



Efficiency and Stability in Euclidean Network Design

Wilhelm Friedemann
Hasso Plattner Institute
University of Potsdam
Potsdam, Germany

Tobias Friedrich
Hasso Plattner Institute
University of Potsdam
Potsdam, Germany

Hans Gawendowicz
Hasso Plattner Institute
University of Potsdam
Potsdam, Germany

Pascal Lenzen
Hasso Plattner Institute
University of Potsdam
Potsdam, Germany

Anna Melnichenko
Hasso Plattner Institute
University of Potsdam
Potsdam, Germany

Jannik Peters
Hasso Plattner Institute
University of Potsdam
Potsdam, Germany

Daniel Stephan
Hasso Plattner Institute
University of Potsdam
Potsdam, Germany

Michael Vaichenker
Hasso Plattner Institute
University of Potsdam
Potsdam, Germany

ABSTRACT

Network Design problems typically ask for a minimum cost sub-network from a given host network. This classical point-of-view assumes a central authority enforcing the optimum solution. But how should networks be designed to cope with selfish agents that own parts of the network? In this setting, minimum cost networks may be very unstable in that agents will deviate from a proposed solution if this decreases their individual cost. Hence, designed networks should be both efficient in terms of total cost and stable in terms of the agents' willingness to accept the network.

We study this novel type of Network Design problem by investigating the creation of (β, γ) -networks, that are in β -approximate Nash equilibrium and have a total cost of at most γ times the optimal cost, for the recently proposed Euclidean Generalized Network Creation Game by Bilò et al. [SPAA 2019]. There, n agents corresponding to points in Euclidean space create costly edges among themselves to optimize their centrality in the created network. Our main result is a simple $O(n^2)$ -time algorithm that computes a (β, β) -network with low β for any given set of points. Moreover, on integer grid point sets or random point sets our algorithm achieves a low constant β . Besides these results for the Euclidean model, we discuss a generalization of our algorithm to instances with arbitrary, even non-metric, edge lengths. Moreover, in contrast to these algorithmic results, we show that no such positive results are possible when focusing on either optimal networks, i.e., $(\beta, 1)$ -networks, or perfectly stable networks, i.e., $(1, \gamma)$ -networks, as in both cases NP-hard problems arise, there exist instances with very unstable

optimal networks, and there are instances for perfectly stable networks with high total cost. Along the way, we significantly improve several results from Bilò et al. and we asymptotically resolve their conjecture about the Price of Anarchy by providing a tight bound.

CCS CONCEPTS

• **Theory of computation** → **Network formation; Quality of equilibria; Algorithmic game theory.**

KEYWORDS

Geometric Network Design, Network Creation Games, Algorithmic Game Theory, Price of Anarchy, Approximate Equilibrium

ACM Reference Format:

Wilhelm Friedemann, Tobias Friedrich, Hans Gawendowicz, Pascal Lenzen, Anna Melnichenko, Jannik Peters, Daniel Stephan, and Michael Vaichenker. 2021. Efficiency and Stability in Euclidean Network Design. In *Proceedings of the 33rd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '21), July 6–8, 2021, Virtual Event, USA*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3409964.3461807>

1 INTRODUCTION

Network Design is a classical and rich research area in Operations Research and Theoretical Computer Science. Core questions in the Network Design literature target how to construct networks with favorable properties like low total cost, high robustness, and good usability. Typically these questions have been addressed as combinatorial optimization problems. Many special cases like the Steiner Tree Problem [38], the Optimum Communication Spanning Tree Problem (ND7 in [32]), the creation of Geometric Spanner Networks [50] or variants of the Network Design Problem [37] have been thoroughly studied [31, 34, 44].

However, almost all previous work on these problems simply assumes that a central authority exists that enforces the optimal network structure obtained by combinatorial optimization. This approach is obviously infeasible for networks with no central governing authority. In such networks the cost of maintaining the network is typically distributed among the network participants, and it can happen that the network with the minimum total cost

This work has been partly supported by DFG Project 442003138 "Geometric Selfish Network Creation", DFG Research Training Group 2434 "Facets of Complexity", and the Gerhard C. Starck Stiftung.



This work is licensed under a Creative Commons Attribution International 4.0 License.

SPAA '21, July 6–8, 2021, Virtual Event, USA.
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8070-6/21/07.
<https://doi.org/10.1145/3409964.3461807>

is not stable, i.e., that selfish participants who own parts of the network prefer to restructure their part. This may be favorable for them individually, but not for the whole network, because this might yield a network with significantly higher total cost. Hence, stability corresponds to reaching an equilibrium in the strategic game that models the interaction of the selfish participants. It can happen that the total cost of any equilibrium of such a game is much higher than the total cost of the optimum network.

Hence, ideally, in the realm of Network Design, we aim for networks that are efficient in terms of total cost and, at the same time, are as stable as possible. This naturally corresponds to a bi-criteria optimization problem, i.e., finding networks that are (β, γ) -approximate solutions, where β is the approximation factor for the cost of the individual agent compared to the induced cost by her best possible strategy and γ is the total cost ratio with the minimum possible total cost. Hence in a (β, γ) -network no agent can improve her cost via a strategy change by more than a factor of β , and the total cost is at most γ times the cost of the social optimum state. Extreme cases are $(1, \gamma)$ -networks, whose analysis yields bounds for the Price of Anarchy (PoA) [40] or the Price of Stability (PoS) [7, 22], and $(\beta, 1)$ -networks that indicate how tolerant the agents have to be to accept the social optimum network [2].

In this work we set out to explore the design of (β, γ) -networks for a natural strategic Network Design setting, where a network between nodes that correspond to points in Euclidean space must be established. Each node represents a selfish agent that strives for centrality in the created network and edges have a cost that is proportional to the Euclidean distance between their endpoints.

1.1 Related Work

We will focus our discussion on game-theoretic network formation models. For an overview over classical Network Design from combinatorial optimization, we refer to the surveys by Magnanti and Wong [44] and Gupta and Könemann [34]. Also related to our work are results about geometric spanners, in particular, Euclidean spanners, that recently have drawn much attention [18, 26, 28, 41]. For a good overview over geometric spanners, we refer to the excellent book by Narasimhan and Smid [50].

Game-theoretic models for network formation can be divided into variants of the *Network Design Game* (NDG) as introduced by Anshelevich et al. [7, 8] and variants of the *Network Creation Game* (NCG) as introduced by Fabrikant et al. [27].

In the NDG a weighted host network is fixed, and agents want to connect subsets of nodes, called terminals, from the network. To do this, the agents decide on payments for the edges of the host network. Thus, the agents buy a sub-network of the host network such that all desired terminal connections are established. If agents can buy arbitrary cost-shares of the edges [8], the PoA is n , where n is the number of agents. Moreover, $(3, 1)$ -networks exist, i.e., the cost of the social optimum network can be shared among the agents to achieve a 3-approximate Nash equilibrium. Similar results have been achieved for tailored cost-sharing protocols [20], for the node weighted version [19], and for a version that guarantees connectivity of the formed network [35]. For the NDG with fair cost sharing [7] the PoA is n , but the PoS is H_n , i.e., the n -th harmonic number, and the latter bound is tight for the version on directed host

networks. Also, Albers and Lenzner [2] showed that instances with only $(\Omega(\log n), 1)$ -networks exist if all agents want to connect to the same terminal. For the general case instances with only $(\Omega(n), 1)$ -networks exist. Hoefer and Krysta [36] analyzed a geometric version of the NDG, where agents correspond to points in the Euclidean plane. They find that the PoA is n and that the PoS for two agents with two terminals each is 1. Recently, also a variant with fair cost sharing and topology dependent edge-cost was studied [13].

