Basic Network Creation Games with Communication Interests*

Andreas Cord-Landwehr, Martina Hüllmann, Peter Kling, and Alexander Setzer

Heinz Nixdorf Institute & Department of Computer Science
University of Paderborn, Germany

Abstract. Network creation games model the creation and usage costs of networks formed by a set of selfish peers. Each peer has the ability to change the network in a limited way, e.g., by creating or deleting incident links. In doing so, a peer can reduce its individual communication cost. Typically, these costs are modeled by the maximum or average distance in the network. We introduce a generalized version of the *basic network creation game* (BNCG). In the BNCG (by Alon et al., SPAA 2010), each peer may replace one of its incident links by a link to an arbitrary peer. This is done in a selfish way in order to minimize either the maximum or average distance to all other peers. That is, each peer works towards a network structure that allows himself to communicate efficiently with *all* other peers. However, participants of large networks are seldom interested in all peers. Rather, they want to communicate efficiently with a small subset only. Our model incorporates these (communication) *interests* explicitly.

Given peers with interests and a communication network forming a tree, we prove several results on the structure and quality of equilibria in our model. We focus on the MAX-version, i.e., each node tries to minimize the maximum distance to nodes it is interested in, and give an upper bound of $\mathcal{O}(\sqrt{n})$ for the private costs in an equilibrium of n peers. Moreover, we give an equilibrium for a circular interest graph where a node has private cost $\Omega(\sqrt{n})$, showing that our bound is tight. This example can be extended such that we get a tight bound of $\Theta(\sqrt{n})$ for the price of anarchy. For the case of general networks we show the price of anarchy to be $\Theta(n)$. Additionally, we prove an interesting connection between a maximum independent set in the interest graph and the private costs of the peers.

1 Introduction

In a network creation game (NCG), several selfish players create a network by egoistic modifications of its edges. One of the most famous NCG models is due to Fabrikant et al. [7]. Their model intends to capture the dynamics in large communication and computer networks built by the individual participants (peers, players) in a selfish way: participants try to ensure a network structure supporting their own communication needs whilst limiting their individual investment into the network. Since the players do not (necessarily) cooperate, the resulting network structure may be suboptimal from

* This work was partially supported by the German Research Foundation (DFG) within the Collaborative Research Centre "On-The-Fly Computing" (SFB 901) and by the Graduate School on Applied Network Science (GSANS).

M. Serna (Ed.): SAGT 2012, LNCS 7615, pp. 72–83, 2012.

a global point of view. The analysis of the resulting structure and its comparison to a (socially) optimal structure is a central aspect in the analysis of network creation games.

In the original model by Fabrikant et al., players may buy (or create) a single edge for a certain (fixed) cost of $\alpha > 0$. Their goal is to improve the network structure with respect to their individual communication needs. There are typically two ways to formalize the corresponding communication cost of a single peer: the maximum distance or the average distance to all other peers in the network. We refer to the different variants by MAX-version and AVG-version. Alon et al. [2] introduce a slightly simpler model, called *basic network creation games* (BNCG), that drops the cost parameter α. Instead, they limit the possible ways in which peers may change the network by restricting them to edge swaps: a peer may only replace one of its incident edges with a new edge to an arbitrary node in the network. Since peers are assumed to be selfish, only edge swaps (including simultaneous swapping of several edges at once) that improve the private communication cost of the corresponding peer are considered. In a *swap equilibrium*, no player can decrease its communication cost by an edge swap. This simpler variant of network creation games has the advantage of polynomially computable best responses of the players. Moreover, it still captures the inherent dynamic character and difficulty of communication networks formed by selfish participants, while avoiding the quite intricate dependence on the parameter α (see related work).

Our work generalizes the BNCG model of Alon et al. by introducing the concept of *interests*. In real communication networks, participants are typically only interested in a small subset of peers rather than the complete network. Thus, instead of trying to minimize the maximal or average distance to *all* other nodes, the individual players consider only the distances to nodes they are interested in. The main part of our analysis focuses on tree networks. Especially, we show that tree networks perform much better than general networks with respect to the price of anarchy. To avoid networks to become disconnected (note that in a BNCG peers want to communicate with all other peers and hence never disconnect the network), we restrict the peers to swaps that preserve connectivity. This restriction is valid from a practical point of view, where a lost network connectivity is to be avoided, since re-connecting a network causes high or even unpredictable costs. Moreover, if you consider that interests of the peers may change over time, it is also important for each single selfish peer to sustain connectivity.

Model and Notions. An instance of the *basic network creation game with interests (I-BNCG)* is given by a set of n players (peers, nodes) $V = \{v_1, v_2, \ldots, v_n\}$, an initial *connection graph* $G = (V,E)$, and an *interest graph* $G_I = (V,I)$. We use $I(v) := \{u \in V \mid \{v,u\} \in I\}$ to refer to the neighborhood of a player v in the interest graph and denote them as the *interests of* v. Both the connection graph and the interest graph are undirected. Thus, interests are always mutual. The connection graph represents the current communication network and can change during the course of the game. We consider only instances where the (initial) connection graph is a tree, whereas the interest graph G_I may be an arbitrary and not necessarily connected graph. Each player is assumed to have at least one interest. We study two different ways to formalize the private communication costs of nodes: the *MAX-version* and the *AVG-version*. In the first, the private cost $c(v) := \max\{d(v,u) \mid u \in I(v)\}$ of a node $v \in V$ is defined as the maximum distance from v to its interests. In the second, we define $c(v) := \sum_{u \in I(v)} \frac{d(v,u)}{|I(v)|}$

as the average distance to its interests. Here, $d(v,u)$ denotes the (shortest path) distance between u and v in the connection graph.

