Greedy Selfish Network Creation

Pascal Lenzner

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

Abstract. We introduce and analyze *greedy equilibria* (GE) for the well-known model of selfish network creation by Fabrikant et al. [PODC'03]. GE are interesting for two reasons: (1) they model outcomes found by agents which prefer smooth adaptations over radical strategy-changes, (2) GE are outcomes found by agents which do not have enough computational resources to play optimally. In the model of Fabrikant et al. agents correspond to Internet Service Providers which buy network links to improve their quality of network usage. It is known that computing a best response in this model is NP-hard. Hence, poly-time agents are likely not to play optimally. But how good are networks created by such agents? We answer this question for very simple agents. Quite surprisingly, naive greedy play suffices to create remarkably stable networks. Specifically, we show that in the SUM version, where agents attempt to minimize their average distance to all other agents, GE capture Nash equilibria (NE) on trees and that any GE is in 3-approximate NE on general networks. For the latter we also provide a lower bound of $\frac{3}{2}$ on the approximation ratio. For the MAX version, where agents attempt to minimize their maximum distance, we show that any GE-star is in 2-approximate NE and any GE-tree having larger diameter is in $\frac{6}{5}$-approximate NE. Both bounds are tight. We contrast these positive results by providing a linear lower bound on the approximation ratio for the MAX version on general networks in GE. This result implies a locality gap of $\Omega(n)$ for the metric min-max facility location problem, where n is the number of clients.

1 Introduction

The area of Network Design is one of the classical and still very active fields in the realm of Theoretical Computer Science and Operations Research. But there is this curious fact: One of the most important networks which is increasingly shaping our everyday life – the Internet – cannot be fully explained by classical Network Design theory. Unlike centrally designed and optimized networks, the Internet was and still is created by a multitude of selfish agents (Internet Service Providers (ISPs)), who control and modify varying sized portions of the network structure ("autonomous systems"). This decentralized nature is an obstacle to approaching the design and analysis of the Internet as a classical optimization problem. Interestingly, each agent does face classical Network Design problems, i.e. minimizing the cost of connecting the *own* network to the rest of the Internet while ensuring a high quality of service. The Internet is the result of the interplay

P.W. Goldberg and M. Guo (Eds.): WINE 2012, LNCS 7695, pp. 142–155, 2012.

of such local strategies and can be considered as an equilibrium state of a game played by selfish agents.

The classical and most popular solution concept of such games is the (pure) Nash equilibrium [15], which is a stable state, where no agent unilaterally wants to change her current (pure) strategy. However, Nash equilibria (NE) have their difficulties. Besides their purely descriptive, non-algorithmic nature, there are two problems: (1) With NE as solution concept agents only care if there is a better strategy and would perform radical strategy-changes even if they only yield a tiny improvement. (2) In some games it is computationally hard to even tell if a stable state is reached because computing the best possible strategy of an agent is hard. Thus, for such games NE only predict stable states found by supernatural agents.

But what solutions are actually found by more realistic players, i.e. by agents who prefer smooth strategy-changes and who can only perform polynomial-time computations? And what impact on the stability has this transition from supernatural to realistic players?

In this paper, we take the first steps towards answering these questions for one of the most popular models of selfish network creation. This model, called the Network Creation Game (NCG), was introduced a decade ago by Fabrikant et al. [9]. In NCGs agents correspond to ISPs who create links towards other ISPs while minimizing cost and maximizing their quality of network usage. It seems reasonable that ISPs prefer greedy refinements of their current strategy (network architecture) over a strategy-change which involves a radical re-design of their infrastructure. Furthermore, computing the best strategy in NCGs is NP-hard. Hence, it seems realistic to assume that agents perform smooth strategy-changes and that they do not play optimally. We take this idea to the extreme by considering very simple agents and by introducing and analyzing a natural solution concept, called *greedy equilibrium*, for which agents can easily compute whether a stable state is reached and which models an ISP's preference for smooth strategy-changes.

1.1 Model and Definitions

In NCGs [9] there is a set of n agents V and each agent $v \in V$ can buy an edge $\{v, u\}$ to any agent $u \in V$ for the price of $\alpha > 0$. Here α is a fixed parameter of the game which specifies the cost of creating any link. The strategy S_v of an agent v is the set of vertices towards which v buys an edge. Let $G = (V, E)$ be the induced network, where an edge $\{x, y\} \in E$ is present if $x \in S_y$ or $y \in S_x$. The network G will depend heavily on the parameter α. To state this explicitly, we let (G, α) denote the network induced by the strategies of all agents V. In a NCG agents selfishly choose strategies to minimize their cost. There are basically two versions of NCGs, depending on the definition of an agent's cost-function. In the Sum *version* [9], agents try to minimize the sum of their shortest path lengths to all other nodes in the network, while in the Max *version* [7], agents try to minimize their maximum shortest path distance to any other network node. The precise definitions are as follows: Let S_v denote agent v's strategy in (G, α), then we have for the Sum version that the cost of agent v is $c_v(G,\alpha) = \alpha|S_v| + \sum_{w \in V(G)} d_G(v,w)$, if G is connected and $c_v(G,\alpha) = \infty$, otherwise. For the Max version we define agent v's cost as $c'_v(G,\alpha) = \alpha|S_v| + \max_{w \in V(G)} d_G(v,w)$, if G is connected and $c'_v(G,\alpha) = \infty$, otherwise. In both cases $d_G(\cdot,\cdot)$ denotes the shortest path distance in the graph G. Note that both cost functions nicely incorporate two conflicting objectives: Agents want to pay as little as possible for being connected to the network while at the same time they want to have good connection quality. For NCGs we are naturally interested in networks where no agent unilaterally wants to change her strategy. Clearly, such outcomes are pure NE and we let Sum-NE denote the set of all pure NE of NCGs for the Sum version and Max-NE denotes the corresponding set for the Max version.