Much closer to our model is research on the NCG [27], where the agents correspond to nodes of a network and any node can establish undirected links to other agents for the cost of α per link, where $\alpha > 0$ is a fixed parameter. The created network then consists of the union of all links created by the agents. The goal of the agents is to minimize the sum of their cost for creating edges and their average distance to all other agents in the created network, i.e., their closeness centrality [51]. A long line of research [1, 4–6, 10, 23, 27, 45, 47] has established that the PoA of the NCG is constant for almost all $\alpha > 0$ and it is conjectured that this holds for all α [3, 27, 45, 46]. Computing the best possible strategy of an agent in the NCG was shown to be NP-hard [27] and this also holds for many NCG variants [16, 17, 21, 46]. However, restricted variants with efficient best response computation also exist [3, 9, 30, 43]. Regarding the dynamics, it has been shown [39, 42] that many NCG variants do not have the *finite improvement property* (FIP) [48]. Thus, natural convergence protocols, such as iterated best response dynamics, have no convergence guarantee.

There are only a few works that investigate NCG variants in a geometric setting. Eidenbenz et al. [25] consider agents corresponding to points in the Euclidean plane that strategically buy incident edges to create a connected network. Another related geometric game was proposed by Moscibroda et al. [49], where the agents pay a fixed price $\alpha > 0$ for each edge and aim at minimizing their total stretch, where the stretch is the ratio of the shortest path length in the network and the geometric distance. Guylás et al. [33] considered a NCG variant where agents correspond to uniformly sampled points in the hyperbolic plane that strive for maximum navigability. Bilò et al. [11] considered a variant with a dynamically changing underlying geometry. Finally, and closest to our work, Bilò et al. [12] recently introduced the *Generalized Network Creation Game* (GNCG) with a given weighted host network. It generalizes the NCG since edges can have weights, and the cost of an edge is defined as α times its weight, for $\alpha > 0$. Besides the version with arbitrary edge weights, variants with arbitrary metric weights, metric weights defined by a tree metric, and metric weights defined by Euclidean distance were proposed. The latter of which is the most natural setting for the creation of communication networks and will be our main focus. The authors of [12] prove an upper bound on the PoA of $\left(\frac{\alpha+2}{2}\right)^2$ for the most general version and give a better PoA upper bound of $\frac{\alpha+2}{2}$ for metric weights that is conjectured to hold for the general case as well. For the version with weights defined by a tree metric the bound of $\frac{\alpha+2}{2}$ was shown to be tight. For the Euclidean version using the 2-norm, the setting that is at the heart of our paper, the shown results are far from being tight: a constant lower bound on the PoA that is slightly above 3 was provided while the upper bound is $\frac{\alpha+2}{2}$. Furthermore, it was shown that the PoA lower bound using the 1-norm approaches $\frac{\alpha+2}{2}$.

if the number of dimensions tends to infinity. Non-trivial bounds on the PoS are given only for the tree metric variant where the PoS is 1. Moreover, it was shown that computing the best possible strategy for an agent is NP-hard for all versions. Regarding the FIP, the status for the Euclidean version with p -norm for $p \geq 2$ was left open. For all the other variants, it was shown that the FIP does not hold. Finally, concerning the existence of pure Nash equilibria, it is shown that equilibria exist for weights induced by a tree metric and for suitable α values if all weights are either 1 or 2. For the other cases with metric weights, only the existence of $3(\alpha + 1)$ -approximate equilibria is claimed.

Besides NCG variants where agents minimize edge costs and distance costs, also models where the cost of an agent depends on her local clustering coefficient [14] or on the edge-connectivity of the created network [24] have been proposed.

1.2 Model and Notation

We consider the recently introduced model by Bilò et al. [12] for the distributed creation of a network by selfish agents with an underlying geometry, called the *Generalized Network Creation Game* (GNCG). Our main focus will be on the natural special case of the GNCG where the underlying geometry is Euclidean, called the *Euclidean Generalized Network Creation Game* (\mathbb{R}^d -GNCG). In this model, a set of n points $P = \{p_1, \dots, p_n\}$ in the d -dimensional Euclidean space \mathbb{R}^d is considered, where each point $p_i = (p_{i,1}, \dots, p_{i,d})$, for $1 \leq i \leq n$, corresponds to a selfish agent. As we will see, all agents will jointly form an undirected weighted network $G = (P, E)$ among themselves, where P is the set of nodes and E is the edge set of G . Hence, we will use point, agent and node interchangeably. Moreover, we will use the shorthand uv for denoting the edge $\{u, v\} \in E$. The weight of an edge uv is denoted by $\|u, v\|$ and it is defined as the metric induced by the 2-norm¹, i.e., for any two points $u = (u_1, \dots, u_d), v = (v_1, \dots, v_d)$, we have $\|u, v\| := \|u - v\|_2 = \sqrt{\left(\sum_{k=1}^d (u_k - v_k)^2\right)}$. We will call $\|u, v\|$ the *length* of the edge uv . Let $w_{max} := \max_{p_i, p_j \in P} \|p_i, p_j\|$ and $w_{min} := \min_{p_i, p_j \in P, p_i \neq p_j} \|p_i, p_j\|$ denote the longest and shortest distance between any two points in the point set P and let $r := \frac{w_{max}}{w_{min}}$ be the *aspect ratio* of P . Furthermore for any network $G = (P, E)$ and any $u, v \in P$, let $\pi_G(u, v) \subseteq E$ be a shortest path from u to v in G .

Each agent strategically decides which incident edges to buy in order to minimize her total distance to all other agents in the created network. More precisely, the strategy of agent u , denoted as S_u , is a subset of nodes in $P \setminus \{u\}$ to which agent u wants to create an undirected edge. For each $v \in S_u$, we say that u is the *owner* of the undirected edge uv . Edges are costly and we assume that the edge cost is proportional to the Euclidean distance between the endpoints of the respective edge. Formally, the price of an edge uv that must be paid by its owner is equal to $\alpha \cdot \|u, v\|$, where $\alpha > 0$ is a fixed parameter of the game.

Any vector $s = (S_{p_1}, \dots, S_{p_n})$, is called a strategy profile. Every strategy profile s uniquely determines a created network $G(s) =$

$(P, E(s))$, where $E(s) = \{p_i p_j \mid p_i \in P, p_j \in S_{p_i}\}$.² We will omit the reference to s when it is clear from the context.

Let $d_G(u, v)$ denote the distance between two nodes u, v in a network G , where $d_G(u, v)$ is the sum of the edge lengths of the edges in the shortest u - v path in G , i.e., $d_G(u, v) := \sum_{x, y \in \pi_G(u, v)} \|x, y\|$. If there is no u - v path in G , then $d_G(u, v) := +\infty$. We use $d_G(u, U)$ to denote the sum of distances from u to all nodes in $U \subseteq P$ in G , and we use $\|u, U\|$ to denote the sum of the lengths of the edges between u and U . We call $d_G(u, P)$ the *distance cost* and $\alpha \cdot \|u, S_u\|$ the *edge cost* of agent u . Each agent u aims at minimizing her cost $cost(u, G(s))$, that is the sum of the agent's edge cost and her distance cost:

$$cost(u, G(s)) := \alpha \cdot \|u, S_u\| + d_{G(s)}(u, P).$$

Note that the parameter α expresses the agents' relative importance of edge costs versus distance costs.

We measure the *efficiency* of a network $G(s)$ by its *social cost* $SC(G(s)) := \sum_{u \in P} cost(u, G(s))$. The strategy profile s^* that minimizes $SC(G(s^*))$ for given points P is called the *social optimum*. We refer to the network $G(s^*)$ also as the *social optimum network for the point set P* , denoted as OPT_P .

An *improving move* for an agent u is a strategy change, that decreases her cost. Agent u plays her *best response* if agent u has no improving move. A strategy profile s is a *pure Nash equilibrium* (NE) if all agents play a best response in $G(s)$. We say that s is a β -*approximate NE* (β -NE), if no agent can change her strategy to improve her cost by more than a factor of β . Moreover, we call strategy profile s a (β, γ) -NE if $SC(G(s)) \leq \gamma \cdot SC(OPT_P)$ and it is a β -NE. Strategy profiles induce networks, thus we call $G(s)$ a (β, γ) -network if s is a (β, γ) -NE. Moreover, if the edge ownership of some network G is not specified, then we assume arbitrary edge ownership when calling G a (β, γ) -network.

