On Selfish Creation of Robust Networks

Ankit Chauhan[1], Pascal Lenzner[1(✉)], Anna Melnichenko[1], and Martin Münn[2]

[1] Algorithm Engineering Group,
Hasso-Plattner-Institute Potsdam, Potsdam, Germany
{ankit.chauhan,pascal.lenzner,anna.melnichenko}@hpi.de
[2] Department of Computer Science, University of Liverpool, Liverpool, UK
M.F.Munn@liverpool.ac.uk

Abstract. Robustness is one of the key properties of nowadays networks. However, robustness cannot be simply enforced by design or regulation since many important networks, most prominently the Internet, are not created and controlled by a central authority. Instead, Internet-like networks emerge from strategic decisions of many selfish agents. Interestingly, although lacking a coordinating authority, such naturally grown networks are surprisingly robust while at the same time having desirable properties like a small diameter. To investigate this phenomenon we present the first simple model for selfish network creation which explicitly incorporates agents striving for a central position in the network while at the same time protecting themselves against random edge-failure. We show that networks in our model are diverse and we prove the versatility of our model by adapting various properties and techniques from the non-robust versions which we then use for establishing bounds on the Price of Anarchy. Moreover, we analyze the computational hardness of finding best possible strategies and investigate the game dynamics of our model.

1 Introduction

Networks are everywhere and we crucially rely on their functionality. Hence it is no surprise that designing networks under various objective functions is a well established research area in the intersection of Operations Research, Computer Science and Economics. However, investigating how to create suitable networks *from scratch* is of limited use for understanding most of nowadays networks. The reason for this is that most of our resource, communication and online social networks have not been created and designed by some central authority. Instead, these critical networks emerged from the interaction of many selfish agents who control and shape parts of the network. This clearly calls for a game-theoretic perspective.

One of the most prominent examples of such a selfishly created network is the Internet, which essentially is a network of sub-networks which are each owned and controlled by Internet service providers (ISP). Each ISP decides selfishly how to connect to other ISPs and thereby balancing the cost for creating links (buying the necessary hardware and/or peering agreement contracts for

M. Gairing and R. Savani (Eds.): SAGT 2016, LNCS 9928, pp. 141–152, 2016.
DOI: 10.1007/978-3-662-53354-3_12

routing traffic) and the obtained service quality for its customers. Interestingly, although the Internet is undoubtedly an important and critical infrastructure, there is no central authority which ensures its functionality if parts of the network fail. Despite this fact, the Internet seems to be robust against node or edge failures, which hints that a socially beneficial property like network robustness may emerge from selfish behavior.

Modeling agents with a desire for creating a *robust* network has long been neglected and was started to be investigated only very recently. This paper contributes to this endeavor by proposing and analyzing a model of selfish network creation, which explicitly incorporates agents who strive for occupying a central position in the network while at the same time ensuring that the overall network remains functional even under edge-failure.

1.1 Related Work

Previous research on game-theoretic models for network creation has either focused on *centrality models*, where the agents' service quality in the created network depends on the distances to other agents, or on *reachability models* where agents only care about being connected to many other agents.

Some prominent examples of centrality models for selfish network creation are [2,7–10,13,18]. They all have in common that agents correspond to nodes in a network and that the edge set of the network is determined by the combination of the agents' strategies. The utility function of an agent contains a service quality term which depends on the distances to all other agents. To the best of our knowledge the very recent paper by Meirom et al. [19] is the only centrality model which incorporates edge-failures. The authors consider two types of agents, major-league and minor-league agents, which maintain that the network remains 2-connected while trying to minimize distances, which are a linear combination of the length of a shortest path and the length of a best possible vertex disjoint backup path. Under some specific assumptions, e.g. that there are significantly more minor-league than major-league agents, they prove various structural properties of equilibrium networks and investigate the corresponding game-dynamics. In contrast to this, we will investigate a much simpler model with homogeneous agents which is more suitable for analyzing networks created by equal peers. Our results can be understood as zooming in on the sub-network formed by the major-league agents (i.e. top tier ISPs).

In reachability models, e.g. [3–5,11,12,15], the service quality of an agent is simply defined as the number of reachable agents and distances are ignored completely. For reachability models the works of Kliemann [15,16] and the very recent paper by Goyal et al. [12] explicitly incorporate a notion of network robustness in the utility function of every agent. All models consider an external adversary who strikes after the network is built. In [15,16] the adversary randomly removes a single edge and the agents try to maximize the expected number of reachable nodes post attack. Two versions of the adversary are analyzed: edge removal uniformly at random or removal of the edge which hurts the society of agents most. For the former adversary a constant Price of Anarchy is shown, whereas for the latter adversary this positive result is only true if edges can be created unilaterally. In [12] nodes are attacked (and killed) and this attack spreads virus-like to neighboring nodes unless they are protected by a firewall. Interestingly also this model has a low Price of Anarchy and the authors prove a tight linear bound on the amount of edge-overbuilding due to the adversary.