To improve its private cost, a player u may perform *edge swaps* in the connection graph: replace an incident edge $\{u,v\}$ with a new edge $\{u,w\}$ to an arbitrary player $w \in V$, written as $u : [v \to w]$. We refer to a single as well as to a series of simultaneously executed edge swaps of a player u as an *improving step* if u's private cost decreases. A player is only allowed to perform an improving step if the connection graph stays connected. If no player can perform an improving step, we say the connection graph is in a *MAX-equilibrium* or *AVG-equilibrium*, respectively. See Figure 1 for an example.

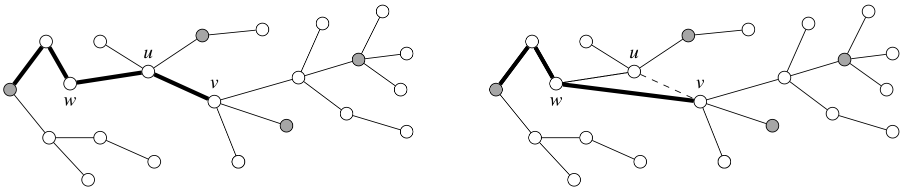

(a) Connection graph with v having $c(v) = 4$. (b) After swap $v : [u \to w]$ with $c(v) = 3$.

Fig. 1. MAX-version example of an improving swap performed by v. The gray nodes denote $I(v)$, the thick lines indicate the largest distance to a node in $I(v)$.

The quality of a connection graph G is measured by the *social cost* $c(G) = \sum_{v \in V} c(v)$ as the sum over all private costs. Our goal is to analyze the structure and social cost of worst case swap equilibria and compare them with a general optimal solution. As usual in algorithmic game theory, we use the ratio of these two values (*price of anarchy*, see Section 2.3) for this comparison [9]. Note that if the interest graph is the complete graph, I-BNCG coincides with the BNCG by Alon et al. [2].

Related Work. Network creation games combine two crucial aspects of modern communication networks: network design and routing. Many such networks consist of autonomous peers and have a highly dynamic character. Thus, it seems natural to use a game theoretic approach to study their evolution and behavior. Given the possibility to change the network structure (buy bandwidth, create new links, etc.), peers typically try to improve their individual communication experience. The question whether this selfish behavior results in an overall good network structure constitutes the central question of the study of network creation games as introduced by Fabrikant et al. [7]. In their model, the authors use a fixed cost parameter $\alpha > 0$ representing the cost of buying a single edge. The players (nodes) in such a game can buy edges to decrease their local communication cost (the average distance to all other nodes in the network). Each player's objective is to minimize the sum of its individual communication cost and the money spent on buying edges. In their seminal work, the authors proved (among other things) an upper bound of $\mathcal{O}\left(\sqrt{\alpha}\right)$ on the price of anarchy (PoA) in the case of $\alpha < n^2$. Albers et al. [1] proved a constant PoA for $\alpha \in \mathcal{O}\left(\sqrt{n}\right)$ and the first sublinear worst case bound of $\mathcal{O}\left(n^{1/3}\right)$ for general α. Demaine et al. [6] were the first to prove an $\mathcal{O}\left(n^{\varepsilon}\right)$ bound for α in the range of $\Omega\left(n\right)$ and $\mathrm{o}\left(n \lg n\right)$. Furthermore, Demaine et al. introduced a new cost measure for the private cost, causing the individual nodes to consider their maximum distance to all remaining nodes instead of the average distance. For this variant they showed that the PoA is at most 2 for $\alpha \geq n$, $\mathcal{O}\left(\min\{4^{\sqrt{\lg n}}, (n/\alpha)^{1/3}\}\right)$ for $2\sqrt{\lg n} \leq \alpha \leq n$, and $\mathcal{O}\left(n^{2/\alpha}\right)$ for $\alpha < 2\sqrt{\lg n}$. Recently, Mihalák and Schlegel [11] could prove that for $\alpha > 273 \cdot n$ all equilibria in the AVG-version are trees (and thus the PoA is constant). The same result applies to the MAX-version if $\alpha > 129$.

While network creation games, as defined by Fabrikant et al., and their variants seem to capture the dynamics and evolution caused by the selfish behavior of peers in an accurate way, there is a major drawback of these models: most of them compute the private communication cost of the peers over the *complete* network. Given the immense size of such communication networks, this seems rather unrealistic. Typically, participants want to communicate only in small groups, with a small subset of participants they know. To the best of our knowledge, the only other work taking this into account is due to Halevi and Mansour [8]. They introduce a concept similar to our interests (see model description). For the objective of minimizing the average distance of a peer to its interests, they proved the existence of pure nash equilibria for $\alpha \leq 1$ and $\alpha \geq 2$ and upper bounded the PoA by $\mathcal{O}(\sqrt{n})$ for general α. In the case of constant α or d (where d denotes the average degree in the interest graph) or $\alpha \in \mathcal{O}(nd)$, Halevi and Mansour upper bounded the PoA by a constant. Furthermore, the authors provided a family of problem instances for which the PoA is lower bounded by $\Omega(\log n / \log\log n)$.