We measure the loss of efficiency due to selfishness via the *Price of Anarchy* (PoA) [40] and the *Price of Stability* (PoS) [7, 22]. Let $worst_P$ (respectively $best_P$) be the highest (respectively the lowest) social cost of any NE on the point set P and let \mathcal{P} be the set of all possible finite point sets in \mathbb{R}^d . Then the PoA is defined as $\sup_{P \in \mathcal{P}} \frac{worst_P}{SC(OPT_P)}$ and the PoS is $\sup_{P \in \mathcal{P}} \frac{best_P}{SC(OPT_P)}$.

1.3 Our Contribution

We explore the design of networks that at the same time should be efficient in terms of total cost and stable in terms of local changes to the network infrastructure by selfish agents. This very natural focus on approximating both efficiency and stability seems to be novel in the literature on variants of Network Creation Games. Moreover, also in the wider Network Design literature, we are only aware of the works on Network Design Games with cost-sharing on the edges [2, 8, 19, 20, 35] that take a similar point-of-view.

We study the creation of (β, γ) -networks for the \mathbb{R}^d -GNCG [12] using the 2-norm. Such networks are β -approximate Nash equilibrium and at the same time have a total cost that is at most γ times the optimal total cost. See Table 1 for a result overview.

Our main result is a simple $\mathcal{O}(n^2)$ algorithm that computes a (β, β) -network for $\alpha \leq n^x$ with $\beta \in \mathcal{O}\left(\alpha^{1-\frac{1}{2x}} + 1\right)$ for $x \geq 1$ and

¹Our results can be adapted to any p -norm. We focus on the 2-norm for the sake of presentation.

²Note that if both $p_i \in S_{p_j}$ and $p_j \in S_{p_i}$ holds, then both agents p_i and p_j would pay $\alpha \|u, v\|$ for the edge uv . As we will see, this cannot happen in any equilibrium.

Table 1: Overview of our results on (β, γ) -networks for the \mathbb{R}^d -GNCG using the 2-norm.

Socially optimal	Apx. optimal and apx. stable	Perfectly stable
$(\beta, 1)$ -networks	(β, γ) -networks with $\beta, \gamma > 1$	$(1, \gamma)$ -networks
← efficiency		stability →
Social optimum computation NP-hard for metric instances [Thm. 2.2]	Simple $O(n^2)$ algorithm for computing (β, β) -networks [Thm. 3.7]:	Stability check via computing best responses NP-hard [12]
Instances with only $(\Omega(\sqrt{\alpha}), 1)$ -networks exist [Thm. 2.1]	$\beta \in O(1)$ for $\alpha \leq \sqrt[3]{n}$	Existence of stable networks open, no FIP [Thm. 3.1]
	$\beta \in O(\alpha^{\frac{2x-1}{2x}} + 1)$ for $\alpha \leq n^x, x \geq 1$	PoS > 1 [Thm. 4.3]
	$\beta \in O(\alpha^{\frac{3x-1}{4x}} + 1)$ for $\alpha \leq n^x, x \leq 1$	PoA $\in \Theta(\alpha), d \rightarrow \infty$ [Thm. 4.1]
	MST is $(n-1, n-1)$ -network [Thm. 3.9]	PoA $\in \Omega(\alpha^{\frac{2}{3}}), d \geq 1$ [Thm. 4.2]
	$(1 + \varepsilon, 1 + \varepsilon)$ -network, a.a.s., for random points in $[0, 1]^2, \alpha \in o(n)$ [Thm. 3.12]	GNCG PoA $\in \Theta(\alpha)$ [Cor. 5.5]
	$(2d, 2d)$ -networks for integer grid point sets in \mathbb{R}^d [Thm. 3.13]	

$\beta \in O\left(\alpha^{\frac{3x-1}{4x}} + 1\right)$ for $0 < x < 1$. See Figure 3 for a graphical illustration of these bounds. For $\alpha \leq \sqrt[3]{n}$, i.e., if edges are cheap or for large networks, this implies that our algorithm constructs a $(O(1), O(1))$ -network. We further demonstrate the power of our algorithm by investigating special instance types: grid point sets and uniform random point sets. For them we obtain particularly low constant values for β and γ . Additionally, we also provide (β, γ) -networks with an even simpler construction: we show that any MST on the point set P is a $(n-1, n-1)$ -network. Using the better outcome of either the MST or the network obtained by our algorithm then yields a $(O(\alpha^{\frac{2}{3}}), O(\alpha^{\frac{2}{3}}))$ -network for arbitrary α . Moreover, we show that the complete network on P is a $(\alpha + 1, \frac{\alpha}{2} + 1)$ -network. This even holds for the GNCG with arbitrary, even non-metric, edge lengths and hence proves that $(\alpha + 1)$ -approximate NEs always exist. This is an improvement over the claimed bound of $3(\alpha + 1)$ in [12] and it resolves an open problem from that paper as no bound for the GNCG was provided.

In contrast to these positive results, we provide negative results for the extreme cases of (β, γ) -networks, i.e., $(\beta, 1)$ -networks with the optimal total cost and $(1, \gamma)$ -networks with perfect stability. For both of these extremes, our results indicate that, unless P=NP, such networks cannot be computed efficiently. In particular, we show that computing a $(\beta, 1)$ -network is NP-hard for a generalization of the model with metric edge lengths and that $(1, \gamma)$ -networks cannot be found via improving response dynamics since the finite improvement property does not hold. The latter was left open in [12]. Also, there it was already shown that computing a best possible strategy is NP-hard, which indicates that also deciding stability is a hard problem. Moreover, we provide an instance having very unstable networks with optimal total cost, i.e., only $(\Omega(\sqrt{\alpha}), 1)$ -networks exist, and this instance also shows that the PoS is larger than 1, i.e., $(1, 1)$ -networks cannot exist in general. Moreover, we

show that $(1, \gamma)$ -networks with $\gamma \in \Omega(\alpha^{\frac{2}{3}})$ exist for all $d \geq 1$, i.e., a $\Omega(\alpha^{\frac{2}{3}})$ lower bound for the PoA that significantly improves over the known constant lower bound and which is close to the known $O(\alpha)$ upper bound. Additionally, we show that for d tending to infinity the PoA is in $\Theta(\alpha)$. Besides these PoA results for the Euclidean version, we also prove an upper bound of $2(\alpha + 1)$ on the PoA for the GNCG with arbitrary, even non-metric, edge lengths. This asymptotically matches the lower bound of $\frac{\alpha+2}{2}$, and it proves a conjecture from [12] up to constant factors.

We refer to [29] for all details which were omitted due to space constraints.

2 SOCIAL OPTIMUM

We show that minimum cost networks can be rather unstable.

THEOREM 2.1. *There exists a set of points where in the unique social optimum network an agent can improve by a factor of at least $\frac{\sqrt{\alpha}}{3}$, i.e., it is a $(\frac{\sqrt{\alpha}}{3}, 1)$ -network. (See Figure 1 (left).)*

PROOFSKETCH. For simplification we allow co-located points. Note that the result still holds asymptotically without co-location, since we can place all co-located points arbitrarily close together.