1.2 Model and Notation

We consider the Network Creation Game (NCG) by Fabrikant et al. [10] augmented with the uniform edge-deletion adversary from Kliemann [15] and we call our model *Adversary NCG* (Adv-NCG). More specifically, in an Adv-NCG there are n selfish agents which correspond to the nodes of an undirected multi-graph $G = (V, E)$ and we will use the terms agent and node interchangeably. A pure strategy S_u of agent $u \in V(G)$ is any multi-set over elements from $V \setminus \{u\}$. If v is contained k times in S_u then this encodes that agent u wants to create k undirected edges to node v. Moreover we say that u is the owner of all edges to the nodes in S_u. We emphasize the edge-ownership in our illustrations by drawing directed edges which point away from their owner - all edges are nonetheless understood to be undirected. Given an n-dimensional vector of pure strategies for all agents, then the union of all the edges encoded in all agents' pure strategies defines the edge set E of the multi-graph G. Since there is a bijection of multi-graphs with edge-ownership information and pure strategy-vectors, we will use networks and strategy-vectors interchangeably, e.g. by saying that a network is in equilibrium.

The agents prepare for an adversarial attack on the network after creation. This attack deletes one edge uniformly at random. Hence, agents try to minimize the attack's impact on themselves by minimizing their *expected cost*. Let $G - e$ denote the network G where edge e is removed. Let $\delta_G(u) = \sum_{v \in V(G)} d_G(u, v)$, where $d_G(u, v)$ is the number of edges of a shortest path from u to v in network G. Let

$$dist_G(u) = \frac{1}{|E|} \sum_{e \in E} \delta_{G-e}(u) = \frac{1}{|E|} \sum_{e \in E} \sum_{v \in V} d_{G-e}(u, v)$$

denote agent u's *expected distance cost* after the adversary has removed some edge uniformly at random from G. The expected cost of an agent u in network $G = (V, E)$ with edge-price α is defined as $cost_u(G, \alpha) = edge_u(G, \alpha) + dist_G(u)$, where $edge_u(G, \alpha) = \alpha \cdot |S_u|$ is the total edge-cost for agent u with strategy S_u in (G, α). Thus, compared to the NCG [10], the distance cost term is replaced by the expected distance cost with respect to uniform edge deletion.

Let (G, α) be any network with edge-ownership information. We call any strategy-change from S_u to S_u' of some agent u a *move*. Specifically, if $|S_u| = |S_u'|$, then such a move is called a *multi-swap*, if $|S_u \cap S_u'| < |S_u| - 2$ and a *swap* if $|S_u \cap S_u'| = |S_u| - 2$. If a move of agent u strictly decreases agent u's cost, then it is called an *improving move*. If no improving move exists, then we say that agent u plays its *best response*. Analogously we call a strategy-change towards a best response a *best response move*. A sequence of best response moves which starts and ends with network (G, α) is called a *best response cycle*. We say that (G, α) is in *Pure Nash Equilibrium* (NE) if all agents play a best response.

We measure the overall quality of a network (G, α) with its *social cost*, which is defined as $cost(G, \alpha) = \sum_{u \in V(G)} cost_u(G, \alpha) = edge(G, \alpha) + dist(G)$, where $edge(G, \alpha) = \sum_{u \in V(G)} edge_u(G, \alpha) = \alpha \cdot |E|$ and $dist(G) = \sum_{u \in V(G)} dist_G(u)$. Let $OPT(n, \alpha)$ be a network on n nodes with edge-price α which minimizes the social cost and we call $\mathrm{OPT}(n, \alpha)$ the *optimum network* for n and α. Let $\mathrm{maxNE}(n, \alpha)$ be the maximum social cost of any NE network on n agents with edge-price α and analogously let $\mathrm{minNE}(n, \alpha)$ be the NE having minimum social cost. Then, the *Price of Anarchy* is the maximum over all n and α of the ratio $\frac{\mathrm{maxNE}(n,\alpha)}{\mathrm{OPT}(n,\alpha)}$, whereas the *Price of Stability* is the maximum over all n and α of the ratio $\frac{\mathrm{minNE}(n,\alpha)}{OPT(n,\alpha)}$.

