Network Creation Games: Think Global – Act Local

Andreas Cord-Landwehr[1](✉) and Pascal Lenzner[2](✉)

[1] Heinz Nixdorf Institute and Department of Computer Science, University of Paderborn, Paderborn, Germany
andreas.cord-landwehr@uni-paderborn.de
[2] Department of Computer Science, Friedrich-Schiller-University Jena, Jena, Germany
pascal.lenzner@uni-jena.de

Abstract. We investigate a non-cooperative game-theoretic model for the formation of communication networks by selfish agents. Each agent aims for a central position at minimum cost for creating edges. In particular, the general model (Fabrikant et al., PODC'03) became popular for studying the structure of the Internet or social networks. Despite its significance, locality in this game was first studied only recently (Bilò et al., SPAA'14), where a worst case locality model was presented, which came with a high efficiency loss in terms of quality of equilibria. Our main contribution is a new and more optimistic view on locality: agents are limited in their knowledge and actions to their local view ranges, but can probe different strategies and finally choose the best. We study the influence of our locality notion on the hardness of computing best responses, convergence to equilibria, and quality of equilibria. Moreover, we compare the strength of local versus non-local strategy changes. Our results address the gap between the original model and the worst case locality variant. On the bright side, our efficiency results are in line with observations from the original model, yet we have a non-constant lower bound on the Price of Anarchy.

1 Introduction

Many of today's networks are formed by selfish and local decisions of their participants. Most prominently, this is true for the Internet, which emerged from the uncoordinated peering decisions of thousands of autonomous subnetworks. Yet, this process can also be observed in social networks, where participants selfishly decide with whom they want to interact and exchange information. *Network Creation Games* (NCGs) are known as a widely adopted model to study the evolution and outcome of such networks. In the last two decades, several such game variants were introduced and analyzed in the fields of economics, e.g. Jackson

A. Cord-Landwehr—This work was partially supported by the German Research Foundation (DFG) within the Collaborative Research Centre "On-The-Fly Computing" (SFB 901).

G.F. Italiano et al. (Eds.): MFCS 2015, Part II, LNCS 9235, pp. 248–260, 2015.
DOI: 10.1007/978-3-662-48054-0_21

and Wolinsky [12], Bala and Goyal [3], and theoretical computer science, e.g. Fabrikant et al. [10], Corbo and Parkes [6].

In all of these models, the acting agents are assumed to have a global knowledge about the network structure on which their decisions are based. Yet, due to the size and dynamics of those networks, this assumption is hard to justify. Only recently, Bilò et al. [5] introduced the first variant of the popular (and mathematically beautiful) model by Fabrikant et al. [10] that explicitly incorporates a locality constraint. In this model, the selfish agents are nodes in a network, which can buy arbitrary incident edges. Every agent strives to maximize her service quality in the resulting network at low personal cost for creating edges. The locality notion by Bilò et al. [5] incorporates a worst case view on the network, which limits agents to know only their neighborhood within a bounded distance. The network structure outside of this view range is assumed to be worst possible. In particular, the assumed resulting cost for any agent's strategy-change is estimated as the worst case over all possible network structures outside of this view. In a follow-up work [4], this locality notion was extended by enabling agents to obtain information about the network structure by traceroute based strategies. Interestingly, in both versions the agents' service quality still depends on the whole network, which is actually a realistic assumption. As their main result, Bilò et al. show a major gap in terms of efficiency loss caused by the selfish behavior when compared to the original non-local model.

In this paper, we extend the investigation of the influence of locality in NCGs by studying a more optimistic but still very natural model of locality. Our model allows us to map the boundary of what locally constrained agents can actually hope for. Thereby, we close the gap between the non-local original version and the worst case locality models by Bilò et al. Besides studying the impact of locality on the outcomes' efficiency, we also analyze the impact on the computation of best response strategies and on the dynamic properties of the induced network creation process and we compare the strength of local versus non-local strategy-changes from an agent's perspective.

Due to space constraints we refer to [7] for all omitted proofs and details.

Our Locality Approach. We assume that an agent u in a network only has complete knowledge of her k-neighborhood. Yet, in contrast to Bilò et al. [4,5], besides knowing the induced subnetwork of all the agents that have distance of at most k to u, the agent can, e.g. by sending messages, judge her actual service quality that would result from a strategy-change. That is, we assume rational agents that "probe" different strategies, get direct feedback on their cost and finally select the best strategy. It is easy to see that allowing the agents to probe all available strategy-changes within their respective k-neighborhood and then selecting the best of them is equivalent to providing the agents with a global view, but restricting their actions to k-local moves. Here, a *k-local move* is any combination of (1) removing an own edge, (2) swapping an own edge towards an agent in the k-neighborhood, and (3) buying an edge towards an agent in the k-neighborhood.