Another important notion is the concept of *approximate Nash equilibria*. Let (G,α) be any network in a NCG. For all $u \in V(G)$ let $c(u)$ and $c^*(u)$ denote agent u's cost induced by her current pure strategy in (G,α) and by her best possible pure strategy, respectively. We say that (G,α) is a β-approximate Nash equilibrium if for all agents $u \in V(G)$ we have $c(u) \leq \beta c^*(u)$, for some $\beta \geq 1$.

1.2 Related Work

The work of Fabrikant et al. [9] did not only introduce the very elegant model described above. Among other results, the authors showed that computing a best possible strategy of an agent is NP-hard.

To remove the quite intricate dependence on the parameter α, Alon et al. [3] recently introduced the Basic Network Creation Game (BNCG), in which a network G is given and agents can only "swap" incident edges to decrease their cost. Here, a swap is the exchange of an incident edge with a non-existing incident edge. The cost of an agent is defined like in NCGs but without the edge-cost term. The authors of [3] proposed the *swap equilibrium* (SE) as solution concept for BNCGs. A network is in SE, if no agent unilaterally wants to swap an edge to decrease her cost. This solution concept has the nice property that agents can check in polynomial time if they can perform an improving strategy-change. The greedy equilibrium, which we analyze later, can be understood as an extension of the swap equilibrium which has similar properties but provides agents more freedom to act. Note, that in BNCGs an agent can swap *any* incident edge, whereas in NCGs only edges which are bought by agent v can be modified by agent v. This problem, first observed by Mihalák and Schlegel [13], can easily be circumvented, as recently proposed by the same authors in [14]: BNCGs are modified such that every edge is owned by exactly one agent and agents can only swap *own* edges. The corresponding stable networks of this modified version are called *asymmetric swap equilibrium* (ASE). However, independent of the ownership, edges are still two-way. These simplified versions of NCGs are an interesting object of study since (asymmetric) swap equilibria model the local weighing of decisions of agents and despite their innocent statement they tend to be quite complicated structures. In [12] it was shown that greedy dynamics in a BNCG converge very quickly to a stable state if played on a tree. The authors of [5] analyzed BNCGs on trees with agents having communication interests.

However, simplifying the model as in [3] is not without its problems. Allowing only edge-swaps implies that the number of edges remains constant. Hence, this model seems too limited to explain the creation of rapidly growing networks.

A part of our work focuses on tree networks. Such topologies are common outcomes of NCGs if edges are expensive, which led the authors of [9] to conjecture that all (non-transient) stable networks of NCGs are trees if α is greater than some constant. The conjecture was later disproved by Albers et al. [1] but it was shown to be true for high edge-cost. In particular, the authors of [13] proved that all stable networks are trees if $\alpha > 273n$ in the SUM version or if $\alpha > 129$ in the MAX version. Experimental evidence suggests that this transition to tree networks already happens at much lower edge-cost and it is an interesting open problem to improve on these bounds.

Demaine et al. [6] investigated NCGs, where agents cannot buy every possible edge. Furthermore, Ehsani et al. [8] recently analyzed a bounded-budget version. Both versions seem realistic, but in the following we do not restrict the set of edges which can be bought or the budget of an agent. Clearly, such restrictions reduce the qualitative gap between simple and arbitrary strategy-changes and would lead to weaker results for our analysis. Note, that this indicates that outcomes found by simple agents in the edge or budget-restricted version may be even more stable than we show in the following sections.

To the best of our knowledge, approximate Nash equilibria have not been studied before in the context of selfish network creation. Closest to our approach here may be the work of Albers et al. [2], which analyzes for a related game how tolerant the agents have to be in order to accept a centrally designed solution. We adopt a similar point of view by asking how tolerant agents have to be to accept a solution found by greedy play.

Guylás et al. [10] recently published a paper having a very similar title to ours. They investigate networks created by agents who use the length of "greedy paths" as communication cost and show that the resulting equilibria are substantially different to the ones we consider here. Their term "greedy" refers to the distances whereas our term "greedy" refers to the behavior of the agents.

1.3 Our Contribution

We introduce and analyze greedy equilibria (GE) as a new solution concept for NCGs. This solution concept is based on the idea that agents (ISPs) prefer greedy refinements of their current strategy (network architecture) over a strategy-change which involves a radical re-design of their infrastructure. Furthermore, GE represent solutions found by very simple agents, which are computationally bounded. We show in Section 2 that such greedy refinements can be computed efficiently and clarify the relation of GE to other known solution concepts for NCGs.

Our main contribution follows in Section 3 and Section 4, where we analyze the stability of solutions found by greedily playing agents. For the SUM version we show the rather surprising result that, despite the fact that greedy strategy-changes may be sub-optimal from an agent's point of view, SUM-GE capture

SUM-NE on trees. That is, in any tree network which is in SUM-GE *no* agent can decrease her cost by performing *any* strategy-change. For general networks we prove that any network in SUM-GE is in 3-approximate SUM-NE and we provide a lower bound of $\frac{3}{2}$ for this approximation ratio. Hence, we are able to show that greedy play almost suffices to create perfectly stable networks.