1.3 Our Contribution

This paper introduces and analyzes an accessible model, the Adv-NCG, for selfish network creation in which agents strive for a central position in the network while protecting against random edge-failures.

We show that optimum networks in the Adv-NCG are much more diverse than without adversary, which also indicates that the same holds true for the Nash equilibria of the game. However, we also show that many techniques and results from the NCG can be adapted to cope with the Adv-NCG, which indicates that the influence of the adversary is limited. In particular, we prove NP-hardness of computing a best possible strategy and W[2]-hardness of computing a best multi-swap. Moreover, we show that the Adv-NCG is not weakly acyclic, which is the strongest possible non-convergence result for any game. On the positive side, we prove that the amount of edge-overbuilding due to the adversary is limited, which is then used for proving that upper bounding the diameter essentially bounds the Price of Anarchy from above. We apply this by adapting two diameter-bounding techniques from the NCG to the adversarial version which then yields an upper bound on the PoA of $\mathcal{O}(1 + \alpha/\sqrt{n})$.

Due to space constraints some proofs and details are omitted. All missing material can be found in the full version [6].

2 Optimal Networks

Clearly, every optimal network must be 2-edge connected. Thus, every optimal network must have at least n edges. We first prove the intuitive fact that if edges get more expensive, then the optimum networks will have fewer edges.

Theorem 1. *Let $(G = (V, E), \alpha)$ and $(G' = (V, E'), \alpha')$ be optimal networks on n nodes in the Adv-NCG for α and α', respectively. If $\alpha' > \alpha$, then $|E| \geq |E'|$.*

In the following, we show that the landscape of optimum networks is much richer in the Adv-NCG, compared to the NCG where the optimum is either a clique

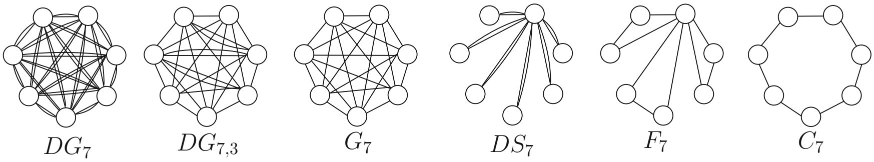

Fig. 1. Different candidates for optimum networks.

or a star, depending on α. In particular, we prove that there are $\Omega(n^2)$ different optimal topologies. We consider the following types of networks: Here DG_n is a clique of n nodes where we have a double edge between all pairs of nodes. Let $DG_{n,k}$ be a n node clique with exactly k pairs of nodes which are connected with double edges. Thus, $DG_{n,0} = G_n$ and $DG_{n,\binom{n}{2}} = DG_n$. Moreover, let F_n denote the fan-graph on n nodes which is a collection of triangles which all share a single node and let DS_n denote a star on n nodes where all connections between the center and the leaves are double edges. Finally, let C_n be a cycle of length n.

Clearly, if $\alpha = 0$, then the optimum network on n nodes must be a DG_n, since in this network no edge deletion of the adversary has any effect, it minimizes the sum of distances of each agent and since edges are for free.

Now consider what happens, if one pair of agents, say u and v, are just connected via a single edge instead of a double edge. The probability that the adversary removes this edge is $\frac{1}{n(n-1)-1}$. The removal would cause an increase in distance cost of 1 between u and v and between v and u. Thus, if $\alpha < \frac{2}{n(n-1)-1}$, then agent u and v would individually be better off buying another edge between each other. Thus, we have the following observation.

Observation 2. *If* $0 \leq \alpha \leq \frac{2}{n(n-1)-1}$, *then* $\mathrm{OPT}(n, \alpha) = DG_n$.

Lemma 1. *If* $\frac{2n(n-1)}{(\binom{n}{2}+k)(\binom{n}{2}+k+1)} \leq \alpha \leq \frac{2n(n-1)}{(\binom{n}{2}+k)(\binom{n}{2}+k-1)}$, *for* $1 \leq k \leq \binom{n}{2} - 1$, *then the network* $DG_{n,k}$ *is optimal.*

Note, that the proof of the above statement implies that the complete graph G_n is an optimum, if $\frac{4}{\binom{n}{2}+1} \leq \alpha < 2 - \frac{2}{\binom{n}{2}}$.

We also remark that we conjecture that Fig. 1 resembles a snapshot of optimum networks for increasing α from left to right. In fact, extensive simulations indicate that the optimum changes from G_n to DS_n and then, for slightly larger α for F_n. After this the cycles in the fan-graph increase and get fewer in number until, finally, for $\alpha \in \Omega(n^3)$ the cycle appears as optimum.