Depending on the size of the neighborhood, probing all k-local moves may be unrealistic since there can be exponentially many such strategy-changes. To address this issue, we will also consider *k-local greedy moves*, which are k-local moves consisting only of exactly one of the options (1)–(3). It is easy to see that the number of such moves is quadratic in the number of vertices in the k-neighborhood and especially for small k and sparse networks this results in a small number of probes.

We essentially investigate the trade-off between the cost for eliciting a good local strategy by repeated probing and the obtained network quality for the agents. For this, we consider the extreme cases where agents either probe all k-local moves or only a polynomial fraction of them. Note that the former sheds light on the networks created by the strongest possible locally constrained agents.

Model and Notation. We consider the NCGs as introduced by Fabrikant et al. [10], where n agents V want to create a connected network among themselves. Each agent selfishly strives for minimizing her cost for creating network links, while maximizing her own service quality in the network. All edges in the network are undirected, have unit length and agents can create any incident edge for the price of $\alpha > 0$, where α is a fixed parameter of the game. (Note that in our illustrations, we depict the edge-ownerships by directing the edges away from their owners, yet still understand them as undirected.) The strategy $S_u \subseteq V \setminus \{u\}$ of an agent u determines which edges are bought by this agent, that is, agent u is willing to create (and pay for) all the edges ux, for all $x \in S_u$. Let $\mathcal{S}$ be the n-dimensional vector of the strategies of all agents, then $\mathcal{S}$ determines an undirected network $G_{\mathcal{S}} = (V, E_{\mathcal{S}})$, where for each edge $uv \in E_{\mathcal{S}}$ we have $v \in S_u$ or $u \in S_v$. If $v \in S_u$, then we say that agent u is the *owner* of edge uv, otherwise, if $u \in S_v$, then agent v owns the edge uv. We assume throughout the paper that each edge in $E_{\mathcal{S}}$ has a unique owner, which is no restriction, since no edge can have two owners in any equilibrium network. In particular, we assume that the cost of an edge cannot be shared and every edge is fully paid by its owner. With this, it follows that there is a bijection between strategy-vectors $\mathcal{S}$ and networks G with edge ownership information. Thus, we will use networks and strategy-vectors interchangeably, i.e., we will say that a network (G, α) is in equilibrium meaning that the corresponding strategy-vector $\mathcal{S}$ with $G = G_{\mathcal{S}}$ is in equilibrium for edge-price α. The edge-price α will heavily influence the equilibria of the game, which is why we emphasize this by using (G, α) to denote a network G with edge ownership information and edge-price α. Let $N_k(u)$ in a network G denote the set of all nodes in G with distance of at most k to u. The subgraph of G that is induced by $N_k(u)$ is called the *k-neighborhood* of u.

There are two versions for the cost function of an agent, which will yield two different games called the SUM-NCG [10] and the MAX-NCG [8]. The cost of an agent u in the network $G_{\mathcal{S}}$ with edge-price α is $\text{cost}_u(G_{\mathcal{S}}, \alpha) = \text{edge}_u(G_{\mathcal{S}}, \alpha) + \text{dist}_u(G_{\mathcal{S}}) = \alpha|S_u| + \text{dist}_u(G_{\mathcal{S}})$, where in the SUM-NCG we have that $\text{dist}_u(G_{\mathcal{S}}) = \sum_{w \in V} d_{G_{\mathcal{S}}}(u, w)$, if $G_{\mathcal{S}}$ is connected and $\text{dist}_u(G_{\mathcal{S}}) = \infty$, otherwise. Here, $d_{G_{\mathcal{S}}}(u, w)$ denotes the length of the shortest path between u and w in $G_{\mathcal{S}}$. Since all edges have unit length, $d_{G_{\mathcal{S}}}(u, w)$ is the hop-distance between u and w. In the MAX-NCG, the sum-operator in $\text{dist}_u(G_{\mathcal{S}})$ is replaced by a max-operator.[1]

A network $(G_{\mathcal{S}}, \alpha)$ is in *Pure Nash Equilibrium* (NE), if no agent can unilaterally change her strategy to strictly decrease her cost. Since the NE has undesirable computational properties, researchers have considered weaker solution concepts for NCGs. $G_{\mathcal{S}}$ is in *Asymmetric Swap Equilibrium* (ASE) [19], if no agent can strictly decrease her cost by swapping one own edge. Here, a swap of agent u is the replacement of one incident edge uv by any other new incident edge uw, abbreviated by $uv \to uw$. Note that this solution concept is independent of the parameter α since the number of edges per agent cannot change. A network $(G_{\mathcal{S}}, \alpha)$ is in *Greedy Equilibrium* (GE) [16], if no agent can buy, swap or delete exactly one own edge to strictly decrease her cost. These solution concepts induce the following k-local solution concepts: (G, α) is in *k-local Nash Equilibrium* (k-NE) if no agent can improve by a k-local move, (G, α) is in *k-local Greedy Equilibrium* (k-GE) if no agent can improve by a k-local greedy move. By slightly abusing notation, we will use the names of the above solution concepts to also denote the sets of instances that satisfy the respective solution concept, i.e., NE denotes the set of all networks (G, α) that are in Pure Nash Equilibrium. The notions GE, k-NE, and k-GE are used respectively. With this, we have the following:

Observation 1. *For $k \geq 1$: NE $\subseteq$ k-NE $\subseteq$ k-GE and NE $\subseteq$ GE $\subseteq$ k-GE.*

We will also consider approximate equilibria and say that a network (G, α) is in *β-approximate Nash Equilibrium*, if no strategy-change of an agent can decrease her cost to less than a β-fraction of her current cost in (G, α). Similarly, we say (G, α) is in *β-approximate Greedy Equilibrium*, if no agent can decrease her cost to less than a β-fraction of her current cost by buying, deleting or swapping exactly one own edge.