For the MAX version we show that these games have a strong non-local flavor which yields diminished stability. Here even GE-trees may be susceptible to non-greedy improving strategy-changes. Interestingly, susceptible trees can be fully characterized and we show that their stability is very close to being perfect. Specifically, we show that any GE-star is in 2-approximate MAX-NE and that any GE-tree having larger diameter is in $\frac{6}{5}$-approximate MAX-NE. We give a matching lower bound for both cases. For non-tree networks in GE the picture changes drastically. We show that for GE-networks having a very small α the approximation ratio is related to their diameter and we provide a lower bound of 4. For $\alpha \geq 1$, we show that there are non-tree networks in MAX-GE, which are only in $\Omega(n)$-approximate MAX-NE. The latter result yields that the locality gap of uncapacitated metric min-max facility location is in $\Omega(n)$.

Regarding the complexity of deciding Nash-stability, we show that there are simple polynomial time algorithms for tree networks in both versions. Furthermore, greedy-stability represents an easy to check certificate for 3-approximate Nash-stability in the SUM version.

2 Greedy Agents and Greedy Equilibria

We consider agents which check three simple ways to improve their current infrastructure. The three operations are

- *greedy augmentation*, which is the creation of *one* new own link,
- *greedy deletion*, which is the removal of *one* own link,
- *greedy swap*, which is a swap of *one* own link.

Computing the best augmentation/deletion/swap for one agent can be done in $\mathcal{O}(n^2(n+m))$ steps by trying all possibilities and re-computing the incurred cost. Observe, that these smooth strategy-changes induce some kind of organic evolution of the whole network which seems highly adequate in modeling the Internet. This greedy behavior naturally leads us to a new solution concept:

Definition 1 (Greedy Equilibrium). (G, α) *is in* greedy equilibrium *if no agent in* G *can decrease her cost by buying, deleting or swapping* one *own edge.*

Note, that GE can be understood as solutions which are obtained by a distributed local search procedure performed by selfish agents.

The next theorem relates GE to other solution concepts in the SUM version. See Fig. 1 for an illustration. Relationships are similar in the MAX version.

Theorem 1. *For the* SUM *version it is true that* $NE \subset GE \subset ASE$ *and that* $SE \subset ASE$*. Furthermore, we have* $NE \backslash SE \neq \emptyset$, $GE \backslash SE \neq \emptyset$, $(GE \backslash SE) \backslash NE \neq \emptyset$, $(GE \backslash NE) \cap SE \neq \emptyset$ *and* $NE \cap GE \cap SE \neq \emptyset$.

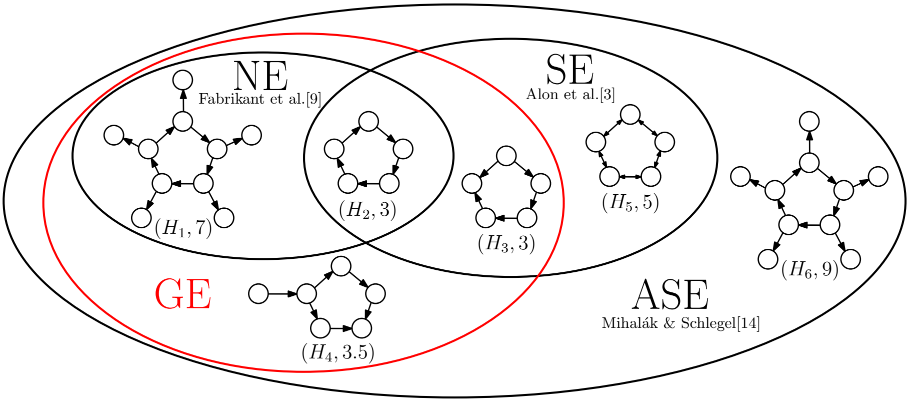

Fig. 1. Relations between solution concepts for NCGs in the SUM version. Edge-directions indicate edge-ownership, edges point away from its owner.

3 The Quality of Sum Greedy Equilibria

This section is devoted to discussing the quality of greedy equilibrium networks in the SUM version. We begin with a simple but very useful property.

Lemma 1. *If an agent v cannot decrease her cost by buying* one *edge in the* SUM *version, then buying $k > 1$ edges cannot decrease agent v's cost.*

3.1 Tree Networks in Sum Greedy Equilibrium

We show that in a NCG all stable trees found by greedily behaving agents are even stable against *any* strategy-change. Hence, in case of a tree equilibrium *no* loss in stability occurs by greedy play. This is a counter-intuitive result, since for each agent alone being greedy is clearly sub-optimal (the network in Fig. 2 with $\alpha = 6$ is an example). Thus, the following theorem shows the emergence of an optimal outcome out of a combination of sub-optimal strategies.

Theorem 2. *If (T, α) is in* SUM*-GE and T is a tree, then (T, α) is in* SUM*-NE.*

Before we prove Theorem 2, we first provide some useful observations. The well-known notion of a 1-*median* [11] is used: A 1-*median* of a connected graph G is a vertex $x \in V(G)$, where $x \in \arg\min_{u \in V(G)} \sum_{w \in V(G)} d(u, w)$.