3 Computing Best Responses and Game Dynamics

In this section we investigate computational aspects of the Adv-NCG. First we analyze the hardness of computing a best response and the hardness of computing a best possible multi-swap. Then we analyze a natural process for finding an equilibrium network by sequentially performing improving moves.

3.1 Hardness of Best Response Computation

We first introduce useful properties for ruling out multi-buy or multi-delete moves. The proof is similar to the proof of Lemma 1 in [17].

Proposition 1. *If an agent cannot decrease its expected cost by buying (deleting) one edge in the Adv-NCG, then buying (deleting) $k > 1$ edges cannot decrease the agent's expected cost.*

Lemma 2. *If $1 - \frac{1}{|E|+1} < \alpha < 1 + \frac{1}{|E|(|E|-1)}$ and if agent u is not an endpoint of any double-edge in the Adv-NCG, then buying the minimum number of edges such that u's expected distance to all nodes in $V \setminus N_u$ is 2 and to nodes in N_u is $1 + \frac{1}{|E|}$ is u's best response.*

Now we show that computing the best possible strategy-change is intractable.

Theorem 3. *1. It is NP-hard to compute the best response of agent u in the Adv-NCG.*
2. It is $W[2]$-hard to compute the best multi-swap of agent u in the Adv-NCG.

3.2 Game Dynamics

We investigate the dynamic properties of the Adv-NCG. That is, we turn the model into a sequential version which starts with some initial network (G, α) and then agents move sequentially in some order and perform improving moves, if possible. One natural question is, if this process is guaranteed to converge to a Nash equilibrium of the game.

For the game dynamics of the Adv-NCG we prove the strongest possible negative result, which essentially shows that there is no hope for convergence if agents stick to performing improving moves only. In particular, we prove that the order of the agents moves or any tie-breaking between different improving moves does not help for achieving convergence. This result is even stronger than the best known non-convergence results for the NCG [14].

Theorem 4. *The $Adv - NCG$ is not weakly acyclic.*

4 Analysis of Networks in Nash Equilibrium

In this section we establish the existence of networks in Nash Equilibrium for almost the whole parameter space and we compare NE networks in the Adv-NCG with NE networks from the NCG [10] and Kliemann's adversarial model [15]. Moreover, we investigate structural properties which allow us to provide bounds on the Price of Stability and the Price of Anarchy.

We start with the existence result:

Theorem 5. *The networks DG_n and DS_n are in pure Nash Equilibrium if $\alpha \leq \frac{1}{n(n-1)-1}$ and $\alpha \geq 1 - \frac{1}{2n-1}$, respectively.*

Next, we show that NE in the Adv-NCG are not comparable with NE from the NCG or Kliemann's model.

Theorem 6. *There is a NE in the Adv-NCG which is not an NE in the NCG and vice versa. The analogous statement also holds for Kliemann's model.*

4.1 Relation Between the Diameter and the Social Cost

We prove a property which relates the diameter of a network with its social cost. With this, we prove that one of the most useful tools for analyzing NE in the NCG [10] can be carried over to the Adv-NCG.

Before we start, we analyze the diameter increase induced by removing a single edge in a 2-edge-connected network.

Lemma 3. *Let $G = (V, E)$ be any 2-edge-connected network having diameter D and let $G - e$ be the network G where some edge $e \in E$ is removed. Then the diameter of $G - e$ is at most $2D$.*

Next, we will focus on edges which are part of cuts of the network of size two. Remember that a bridge is an edge whose removal from a network increases the number of connected components of that network. Let $G = (V, E)$ be any 2-edge-connected network. We say that an edge $e \in E$ is a 2-*cut-edge* if there exists a cut of G of size 2 which contains edge e. Equivalently, e is a 2-cut-edge of G if its removal from G creates at least one bridge in $G-e$. We now bound the number of 2-cut-edges in any 2-edge-connected network G. This is an important structural result, since this proves that the amount of edge-overbuilding due to the adversary is sharply limited.

Lemma 4. *Any 2-edge-connected network G with n nodes can have at most $2(n-1)$ edges which are 2-cut-edges.*

Proof. Let e be any 2-cut-edge in network G. By definition, the removal of e creates one or more bridges in $G - e$. Let $b_1, \ldots, b_l$ denote those bridges. Note, that $b_1, \ldots, b_l$ also must be 2-cut-edges in G. Moreover, it follows that there must be a shortest cycle C in G which contains all the edges $e, b_1, \ldots, b_l$. If there are more than one such cycles, then fix one of them. We call the fixed cycle C a cut-cycle.