The *social cost* of a network $(G_{\mathcal{S}}, \alpha)$ is $\text{cost}(G_{\mathcal{S}}, \alpha) = \sum_{u \in V} \text{cost}_u(G_{\mathcal{S}})$. Let OPT_n be the minimum social cost of an n agent network. Let maxNE_n be the maximum social cost of any NE network on n agents and let minNE_n be the minimum social cost of any NE network on n agents. Then, the *Price of Anarchy* (PoA) [14] is the maximum over all n of the ratio $\frac{\text{maxNE}_n}{\text{OPT}_n}$, whereas the *Price of Stability* (PoS) [2] is the maximum over all n of the ratio $\frac{\text{minNE}_n}{\text{OPT}_n}$.

Known Results. SUM-NCGs were introduced in [10], where the authors proved the first PoA upper bounds, among them a constant bound for $\alpha \geq n^2$ and for trees. Later, by different authors and papers [1,8,17,18], the initial PoA bounds were improved for several ranges of α, resulting in the currently best known bounds of PoA $= \mathcal{O}(1)$, in the range of $\alpha = \mathcal{O}(n^{1-\varepsilon})$, for any fixed $\varepsilon \geq \frac{1}{\log n}$, and the range of $\alpha \geq 65n$, as well as the bound $o(n^{\epsilon})$, for α between $\Omega(n)$ and $\mathcal{O}(n \log n)$. Fabrikant et al. [10] also showed that for $\alpha < 2$ the network having minimum social cost is the complete network and for $\alpha \geq 2$ it is the

[1] Throughout this paper, we will only consider connected networks as they are the only ones which induce finite costs.

spanning star which yields a constant PoS for all α. Since it is known that computing the optimal strategy change is NP-hard [10], Lenzner [16] studied the effect of allowing only single buy/delete/swap operations, leading to efficiently computable best responses and 3-approximate NEs. NCG versions where the cost of edges can be shared have been studied by Corbo and Parkes [6] and Albers et al. [1].

For the Max-NCG, Demaine et al. [8] showed that the PoA is at most 2 for $\alpha \geq n$, for α in range $2\sqrt{\log n} \leq \alpha \leq n$ it is $\mathcal{O}(\min\{4^{\sqrt{\log n}}, (n/\alpha)^{1/3}\})$, and $\mathcal{O}(n^{2/\alpha})$ for $\alpha < 2\sqrt{\log n}$. For $\alpha > 129$, Mihalák and Schlegel [18] showed, similarly to the Sum-NCG, that all equilibria are trees and the PoA is constant.

Kawald and Lenzner [13,15] studied convergence properties of the sequential versions of many NCG-variants and provided mostly negative results. The agents in these variants are myopic in the sense that they only optimize their next step which is orthogonal to our locality constraint.

To the best of our knowledge, the only models in the realm of NCGs that consider locality constraints are [4,5]. As discussed above, both model the local knowledge in a very pessimistic way. Hence, it is not surprising that in [5] the authors lower bound the PoA by $\Omega(n/(1+\alpha))$ for Max-NCG and by $\Omega(n/k)$ for Sum-NCG, when $k = o(\sqrt[3]{\alpha})$. In particular, for Max-NCG they show that their lower bound is still $\Omega(n^{1-\varepsilon})$ for every $\varepsilon > 0$, even if k is poly-logarithmic and $\alpha = \mathcal{O}(\log n)$. On the bright side, they provide several PoA upper bounds that match with their lower bounds for different parameter combinations. In their follow-up paper [4], they equip agents with knowledge about either all distances to other nodes, a shortest path tree, or a set of all shortest path trees. However, the Max-NCG PoA bounds are still $\Theta(n)$ for $\alpha > 1$, while the bounds for Sum-NCG improve to $\Theta(\min\{1 + \alpha, n\})$. Note, that this is in stark contrast to the known upper bounds for the non-local version.

Apart from NCGs the influence of locality has already been studied for different game-theoretic settings, e.g. in the local matching model by Hoefer [11], where agents only know their 2-neighborhood and choose their matching partners from this set.