Lemma 2. *Let (T, α) be a tree network in* SUM*-GE. If agent u owns edge $\{u, w\}$ in (T, α), then w must be a 1-median of its tree in the forest $T - \{u\}$.*

Let (T, α) be any tree network in SUM-GE and let T^u be the forest induced by removing all edges owned by agent u from T. Let F^u be the forest T^u without the tree containing vertex u. The above lemma directly implies the following:

Corollary 1. *Let (T, α) be in* SUM*-GE, and let F^u be defined as above. Agent u's strategy in (T, α) is the optimal strategy among all strategies that buy exactly one edge into each tree of F^u.*

Let $x \in V(T)$ be a 1-median of the tree T. Let $u \notin V(T)$ be a special vertex. We consider the network (G_T^u, α), which is obtained by adding vertex u and inserting edge $\{u, x\}$, which is owned by u, in T and by assigning the ownership of all other edges arbitrarily among the respective endpoints of any other edge in G_T^u. Furthermore, let $y_1, \ldots, y_l$ denote the neighbors of vertex x in T and let T_{y_i}, for $1 \leq i \leq l$, denote the maximal subtree of T which is rooted at y_i and which does not contain vertex x. See Fig. 2 (left) for an illustration. We consider a special

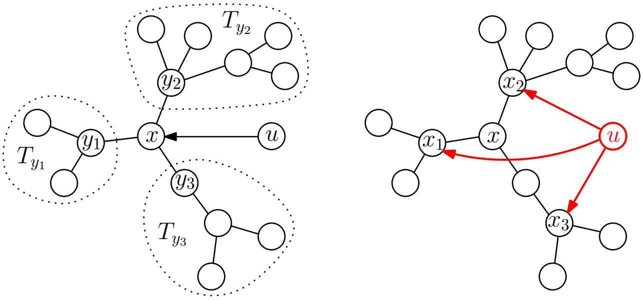

Fig. 2. The network (G_T^u, α) before and after agent u changes her strategy to S_u^*

strategy of agent u in (G_T^u, α): Let $S_u^* = \{x_1, \ldots, x_k\}$ be the best strategy of agent u which purchases *at least two* edges. The situation with agent u playing strategy S_u^* is depicted in Fig. 2 (right).

Lemma 3. *Let (G_T^u, α), $S_u^* = \{x_1, \ldots, x_k\}$ and the subtrees T_{y_i}, for $1 \leq i \leq l$ be specified as above. There is no subtree T_{y_i}, which contains all vertices $x_1, \ldots, x_k$.*

Next, let us consider two special strategies of agent u. Let S_u^1 be agent u's best strategy, which buys at least two edges including one edge towards vertex x. Furthermore, let S_u^2 be agent u's best strategy, which buys at least two edges, but no edge towards vertex x.

Lemma 4. *Let (G_T^u, α), S_u^1, S_u^2 and vertex x be specified as above. Let $x_j \in S_u^2$ be a vertex which has minimum distance to x among all vertices in S_u^2. If strategy S_u^2 yields less cost for agent u than strategy S_u^1, then x_j cannot be a leaf of G_T^u.*

Now we have all the tools we need to prove Theorem 2.

Proof (of Theorem 2). We will prove the contra-positive statement of Theorem 2. We show that if an agent u can decrease her cost by performing a strategy-change in a tree network (T, α) which is in Sum-GE, then there is an agent z in $V(T)$ who can decrease her cost by performing a *greedy* strategy-change. In that case we have a contradiction to (T, α) being in Sum-GE.

If agent u can decrease her cost by buying, deleting or swapping one own edge, then we have $u = z$ and we are done. Hence, we assume that agent u cannot decrease her cost by a greedy strategy-change but by performing an arbitrary strategy-change. We consider agent u's strategy-change towards the best possible *arbitrary* strategy S^* (if u has more than one such strategy, then we choose the one which buys the least number of edges). Clearly, agent u cannot remove any owned edge without purchasing edges, since T is a tree and the removal would disconnect T. Furthermore, since (T, α) is in SUM-GE and by Lemma 1, agent u cannot decrease her cost by purchasing $k > 0$ *additional* edges. Hence, the only way agent u can possibly decrease her cost is by removing j own edges and building k edges simultaneously. Clearly, $k \geq j$ must hold. Furthermore, by Corollary 1, it follows that $k > j$. Let F^u be the forest obtained by removing the j edges owned by agent u from T and let T^* be the tree in F^u which contains vertex u. Observe that among the k new edges, there cannot be edges having an endpoint in T^*. This is true because (T, α) is in SUM-GE and by Lemma 1. Any such edge would be a possible greedy augmentation which we assume not to exist. Hence, by the pigeonhole principle, we have that there must be at least one tree T_q in F^u into which agent u buys at least *two* edges with strategy S^*. We focus on T_q and will find agent z within.

Let $\{u, x\}$, with $x \in V(T_q)$, be the unique edge of T which connects u to the subtree T_q. Hence, agent u's strategy-change to S^* removes edge $\{u, x\}$ and buys $k_q > 1$ edges $\{u, x_1\}, \dots, \{u, x_{k_q}\}$, with $x_j \in V(T_q)$ for $1 \leq j \leq k_q$. Let $X = \{x_1, \dots, x_{k_q}\}$. By Lemma 1, we have $x_j \neq x$, for $x_j \in X$. Let $y_1, \dots, y_l$ denote the neighbors of vertex x in T_q and let $T_{y_1}, \dots, T_{y_l}$ be the maximal subtrees of T_q not containing vertex x, which are rooted at vertex $y_1, \dots, y_l$, respectively. Let $x_a \in X$ be a vertex of X which has minimum distance to vertex x. Let $T_a \in \{T_{y_1}, \dots, T_{y_l}\}$ be the subtree containing x_a. By Lemma 3, we have that there is a subtree $T_b \in \{T_{y_1}, \dots, T_{y_l}\}$, with $T_b \neq T_a$, which contains at least one vertex of X. Let $B = \{x_{b_1}, \dots, x_{b_p}\} = X \cap V(T_b)$. Furthermore, since no strategy which buys at least two edges including an edge towards x into T_q outperforms u's greedy strategy within T_q and by Lemma 4, we have that vertex x_a cannot be a leaf. That is, there is a vertex $z \in V(T_q)$, which is a neighbor of x_a, such that $d(z, x) > d(x_a, x)$. We show that agent z can decrease her cost by buying *one* edge in (T, α).