Notice that any 2-cut-edge corresponds to exactly one cut-cycle in the network and that every cut-cycle contains at least two 2-cut-edges. We show in the following that if any cut-cycle in the network contains at least three 2-cut-edges, then we can modify the network to obtain strictly more 2-cut-edges and strictly more cut-cycles. This implies that the number of 2-cut-edges is maximized if the number of cut-cycles is maximized and every cut-cycle contains exactly two 2-cut-edges.

Now we describe the procedure which converts any network with at least one cut-cycle containing at least three 2-cut-edges into a modified network with a strictly increased number of cut-cycles and 2-cut-edges (see Fig. 2). Let G be

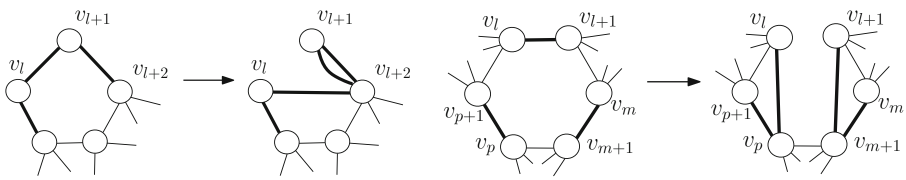

a) Conversion of adjacent 2-cut-edges b) Conversion of non-adjacent 2-cut-edges

Fig. 2. Increasing the number of 2-cut-edges by splitting up a cut-cycle.

any network with at least one cut-cycle $C = v_1, \ldots, v_k, v_1$ containing at least three 2-cut-edges. If there are two adjacent 2-cut-edges $\{v_l, v_{l+1}\}, \{v_{l+1}, v_{l+2}\}$ in cycle C, then delete the 2-cut-edge $\{v_l, v_{l+1}\}$ and insert two new edges $\{v_l, v_{l+2}\}$ and $\{v_{l+1}, v_{l+2}\}$. First of all, note that these new edges have not been present in network G before the insertion since otherwise $\{\{v_l, v_{l+1}\}, \{v_{l+1}, v_{l+2}\}\}$ cannot be a cut of G. We claim that both new edges are 2-cut-edges and that the cycle C is divided into two new cut-cycles $v_1, \ldots, v_l, v_{l+2}, \ldots, v_k, v_1$ and $v_{l+1}, v_{l+2}, v_{l+1}$. Indeed, there are at least two bridges $\{v_{l+1}, v_{l+2}\}$ and $\{v_k, v_{k+1}\}$ in the cut-cycle C after deleting $\{v_l, v_{l+1}\}$, and both of them end up in different new cut-cycles. Hence, deleting any of the newly inserted edges $\{v_l, v_{l+2}\}$ or $\{v_{l+1}, v_{l+2}\}$ implies that $\{v_k, v_{k+1}\}$ or $\{v_{l+1}, v_{l+2}\}$ becomes a bridge. Thus, both new edges are 2-cut-edges and both of new cycles are cut-cycles.

If there are three pairwise non-adjacent 2-cut-edges $\{v_l, v_{l+1}\}$, $\{v_m, v_{m+1}\}$, $\{v_p, v_{p+1}\}$ in cycle C, then delete one 2-cut-edge $\{v_l, v_{l+1}\}$ and insert two new edges $\{v_l, v_{p+2}\}$ and $\{v_{l+1}, v_{m+1}\}$. Analogous to above, both new edges cannot be already present in G and both are 2-cut-edges because deleting any of them renders edge $\{v_m, v_{m+1}\}$ or $\{v_p, v_{p+1}\}$ a bridge. Moreover, cut-cycle C is divided into two new cut-cycles.

Finally, we claim that the maximum number of cut-cycles in any n-vertex network G is at most $n-1$. Since we know that every such cut-cycle contains exactly two 2-cut-edges this then implies that there can be at most $2(n-1)$ 2-cut-edges in any network G.

Now we prove the above claim. Note that applying our transformation does not disconnect the network. Thus, we know that network G after all transformations is connected. Now we iteratively choose any cut-cycle C in G and we delete the two 2-cut-edges contained in C. This deletion increases the number of connected components of the current network by exactly 1. We repeat this process until we have destroyed all cut-cycles in G. Note that deleting edges from G may create new cut-cycles, but we never destroy more than one of them at a time. Thus, since each iteration increases the number of connected components of the network by 1, it follows that there can be at most $n-1$ iterations since network G with n vertices cannot have more than n connected components. $\square$

Remark 1. Lemma 4 is tight, since a path of length $n-1$, where all neighboring nodes are connected via double edges, has exactly $2(n-1)$ 2-cut-edges.

