *If $p \geq 1$ edges can be swapped at a time, then the best response of a player v can be computed in linear time if G is a tree.*

Proof. Let v be a degree d vertex in G. Clearly, player v can swap at most $\min\{p, d\}$ edges and the task is to determine the $k \leq \min\{p, d\}$ edge swaps that decrease player v's connection cost most. Let $v_1, \ldots, v_d$ denote the neighbors of v in G.

Let $F = T_1 \cup T_2 \cup \cdots \cup T_d$ be the forest obtained by deleting vertex v. Let $c_i = |T_i| + \sum_{w \in T_i} d_G(v_i, w)$ denote player v's connection cost to vertices in T_i, where $1 \leq i \leq d$. And let $c_i(v_i) = c_i - |T_i|$ be vertex v_i's connection cost in tree T_i. By Lemma 4 we have that every swap in player v's best response is of the form (v_i, w_i), where w_i is a center-vertex of T_i. Let $c_i(w_i)$ denote vertex w_i's connection cost in tree T_i. Let $z_i = c_i(w_i) - c_i(v_i)$ denote player v's change in costs after the swap (v_i, w_i). Clearly, if $z_i \geq 0$ then the swap (v_i, w_i) will not be part of v's best response, since it does not yield a cost reduction. If $z_i < 0$, then we call the corresponding swap (v_i, w_i) *attractive.* If there are l attractive swaps for player v, then we have that v's best response will consist of the $\min\{k, l\}$ attractive swaps having the smallest z_i values.

Thus, computing player v's best response reduces to finding a center-vertex in each tree T_i and to computing the corresponding $c_i(w_i)$-value. By Lemma 11 we have that both tasks can be done in time linear to the number of vertices in each T_i. Observe that all negative z_i-values are in the range $[-n^2, 0]$. Hence we can use radixsort, to find the $\min\{k, l\}$ attractive swaps having the smallest z_i values in linear time. $\square$

3 Playing on General Graphs

3.1 Best Response Dynamics on General Graphs

Definition 3. *A cycle $x_1, \ldots, x_l$ is a* best response cycle, *if $x_1 = x_l$ and each x_i is a pure strategy profile in the* SUM BASIC NETWORK CREATION GAME *and for all $1 \leq k \leq l-1$ there is a player p_k whose best response move transforms the profile x_k into x_{k+1}.*

Theorem 7. *The* SUM BASIC NETWORK CREATION GAME *allows best response cycles.*

Proof. Consider the graph G depicted left in Figure 3 and let x_1 denote the corresponding strategy profile. Player a can decrease its connection cost and one of its best responses is to

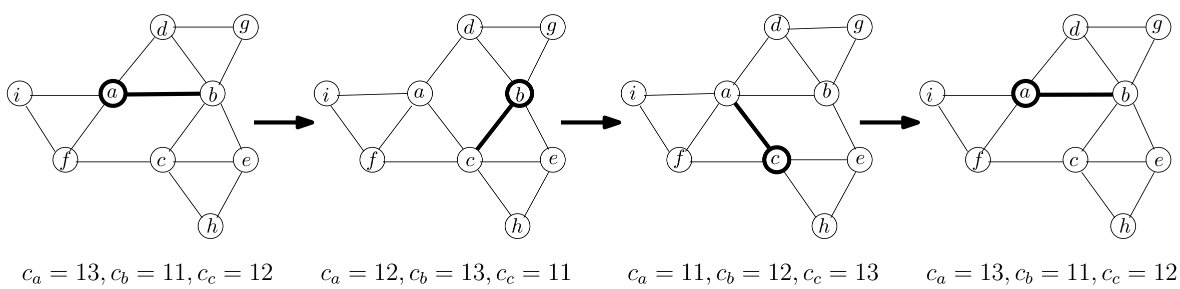

Figure 3: Example of a graph, where the SUM BASIC NETWORK CREATION GAME allows a best response cycle. The steps of the cycle are shown.

swap edge ab with edge ac. This leads to the second graph depicted in Figure 3. Call the corresponding strategy profile x_2. Now, player b has the swap bc to ba as its best response, which leads to the third graph depicted in the illustration, with x_3 as its strategy profile. Finally, player c can perform the swap ca to cb as its best response, which leads to profile $x_4 = x_1$. Thus, x_1, x_2, x_3, x_4 is a best response cycle in the SUM BASIC NETWORK CREATION GAME on graph G. $\square$

Voorneveld [14] introduced the class of *best-response potential games*, which is a super-class of ordinal potential games. Furthermore he proves, that if the strategy space is countable, then a strategic game is a best-response potential game if and only if there is no best response cycle. This implies the following Corollary.