Consider the three corners of an equilateral triangle with side length 1, see Figure 1 (left). We place $\frac{n}{3}$ points on each corner. Note that it is without cost to buy all the edges of length 0 and that two length-1-edges need to be bought or the network will be disconnected. We observe that buying all three length-1-edges gives us a social optimum if $2 \cdot 2 \cdot (\frac{n}{3})^2 > \alpha + 2 \cdot (\frac{n}{3})^2$, which is equivalent to $\alpha < 2 \cdot (\frac{n}{3})^2$. Given this, we set $n = 3 \lfloor \sqrt{\alpha} + 1 \rfloor$ and consider the social optimum where every agent buys at most one length-1-edge. Finally, we compute the improvement factor for one

agent by selling a length-1-edge resulting in

$$\frac{\alpha + 2\frac{n}{3}}{3\frac{n}{3}} \geq \frac{\alpha + 2\sqrt{\alpha}}{3\sqrt{\alpha} + 3} \geq \frac{\sqrt{\alpha}}{3}. \quad \square$$

Next, we show that computing the social optimum network is NP-hard for any fixed α in the more general metric version of the GNCG (M-GNCG) [12]. There, agents are nodes of a given complete weighted host network $H = (V, E(H))$ with edge weights $w: V \times V \rightarrow \mathbb{R}_+$ satisfying the triangle inequality. Hence, the edge price of $uv \in E(H)$ is $\alpha \cdot w(u, v)$.

THEOREM 2.2. *For any $\alpha > 0$, computing a social optimum in the M-GNCG is NP-hard.*

Clearly, hardness for a problem on metric instances does not imply hardness for Euclidean instances. However, given that many variants of minimum weight Euclidean t-spanner problems are also NP-hard (e.g., see [15] and the references therein), and since these problems seem to be very close to computing a social optimum network, we think it could be possible to either adapt our reduction for the metric case via suitable gadgets or to reduce from hard minimum weight t-spanner problems directly. Therefore, we conjecture the following:

CONJECTURE 1. *Computing a social optimum network in the \mathbb{R}^d -GNCG is NP-hard.*

3 NASH EQUILIBRIUM

We investigate the existence and the computation of (approximate) NEs. We show that iteratively playing best responses is not guaranteed to lead to a NE and give sufficient conditions for their existence. Then follows the main result of our paper: a simple and efficient algorithm for computing a $(\mathcal{O}(\alpha^{2/3}), \mathcal{O}(\alpha^{2/3}))$ -network.

3.1 Existence

An obvious way of finding a NE would be to iteratively play best responses. However, the following theorem shows that doing so does not necessarily lead to a NE.

THEOREM 3.1. *The \mathbb{R}^d -GNCG with $d \geq 2$ does not have the finite improvement property.*

PROOF SKETCH. We prove this statement by providing a best response cycle, i.e., a cyclic sequence of networks obtained by iterative strategy changes to best responses, in \mathbb{R}^2 for $\alpha = 1$. See Figure 1 (right) for illustrations of the steps of the cycle. \square

As a first step towards showing that in some cases Nash equilibria exist, we show that, if α is large enough, any center sponsored star, where the center buys all edges, is a NE.

LEMMA 3.2. *Let T be the star of a point set P centered at the node n_0 , such that n_0 owns all edges. If $\alpha \geq \max_{u,v \in P, u \neq v} \frac{\|u, n_0\| + \|n_0, v\|}{\|u, v\|} - 1$, then T is a NE.*

PROOF. Since n_0 owns all edges, she does not want to buy, sell, or swap any additional edges. Therefore we only need to check, whether any of the non-center nodes wants to buy an edge.

We consider the case, where node v buys an edge towards node u . Due to the triangle inequality, this would only improve the distance




Figure 1: Left: Sketch of the optimum network from Theorem 2.1. In each corner there is a node-cluster of size $\frac{n}{3}$ arbitrarily close together. The improving move removes the dotted orange length-1-edge. Right: Best response cycle for the \mathbb{R}^d -GNCG for $d = 2$ and $\alpha = 1$. In each step the marked agent may buy edges (dashed green line) and/or sell edges (dotted orange line). The network after step 3 is the same as the initial network, but mirrored. Edge directions indicate ownership. Edges point away from their owners.

towards u and it would not change any of the other distances. However, since, by assumption, $(\alpha + 1)\|u, v\| \geq \|u, n_0\| + \|n_0, v\|$ buying this edge does not decrease the cost. Therefore T is a NE. \square

COROLLARY 3.3. *If $\alpha \geq 2r - 1$, for aspect ratio r , then every center sponsored star is a NE.*

Using this corollary we can now show that a uniform random point set asymptotically almost surely has a Nash equilibrium if α is asymptotically larger than n .

THEOREM 3.4. *Let $P_n = \{v_1, \dots, v_n\} \subseteq [0, 1] \times [0, 1]$ be n points chosen uniformly at random and let $(\alpha_n)_{n \in \mathbb{N}}$ be a sequence of positive real numbers. If $\alpha_n \in \omega(n)$, then P_n asymptotically almost surely (a.a.s.) has a NE for any $\alpha \geq \alpha_n$.*

PROOF. By Corollary 3.3 it suffices to show that the aspect ratio of P_n is upper bounded by $\frac{\alpha+1}{2}$. Since the maximum possible distance is $\sqrt{2}$, it is enough to show that the closest pair of points are at least $\frac{2\sqrt{2}}{\alpha+1} =: d_\alpha$ apart. For any $i = 1, \dots, n$ we define the random variables

$$X_i = \begin{cases} 1, & \text{if } \min_{v_j \in P_n \setminus \{v_i\}} \|v_i, v_j\| \leq d_\alpha \\ 0, & \text{otherwise.} \end{cases}$$

Also, let $X = \sum_{i=1}^n X_i$. We thus want to show that a.a.s. $X = 0$.

First, we observe that the probability of a single point falling into the d_α neighbourhood of v_i is at most πd_α^2 . Thus, $\mathbb{P}[X_i = 0] \geq (1 - \pi d_\alpha^2)^n$ and $\mathbb{E}[X_i] = \mathbb{P}[X_i = 1] \leq 1 - (1 - \pi d_\alpha^2)^n$.

Now we can apply Markov's and Bernoulli's inequalities to bound the probability:

$$\begin{aligned} \mathbb{P}[X \neq 0] &= \mathbb{P}[X \geq 1] \leq \mathbb{E}[X] \leq n(1 - (1 - \pi d_\alpha^2)^n) \\ &\leq n^2 \pi d_\alpha^2 = \frac{8\pi n^2}{(\alpha + 1)^2}. \end{aligned}$$

Since $\alpha \geq \alpha_n$ and $\alpha_n \in \omega(n)$, this probability approaches zero as n tends to infinity. Therefore, P_n asymptotically almost surely has a Nash equilibrium. \square

3.2 Approximation

We show how to construct (β, γ) -networks for different values of α . First, we prove a general upper bound on the approximation factor.

THEOREM 3.5. *Let P be a set of points in \mathbb{R}^d . Any complete network $K = (P, P \times P)$ is a $(\alpha + 1, \frac{\alpha}{2} + 1)$ -network in the \mathbb{R}^d -GNCG.*

PROOF. Since K is a complete network, every agent u can improve its strategy only by deleting its edges. Let K' be a network obtained after a strategy change. Since the deletion of edges increases the distance cost, $\text{cost}(u, K') \geq d_{K'}(u, P) \geq d_K(u, P)$. In the worst case, u owns all its incident edges in K . Hence, $\text{cost}(u, K) \leq \alpha \cdot d_K(u, P) + d_K(u, P)$, and we get

$$\frac{\text{cost}(u, K)}{\text{cost}(u, K')} \leq \frac{\alpha \cdot d_K(u, P) + d_K(u, P)}{d_K(u, P)} = \alpha + 1.$$

In a similar way we can prove that the social cost of K is at most $(\frac{1}{2}\alpha + 1)$ times the social cost of OPT_P . By the triangle inequality, for any edge uv in K , $\|u, v\| \leq d_{OPT_P}(u, v)$. Hence, $\text{SC}(K) = \frac{1}{2}\alpha \sum_{u, v \in V} \|u, v\| + \sum_{u, v \in V} \|u, v\| \leq (\frac{1}{2}\alpha + 1) \sum_{u \in P} d_{OPT_P}(u, P)$, while the social cost of the social optimum is at least its distance cost $\sum_{u \in P} d_{OPT_P}(u, P)$. Therefore, $\text{SC}(K) \leq (\frac{1}{2}\alpha + 1) \text{SC}(OPT_P)$. \square

In the following we show that it is possible to construct a (β, γ) -network with $\beta, \gamma \in o(\alpha)$. We call S a k -degree t -spanner on P if all its nodes have degree at most k and for any two points $p_1, p_2 \in P$, $d_S(p_1, p_2) \leq t \cdot \|p_1, p_2\|$.³ Consider Algorithm 1. The idea of this algorithm is simple: if P has a large cluster of closely located points, a spanner for the cluster points (set B_v) and all close points (set C_v) is built. The rest of the points is connected with the shortest edges to C_v . (See Figure 2 (left).) If the set of points P is sparsely distributed, i.e., there is no cluster, a spanner for the entire set of points is constructed. (See Figure 2 (right).)