First of all, notice that by definition of S^*, we have that *each* edge $\{u, x_j\}$, with $x_j \in X$, must independently of the other bought edges yield a distance decrease of *more than* α for agent u. Otherwise agent u could remove this edge and obtain a strictly better (or smaller) strategy, which contradicts the fact that S^* is the best possible strategy (buying the least number of edges). Let $D_j \subset V(T_q)$ be the set of vertices to which edge $\{u, x_j\}$ is the first edge on agent u's unique shortest path. Since x_a has minimum distance to x, it follows that $D_r \subseteq V(T_b)$ for $r \in \{b_1, \dots, b_p\}$. The main observation is that agent z faces in some sense the same situation as agent u with strategy S^* but without all edges $\{u, y\}$, where $y \in B$: Both have vertex x_a as neighbor and their shortest paths to any vertex in T_b all traverse x_a and x. Remember, that each edge $\{u, y\}$, for all $y \in B$, yields a distance decrease of more than α for agent u and that $D_r \subseteq V(T_b)$, for $r \in \{b_1, \dots, b_p\}$. Furthermore, removing all those edges from S^* yields a strict cost increase for agent u. This implies that agent z can decrease her cost by buying all edges $\{z, y\}$, for $y \in B$, simultaneously. If $|B| = 1$, then this strategy-change is a greedy move by agent z which decreases z's cost. If $|B| > 1$, then, by the contra-positive statement of Lemma 1, it follows that there exists *one* edge $\{z, y^*\}$, with $y^* \in B$, which agent z can greedily buy to decrease her cost. □

3.2 Non-tree Networks in Sum Greedy Equilibrium

There exist non-tree networks in SUM-GE, since, as shown by Albers et al. [1], there exist non-tree networks in SUM-NE and we have SUM-NE $\subseteq$ SUM-GE. Having Theorem 2 at hand, one might hope that this nice property carries over to non-tree greedy equilibria. Unfortunately, this is not true.

Theorem 3. *There is a network in* SUM-*GE which is not in β-approximate* SUM-*NE for $\beta < \frac{3}{2}$.*

Fig. 3 shows the construction of a critical greedy equilibrium network.

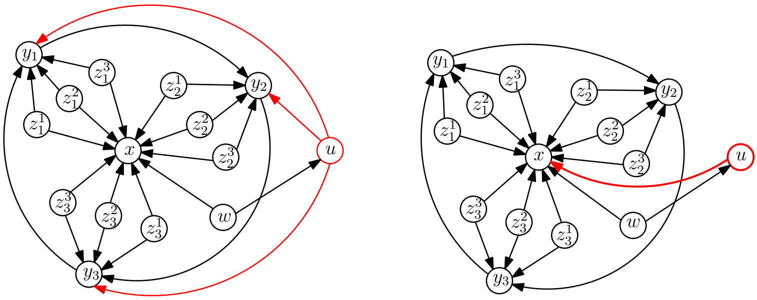

Fig. 3. The network $(G_k, k+1)$ for $k = 3$ and agent u's best response. Edges point away from its owner. For $k \to \infty$ agent u's improvement approaches a factor of $\frac{3}{2}$.

Now let us turn to the good news. We show that SUM-GEs cannot be arbitrarily unstable. On the contrary, they are very close to SUM-NEs in terms of stability.

Theorem 4. *Every network in* SUM-*GE is in* 3*-approximate* SUM-*NE.*

Proof. We prove Theorem 4 by providing a "locality gap preserving" reduction to the UNCAPACITATED METRIC FACILITY LOCATION problem (UMFL) [16]. Let u be an agent in (G, α) and let Z be the set of vertices in $V(G)$ which own an edge towards u. Consider the network (G', α), where all edges owned by agent u are removed. Observe, that the set Z is the same in (G, α) and (G', α). Let $\mathcal{S} = \{U \mid U \subseteq (V(G') \setminus \{u\}) \wedge U \cap Z = \emptyset\}$ denote the set of agent u's pure strategies in (G', α) which do not induce multi-edges or a self-loop. We transform (G', α) into an instance $I(G')$ for UMFL as follows:

Let $V(G') \setminus \{u\} = F = C$, where F is the set of facilities and C is the set of clients. For all facilities $f \in Z \cap F$ we define the opening cost to be 0, all other facilities have opening cost α. Thus, Z is exactly the set of cost 0 facilities in $I(G')$. For every $i, j \in F \cup C$ we define $d_{ij} = d_{G'}(i, j) + 1$. If there is no path between i and j in G', then we define $d_{ij} = \infty$. Clearly, since the distance in G' is metric we have that all distances d_{ij} in $I(G')$ are metric as well. See Fig. 4 for an example.
Now, observe that any strategy $S \in \mathcal{S}$ of agent u in (G', α) corresponds to the solution of the UMFL instance $I(G')$, where exactly the facilities in $F_S = S \cup Z$ are opened and where all clients are assigned to their nearest open facility. Moreover, every solution $F' = X \cup Z$, where $X \subseteq F \setminus Z$, for instance $I(G')$ corresponds to agent u's strategy $X \in \mathcal{S}$ in (G', α). Let $\mathcal{S}_{\text{UMFL}} = \{W \subseteq F \mid Z \subseteq$