Note that all these results largely depend on the cost parameter α. Moreover, as has been stated in [7], computing a player's best response for these models is NP-hard. This observation leads to a new, simplified formalization by Alon et al. [2], trying to capture the crux of the problem without the burden of this additional parameter. They introduce basic network creation games (BNCG), where players no longer have to pay for edges. Instead, possible actions are limited to *improving edge swaps*: replacing a single, incident edge by an edge to some arbitrary node which improves the node's private cost. Other than that, the general problem stays mostly untouched, especially the private cost function (average distance or maximum distance to all other nodes). Best responses in this game turn out to be polynomially computable. Restricting the initial network to trees, they show that the only equilibrium in the AVG-version is a star graph. Without restrictions, all swap equilibria are proven to have a diameter of $2^{\mathcal{O}(\sqrt{\lg n})}$. For the MAX-version, the authors prove a maximum diameter of three if the resulting equilibrium is a tree. Furthermore, the authors construct an equilibrium of diameter $\Theta(\sqrt{n})$. Our model is a direct generalization of these BNCGs, introducing the concept of interests. Up to now, the only other work on BNCGs we are aware of is due to Lenzner [10]. He studies the dynamics of the AVG-version of BNCGs and proves for the case of tree connection graphs a convergence to pure equilibria. Moreover, he proves that any sequence of improving edge swaps converges in at most $\mathcal{O}(n^3)$ steps to a star equilibrium.

Our Contribution. We introduce a generalized class of the BNCG by taking the different interests of individual peers into account. We analyze the structure and quality for the case that the initial connection graph is a tree. For the MAX-version, we derive a worst case upper bound of $\mathcal{O}(\sqrt{n})$ for the private costs of the individual players in an equilibrium. Thereto, we introduce and apply a novel combinatorial technique that captures the structural properties of our equilibria (see MAX-arrangement, Definition 1). Furthermore, for interest graphs with a maximum independent set of size $M \leq \sqrt{n}$ (e.g., the clique graph with $M = 1$), we can improve the private cost upper bound to $\mathcal{O}(M)$. Using a circular interest graph, we construct an equilibrium with a player having private cost $\Omega(\sqrt{n})$, showing that our bound is tight. By extending this construction, we are able to prove a tight bound of $\Theta(\sqrt{n})$ on the price of anarchy (ratio between the social cost of a worst case equilibrium and an optimum [9]). Using a star-like connection graph, we show the existence of a MAX-equilibrium with small social cost, yielding a *price of stability* (ratio between the social cost of a best case equilibrium and an optimum [3,4]) of at most two for an I-BNCG. For the case of an I-BNCG featuring a general connection graph (instead of a tree), we show that the price of anarchy is $\Theta(n)$.

2 Quality of Equilibria in I-BNCGs

In this section we show a tight worst case private cost upper bound of $\Theta\left(\sqrt{n}\right)$ for every MAX-equilibrium on trees as well as the same bound for the price of anarchy. The price of stability we can limit to be at most two. For general connection graphs we provide an instance with social cost $\Omega\left(n^2\right)$, yielding a price of anarchy of $\Theta(n)$.

2.1 Private Cost Upper Bound

In the following we prove the private cost upper bound as stated below:

Theorem 1. *Let I be a set of interests and $G = (V,E)$ a corresponding tree in a MAX-equilibrium, $n := |V|$. Then, for all $v \in V$ we have $c(v) \in \mathcal{O}(\sqrt{n})$.*

Outline of the proof: We consider a tree network in a MAX-equilibrium and take any node with maximal private cost. Starting with this node, we define a special node sequence, called MAX-arrangement, that will contribute the following properties: each two successive nodes are interested in each other and every node is "far away" from all previous nodes of the sequence. We will prove that such a sequence necessarily exists and that its length is proportional to the private cost of the starting node.

In detail, we prove with Lemma 3 and Lemma 4 that a shortest path traversal of a MAX-arrangement in the connection graph uses each edge at most twice and by this limits its length. Lemma 5 constructively shows that given a node with maximal private cost, there always exists a MAX-arrangement starting with this node and ending with a node with a private cost of 3. Lemma 2 gives us that the number of nodes in this MAX-arrangement is proportional to the maximal private cost of the first node. Comparing the maximum private cost of a node with the length of a shortest path traversal of any corresponding MAX-arrangement gives us the upper bound.

Remark 1. Note that in a MAX-equilibrium, each node v with $|I(v)| = 1$ has $c(v) = 1$. Hence, for a node v' with $c(v') > 1$, it holds $|I(v')| > 1$.

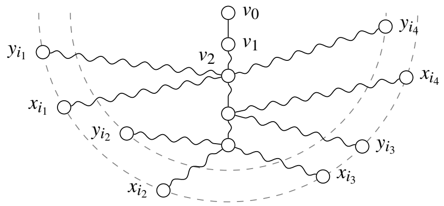

Fig. 2. Visualization for Lemma 1. v_0 can perform improving swap $v_0 : [v_1 \to v_2]$.