Now we relate the diameter with the social cost.

Theorem 7. *Let (G, α) be any NE network on n nodes having diameter D and let $OPT(n, \alpha)$ be the optimum network on n nodes for the same edge-cost α. Then we have that*

$$\frac{cost(G, \alpha)}{cost(OPT(n, \alpha))} \in \mathcal{O}(\mathcal{D}).$$

Proof. Since $OPT(n, \alpha)$ must be 2-edge-connected, it must have at least n edges. Moreover, the minimum expected distance between each pair of vertices in $OPT(n, \alpha)$ is at least 1. Thus, we have that $cost(OPT(n, \alpha)) \in \Omega(\alpha \cdot n + n^2)$.

Now we analyze the social cost of the NE network (G, α), where $G = (V, E)$. We have $cost(G, \alpha) = edge(G, \alpha) + dist(G)$ and we will analyze both terms separately. We start with an upper bound on $dist(G)$.

Since (G, α) has diameter D and since (G, α) is 2-edge-connected, Lemma 3 implies that the expected distance between each pair of vertices in (G, α) is at most $2D$. Thus, we have that $dist(G) \in \mathcal{O}(n^2 \cdot D)$.

Now we analyze $edge(G, \alpha)$. By Lemma 4 we have at most $2n$ many 2-cut-edges in G. Buying all those edges yields cost of at most $2n \cdot \alpha$.

We proceed with bounding the number of non-2-cut-edges in G. We consider any agent v and analyze how many non-2-cut-edges agent v can have bought. We claim that this number is in $\mathcal{O}\left(\frac{nD}{\alpha}\right)$, which yields total edge-cost of $\mathcal{O}(nD)$ for agent v. Summing up over all n agents, this yields total edge-cost of $\mathcal{O}(n^2 D)$ for all non-2-cut-edges of G. This then implies an upper bound of $\mathcal{O}(\alpha \cdot n + n^2 D)$ on the social cost of G which finishes the proof.

Now we prove our claim. Fix any non-2-cut-edge $e = \{v, w\}$ of G which is owned by agent v. Let $V_e \subset V$ be the set of nodes of G to which all shortest paths from v traverse the edge e.

We first show that removing the edge e increases agent v's expected distance to any node in V_e to at most $4D$. By Lemma 3, removing edge e increases the diameter of G from D to at most $2D$. Since e is a non-2-cut-edge, we have that $G - e$ is still 2-edge-connected. Thus, again by Lemma 3, it follows that agent v's expected distance to any other node in $G - e$ is at most $4D$.

However, removing edge e not only increases v's expected distance towards all nodes in V_e, instead, since $G - e$ has a less many edges than G, agent v's expected distance to *all* other nodes in $V \setminus (V_e \cup \{v\})$ increases as well. We now proceed to bound this increase in expected distance cost.

We compare agent v's expected distance cost in network G and in network $G - e$. Let m denote the number of edges in G. Thus, $G - e$ has $m - 1$ many edges. For network G agent v's expected distance cost is

$$dist_G(v) = \frac{1}{m} \sum_{f \in E} \delta_{G-f}(v) = \frac{1}{m} \sum_{f \in E \setminus \{e\}} \delta_{G-f}(v) + \frac{\delta_{G-e}(v)}{m}.$$

In network $G - e$, we have $dist_{G-e}(v) = \frac{1}{m-1} \sum_{f \in E \setminus \{e\}} \delta_{G-e-f}(v)$. Now we upper bound the increase in expected distance cost for agent v due to removal of edge e from G. $dist_{G-e}(v) - dist_G(v)$ is

$$\frac{1}{m-1}\sum_{f\in E\setminus\{e\}}\delta_{G-e-f}(v)-\left(\frac{1}{m}\sum_{f\in E\setminus\{e\}}\delta_{G-f}(v)+\frac{\delta_{G-e}(v)}{m}\right)$$
$$=\sum_{f\in E\setminus\{e\}}\left(\frac{\delta_{G-e-f}(v)}{m-1}-\frac{\delta_{G-f}(v)}{m}\right)-\frac{\delta_{G-e}(v)}{m}.$$