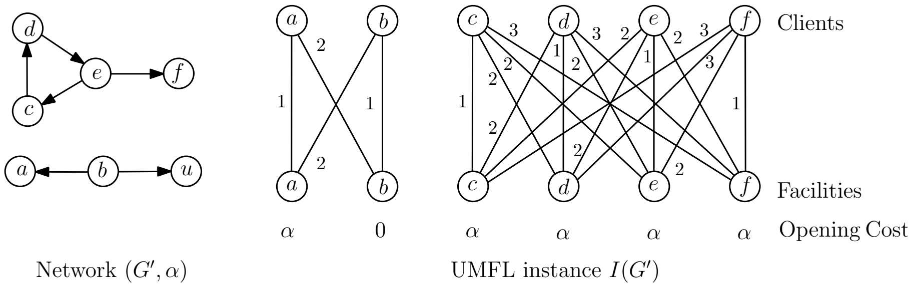

Fig. 4. Network (G', α) and its corresponding UMFL instance $I(G')$. Edges between clients and between facilities are omitted. All other omitted edges have length ∞.

$W\}$ denote the set of all solutions to instance $I(G')$, which open at least all cost 0 facilities. Hence, we have a bijection $\pi : \mathcal{S} \to \mathcal{S}_{\text{UMFL}}$, with $\pi(S) = S \cup Z$ and $\pi^{-1}(X) = X \setminus Z$. Let $\pi(S) = F_S$ and let (G_S, α) denote the network (G', α), where agent u has bought all edges towards vertices in S. Let $cost(F_S)$ denote the cost of the solution F_S to instance $I(G')$. We have that agent u's cost in (G_S, α) is equal to the cost of the corresponding UMFL solution F_S, since

$$\begin{aligned} c_u(G_S, \alpha) &= \alpha|S| + \sum_{w \in V(G_S)\setminus\{u\}} \left(1 + \min_{x \in S \cup Z} d_{G'}(x, w)\right) \\ &= \alpha|S| + 0|Z| + \sum_{w \in V(G_S)\setminus\{u\}} \min_{x \in S \cup Z} d_{xw} \\ &= \alpha|F_S \setminus Z| + 0|Z| + \sum_{w \in C} \min_{x \in F_S} d_{xw} \quad = cost(F_S). \end{aligned}$$

We claim the following: If agent u plays strategy $S \in \mathcal{S}$ and cannot decrease her cost by buying, deleting or swapping *one* edge in (G_S, α), then we have that the cost of the corresponding solution $F_S \in \mathcal{S}_{\text{UMFL}}$ to instance $I(G')$ cannot be strictly decreased by opening, closing or swapping *one* facility.

Proving the above claim suffices to prove Theorem 4. This can be seen as follows: For UMFL, Arya et al. [4] have already shown that the locality gap of UMFL is 3, that is, that any UMFL solution in which clients are assigned to their nearest open facility and which cannot be improved by opening, closing or swapping *one* facility is a 3-approximation of the optimum solution. By construction of $I(G')$, we have that every facility $z \in Z$ is the unique facility which is nearest to some client $w \in C$. Thus, we have that in any locally optimal and any globally optimal UMFL solution to $I(G')$ all cost 0 facilities must be open, since otherwise such a solution can be improved by opening a cost 0 facility. Hence, every locally or globally optimal solution to $I(G')$ has a corresponding strategy of agent u which yields the same cost. Using the claim and the result by Arya et al. [4], it follows that if agent u cannot decrease her cost by buying, deleting or swapping an edge in (G_S, α) then we have $c_u(G_S, \alpha) \leq 3c_u(G_{S^*}, \alpha)$, where S^* is agent u's optimal (non-greedy) strategy in (G', α) and (G_{S^*}, α) the network induced by S^*.

Now we prove the claim. Let $\pi(S) = F_S$. We have already shown that $c_u(G_S, \alpha) = cost(F_S)$. Furthermore, we have $Z \subseteq F_S$. We prove the contra-positive statement of the claim. Assume that solution F_S can be improved by opening, closing or swapping *one* facility. Let F'_S be this locally improved solution and let $cost(F'_S) < cost(F_S)$. Note, that $Z \subseteq F'_S$ must hold. This is true, since by construction of $I(G')$ closing a cost 0 facility increases the cost of any solution to $I(G')$. Hence, no facility $z \in Z$ can be included in a closing or swapping operation. It follows that the strategy $S' := \pi^{-1}(F'_S)$ exists. Observe, that $S = F_S \setminus Z$ and $S' = F'_S \setminus Z$ must differ by one element. Furthermore, by cost-equality, we have that $c_u(G_{S'}, \alpha) = cost(F'_S) < cost(F_S) = c_u(G_S, \alpha)$. Hence, agent u can buy, delete or swap one edge in (G_S, α) to decrease her cost. □

4 The Quality of Max Greedy Equilibria

In this section, we discuss the stability of networks in Max-GE. We will start by showing that operations of buying, deleting and swapping edges each may have a strong non-local flavor. See Fig. 5 for an illustration.