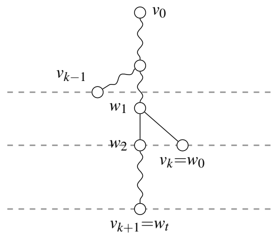

Fig. 3. Visualization for proof of Lemma 4. Edge $\{w_0, w_1\}$ is used twice.

Lemma 1 (T-configuration). *Let I be a set of interests, $G = (V, E)$ a corresponding tree in a MAX-equilibrium and $v \in V$ with $|I(v)| \geq 2$. Then there exist nodes $x, y \in I(v)$ such that $|d(x, v) - d(v, y)| \leq 1$ and v is connected by at most one edge to the shortest path from x to y and $c(v) = d(v, x)$.*

Proof. Let $v \in V$ with $|I(v)| \geq 2$ and $x \in I(v)$ with $d(v,x) = c(v)$. Assume that all $x' \in I(v) \setminus \{x\}$ are at distance $d(x', v) \leq c(v) - 2$ from v. Consider the shortest path $v \to v_1 \to v_2 \to \ldots \to x$ to x. In this case v can reduce its private cost by $v : [v_1 \to v_2]$ since this swap improves v's distance to x by 1 but increases the distances to every node in $I(v) \setminus \{x\}$ by at most 1. But this contradicts G being in a MAX-equilibrium.

We now consider all pairs $(x_i, y_i) \in I(v) \times I(v)$ for that hold $d(v, x_i) = c(v)$ and $d(v, y_i) \geq c(v) - 1$. Let us assume that v is connected to each shortest path from x_i to y_i by at least two edges that do not lie on that path. (See Figure 2 for a visualization.) Thus, v is not located on the shortest path from x_i to y_i for all i. This implies that in the graph $G \setminus \{v\}$ for each pair (x_i, y_i) there exists a connected component containing both nodes x_i, y_i. Since each two nodes at distance exactly $c(v)$ form such a pair, all nodes of $I(v)$ at distance exactly $c(v)$ must be located in the same connected component, which gives for every pair (x_i, y_i) that both nodes are contained in the same component. Hence, all nodes $x' \in I(v)$ at distance $d(x', v) \geq c(v) - 1$ from v are in the same connected component and by the two edges distance constraint, there must be a path $v \to v_1 \to v_2$ that is a subpath of every path from v to every node x_i and y_i. Hence, v can perform the improving swap $v : [v_1 \to v_2]$ (cf. Figure 2). This swap decreases the distance to all nodes x_i, y_i by one and increases each distance to other nodes (i.e., nodes $w \in I(v)$ with $d(w, v) \leq c(v) - 2$) by at most one and hence contradicts G being in a MAX-equilibrium. □

Definition 1 (MAX-arrangement). *Let $v_0 \in V$ and $v_1 \in I(v_0)$ such that $d(v_0, v_1) = c(v_0)$. Consider a sequence of nodes $v_0, \ldots, v_m$ with $v_i \in I(v_{i-1})$, $i = 1, \ldots, m$, with private costs $c(v_i) > 3$ for $i = 0, \ldots, m-1$ and $c(v_m) = 3$. We call this sequence a* MAX-arrangement *if for all $i = 2, \ldots, m$ it holds (see Figure 4 for a visualization):*

$$v_i = \operatorname*{argmax}_{v_i \in I(v_{i-1})} \left\{ d(v_{i-2}, v_i) \;\middle|\; \begin{array}{l} v_{i-1} \text{ is connected by} \leq 1 \text{ edge to the} \\ \text{shortest path from } v_{i-2} \text{ to } v_i \end{array} \right\}$$

The key property of a MAX-arrangement is stated by the following two lemmas: consider a node v_i in a MAX-arrangement, then (1) its MAX-arrangement successor node

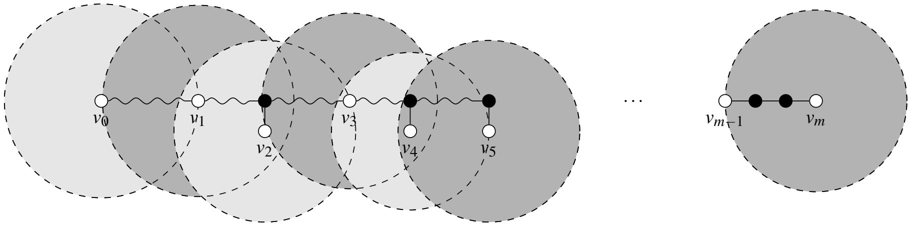

Fig. 4. Visualization of a MAX-arrangement. The radius of a circle around a node corresponds to the node's private cost. Curled lines denote shortest paths.

v_{i+1} cannot have a "much lower" private cost than v_i and (2) in the connection graph the shortest path from v_i to v_{i+1} can overlap by at most one edge with the shortest path to v_i's MAX-arrangement predecessor node.