We now prove our main results, i.e., the approximation bounds achieved by constructing the network via Algorithm 1.

THEOREM 3.6. *Let $b \in \mathbb{R}_{\geq 1}, 0 \leq c \leq n - 1$, and let $k \in \mathbb{N}, t \in \mathbb{R}_{>1}$ be parameters such that we can construct a k -degree t -spanner for P . Algorithm 1 computes a β -NE with*

$$\beta = \max \left\{ \frac{kb}{c} \alpha + t, \frac{4k}{b} \alpha + 2t + 1, \frac{2\alpha}{n-c} + 2, \frac{4c(b+2t)}{n-c} + 6t \right\}.$$

PROOF. We divide the proof into two parts corresponding to each if-case of the algorithm. First, we prove that if all nodes have at least c nodes at distance of at least $\frac{w_{max}}{b}$, i.e., for all $v \in P$, $|P \setminus B_v| \geq c$, then the k -degree t -spanner G is a $(\frac{kb\alpha}{c} + t)$ -NE.

Consider an agent $u \in P$ playing a strategy S_u in G . We need to evaluate the maximal improvement of the cost function that can be

³ Algorithm 1 and all the following results hold for a more general setting. We call S a k -distributable t -spanner if it is possible to assign all edges of S to agents such that each agent owns at most k edges. Then we can generalize Algorithm 1 by constructing a k -distributable t -spanner in step 4 and 9.

Algorithm 1: $O(n^2)$ algorithm for computing a (β, β) -network in the \mathbb{R}^d -GNCG.

input: n points P in \mathbb{R}^d , parameters
 $k \in \mathbb{N}, t \in \mathbb{R}_{>1}, b \in \mathbb{R}_{\geq 1}, 0 \leq c \leq n - 1$

- 1 **for** $v \in P$ **do**
- 2 $B_v \leftarrow \{u \in P \mid \|u, v\| \leq \frac{w_{max}}{b}\}$ and
 $C_v \leftarrow \{u \in P \mid \|u, v\| \leq 2 \frac{w_{max}}{b}\};$
- 3 **if** there is a node $v \in P$ with $|P \setminus B_v| < c$ **then**
- 4 construct a k -degree t -spanner G on C_v ;
- 5 assign the edge ownership such that each agent owns at most k edges in G ;
- 6 **for** $u \in P \setminus C_v$ **do**
- 7 Find node $u' \in C_v$ closest to u and add u' to the strategy of agent u ;
- 8 **else**
- 9 construct a k -degree t -spanner G on P ;
- 10 assign the edge ownership such that each agent owns at most k edges in G ;




Figure 2: Illustration of the network computed by Algorithm 1. Left: The network if there exists a node v with at most c points with distance at least $\frac{w_{max}}{b}$. Right: a bounded degree spanner produced by the algorithm in case all points are sparsely distributed. Edges point away from their owners. Undirected edges can be assigned arbitrarily as long as every agent owns at most k edges.

made by u after changing its strategy from S_u to S'_u . Let $G' = (P, E')$ be the network obtained after the improving move. Since G is a t -spanner and u owns at most k edges of weight w_{max} , the cost of agent u before the move is

$$\begin{aligned} \text{cost}(u, G) &= \alpha \cdot \|u, S_u\| + d_G(u, P) \leq \alpha k \cdot w_{max} + t \cdot \|u, P\| \\ &\leq \alpha k \cdot w_{max} + t \cdot d_{G'}(u, P). \end{aligned}$$

By construction, u has at least c nodes at distance at least $\frac{w_{max}}{b}$. Hence, $\frac{b}{c \cdot w_{max}} d_{G'}(u, P) \geq 1$. Combining with the inequality above, we get $\text{cost}(u, G) \leq \frac{\alpha k b}{c} \cdot d_{G'}(u, P) + t \cdot d_{G'}(u, P)$. With this we can

evaluate the maximal improvement made by u :

$$\begin{aligned} \frac{\text{cost}(u, G)}{\text{cost}(u, G')} &\leq \frac{\frac{\alpha kb}{c} \cdot d_{G'}(u, P) + t \cdot d_{G'}(u, P)}{\alpha \cdot \|u, S'_u\| + d_{G'}(u, P)} \\ &\leq \frac{\left(\frac{\alpha kb}{c} + t\right) \cdot d_{G'}(u, P)}{d_{G'}(u, P)} = \frac{\alpha kb}{c} + t. \end{aligned} \quad (1)$$

Now we analyze the case when there is a point $v \in P$ such that $|P \setminus B_v| < c$. In this situation, the algorithm computes a network $G = (P, E)$ that contains a spanner on C_v with attached leaf nodes from $P \setminus C_v$ (see Figure 2 (left) for an illustration). In the following part of the proof we show that the improvement factor of every agent $u \in P$ is bounded as well. We distinguish two cases depending on whether $u \in C_v$ or $u \in P \setminus C_v$.

If $u \in C_v$, we observe that, by construction, u owns at most k edges in G . Let S_u be the strategy of u in G , let S'_u be the new improving strategy, and let G' be the network obtained after the improving move. Since u only buys edges to the nodes from C_v in G , each of her edges has length of at most $\frac{4w_{max}}{b}$. Then the total edge cost of u is at most $4\alpha k \frac{w_{max}}{b}$. Note that $d_{G'}(u, P) \geq \sum_{w \in P} \|u, w\| \geq w_{max}$. Thus, the edge cost of u is at most $\frac{4\alpha k}{b} \cdot d_{G'}(u, P)$.

The distance cost for u is $d_G(u, P) = d_G(u, C_v) + d_G(u, P \setminus C_v)$. Clearly, $d_G(u, C_v) \leq t \cdot \|u, C_v\| \leq t \cdot d_{G'}(u, C_v)$. To analyze the second term $d_G(u, P \setminus C_v)$, consider a node $x \in P \setminus C_v$. Let $y \in C_v$ be a node to which x buys an edge. The existence of y follows from the construction of G . By the triangle inequality, $d_G(u, x) \leq d_G(u, y) + \|y, x\|$. Since G is a t -spanner for C_v , this implies that $d_G(u, y) \leq t \cdot \|u, y\| \leq t(\|u, x\| + \|y, x\|) \leq 2t \cdot \|u, x\|$, where the last inequality holds since, by construction, $\|x, y\| \leq \|x, u'\|$ for any $u' \in C_v$. Therefore, $d_G(u, x) \leq 2t \cdot \|u, x\| + \|y, x\| \leq (2t+1) \cdot \|u, x\|$. Finally, we can evaluate the maximum improvement that can be made by agent u :

$$\begin{aligned} \frac{\text{cost}(u, G)}{\text{cost}(u, G')} &\leq \frac{\alpha \|u, S_u\| + d_G(u, P)}{\alpha \|u, S'_u\| + d_{G'}(u, P)} \\ &\leq \frac{\frac{4\alpha k}{b} d_{G'}(u, P) + (2t+1)d_{G'}(u, P)}{d_{G'}(u, P)} \\ &= \frac{4\alpha k}{b} + 2t + 1. \end{aligned} \quad (2)$$