Lemma 5. *For $k \geq 2$ there is a network (G, α), where an agent can decrease her cost by buying/deleting/swapping k edges but not by buying/deleting/swapping $j < k$ edges.*

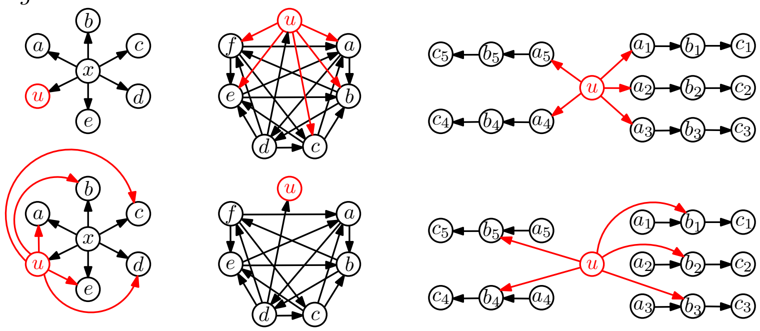

Fig. 5. The networks and strategy-changes for $k = 5$

Having seen Lemma 5, it should not come as a surprise that greedy local optimization may get stuck at sub-optimal states of the game.

4.1 Tree Networks in Max Greedy Equilibrium

The examples on the left and right side of Fig. 5 already show that there are tree networks, which are in Max-GE but not in Max-NE. In the following we show that this undesired behavior is restricted only to two families of tree networks in Max-GE. That is, we provide a characterization of all tree networks in Max-GE which are not in Max-NE. Furthermore, we show tight bounds on the stability for both mentioned families which are very close to the optimum. We start by introducing the main actors: *Cheap Stars* and *Badly Connected Trees.*

Definition 2 (Cheap Star). *A network (T, α) in* Max-*GE is called a* Cheap Star*, if T is a star having at least $n \geq 4$ vertices and $\alpha < \frac{1}{n-2}$. Furthermore, the ownership of all edges in T is arbitrary.*

Definition 3 (Badly Connected Tree). *A tree network* (T, α) *in* Max-*GE is a* Badly Connected Tree *if there is an agent* $u \in V(T)$ *who can decrease her cost by swapping* $k > 1$ *own edges simultaneously.*

Intuitively, Cheap Stars owe their instability to a multi-buy operation, whereas Badly Connected Trees owe their instability to a multi-swap operation. Observe that Cheap Stars have diameter 2 and that Badly Connected Trees have diameter at least 3. Hence, these families are disjunct. The following theorem shows that Cheap Stars and Badly Connected Trees are the *only* tree networks in Max-GE which are not in Max-NE.

Theorem 5. *Let* (T, α) *be a network in* Max-*GE, where* T *is a tree. The network* (T, α) *is in* Max-*NE if and only if it is not a Cheap Star or a Badly Connected Tree.*

We can use the characterization provided by Theorem 5 to "circumvent" the hardness of deciding whether a tree network is in Max-NE.

Theorem 6. *For every tree network* (T, α) *it can be checked in* $\mathcal{O}(n^4)$ *many steps whether* (T, α) *is in* Max-*NE.*

We are interested in the stability of tree networks in Max-GE. By Theorem 5, we only have to analyze the stability of Cheap Stars and Badly Connected Trees to get bounds on the stability on any tree network in Max-GE.

Lemma 6. *Every Cheap Star is in* 2*-approximate* Max-*NE. Furthermore, this bound is tight.*

Lemma 7. *Every Badly Connected Tree is in* $\frac{6}{5}$*-approximate* Max-*NE. Furthermore, this bound is tight.*

Combining Theorem 5 with Lemma 6 and Lemma 7 we arrive at the following:

Theorem 7. *Let* (T, α) *be a tree network in* Max-*GE. If* T *has diameter at most* 2*, then* (T, α) *is in* 2*-approximate* Max-*NE. If* T *has diameter at least* 3*, then* (T, α) *is in* $\frac{6}{5}$*-approximate* Max-*NE. Moreover, both bounds are tight.*

4.2 Non-tree Networks in Max Greedy Equilibrium

Fig. 5 (middle) shows that there are non-tree networks in Max-GE, which are not in Max-NE. We want to quantify the loss in stability of Max-GEs versus Max-NEs. For tree networks we have that Cheap Stars play a crucial role. These networks owe their instability to a multi-buy operation and to the fact that they are in Max-GE for arbitrarily small α. We generalize this property of Cheap Stars to non-tree networks.

Definition 4 (Cheap Network). *A network* (G, α) *in* Max-*GE, is called a* Cheap Network*, if* (G, α) *remains in* Max-*GE when* α *tends to* 0*.*

Cheap Stars yield a lower bound on the stability approximation ratio which equals their diameter. We can generalize this observation:

Theorem 8. *If there is Cheap Network* (G, α) *having diameter* d*, then there is an* α^* *such that the network* (G, α^*) *is in* MAX-*GE but not in* β*-approximate* MAX-*NE for any* $\beta < d$.

Lemma 8. *There is a Cheap Network having diameter* 4.

Corollary 2. *For* $\alpha < 1$ *there is a network* (G, α) *in* MAX-*GE, which is not in* β*-approximate* MAX-*NE for any* $\beta < 4$.

Now we consider the case, where $\alpha \geq 1$. Quite surprisingly, it turns out that this case yields a very high lower bound on the approximation ratio.

Theorem 9. *For* $\alpha \geq 1$ *there is a* MAX-*GE network* (G, α) *having* n *vertices, which is not in* β*-approximate* MAX-*NE for any* $\beta < \frac{n-1}{5}$.

We give a family of networks in MAX-GE each having an agent u who can decrease her cost by a factor of $\frac{n-1}{5}$ by a non-greedy strategy-change. The network (G_1, α) can be obtained as follows: $V(G_1) = \{u, v, l_1, l_2, a_1, a_2, b_1, b_2, x_1, y_1\}$ and agent u owns edges to a_1, a_2 and x_1. For $i \in \{1, 2\}$, agent b_i owns an edge to v and to a_i and agent l_i owns an edge to b_i. Finally, agent y_1 owns an edge to x_1 and to v. Fig. 6 (left) provides an illustration. To get the k-th member of the family, for $k \geq 2$, we simply add the vertices x_j, y_j, for $2 \leq j \leq k$, and let agent y_j own edges towards x_j and v. See Fig. 6 (right).