Lemma 2. *For each two successive nodes v_i, v_{i+1} ($0 \leq i < m$) in a MAX-arrangement $v_0, \dots, v_m$ it holds $d(v_i, v_{i+1}) \geq c(v_i) - 1$ and hence $c(v_{i+1}) \geq c(v_i) - 1$.*

Proof. Consider a node v_i, $0 \leq i < m$, in the MAX-arrangement. Then by Lemma 1, there exist $x, y \in I(v_i)$ with $d(v_i, x) = c(v_i)$ and $c(v_i) \geq d(v_i, y) \geq c(v_i) - 1$ such that v_i is connected by at most one edge to the shortest path from x to y. At least one of these nodes is a valid candidate for the next MAX-arrangement node v_{i+1} (even if neither x nor y is v_{i+1}, we get a distance lower bound) and we get $d(v_i, v_{i+1}) \geq \min\{d(v_i, x), d(v_i, y)\} \geq c(v_i) - 1$. This gives, $c(v_{i+1}) \geq c(v_i) - 1$. □

Lemma 3 (Increasing Distance). *Let I be a set of interests and $G = (V, E)$ a corresponding tree in a MAX-equilibrium with $v_0, \dots, v_k$ a MAX-arrangement. Then the distances to v_0 are monotonically increasing, i.e., $d(v_0, v_i) \leq d(v_0, v_{i+1})$ for $i = 1, \dots, k-1$.*

Proof. By $c(v_1) \geq 3$ we get with Remark 1 that $|I(v_1)| \geq 2$. Hence by Lemma 1, there exists a node v_2, such that the paths v_1 to v_0 and v_1 to v_2 overlap by at most one edge. By construction of the MAX-arrangement the distance $d(v_0, v_2)$ is maximal among all distances from v_0 to nodes $v \in I(v_1)$ and hence we get $d(v_0, v_1) \leq d(v_0, v_2)$.

Assume that there is a node v_i with smallest index $i \geq 2$ in the MAX-arrangement for which the claim does not hold. That is $d(v_0, v_{i-1}) \leq d(v_0, v_i) > d(v_0, v_{i+1})$. Denote by x the most distant node from v_0 that is on all shortest paths from v_0 to v_{i-1}, v_0 to v_i, and v_0 to v_{i+1}. (Such a node x exists since especially v_0 fulfills the restrictions.) By the choice of i and since all these paths cross node x, we get:

$$d(x, v_{i-1}) \leq d(x, v_i) > d(x, v_{i+1}) \tag{1}$$

By definition of the MAX-arrangement, v_i is connected by at most one edge to the shortest path from v_{i-1} to v_{i+1}. Hence, x must be a node on the path from v_{i-1} to v_{i+1}. First note that x cannot be v_i or a neighbor of v_i, since for those cases with (1) we get $d(x, v_{i+1}) < d(x, v_i) \leq 1$. Further, x must lie on the shortest path from v_{i-1} to v_i, since otherwise x would lie on the shortest path from v_i to v_{i+1} which implies by $d(v_{i-1}, v_i) \geq 3$ that $d(x, v_i) < d(x, v_{i-1})$. This gives $d(x, v_i) \leq d(x, v_{i+1})$ and is a contradiction. □

Lemma 4. *Let I be a set of interests and $G = (V,E)$ a corresponding tree in a MAX-equilibrium. Consider a MAX-arrangement $v_0,\ldots,v_m$. Then, no edge in E is used more than two times by the shortest path visiting the nodes $v_0,\ldots,v_m$ in the given order.*

Proof. We label the nodes of G by their distances to v_0. This is, for every $v \in V$ we define a *level* by $\text{level}(v) := d(v_0,v)$. We consider an arbitrary node v_k with $k \in \{1,\ldots,m-1\}$ and the corresponding shortest path $v_k =: w_0 \to w_1 \to \ldots \to w_t := v_{k+1}$ to node v_{k+1}. By definition, v_k is connected by at most one edge to the shortest path from v_{k-1} to v_{k+1} (see Figure 3). By Lemma 3 we have $\text{level}(v_{k-1}) \leq \text{level}(v_k) \leq \text{level}(v_{k+1})$. Hence, for $i = 2,\ldots,t-1$ we get $\text{level}(w_i) < \text{level}(w_{i+1})$. This is, at most one edge (explicitly edge $\{w_0,w_1\}$) of the shortest path from v_0 to v_k is used a second time by the shortest path traversal from v_k to v_{k+1}. By Lemma 2 we have $t \geq c(v_k) - 1 \geq 3$ and get $\text{level}(v_k) < \text{level}(v_{k+1})$. □

Now we prove that given a node v_0 of a MAX-equilibrium tree with $c(v_0) > 3$, there exists a MAX-arrangement starting with v_0 and closing with a node with private cost 3. With the previous results about MAX-arrangements, this leads to the upper bound.