Our Contribution. We introduce a new locality model for Network Creation Games which allows to explore the intrinsic limits induced by a locality constraint and apply it to one of the mostly studied NCG versions, namely the Sum-NCG introduced by Fabrikant et al. [10]. In Sect. 2, we prove that constraining the agents' actions to their k-neighborhood has no influence on the hardness of computing good strategies and on the game dynamics, even for very small k. In Sect. 3, we explore the impact of locality from the agents' perspective by studying the strength of local versus global strategy-changes obtaining an almost tight general approximation gap of $\Omega(\frac{\log n}{k})$, which is tight for trees. Finally, in Sect. 4 we provide drastically improved PoA upper bounds compared to [4], which are in line with the known results about the non-local Sum-NCG. In contrast to this, we also prove a non-constant lower bound on the PoA, which proves that even in the most optimistic locality model the social efficiency deteriorates significantly.

2 Computational Hardness and Game Dynamics

In this section, we study the effect of restricting agents to k-local (greedy) moves on the hardness of computing a best response and the convergence to equilibria. We start with the observation that for any $k \geq 1$ computing a best possible k-local move is not easier than in the general setting. See [7] for omitted proofs.

Theorem 2. *Computing a best possible k-local move is NP-hard for all $k \geq 1$.*

Proof (Sketch). We reduce from Dominating Set [22] and we focus on an agent u who owns edges to any other agent, thus who can only delete edges to improve. For $1 < \alpha < 2$ this yields that u's best possible k-local move only keeps edges to nodes belonging to a minimum dominating set in the original instance. □

Clearly, the best possible k-local greedy move of an agent can be computed in polynomial time by simply trying all possibilities. Similarly to [16], it is true that the best k-local greedy move is a 3-approximation of the best k-local move. This yields:

Theorem 3. *For any $k \geq 1$, every network in k-GE is in 3-approximate k-NE.*

The same construction as in [16] yields:

Corollary 1. *For $k \geq 2$ there exist k-GE networks that are in $\frac{3}{2}$-approximate k-NE.*

In the following, we analyze the influence of k-locality on the dynamics of the network creation process. For this, we consider several sequential versions of the Sum-NCG, which were introduced in [13,15], and refer to the corresponding papers for further details. Our results only cover the Sum-NCG but we suspect similar results for corresponding versions of Max-NCG, which would be in line with the non-local version [13].

In short, we consider the following network creation process: Starting with any connected network having n agents, edge price α, and arbitrary edge ownerships. Now agents move sequentially, that is, at any time exactly one agent is active and checks if an improvement of her current cost is possible. If this is the case, then this agent will perform the strategy-change towards her best possible new strategy. After every such move this process is iterated until there is no agent who can improve by changing her current strategy, i.e., the network is in equilibrium. If there is a cyclic sequence of such moves, then clearly this process is not guaranteed to stop and we call such a situation a *best response cycle* (BR-cycle). Note that the existence of a BR-cycle, even if the participating active agents are chosen by an adversary, implies the strong negative result that no ordinal potential function [20] can exist.

In Table 1, we summarize our results for four types of possible strategy-changes: If agents can only swap own edges in their k-neighborhood, then this is called the *k-local Asymmetric Swap Game* (k-ASG). If agents are allowed to swap any incident edge within their k-neighborhood, then we have the *k-local Swap Game* (k-SG). If agents are allowed to perform k-local greedy moves only, then we have the *k-local Greedy Buy Game* (k-GBG), and if any k-local move is allowed, then we have the *k-local Buy Game* (k-BG).

Table 1. Overview of convergence speeds in the sequential versions. Note that the existence of a BR-cycle implies that there is no convergence guarantee. Here a move is a strategy-change by one agent.

k	Sum k-SG	Sum k-ASG	Sum k-GBG	Sum k-BG
$k = 1$	no moves	no moves	$\Theta(n^2)$ moves	$\Theta(n)$ moves
$k = 2$	BR-cycle	OPEN	BR-cycle	BR-cycle
$k \geq 3$	BR-cycle	BR-cycle	BR-cycle	BR-cycle
on trees	$\mathcal{O}(n^3)$ moves	$\mathcal{O}(n^3)$ moves	BR-cycle for $k \geq 2$	BR-cycle for $k \geq 2$

3 Local versus Global Moves

In this section, we investigate the agents' perspective. We ask how much agents lose by being forced to act locally. That is, we compare k-local moves to arbitrary non-local strategy-changes. We have already shown that the best possible k-local greedy move is a 3-approximation of the best possible k-local move. The same was shown to be true for greedy moves versus arbitrary strategy-changes [16]. Thus, if we ignore small constant factors, it suffices to compare k-local greedy moves with arbitrary greedy moves. All omitted details of this section can be found in [7].

We start by providing a high lower bound on the approximation ratio for k-local greedy moves versus arbitrary greedy moves. Corresponding to this lower bound, we provide two different upper bounds. The first one is a tight upper bound for tree networks. The second upper bound holds for arbitrary graphs, but is only tight for constant k. Here, the structural difference between trees and arbitrary graphs is captured in the difference of Lemmas 1 and 2. Whereas for tree networks only edge-purchases have to be considered, for arbitrary networks edge-swaps are critical as well.