We have that $\delta_{G-e-f}(v) \leq \delta_{G-f}(v) + |V_e| \cdot 4D$, since in $G-e-f$ only the distances to nodes in V_e increase, compared to the network $G-f$ and since e is a non-2-cut-edge in G. Moreover, by Lemma 3, the distances to nodes in V_e in $G-e-f$ increase to at most $4D$ for each node in V_e. Thus, we have that $dist_{G-e}(v) - dist_G(v)$ is

$$\sum_{f\in E\setminus\{e\}}\left(\frac{\delta_{G-e-f}(v)}{m-1}-\frac{\delta_{G-f}(v)}{m}\right)-\frac{\delta_{G-e}(v)}{m}$$
$$\leq \sum_{f\in E\setminus\{e\}}\left(\frac{\delta_{G-f}(v)+|V_e|4D}{m-1}-\frac{\delta_{G-f}(v)}{m}\right)$$
$$=|V_e|4D+\sum_{f\in E\setminus\{e\}}\left(\frac{\delta_{G-f}(v)}{m(m-1)}\right)\leq |V_e|4D+\sum_{f\in E\setminus\{e\}}\left(\frac{2D\cdot n}{n(m-1)}\right)$$
$$\leq |V_e|4D+4D=(|V_e|+1)4D.$$

Since G is in Nash Equilibrium, we know that removing edge e is not an improving move for agent v. Thus, we have that

$$\alpha \leq (|V_e|+1)4D \iff |V_e| \geq \frac{\alpha}{4D}-1.$$

Thus, for all non-2-cut-edges e which are bought by agent v, we have that $|V_e| \in \Omega(\frac{\alpha}{D})$. Since all these sets V_e are disjoint, it follows that v can have bought at most $\frac{n}{\Omega(\frac{\alpha}{D})} \in \mathcal{O}(\frac{nD}{\alpha})$ many non-2-cut-edges. □

4.2 Price of Stability and Price of Anarchy

Theorem 8. *If $\alpha \leq \frac{1}{n(n-1)-1}$, then the PoS is 1. If $\frac{1}{n(n-1)-1} < \alpha < \frac{2}{n(n-1)-1}$, then PoS is strictly larger than 1, if $\alpha > 1-\frac{1}{2n-1}$, then the PoS is at most 2.*

We now show how to adapt two techniques from the NCG for bounding the diameter of equilibrium networks to our adversarial version. This can be understood as a proof of concept showing that the Adv-NCG can be analyzed as rigorously as the NCG. However, carrying over the currently strongest general diameter bound of $2^{\mathcal{O}(\sqrt{\log n})}$ due to Demaine et al. [8], which is based on interleaved region-growing arguments seems challenging due to the fact that we can only work with expected distances.

We start with a simple diameter upper bound based on [10].

Theorem 9. *The diameter of any NE network* (G, α) *is in* $\mathcal{O}(\sqrt{\alpha})$.

Proof. We prove the statement by contradiction. Assume that there are agents u and v in network G with $d_G(u,v) \geq 4\ell$, for some ℓ. Since expected distances cannot be shorter than distances in G, it follows that u's expected distance to v is at least 4ℓ. If u buys an edge to v for the price of α then u's decrease in expected distance cost is at least $\frac{|E|}{|E|+1}(4\ell - 1 + 4\ell - 3 + \cdots + 1) = \frac{|E|}{|E|+1}2\ell^2$.

Thus, if $d_G(u,v) > 4\sqrt{\alpha}$, then $u's$ decrease in expected distance cost by buying the edge uv is at least $\frac{|E|}{|E|+1}2\alpha > \alpha$. Thus, if the diameter of G is at least $4\sqrt{\alpha}$ then there is some agent who has an improving move. □

Together with Theorem 7 this yields the following statement:

Corollary 1. *The Price of Anarchy of the Adv-NCG is in* $\mathcal{O}(\sqrt{\alpha})$.

Adapting a technique by Albers et al. [1] yields a stronger statement, which implies constant PoA for $\alpha \in \mathcal{O}(\sqrt{n})$.

Theorem 10. *The Price of Anarchy of the Adv-NCG is in* $\mathcal{O}\left(1 + \frac{\alpha}{\sqrt{n}}\right)$.

We conclude with a weak lower bound on the PoA, which is tight for high α.

Theorem 11. *The Price of Anarchy of the Adv-NCG is at least* 2 *and for very large* α *this bound is tight.*

5 Conclusion

Our work is the first step towards incorporating both centrality and robustness aspects in a simple and accessible model for selfish network creation. In essence we proved that many properties and techniques can be carried over from the non-adversarial NCG and we indicated that the landscape of optimum and equilibrium networks in the Adv-NCG is much more diverse than without adversary. As for the NCG, proving strong upper or lower bounds on the PoA is very challenging. Especially surprising is the hardness of constructing higher lower bounds than in the NCG since by introducing suitable gadgets it is always possible to enforce that no agent wants to swap or delete edges. A non-constant lower bound on the PoA seems possible if α is linear in n.