Lemma 5. *Let I be a set of interests and $G = (V,E)$ a corresponding tree in a MAX-equilibrium. Then for $v_0,v_1 \in V$ with $d(v_0,v_1) > 3$ and $v_1 \in I(v_0)$ there exists a MAX-arrangement starting with v_0. And for every such MAX-arrangement it holds that the shortest path that visits all nodes of the MAX-arrangement in the given order uses at least $(c(v_0)^2 + c(v_0) - 6)/4$ different edges of G.*

Proof. Existence: v_0,v_1 obviously fulfill the conditions of a MAX-arrangement. Thus, it suffices to show that, given the beginning of a MAX-arrangement $v_0,\ldots,v_i$ with $c(v_j) > 3, j = 0,\ldots,i-1$, we can either find a next node v_{i+1} that suffices the conditions or otherwise $c(v_i) = 3$. Assume $c(v_i) > 3$. Then, by Lemma 1 there exist $x,y \in I(v_i)$ with $d(v_i,x) = c(v_i)$ and $c(v_i) \geq d(v_i,y) \geq c(v_i) - 1$ such that v_i is connected by at most one edge to the shortest path from x to y. Since $c(v_i) > 3$, also $c(x) \geq 3$ and $c(y) \geq 3$ hold. Now, for at least one node (x or y) we have that this node is most distant to v_{i-1}, it is not v_{i-2}, and thus it fulfills the conditions for a MAX-arrangement.

Traversal: We now can apply the previous lemmas for providing the minimal length of such a MAX-arrangement: Lemma 2 states that by construction we always have $c(v_{i+1}) \geq c(v_i) - 1$. Lemma 3 implies that no node can be contained more than once in a MAX-arrangement. By the arguments above, we can always find a new node for the MAX-arrangement until we reach a node w with $c(w) = 3$. Hence, the MAX-arrangement contains at least $c(v_0) - 2$ nodes. Since the distance between two succeeding nodes in the MAX-arrangement decreases by at most one per node, a traversal of this MAX-arrangement consists of at least $\sum_{i=3}^{c(v_0)} i = (c(v_0)^2 + c(v_0) - 6)/2$ edges. From these, by Lemma 4, at least $(c(v_0)^2 + c(v_0) - 6)/4$ edges are different. □

Theorem 1 (Restated) *Let I be a set of interests and $G = (V,E)$ a corresponding tree in a MAX-equilibrium, $n := |V|$. Then, for all $v \in V$ we have $c(v) \in \mathcal{O}(\sqrt{n})$.*

Proof. W.l.o.g. we may assume that there is at least one $\{v,v'\} \in I$ with $d(v,v') \geq 3$. Let nodes $v_0,v_1 \in V$, $v_1 \in I(v_0)$ have maximal distance among all nodes,

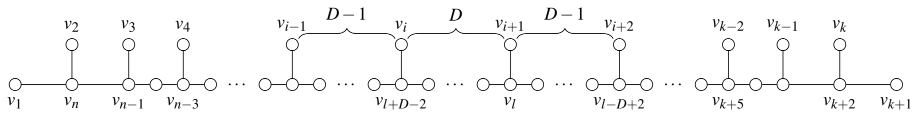

Fig. 5. Tree $G = (V,E)$ in a MAX-equilibrium with private cost $\Omega(D)$ for node v_i, with $D := \sqrt{|V|-2}+1$, $k := 2D-3$, and $l = n - \sum_{i=1}^{D} i$.

$D := d(v_0, v_1) = c(v_0)$. Then, by Lemma 5 we can find a MAX-arrangement $v_0, \ldots, v_m$ whose traversal uses at least $(D^2 + D - 6)/4$ different edges. Since our tree has exactly $n-1$ edges, we get $(D^2 + D - 6)/4 \le n - 1$ as an upper bound for the size of every MAX-arrangement and hence the private cost upper bound is $D \in \mathcal{O}(\sqrt{n})$. □

2.2 The Private Cost Upper Bound Is Tight

Next, we show that the upper bound of $\mathcal{O}(\sqrt{|V|})$ for the private costs is tight by constructing a MAX-equilibrium instance with one player having private cost $\Omega(\sqrt{|V|})$.

Remark 2. For a connection graph with nodes $V := \{v_1, \ldots, v_n\}$, let $I := \{\{v_i, v_{i+1}\} | i = 1, \ldots, n-1\} \cup \{\{v_n, v_1\}\}$ be interests such that (V, I) is a circle. Then, a node v_i with degree one in G cannot perform any swap if and only if it holds $|d(v_{i-1}, v_i) - d(v_i, v_{i+1})| \le 1$ and v_i is connected by one edge to the shortest path from v_{i-1} to v_{i+1} (cf. Lemma 1).

Theorem 2. *There exists a set of interests and a corresponding tree $G = (V,E)$ in a MAX-equilibrium with a node $v_i \in V$ that has private cost $c(v_i) \in \Omega\left(\sqrt{|V|}\right)$.*

Sketch. We consider interests $I := \{\{v_i, v_{i+1}\} | i = 1, \ldots, n-1\} \cup \{\{v_n, v_1\}\}$ and the connection graph as stated in Figure 5. (For the proof see the full version [5].) □

2.3 Existence of MAX-equilibria and the Price of Anarchy

In this section we compute the *price of stability* (PoS) and the *price of anarchy* (PoA). Let the *social optimum* represent an instance with the smallest social cost of any tree over all nodes (which is not necessarily in a MAX-equilibrium). Then, the PoS denotes the ratio between the minimum social cost of a MAX-equilibrium and the cost of a social optimum. Whereas the PoA denotes the ratio between the worst social cost of a MAX-equilibrium and the cost of a social optimum.

In the full version [5], we provide a simple approximation algorithm that generates for any interest graph a MAX-equilibrium tree whose social cost is at most twice as high as an optimal solution, yielding the following lemma.