Theorem 4 (Locality Lower Bound). *For any n' there exist tree networks on $n \geq n'$ vertices having diameter $\Theta(\log n)$ which are in k-GE but only in $\Omega\left(\frac{\log n}{k}\right)$-approximate GE.*

For the first upper bound, let (T, α) be a tree network in k-GE and assume that it is not in GE, that is, we assume there is an agent u, who can decrease her cost by buying or swapping an edge. Note that we can rule out single edge-deletions, since these are 1-local greedy moves and assume that no improving k-local greedy move is possible for any agent. First, we will show that we only have to analyze the cost decrease achieved by single edge-purchases.

Lemma 1. *Let T be any tree network. If an agent can decrease her cost by performing any single edge-swap in T, then this agent also can decrease her cost by performing a k-local single edge-swap in T, for any $k \geq 2$.*

Now we analyze the maximum cost decrease achievable by buying one edge in T. We show that our obtained lower bound from Theorem 4 is actually tight.

Theorem 5 (Tree Network Locality Upper Bound). *Any tree network on n vertices in k-GE is in $\mathcal{O}\left(\frac{\log n}{k}\right)$-approximate GE.*

We now prove a slightly inferior approximation upper bound for arbitrary networks. For this, we show that for any k the diameter of any network is an upper bound on the approximation ratio. For constant k this matches the $\Omega(\mathrm{diam}(G)/k)$ lower bound from Theorem 4 and is almost tight otherwise.

Lemma 2. *For any $k \geq 2$ there is a non-tree network in which some agent cannot improve by performing a k-local swap but by a $k+1$-local swap.*

It turns out that both buying and swapping operations yield the same upper bound on the approximation ratio, which is even independent of k.

Theorem 6 (General Locality Upper Bound). *Let G' be the network after some agent u has performed a single improving edge-swap or -purchase in a network G, then $\frac{\mathrm{cost}_u(G)}{\mathrm{cost}_u(G')} \leq \mathrm{diam}(G)$. Thus, for any k any network G in k-GE is in $\mathcal{O}(\mathrm{diam}(G))$-approximate GE.*

4 The Quality of k-Local Equilibria

In this section we prove bounds on the Price of Anarchy and on the diameter of equilibrium networks. All omitted proofs can be found in [7].

We start with a theorem about the relationship of β-approximate Nash equilibria and the social cost. This may be of independent interest, since it yields that the PoA can be bounded by using upper bounds on approximate best responses. In fact, we generalize an argument by Albers et al. ([1], Proof of Theorem 3.2).

Theorem 7. *If for $\alpha \geq 2$, (G, α) is in β-approximate NE, then the ratio of social cost and social optimum is at most $\beta(3 + \mathrm{diam}(G))$.*

However, for our bounds we will use another theorem, which directly carries over from the original non-local model.

Theorem 8 ([10,21] Lemma 19.4). *For any $k \geq 1$ and $\alpha \geq 2$, if a k-NE network (G, α) has diameter D, then its social cost is $\mathcal{O}(D)$ times the optimum social cost.*

We obtain several PoA bounds by applying Theorem 8 and the diameter bounds shown later in this section. Summarizing them, the next theorem provides our main insight into the PoA: The larger k grows, the better the bounds get, providing a remarkable gap between $k = 1$ and $k > 1$. In other words, the theorem tells the "price" in terms of minimal required k-locality to obtain a specific PoA upper bound.

Theorem 9. *For k-NE with $k \geq 1$ the following PoA upper bounds hold (cf. Fig. 1):*

$k = 1$: PoA $= \Theta(n)$, *which can be seen by considering a line graph.*

$k \geq 2$: *Theorem 12 gives* PoA $= \begin{cases} \mathcal{O}((\alpha/k)+k) & k < 2\sqrt{\alpha}, \\ \mathcal{O}(\sqrt{\alpha}) & k \geq 2\sqrt{\alpha}. \end{cases}$

$k \geq 6$: *Theorem 13 gives* PoA $= \mathcal{O}\left(n^{1-\varepsilon(\log(k-3)-1))}\right)$, *provided* $1 \leq \alpha \leq n^{1-\varepsilon}$ *with* $\varepsilon \geq 1/\log n$. *Note, this yields a constant PoA bound for* $\varepsilon \geq 1/(\log(k-3)-1)$.

For tree networks and $k \geq 2$, *Corollary 2 gives* PoA $= \mathcal{O}(\log n)$.

4.1 Conformity of *k*-Local and Nash Equilibria

We stated in Observation 1 that k-$NE \subseteq NE$ holds. In the following, we discuss the limits of known proof techniques to identify the parameters for which both equilibria concepts coincide, i.e., k-$NE = NE$.