In case $u \in P \setminus C_v$, consider the vertex $y \in C_v$ to which u buys its only edge. Since the edge is the only one owned by u , the agent's edge cost in G is $\alpha \cdot \|u, y\|$. Next we evaluate the distance cost. Every path from u to a node $x \in P$ goes over y in G . If $x \in B_v$, by the triangle inequality and since G is a t -spanner on C_v ,

$$\begin{aligned} d_G(u, x) &= \|u, y\| + d_G(y, x) \leq \|u, y\| + t \cdot \|y, x\| \\ &\leq \|u, y\| + t \cdot (\|y, v\| + \|v, x\|) \leq \|u, y\| + 3t \frac{w_{max}}{b}. \end{aligned}$$

If $x \notin B_v$, the shortest u - x path goes over the t -spanner and contains at most two edges to nodes outside of C_v . Since the length of any shortest path in C_v is at most $2 \cdot \frac{2w_{max}}{b} \cdot t$, then $d_G(u, x) \leq 2w_{max} + 4t \cdot \frac{w_{max}}{b}$. So, we can evaluate the maximum improvement

for u as follows:

$$\begin{aligned} \frac{\text{cost}(u, G)}{\text{cost}(u, G')} &\leq \frac{\alpha \cdot \|u, y\| + d_G(u, B_v) + d_G(u, P \setminus B_v)}{\alpha \cdot \|u, S'_u\| + d_{G'}(u, P)} \\ &\leq \frac{\alpha \|u, y\| + |B_v| \left(\|u, y\| + \frac{3tw_{max}}{b} \right) + |P \setminus B_v| \left(2w_{max} + \frac{4tw_{max}}{b} \right)}{d_{G'}(u, B_v)} \\ &\leq \frac{\alpha \cdot \|u, y\| + |B_v| \left(\|u, y\| + \frac{3tw_{max}}{b} \right) + |P \setminus B_v| \left(2w_{max} + \frac{4tw_{max}}{b} \right)}{\frac{1}{2} \|u, y\| \cdot |B_v| + \frac{w_{max}}{2b} \cdot |B_v|} \end{aligned} \quad (3)$$

Since $|B_v| > n - c$, we obtain $\frac{(\alpha + |B_v|) \|u, y\|}{\frac{1}{2} \|u, y\| |B_v|} \leq \frac{2\alpha}{n-c} + 2$. For the remaining part we get:

$$\begin{aligned} \frac{w_{max} \left(|P \setminus B_v| \left(2 + \frac{4t}{b} \right) + |B_v| \frac{3t}{b} \right)}{\frac{w_{max}}{2b} \cdot |B_v|} &\leq \frac{\left(2 + \frac{4t}{b} \right) |P \setminus B_v|}{\frac{1}{2b} \cdot |B_v|} + 6t \\ &\leq \frac{4c(b+2t)}{n-c} + 6t. \end{aligned} \quad (4)$$

Combining the two inequalities above we obtain⁴ an upper bound for Inequality (3) equal to $\max \left\{ \frac{2\alpha}{n-c} + 2, \frac{4c(b+2t)}{n-c} + 6t \right\}$. Together with Inequality (1) and Inequality (2) we obtain the upper bound of

$$\max \left\{ \frac{kb}{c} \alpha + t, \frac{4k}{b} \alpha + 2t + 1, \frac{2\alpha}{n-c} + 2, \frac{4c(b+2t)}{n-c} + 6t \right\}. \quad \square$$

The next result shows that the proof of Theorem 3.6 also provides an upper bound for the social cost of the network computed by Algorithm 1. This yields a (β, β) -network.

THEOREM 3.7. *Let $b \in \mathbb{R}_{\geq 1}, k, t \in \mathbb{R}_{> 1}, c \in \{0, \dots, n-1\}$. Algorithm 1 computes a (β, β) -network with*

$$\beta = \max \left\{ \frac{kb}{c} \alpha + t, \frac{4k}{b} \alpha + 2t + 1, \frac{2\alpha}{n-c} + 2, \frac{4c(b+2t)}{n-c} + 6t \right\}.$$

PROOF. Let $G = G(s)$ be the network corresponding to the strategy profile computed by Algorithm 1. By Theorem 3.6, G is a β -NE. To complete the proof we need to show that the social cost of G is at most β times the social cost of OPT_P . We evaluate the social cost of G with respect to the social cost of the social optimum network G^* . Clearly, $\frac{SC(G)}{SC(G^*)} = \frac{\sum_{v \in P} \text{cost}(v, G)}{\sum_{v \in P} \text{cost}(v, G^*)} \leq \max_{v \in P} \frac{\text{cost}(v, G)}{\text{cost}(v, G^*)}$. Therefore, we can repeat the analysis from the proof of Theorem 3.6. All upper bounds for the worst-case agent's improvement from the proof of Theorem 3.6 provide the upper bounds for the ratio between the cost of the agent in G and G^* . It holds because we did not assume for G' that strategies of all other agents are the same as in G , i.e., we can replace G' with G^* in all inequalities. Hence, $\max_{v \in P} \frac{\text{cost}(v, G)}{\text{cost}(v, G^*)} \leq \beta$, and the statement follows. \square

⁴The implication follows from the following observation: for any $a, b, c, d \in \mathbb{R}_{> 0}$, $\frac{a+b}{c+d} \leq \max \left\{ \frac{a}{c}, \frac{b}{d} \right\}$

COROLLARY 3.8. *Let $\alpha \leq n^x$ for some $x \in \mathbb{R}_{>0}$. Then we can construct in $O(n^2)$ time a (β, β) -network with $\beta \in O\left(\alpha^{1-\frac{1}{2x}} + 1\right)$ for $x \geq 1$ and $\beta \in O\left(\alpha^{\frac{3x-1}{4x}} + 1\right)$ for $0 < x < 1$.*

PROOF. Consider a (β, β) -network constructed by Algorithm 1 with parameters $b \leq \sqrt{2(n-1)}$ and $c = \frac{b^2}{2}$, some real number $t > 1$ and $k \in O\left((t-1)^{1-2d}\right)$. A k -degree t -spanner can be constructed in $O(n \log n)$ time ([50], Section 10.1). Hence, our algorithm outputs the (β, β) -NE in $O(n^2)$ time with

$$\beta = \max \left\{ \frac{4k}{b} \alpha + 2t + 1, \frac{2\alpha}{n - \frac{b^2}{2}} + 2, \frac{2b^2(b+2t)}{n - \frac{b^2}{2}} + 6t \right\}.$$

We make a case distinction for x when choosing b . If $x \geq 1$, we choose $b = \alpha^{\frac{1}{2x}}$. Then we get $\frac{4k}{b} \alpha + 2t + 1 = 4k\alpha^{1-\frac{1}{2x}} + 2t + 1 \in O(\alpha^{1-\frac{1}{2x}})$. Since $\alpha \leq n^x$, we have that $n - \frac{1}{2}b^2 \geq \alpha^{\frac{1}{x}} - \frac{1}{2}\alpha^{\frac{1}{x}} \geq \frac{1}{2}\alpha^{\frac{1}{x}}$. Therefore, $\max \left\{ \frac{2\alpha}{n - \frac{1}{2}b^2} + 2, \frac{2b^2(b+2t)}{n - \frac{1}{2}b^2} + 6t \right\} \in O\left(\max\{\alpha^{1-\frac{1}{x}}, \alpha^{\frac{1}{2x}}\}\right)$. Since we assume $x \geq 1$, we get

$$\beta \in O\left(\max\{\alpha^{1-\frac{1}{x}}, \alpha^{\frac{1}{2x}}, \alpha^{1-\frac{1}{2x}}\}\right) = O\left(\alpha^{1-\frac{1}{2x}}\right).$$

In case $0 < x < 1$, we choose $b = \alpha^{\frac{x+1}{4x}}$. Then we get $n - \frac{1}{2}b^2 \geq \alpha^{\frac{1}{x}} - \frac{1}{2}\alpha^{\frac{x+1}{2x}} \geq \alpha^{\frac{1}{x}} \left(1 - \frac{1}{2}\alpha^{\frac{x-1}{2x}}\right) \geq \frac{1}{2}\alpha^{\frac{1}{x}}$ since $\alpha^{\frac{x-1}{2x}} < 1$ for $0 < x < 1 < \alpha$. Hence, we get for the approximation factor β that

$$\beta \in O\left(\max\left\{\alpha^{1-\frac{x+1}{4x}}, \alpha^{1-\frac{1}{x}}, \alpha^{\frac{3(x+1)}{4x}-\frac{1}{x}}\right\} + 1\right) = O\left(\alpha^{\frac{3x-1}{4x}} + 1\right). \quad \square$$

Corollary 3.8 claims that for $1 \leq \alpha \leq \sqrt[n]{n}$, there is a (β, β) -NE with constant β , while for the other values of α , the approximation is better than the one obtained for a clique (see Theorem 3.5). However, when x tends to infinity, the value of β approaches α . For this case we show that a minimum spanning tree provides a better approximation.