Lemma 6. *For every set of interests I there exists a corresponding tree $G = (V,E)$ in a MAX-equilibrium that causes social cost $c(G) \le 2n$, $n := |V|$.*

Theorem 3. *The price of stability for I-BNCG is at most 2.*

Proof. Let I be a set of interests over nodes V, $n := |V|$. Then each connection graph that is a tree induces social cost of at least n. By Lemma 6 there exists a connection graph in a MAX-equilibrium with social cost of at most $2n$. Thus, the price of stability is at most $2n/n = 2$. □

Lemma 7. *There exist interest graphs over n nodes with a corresponding MAX-equilibrium tree that causes social cost of* $\Omega(n \cdot \sqrt{n})$.

Sketch. Consider $I := \{\{v_i, v_{i+1}\} | i = 1, \ldots, \lfloor n/2 \rfloor - 1\} \cup \{\{v_i, v_1\} | i = \lfloor n/2 \rfloor, \ldots, n\} \cup \{\{v_{n/2-1}, v_i\} | i = \lfloor n/2 \rfloor, \ldots, n\}$ and construct a similar graph as in Figure 5 but with $\lfloor n/2 \rfloor$ nodes at position of v_i, each with private cost $\Omega(\sqrt{n})$. (See full version [5].) □

Theorem 4. *The price of anarchy for I-BNCG is* $\Theta(\sqrt{n})$.

Proof. Theorem 1 provides an upper bound of $\mathcal{O}(\sqrt{n})$ for the private cost of every node in a tree in a MAX-equilibrium with n nodes. By this, $\mathcal{O}(n \cdot \sqrt{n})$ is an upper bound for the social cost of every MAX-equilibrium. Further, by Lemma 7 we get $\Omega(n \cdot \sqrt{n})$ as a worst case lower bound for the social cost of a graph in a MAX-equilibrium. For the cost of a social optimum, we get $\Theta(n)$. (Each social optimum incurs cost of at least n and at most $2n$.) Hence, we get $\Theta(n \cdot \sqrt{n}/n) = \Theta(\sqrt{n})$ for the price of anarchy. □

2.4 The Price of Anarchy for I-BNCG on General Graphs

Theorem 5. *The price of anarchy for I-BNCG with general connection graphs is* $\Theta(n)$.

Proof. First, note that the social cost of every instance are upper bounded by n^2 and lower bounded by n. Second, we provide an interest graph over n nodes ($n \equiv 0 \mod 6$) and a corresponding MAX-equilibrium graph $G = (V, E)$ with social cost $\Omega(n^2)$ (see Figure 6). To this, we connect $n/2$ nodes to a ring (ring nodes) and connect one additional (satellite) node to each of them. Each of the ring nodes is interested in its three adjacent nodes in G, whereas each satellite node is interested in its neighbor at the ring and in both satellite nodes at distance exactly $n/6 + 2$. This is an equilibrium and all $n/2$ satellite nodes have a private cost of $n/6 + 2$, i.e., the price of anarchy is $\Omega(n)$. □

3 Further Structural Properties of Equilibria

By Lemma 5 we achieved an upper bound for any MAX-arrangement (see Definition 1) contributed only by the property that the network is connected. Here, we introduce a second upper bound for a MAX-arrangement that is given by the size of a *maximum independent set* (MIS) in the interest graph. Having such an MIS of size M, we can bound the maximum private costs by $\mathcal{O}(M)$, which yields improved bounds for specific families of interest graphs. Particularly, this gives asymptotically same upper bounds for complete interest graphs on trees as those explicitly constructed by Alon et al. [2].

Theorem 6. *Let I be a set of interests and $G = (V, E)$ a corresponding tree in a MAX-equilibrium. Let M be the size of a maximum independent set in (V, I). Then for every MAX-arrangement $v_0, \ldots, v_m$: The length of this MAX-arrangement is at most $2 \cdot M$.*

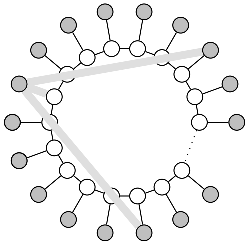

Fig. 6. MAX-equilibrium graph with social cost of $\Omega\left(n^2\right)$. Each white node is interested in its three neighbors. Each gray node is interested in its white neighbor and the two gray nodes at distance $n/6+2$.

Proof. We prove that the nodes of $v_0, \ldots, v_{m-1}$ with even index form an independent set in the interest graph (V,I). Consider an even index i and assume for contradiction that there is an even index $k < i$ such that $v_k \in I(v_i)$. By Lemma 3 we get $d(v_k, v_{k+1}) \leq d(v_k, v_{k+2})$. If $v_{k+2} \neq v_i$ with Lemma 2 and $c(v_j) > 3$ for all v_j in the MAX-arrangement we get $d(v_k, v_i) > d(v_k, v_{k+2}) + 1 \geq c(v_k)$. But this is a contradiction.