Theorem 10. *The equilibrium concepts k-NE and NE coincide for the following parameter combinations and yield the corresponding PoA results (cf. Fig. 1):*

1. *For* $0 < \alpha < 1$ *and* $2 \leq k$, k-$NE = NE$ *and* PoA $= \mathcal{O}(1)$.
2. *For* $1 \leq \alpha \leq \sqrt{n/2}$ *and* $6 \leq k$, k-$NE = NE$ *and* PoA $= \mathcal{O}(1)$.
3. *For* $1 \leq \alpha \leq n^{1-\varepsilon}$, $\varepsilon \geq 1/\log(n)$ *and* $4.667 \cdot 3^{\lceil 1/\varepsilon \rceil} + 8 \leq k$, k-$NE = NE$ *and* PoA $= \mathcal{O}(3^{\lceil 1/\varepsilon \rceil})$.
4. *For* $1 \leq \alpha \leq 12n \log n$ *and* $2 \cdot 5^{1+\sqrt{\log n}} + 24\log(n) + 3 \leq k$, k-$NE = NE$ *and* PoA $= \mathcal{O}(5^{\sqrt{\log n}} \log n)$.
5. *For* $12n \log n \leq \alpha$ *and* $2 \leq k$, k-$NE = NE$ *and* PoA $= \mathcal{O}(1)$.

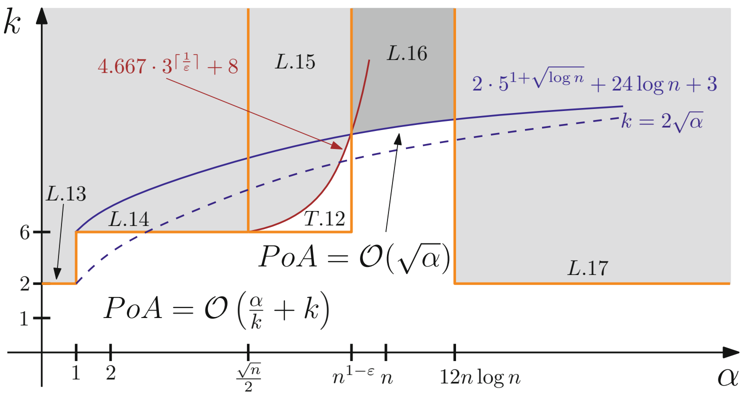

Fig. 1. Overview of our results for both parameters k and α. The shaded areas show where NE and k-NE coincide. In the light gray areas we have a constant PoA, whereas in the dark gray area the PoA is bounded by $\mathcal{O}(5^{\sqrt{\log n}} \log n)$. (L=Lemma, T=Theorem) (Color figure online)

4.2 The Price of Anarchy for Tree Networks

We consider the quality of tree networks that are in k-NE. First of all, by Lemma 1, we have that any tree network in k-GE must be in ASE and therefore, as shown in [9,19], has diameter $\mathcal{O}(\log n)$. Together with Theorem 8 this yields the following upper bound.

Corollary 2. *For k-NE tree networks with* $2 \leq k \leq \log n$*:* $\mathrm{PoA} = \mathcal{O}(\log n)$.

Next, we show a non-constant lower bound, which is tight if $k \in \mathcal{O}(1)$ or if $k \in \Omega(\log n)$. The latter is true, since trees in k-NE have diameter $\mathcal{O}(\log n)$ and thus the technique in [10] yields constant PoA. In particular, this bound states that even with the most optimistic model of locality, there still exists a fundamental difference in the efficiency of the local versus the non-local NCG.

Theorem 11. *For k-NE tree networks with* $2 \leq k \leq \log n$*:* $\mathrm{PoA} = \Omega\left(\frac{\log n}{k}\right)$.

Proof. For any $2 \leq k \leq d$ consider a complete binary tree T_d of depth d with root r, where every edge is owned by the incident agent that is closer to the root r of T. Clearly, independent of α no agent in T_d can delete or swap edges to improve. Thus, towards proving that T_d is in k-NE, we only have to choose α high enough so that no agent can improve by buying any number of edges in her k-neighborhood. Any edge-purchase within the k-neighborhood of some agent u can decrease u's distances to other agents by at most $k-1$. Thus, if $\alpha = (k-1)n$, then no agent can buy one single edge to decrease her cost. Moreover, by buying more than one edge, agent u cannot decrease her distance cost by more than $(k-1)n$, which implies that (T_d, α) is in k-NE for $\alpha = (k-1)n$.