It would also be interesting to consider different adversaries. An obvious candidate for this is node-removal at random. Another promising choice is a local adversary, where every agent considers that some of its incident edges may fail. This local perspective combined with a centrality aspect could explain why many selfishly built networks have a high clustering coefficient.

Another direction is to consider the swap version [9,20] of the Adv-NCG, especially in the case where all agents own at least 2 edges. We note in passing, that the swap-version of the Adv-NCG is not a potential game (an improving move cycle can be found in the full version [6]) and that creating equilibrium networks having diameter 4 is already very challenging.

References

1. Albers, S., Eilts, S., Even-Dar, E., Mansour, Y., Roditty, L.: On nash equilibria for a network creation game. ACM TEAC **2**(1), 2 (2014)
2. Alon, N., Demaine, E.D., Hajiaghayi, M.T., Leighton, T.: Basic network creation games. SIAM J. Discrete Math. **27**(2), 656–668 (2013)
3. Anshelevich, E., Bhardwaj, O., Usher, M.: Friend of my friend: Network formation with two-hop benefit. Theor. Comput. Syst. **57**(3), 711–752 (2015)
4. Bala, V., Goyal, S.: A noncooperative model of network formation. Econometrica **68**(5), 1181–1229 (2000)
5. Bilò, D., Gualà, L., Proietti, G.: Bounded-distance network creation games. ACM Trans. Econ. Comput. **3**(3), 16 (2015)
6. Chauhan, A., Lenzner, P., Melnichenko, A., Münn, M.: On selfish creation of robust networks. arXiv preprint (2016). arXiv:1607.02071
7. Corbo, J., Parkes, D.: The price of selfish behavior in bilateral network formation. In: Proceedings of the twenty-fourth annual ACM symposium on Principles of distributed computing, PODC 2005, pp. 99–107. ACM, New York, NY, USA (2005)
8. Demaine, E.D., Hajiaghayi, M.T., Mahini, H., Zadimoghaddam, M.: The price of anarchy in network creation games. ACM Trans. Algorithms **8**(2), 13 (2012)
9. Ehsani, S., Fadaee, S.S., Fazli, M., Mehrabian, A., Sadeghabad, S.S., Safari, M.A., Saghafian, M.: A bounded budget network creation game. ACM Trans. Algorithms **11**(4), 34 (2015)
10. Fabrikant, A., Luthra, A., Maneva, E., Papadimitriou, C.H., Shenker, S.: On a network creation game. In: PODC 2003, pp. 347–351. ACM, New York, USA (2003)
11. Flammini, M., Gallotti, V., Melideo, G., Monaco, G., Moscardelli, L.: Mobile network creation games. In: Even, G., Halldórsson, M.M. (eds.) SIROCCO 2012. LNCS, vol. 7355, pp. 159–170. Springer, Heidelberg (2012)
12. Goyal, S., Jabbari, S., Kearns, M., Khanna, S., Morgenstern, J.: Strategic network formation with attack and immunization. arXiv preprint 2015. arXiv:1511.05196
13. Jackson, M.O., Wolinsky, A.: A strategic model of social and economic networks. J. Econ. Theor. **71**(1), 44–74 (1996)
14. Kawald, B., Lenzner, P.: On dynamics in selfish network creation. In: SPAA 2013, pp. 83–92. ACM, New York, NY, USA (2013)
15. Kliemann, L.: The price of anarchy for network formation in an adversary model. Games **2**(3), 302–332 (2011)
16. Kliemann, L.: The price of anarchy in bilateral network formation in an adversary model. arXiv preprint (2013). arXiv:1308.1832
17. Lenzner, P.: Greedy selfish network creation. In: Goldberg, P.W. (ed.) WINE 2012. LNCS, vol. 7695, pp. 142–155. Springer, Heidelberg (2012)
18. Meirom, E.A., Mannor, S., Orda, A.: Network formation games with heterogeneous players and the internet structure. In: Proceedings of the 15th ACM Conference on Economics and Computation, pp. 735–752. ACM (2014)
19. Meirom, E.A., Mannor, S., Orda, A.: Formation games of reliable networks. In: INFOCOM 2015, pp. 1760–1768. IEEE (2015)
20. 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)