Corollary 2. *There cannot exist an ordinal potential function for the* SUM BASIC NETWORK CREATION GAME *on graphs containing cycles.*

3.2 Computing a Best Response in General Graphs

Given an undirected, connected graph G, then the best response for player v can be computed in $\mathcal{O}(n^2)$ time, since $|S_G(v)| < n^2$ and we can try all pure strategies to find the best one.

Quite surprisingly, computing the best response is hard if we allow a player to swap $p > 1$ edges at a time.

Theorem 8. *If players are allowed to swap $p > 1$ edges at a time, then computing the best response is NP-hard even if G is planar and has maximum degree 3.*

Proof. We reduce from the p-Median problem [7], which is defined as follows: Given a connected undirected graph $G = (V, E)$ with non-negative weights $w(v)$ for every vertex $v \in V$ and non-negative lengths $l(e)$ for every edge $e \in E$ and given an integer $p > 1$. The task is to find a subset $V' \subseteq V$ with $|V'| = p$ such that $\sum_{v \in V} \min_{u \in V'} w(v) d_G(v, u)$ is minimized. Here $d_G(v, u)$ denotes the length of the shortest path from v to u in G.

The p-Median problem is known [7] to be NP-hard for $p > 1$ even if all vertex weights and edge lengths are one, G is planar and has maximum degree 3.

The reduction works as follows: Let G be an instance of the p-Median problem, where G is planar, has unit vertex weights and edge lengths and has maximum degree 3. We introduce a new vertex v^* and connect v^* with p new edges to G which induces the graph $G' = (V', E')$. Now, let v^* play a best response and let $X \subseteq V$ be the set of vertices incident to v^* after the best response move. We claim that X is the solution to the p-Median problem in G, which implies NP-hardness of computing the best response if p edges can be swapped at a time.

Clearly, we have $|X| = p$, since no best response of v^* will allow multiple edges connecting to the same vertex. By definition of a best response, we have that building edges to vertices in X minimizes the connection cost $c(v^*)$ of player v^*. Thus, we have

$$c(v^*) = \sum_{u \in V'} d_{G'}(v^*, u) = \sum_{u \in V} \left(1 + \min_{x \in X} d_G(x, u)\right) = |V| + \sum_{u \in V} \min_{x \in X} d_G(x, u),$$

which yields that $c(v^*)$ is minimized if and only if the set X minimizes $\sum_{u \in V} \min_{x \in X} d_G(x, u)$. ☐

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 algorithm*, 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] N. Baumann and S. Stiller. The price of anarchy of a network creation game with exponential payoff. In *SAGT*, pages 218–229, 2008.

[4] E. D. Demaine, M. Hajiaghayi, H. Mahini, and M. Zadimoghaddam. The price of anarchy in network creation games. In *Proceedings of the twenty-sixth annual ACM symposium on Principles of distributed computing*, PODC '07, pages 292–298, New York, NY, USA, 2007. ACM.

[5] 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.

[6] M. O. Jackson. A survey of models of network formation: Stability and efficiency. *Group Formation in Economics: Networks, Clubs and Coalitions*, 2003.

[7] O. Kariv and S. L. Hakimi. An algorithmic approach to network location problems. ii: The p-medians. *SIAM Journal on Applied Mathematics*, 37(3):pp. 539–560, 1979.

[8] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th annual conference on Theoretical aspects of computer science*, STACS'99, pages 404–413, Berlin, Heidelberg, 1999. Springer-Verlag.

[9] 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.

[10] H. Lin. On the price of anarchy of a network creation game. Class final project. 2003.

[11] 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.

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

[13] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.

[14] M. Voorneveld. Best-response potential games. *Economics Letters*, 66(3):289 – 295, 2000.