Now we consider the social cost ratio of T_d and the spanning star OPT on $n = 2^{d+1} - 1$ agents, which is the social cost minimizing network, since $\alpha \geq 2$. We will use $\mathrm{cost}(T_d) \geq \alpha(n-1) + n\,\mathrm{dist}_r(T_d)$, which is true, since the root r has minimum distance cost among all agents in T_d. Note that $\mathrm{dist}_r(T_d) > \frac{n}{4}\log n$, since r has at least $\frac{n}{2}$ many agents in distance $\frac{\log n}{2}$. Moreover, the spanning star on n vertices has a total distance cost of less than $2n^2$, since it has diameter 2. With $\alpha = (k-1)n$ this yields:

$$\frac{\mathrm{cost}(T_d)}{\mathrm{cost}(\mathrm{OPT})} > \frac{(k-1)\frac{n^2}{4} + \frac{n^2}{4}\log n}{(k-1)n^2 + 2n^2} > \frac{\frac{n^2}{4}((k-1) + \log n)}{n^2(k+1)} \in \Omega\left(\frac{\log n}{k}\right). \quad \square$$

4.3 The Price of Anarchy for Non-tree Networks

We first provide a simple diameter upper bound that holds for any combination of $k \geq 2$ and $\alpha > 0$.

Theorem 12. *Given a k-NE network* (G, α) *for* $k \geq 2$*, then it holds:*

$$\mathrm{diam}(G) \leq \begin{cases} \alpha/(k-1) + k\frac{3}{2} + 1 & k < 2\sqrt{\alpha}, \\ 2\sqrt{\alpha} & k \geq 2\sqrt{\alpha}. \end{cases}$$

In the following, we provide a more involved upper bound for $1 \leq \alpha < n^{1-\varepsilon}$ by modifying an approach by [8]. First, we give three lemmas that lower bound the number of agents in specific sized neighborhoods for k-NE networks. Then, considering the locality parameter k, we look at the maximal neighborhoods for which these lower bounds apply and present an estimation on the network's diameter.

First, we restate Lemma 3 from [8], which holds for any $k \geq 2$ since only 2-local operations are considered.

Lemma 3 (Lemma 3 from [8]). *For any $k \geq 2$ and any k-NE network G with $\alpha \geq 0$ it holds $|N_2(v)| > n/(2\alpha)$ for every agent $v \in V$.*

Lemma 4. *For $k \geq 6$, let G be a k-NE network and $d \leq k/3 - 1$ an integer. If there is a constant $\lambda > 0$ such that $|N_d(u)| > \lambda$ holds for every $u \in V$, then*

(1) either $|N_{2d+3}(u)| > n/2$ for some *agent $u \in V$*
(2) or $|N_{3d+3}(v)| > \lambda \frac{n}{\alpha}$ for every *agent $v \in V$.*

Lemma 5. *For $k \geq 4$, let G be a k-NE network with $\alpha < n/2$ and $d \leq k/2 - 1$ an integer. If there is an agent $u \in V$ with $|N_d(u)| \geq n/2$, then $|N_{2d+1}(u)| \geq n$.*

Theorem 13. *For $k \geq 6$, $n \geq 4$, and $1 \leq \alpha \leq n^{1-\varepsilon}$ with $\varepsilon \geq 1/\log n$, the maximal diameter of any k-NE network is $\mathcal{O}\left(n^{1-\varepsilon(\log(k-3)-1)}\right)$.*

Proof. Let G be a k-NE network. We define a sequence $(a_i)_{i \in \mathbb{N}}$ by $a_1 := 2$ and for any $i \geq 2$ with $a_i := 3a_{i-1} + 3$. We want to apply Lemma 4 iteratively with $\lambda_i := (n/\alpha)^i/2$. Lemma 3 ensures that with $|N_2(v)| > n/(2\alpha) = \lambda_1$ for all $v \in V$, we have a start for this.

Let m be the highest sequence index with $a_m \leq k/4 - 3$. If there is a $j \leq m$ such that case (1) of Lemma 4 applies, then there is an agent $u \in V$ with $|N_{2a_m+3}(u)| > n/2$. With $\alpha \leq n^{1-\varepsilon} < n/2$ we get with Lemma 5 that $|N_{4a_m+7}| \geq n$ holds and hence the diameter is at most $a_m < k$. Else, case (2) applies for all $i \leq m$ and we know that for every $v \in V$ it holds $|N_{3a_m+3}(v)| > (n/\alpha)^{m-1}/2$. Using $a_i = \frac{7}{6}3^i - 3/2$ and $a_m \leq k/4 - 3$, we get $m \geq \log(k-3) - 1$.

Let D be the diameter of G and p a longest shortest path. We define a set C by selecting the first agent of p as c_1 and than along the path selecting every further agent with distance of $2k$ to the last selected agent. Now consider the operation of c_1 buying an edge to a agent at distance k in the direction of c_2. Using $k \geq 3a_m + 3$, $|N_{3a_m+3}(c)| > (n/\alpha)^{m-1}/2$ for all $c \in C$ and that G forms an equilibrium:

$$\alpha \geq (k-1)(|C|-1)(n/\alpha)^{\log(k-3)-2}/2 \geq \frac{k-1}{2}\left(\frac{D}{2k} - 1\right)(n/\alpha)^{\log(k-3)-2}$$

This gives: $D \leq \frac{4k}{k-1}\alpha\left(\frac{\alpha}{n}\right)^{\log(k-3)-2} + 2k \leq 5n^{1-\varepsilon(\log(k-3)-1)} + 2k$. By using Theorem 10 from [8], we get the claimed diameter upper bound for any $k \geq 6$. □