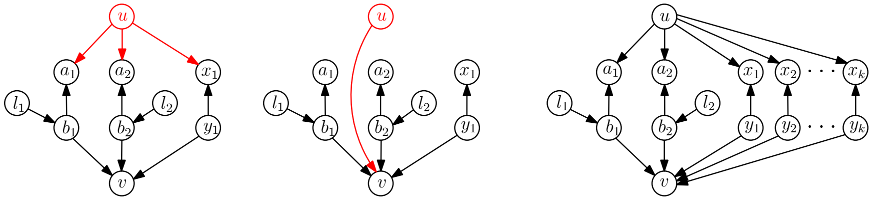

Fig. 6. (G_1, α) before (left) and after (middle) agent u's non-greedy strategy change and the network (G_k, α) (right)

Lemma 9. *Each of the networks* (G_i, α)*, as described above, is in* MAX-*GE for* $1 \leq \alpha \leq 2$.

Proof (of Theorem 9). We focus on agent u in the network (G_k, α) and show that this agent can change her strategy in a non-greedy way and thereby decrease her cost by a factor of $\frac{n-1}{5}$, where n is the number of vertices of G_k. Let S_u be agent u's current strategy in (G_k, α) and let S_u^* be u's strategy which only buys one edge towards vertex v. See Fig 6 (left and middle). Let $cost(u)$ and $cost^*(u)$ denote agent u's cost induced by strategy S_u and S_u^*, respectively. For $\alpha = 2$, we have

$$\frac{cost(u)}{cost^*(u)} = \frac{\alpha(2 + k) + 3}{\alpha + 3} = \frac{7}{5} + \frac{2k}{5} = \frac{n-1}{5},$$

where the last equality follows since $k = \frac{n-8}{2}$, by construction. □

Corollary 3. *Uncapacitated Metric Min-Max Facility Location has a locality gap of* $\frac{n-1}{5}$*, where* n *is the number of clients.*

Acknowledgments. I am grateful to Achim Passen for many interesting discussions and helpful comments. Furthermore, I thank an anonymous referee for pointing to [10].

References

1. Albers, S., Eilts, S., Even-Dar, E., Mansour, Y., Roditty, L.: On nash equilibria for a network creation game. In: Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithm, SODA 2006, pp. 89–98. ACM, New York (2006)
2. Albers, S., Lenzner, P.: On Approximate Nash Equilibria in Network Design. In: Saberi, A. (ed.) WINE 2010. LNCS, vol. 6484, pp. 14–25. Springer, Heidelberg (2010)
3. Alon, N., Demaine, E.D., Hajiaghayi, M., Leighton, T.: Basic network creation games. In: Proceedings of the 22nd ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2010, pp. 106–113. ACM, New York (2010)
4. Arya, V., Garg, N., Khandekar, R., Meyerson, A., Munagala, K., Pandit, V.: Local search heuristics for k-median and facility location problems. SIAM J. Comput. 33(3), 544–562 (2004)
5. Cord-Landwehr, A., Hüllmann, M., Kling, P., Setzer, A.: Basic Network Creation Games with Communication Interests. In: Serna, M. (ed.) SAGT 2012. LNCS, vol. 7615, pp. 72–83. Springer, Heidelberg (2012)
6. Demaine, E.D., Hajiaghayi, M., Mahini, H., Zadimoghaddam, M.: The price of anarchy in cooperative network creation games. SIGecom Exch. 8(2), 2:1–2:20 (2009)
7. Demaine, E.D., Hajiaghayi, M.T., Mahini, H., Zadimoghaddam, M.: The price of anarchy in network creation games. ACM Trans. on Algorithms 8(2), 13 (2012)
8. Ehsani, S., Fazli, M., Mehrabian, A., Sadeghian Sadeghabad, S., Safari, M., Saghafian, M., ShokatFadaee, S.: On a bounded budget network creation game. In: Proceedings of the 23rd ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2011, pp. 207–214. ACM, New York (2011)
9. Fabrikant, A., Luthra, A., Maneva, E., Papadimitriou, C.H., Shenker, S.: On a network creation game. In: Proc. of the 22nd Annual Symp. on Principles of Distributed Computing, PODC 2003, pp. 347–351. ACM, New York (2003)
10. Gulyás, A., Kõrösi, A., Szabó, D., Biczók, G.: On greedy network formation. In: Proceedings of ACM SIGMETRICS Performance W-PIN (2012)
11. Kariv, O., Hakimi, S.L.: An algorithmic approach to network location problems. ii: The p-medians. SIAM J. on Appl. Math. 37(3), 539–560 (1979)
12. Lenzner, P.: On Dynamics in Basic Network Creation Games. In: Persiano, G. (ed.) SAGT 2011. LNCS, vol. 6982, pp. 254–265. Springer, Heidelberg (2011)
13. 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)
14. Mihalák, M., Schlegel, J.: Asymmetric Swap-Equilibrium: A Unifying Equilibrium Concept for Network Creation Games. In: Rovan, B., Sassone, V., Widmayer, P. (eds.) MFCS 2012. LNCS, vol. 7464, pp. 693–704. Springer, Heidelberg (2012)
15. Nash, J.F.: Equilibrium points in n-person games. PNAS 36(1), 48–49 (1950)
16. Vazirani, V.V.: Approximation algorithms. Springer (2001)