THEOREM 3.9. *Any minimum spanning tree $MST(P)$ on a set of points P is a $(n-1, n-1)$ -network.*

Finally, we show that using the better of the networks obtained by Algorithm 1 and the MST yields a $(O(\alpha^{\frac{2}{3}}), O(\alpha^{\frac{2}{3}}))$ -network. See Figure 3 for an illustration.

COROLLARY 3.10. *A (β, β) -network with $\beta \in O(\alpha^{\frac{2}{3}})$ can be constructed in $O(n^2)$ time.*

PROOF. Let $x \in \mathbb{R}^+$ such that $\alpha = n^x$. Then $\alpha^{\frac{1}{x}} = n$. Applying Theorem 3.9 for $x \geq \frac{3}{2}$ and Corollary 3.8 for $x < \frac{3}{2}$ yields the result. \square

We now use Algorithm 1 to obtain a $(1+\epsilon, 1+\epsilon)$ -network for $\epsilon > 0$ if P is chosen uniformly at random from the unit square. We partition the unit square into four quadrants $C = \{a, b, c, d\}$ each containing a length- $\frac{1}{4}$ -square from $C' = \{a', b', c', d'\}$, see Figure 4.

The following Lemma shows that in this case with high probability the second case of Algorithm 1 occurs.




Figure 3: Approximation factor β obtained by Corollary 3.8 and Theorem 3.9 dependent on the relation between α and n . For $x \in \mathbb{R}_{>0}$ such that $\alpha = n^x$, we can construct a (β, β) -network with $\beta \in O(\alpha^y)$.




Figure 4: Illustration of the partition of $[0, 1]^2$ used in Lemma 3.11 and Theorem 3.12.

LEMMA 3.11. *Let $P_n = \{v_1, \dots, v_n\} \subseteq [0, 1]^2$ be n points chosen uniformly at random. Then*

$$\mathbb{P}\left(\bigcup_{\rho' \in C'} \left\{ |P_n \cap \rho'| < (1-\delta) \frac{n}{16} \right\}\right) \leq 4 \exp\left(-\frac{\delta^2 n}{32}\right).$$

By Theorem 10.1.3 from [50], for any $\epsilon > 0$, there is a $(1+\epsilon)$ -spanner with maximum degree only depending on ϵ . We use this construction to provide a $(1+\epsilon, 1+\epsilon)$ -NE.

THEOREM 3.12. *Let $\epsilon > 0$ and $P_n = \{v_1, \dots, v_n\} \subseteq [0, 1]^2$ be a set of n points chosen uniformly at random. Then if $\alpha_n \in o(n)$, there exists a $(1+\epsilon, 1+\epsilon)$ -network for P_n , for any $\alpha < \alpha_n$ asymptotically almost surely.*

PROOF. Consider a network computed by Algorithm 1 with parameters $b = 4$, $c = 2k_\epsilon b \frac{\alpha}{\epsilon}$, and k_ϵ -degree $(1 + \frac{\epsilon}{2})$ -spanner as constructed in [50].

Because $\alpha_n \in o(n)$, we can assume that $\alpha_n \leq \frac{\epsilon}{16k_\epsilon} \frac{n}{16}$, because it holds for all but finitely many n . Since

$$c = 2k_\epsilon b \frac{\alpha}{\epsilon} \leq 2k_\epsilon b \frac{1}{\epsilon} \cdot \frac{\epsilon}{16k_\epsilon} \cdot \frac{n}{16} = \left(1 - \frac{1}{2}\right) \frac{n}{16},$$

then, by Lemma 3.11, each $\rho' \in C'$ has at least c points with probability $1 - 4 \exp(-\frac{n}{128})$. Thus, for any point which is part of a quadrant $\rho \in C$, there are at least c points in ρ' within the distance at least

$\frac{w_{\max}}{b} = \frac{1}{4}$. Because we have a k_ε -degree $(1 + \frac{\varepsilon}{2})$ -spanner, by Theorem 3.6 and Theorem 3.7, we get, with probability $1 - 4 \exp(-\frac{n}{128})$, a (β, β) -network with

$$\beta = \frac{k_\varepsilon b}{c} \alpha + 1 + \frac{\varepsilon}{2} = k_\varepsilon b \frac{\varepsilon}{2k_\varepsilon b \alpha} \alpha + 1 + \frac{\varepsilon}{2} = 1 + \varepsilon. \quad \square$$

Finally, we study integer grids in \mathbb{R}^d and show that Algorithm 1 computes a $(2d, 2d)$ -network, if the grid itself is selected as spanner.

THEOREM 3.13. *Let $(b_1, \dots, b_d) \in \mathbb{N}^d$ and $B = [0, b_1] \times \dots \times [0, b_d]$ the corresponding hyperrectangle. Let $P := \mathbb{Z}^d \cap B$. Then, there exists a $(2d, 2d)$ -network for the nodes in P .*

4 PRICE OF ANARCHY AND PRICE OF STABILITY

In this section we will provide lower bounds on the Price of Anarchy and on the Price of Stability in the Euclidean case. It was shown that in the 1-norm space, the lower bound for the PoA approaches the upper bound of $\frac{\alpha+2}{2}$ when the dimension d tends to infinity [12]. We show that also in Euclidean space, there is an instance that asymptotically almost meets the upper bound.

THEOREM 4.1. *The Price of Anarchy in the \mathbb{R}^d -GNCG is at least $\min \left\{ \frac{\alpha+1}{\sqrt{2}}, \frac{\alpha^2+2\alpha+2}{2\alpha+2} \right\}$ as $d \rightarrow \infty$.*

Next, we now show that the PoA is super-constant in α , even when the underlying space is the \mathbb{R}^1 . For this we significantly improve the analysis of a construction of Bilò et al. ([12], Theorem 3.27).

THEOREM 4.2. *The PoA in the \mathbb{R}^1 and thus in the \mathbb{R}^d is lower bounded by $\frac{3}{5}\alpha^{\frac{2}{3}} \pm o(\alpha^{\frac{2}{3}})$.*

Finally, we show that the PoS is strictly larger than 1.

THEOREM 4.3. *The PoS in the \mathbb{R}^d with $d \geq 2$ is greater than 1 if $\alpha > 2$.*

PROOFSKETCH. Consider the construction from the proof of Theorem 2.1. After connecting the nodes within each cluster there are only two networks that connect the network; one with two length-1-edges and one with three length-1-edges. We set $n = 3(\lceil \alpha \rceil - 1)$ and observe that buying all three length-1-edges results in a social optimum network, as $2 \cdot 2 \cdot (\frac{n}{3})^2 > \alpha + 2 \cdot (\frac{n}{3})^2$. Selling a length-1-edge leads to a NE as $2 \cdot \frac{n}{3} < \alpha + 1 \cdot \frac{n}{3}$. This means that the social optimum network is not a NE, showing that PoS > 1 . \square

5 OUTLOOK: EFFICIENCY AND STABILITY ON A HOST NETWORK

We consider a more general model, i.e., the *Generalized Network Creation Game (GNCG)* by Bilò et al. [12], where a complete host network $H = (V, E)$ with arbitrary edge weights $w: V \times V \rightarrow \mathbb{R}_+$ is given. The price of an edge $uv \in E$ then is $\alpha \cdot w(u, v)$. For an edge subset $E' \subseteq E(H)$, we denote $w(E') := \sum_{uv \in E'} w(u, v)$.