5 Conclusion and Open Problems

Our results show a major gap in terms of social efficiency between the worst case locality model by Bilò et al. [5] and our more optimistic locality assumption. This gap is to be expected since agents in our model can base their decisions on more information. Interestingly, since most of our upper bounds on the PoA are close to the non-local model, this shows that the natural approach of probing different local strategies is quite convenient for creating socially efficient networks. On the other hand, the non-constant lower bound on the PoA and our negative results concerning the approximation of non-local strategies by local strategies show that the locality constraint does have a significant impact on the game, even in the most optimistic setting. Moreover, our negative results on the hardness of computing best responses and on the dynamic properties show that these problems seem to be intrinsically hard and mostly independent of locality assumptions.

The exact choice of the distance cost function seems to have a strong impact in our model. For the Max-NCG, which was extensively studied by Bilò et al. [5], it is easy to see that a cycle of odd length, where every agent owns exactly one edge, yields even in our model a lower bound of $\Omega(n)$ on the PoA for $k = 2$ and $\alpha > 1$. This seems to be an interesting contrast between the Sum-NCG and the Max-NCG, which should be further explored. It might also be interesting to further study what happens if agents are limited to probe only a certain number of strategies. So far, we only considered the cases of probing all strategies and of probing the quadratic number of greedy strategies. Also, applying our locality approach to other models, in particular those motivated from the economics perspective, like [3,12], seems a natural next step of inquiry.

References

1. Albers, S., Eilts, S., Even-Dar, E., Mansour, Y., Roditty, L.: On nash equilibria for a network creation game. ACM Trans. Econ. Comput. **2**(1), 2 (2014)
2. Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, E., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. SIAM J. Comput. **38**(4), 1602–1623 (2008)
3. Bala, V., Goyal, S.: A noncooperative model of network formation. Econometrica **68**(5), 1181–1229 (2000)
4. Bilò, D., Gualà, L., Leucci, S., Proietti, G.: Network creation games with traceroute-based strategies. In: Halldórsson, M.M. (ed.) SIROCCO 2014. LNCS, vol. 8576, pp. 210–223. Springer, Heidelberg (2014)
5. Bilò, D., Gualà, L., Leucci, S., Proietti, G.: Locality-based network creation games. In: SPAA 2014, pp. 277–286. ACM, New York (2014)
6. Corbo, J., Parkes, D.: The price of selfish behavior in bilateral network formation. In: PODC 2005 Proceedings, pp. 99–107. ACM, New York (2005)
7. Cord-Landwehr, A., Lenzner, P.: Network creation games: Think global - act local. CoRR, abs/1506.02616 (2015)
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., Fazli, M., Mehrabian, A., Sadeghabad, S.S., Safari, M., Saghafian, M., ShokatFadaee, S.: On a bounded budget network creation game. In: SPAA 2011, pp. 207–214. ACM (2011)
10. Fabrikant, A., Luthra, A., Maneva, E., Papadimitriou, C.H., Shenker, S.: On a network creation game. In: PODC 2003 Proceedings, pp. 347–351. ACM (2003)
11. Hoefer, M.: Local matching dynamics in social networks. In: Aceto, L., Henzinger, M., Sgall, J. (eds.) ICALP 2011, Part II. LNCS, vol. 6756, pp. 113–124. Springer, Heidelberg (2011)
12. Jackson, M.O., Wolinsky, A.: A strategic model of social and economic networks. J. Econ. Theory **71**(1), 44–74 (1996)
13. Kawald, B., Lenzner, P.: On dynamics in selfish network creation. In: Proceedings of the 25th ACM Symposium on Parallelism in Algorithms and Architectures, pp. 83–92. ACM (2013)
14. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, p. 404. Springer, Heidelberg (1999)
15. Lenzner, P.: On dynamics in basic network creation games. In: Persiano, G. (ed.) SAGT 2011. LNCS, vol. 6982, pp. 254–265. Springer, Heidelberg (2011)
16. Lenzner, P.: Greedy selfish network creation. In: Goldberg, P.W. (ed.) WINE 2012. LNCS, vol. 7695, pp. 142–155. Springer, Heidelberg (2012)
17. Mamageishvili, A., Mihalák, M., Müller, D.: Tree nash equilibria in the network creation game. In: Bonato, A., Mitzenmacher, M., Prałat, P. (eds.) WAW 2013. LNCS, vol. 8305, pp. 118–129. Springer, Heidelberg (2013)
18. 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)
19. Mihalák, M., Schlegel, J.C.: 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)
20. Monderer, D., Shapley, L.S.: Potential games. Games Econ. Behav. **14**(1), 124–143 (1996)
21. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: Algorithmic Game Theory. Cambridge University Press, New York (2007)
22. Williamson, D.P., Shmoys, D.B.: The Design of Approximation Algorithms. Cambridge University Press, New York (2011)