Thus, consider the case $v_{k+2} = v_i$. Since v_{k+1} is connected by at most one edge to the shortest path from v_k to v_{k+2} and $d(v_{k+1}, v_{k+2}) \geq 3$ we get $v_{k+2} \notin I(v_k)$. Otherwise we either get the same contradiction as before, or v_{k+1} would contradict to be the most distant node in $I(v_k)$ that fulfills the MAX-arrangement conditions. Hence, the nodes with even index form an independent set in (V,I), which gives $m \leq 2 \cdot M$ □

Corollary 1. *Let I be a set of interests and $G = (V,E)$ a corresponding tree in a MAX-equilibrium, $n := |V|$. Let the size M of any MIS in (V,I) be limited by $\sqrt{n}$. Then, for $v \in V$ we have $c(v) \in \mathcal{O}(M)$.*

Proof. W.l.o.g. we assume that there is a node with private cost greater than 3. Hence, there is a MAX-arrangement $v_0, \ldots, v_m, v_{m+1}$ with $c(v_i) > 3, i = 1, \ldots, m$ and $c(v_{m+1}) = 3$. By Theorem 6 we get $m \leq 2M$. Analog to Theorem 1, we get the upper bound. □

Corollary 2. *Let I be a set of interests and $G = (V,E)$ a corresponding tree in a MAX-equilibrium. If (V,I) is a complete graph, then $c(v) \in \mathcal{O}(1)$ for all $v \in V$.*

In the full version [5], we provide a scenario and a corresponding cyclic invocation sequence over all nodes, with each node performing a best-response improving swap (if possible), such that the nodes never reach a MAX-equilibrium. This gives:

Remark 3. I-BNCG is no potential game as defined by Monderer and Shapley [12].

4 Outlook and Future Work

In this paper, we presented tight worst case bounds for the private costs as well as for the social cost of any MAX-equilibrium on tree networks. Furthermore, we drew an interesting connection between the size of an MIS in the interest graph and upper bounds on the private/social costs. In comparison with MAX-equilibria on general graphs, we could show that the price of anarchy can perform much worse if the connection graph is not acyclic. However, it remains an open question whether the price of anarchy on general connection graphs with complete interests could perform better than $\mathcal{O}(n)$. For this, so far there is only a worst case lower bound of $\Omega(\sqrt{n})$ (by Alon et al. [2]) for the graph diameter in a MAX-equilibrium, yielding a lower bound for the price of anarchy. Techniques similar to our MAX-arrangement-technique may allow deeper insights into the nature of MAX-equilibria in that scenario. Apart from this, finding good upper bounds on the social cost of an AVG-equilibrium remains a challenging problem (in the full version [5], we give a lower bound of $\Omega(n)$ for the private costs).

Even if the existence of a MAX-equilibrium is always ensured (which we proved for trees), it remains an open question whether the dynamics ever reaches an equilibrium. We could state examples, where the network *never* converges to a MAX-equilibrium. It seems an interesting question whether we can guarantee the convergence by additional policies, e.g., by restricting the order in which nodes perform their swaps. And in case of a guaranteed convergence, how many swaps would it take to reach an equilibrium?

Currently, we only considered static interest graphs. But in practice, interests of network participants might change over time. Introducing a time model and considering certain (possibly restricted) changes of the interest graph seems a natural way to generalize our model, yielding an interesting online problem.

References

1. Albers, S., Eilts, S., Even-Dar, E., Mansour, Y., Roditty, L.: On nash equilibria for a network creation game. In: 17th SODA, pp. 89–98. ACM (2006)
2. Alon, N., Demaine, E.D., Hajiaghayi, M.T., Leighton, T.: Basic network creation games. In: 22nd SPAA, pp. 106–113. ACM (2010)
3. Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, É., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. In: 45th FOCS, pp. 295–304. IEEE (2004)
4. Anshelevich, E., Dasgupta, A., Tardos, É., Wexler, T.: Near-optimal network design with selfish agents. In: 35th STOC, pp. 511–520. ACM (2003)
5. Cord-Landwehr, A., Hüllmann, M., Kling, P., Setzer, A.: Basic network creation games with communication interests. CoRR, arxiv.org/abs/1207.5419 (2012)
6. Demaine, E.D., Hajiaghayi, M.T., Mahini, H., Zadimoghaddam, M.: The price of anarchy in network creation games. In: 26th PODC, pp. 292–298. ACM (2007)
7. Fabrikant, A., Luthra, A., Maneva, E., Papadimitriou, C.H., Shenker, S.: On a network creation game. In: 22nd PODC, pp. 347–351. ACM (2003)
8. Halevi, Y., Mansour, Y.: A Network Creation Game with Nonuniform Interests. In: Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 287–292. Springer, Heidelberg (2007)
9. Koutsoupias, E., Papadimitriou, C.: Worst-Case Equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999)
10. Lenzner, P.: On Dynamics in Basic Network Creation Games. In: Persiano, G. (ed.) SAGT 2011. LNCS, vol. 6982, pp. 254–265. Springer, Heidelberg (2011)
11. Mihalák, M., Schlegel, J.C.: The Price of Anarchy in Network Creation Games Is (Mostly) Constant. In: Kontogiannis, S., Koutsoupias, E., Spirakis, P.G. (eds.) SAGT 2010. LNCS, vol. 6386, pp. 276–287. Springer, Heidelberg (2010)
12. Monderer, D., Shapley, L.S.: Potential Games. Games and Economic Behavior 14(1), 124–143 (1996)