Our main contribution of this section indicates that the geometric and non-geometric versions of the GNCG behave very similarly. Clearly, the hardness results carry over from the special case to the more general case. But we also extend the approximation results for stable and optimum networks to the GNCG, and we show that

the PoA is linear in α , as conjectured in [12]. Thus it matches the bounds for the metric version.

5.1 Approximation

Theorem 3.5 and Theorem 3.9 can be directly generalized to the case with arbitrary, even non-metric, edge weights if we consider a spanning sub-network H' of the host network H such that each edge in H' participates in at least one shortest path.

COROLLARY 5.1. *Let $H = (V, E)$ be a host network in the GNCG. A spanning sub-network $H' = (V, E')$ of H , where $E' = \{(u, v) \mid w(u, v) = d_H(u, v)\}$, is a $(\alpha + 1, \frac{\alpha}{2} + 1)$ -NE.*

COROLLARY 5.2. *Consider a host network $H = (V, E)$. Any minimum spanning tree $MST(H)$ of H is a $(n - 1, n - 1)$ -network.*

The key idea of the two results above is to remove long edges from H to obtain an approximation similar to the case with metric edge weights. We can extend this idea to make our Algorithm 1 work for the GNCG with arbitrary edge weights as follows.

Consider a host network $H = (V, E)$. Update H as follows: starting from the longest edge $uv \in E$, if $d_H(u, v) < w(u, v)$ remove uv from H . Repeat the procedure until all edges are checked. Denote the final network as $H_M = (V, E_M)$. Note that H_M is connected and has metric edge weights, i.e., for any edge $uv \in E_M$ we have $w(u, v) = d_{H_M}(u, v)$. Now we can apply Algorithm 1 to H_M with the only modification on step 7: connect node u with the closest node $u' \in C_v$ by the shortest path $\pi_{H_M}(u, u')$. We assume w_{\max} is the length of the longest shortest path in H_M .

COROLLARY 5.3. *Consider a host network $H = (V, E)$, and let $b \in \mathbb{R}_{\geq 1}, k, t \in \mathbb{R}_{> 1}, c \in \{0, \dots, n - 1\}$ such that a k -distributable t -spanner⁵ exists for H_M . Algorithm 1 computes a (β, β) -network for H with*

$$\beta = \max \left\{ \frac{kb}{c} \alpha + t, \frac{4k}{b} \alpha + 2t + 1, \frac{2\alpha}{n - c} + 2, \frac{4c(b + 2t)}{n - c} + 6t \right\}.$$

However, the above statement relies on the existence of a k -distributable t -spanner for an incomplete weighted host network, which is, to the best of our knowledge, still open. Hence, finding an efficient algorithm that computes a k -distributable t -spanner with low k and t would enable Algorithm 1 to obtain (β, β) -networks with low β for the GNCG.

5.2 Price of Anarchy

We show a $O(\alpha)$ upper bound on the PoA. This asymptotically matches the $\Omega(\alpha)$ lower bound from Bilò et al. [12].

THEOREM 5.4. *In the GNCG the PoA is at most $2(\alpha + 1)$.*

PROOF. Consider a host network $H = (V, E(H))$, a stable network $G = (V, E)$ and an optimum network $G^* = (V, E^*)$. Since every NE is a $(\alpha + 1)$ -spanner (by Lemma 2.2 in [12]) the distance cost of G is

$$\sum_{u, v \in V} d_G(u, v) \leq (\alpha + 1) \sum_{u, v \in V} d_H(u, v) \leq (\alpha + 1) \sum_{u, v \in V} d_{G^*}(u, v). \quad (5)$$

⁵See Footnote 3 for a definition of a k -distributable t -spanner.

Now we evaluate the edge cost of G . We partition the edges in E concerning the edges in the optimum and analyze each set's total cost separately. Let $B := \bigcup_{uv \in E^*} \{\pi_G(u, v)\}$ be a set of edges in G appearing in some shortest u - v path in G for each $uv \in E^*$, i.e., for every edge uv in the social optimum, B contains all edges from a shortest path between u and v in G . We denote the rest of the edges in G as $R := E \setminus B$.

Since G is a $(\alpha + 1)$ -spanner, we can evaluate the cost of the edges in B as follows:

$$\begin{aligned} \alpha \cdot w(B) &= \alpha \cdot w \left(\bigcup_{uv \in E^*} \pi_G(u, v) \right) \leq \alpha \sum_{uv \in E^*} d_G(u, v) \\ &\leq \alpha(\alpha + 1) \sum_{uv \in E^*} d_H(u, v) \leq \alpha(\alpha + 1)w(E^*). \end{aligned} \quad (6)$$

Next, we compute the cost of edges in R . Consider an agent $u \in V$. Let R_u be a set of edges from R that u buys in G . Since G is a NE, deleting all edges from R_u is not an improving move for u . Hence,

$$\alpha w(R_u) + \sum_{v \in V} d_G(u, v) \leq \sum_{v \in V} d_{G-R_u}(u, v),$$

where $G - R_u$ is the network obtained after the deletion. Thus,

$$\alpha w(R_u) \leq \sum_{v \in V} d_{G-R_u}(u, v) - d_G(u, v) \leq d_{G-R_u}(u, V) \leq d_{G-R}(u, V)$$

To evaluate the distance $d_{G-R}(u, v)$, note that it is equal to the distance between u and v in the network G restricted on the edge set B . Consider a shortest path $\pi_{G^*}(u, v)$ in the optimum network G^* . By definition, for each edge $xy \in \pi_{G^*}(u, v)$, set B contains a shortest path $\pi_G(x, y)$ of length

$$d_{G-R}(x, y) \leq (\alpha + 1)d_H(x, y) \leq (\alpha + 1)d_{G^*}(x, y).$$

Thus,

$$\begin{aligned} d_{G-R}(u, v) &\leq \sum_{xy \in \pi_{G^*}(u, v)} d_{G-R}(x, y) \\ &\leq (\alpha + 1) \cdot \sum_{xy \in \pi_{G^*}(u, v)} d_{G^*}(x, y) = (\alpha + 1)d_{G^*}(u, v). \end{aligned}$$

The total cost of set R then is

$$\alpha w(R) = \alpha \sum_{u \in V} w(R_u) \leq \sum_{u \in V} d_{G-R}(u, V) \leq (\alpha + 1) \sum_{u \in V} d_{G^*}(u, V).$$

In combination with Inequality (6) and the upper bound for the distance cost in Inequality (5), we get

$$\begin{aligned} \frac{SC(G)}{SC(G^*)} &= \frac{\alpha w(B) + \alpha w(R) + \sum_{u \in V} d_G(u, V)}{\alpha w(E^*) + \sum_{u \in V} d_{G^*}(u, V)} \\ &\leq \frac{\alpha(\alpha + 1)w(E^*) + 2(\alpha + 1) \sum_{u \in V} d_{G^*}(u, V)}{\alpha w(E^*) + \sum_{u \in V} d_{G^*}(u, V)} \\ &\leq \frac{2(\alpha + 1)(\alpha w(E^*) + \sum_{u \in V} d_{G^*}(u, V))}{\alpha w(E^*) + \sum_{u \in V} d_{G^*}(u, V)} = 2(\alpha + 1). \quad \square \end{aligned}$$

COROLLARY 5.5. *PoA* $\in \Theta(\alpha)$ in the GNCG.

6 CONCLUSION

We studied the problem of designing networks that are both efficient in terms of social cost and stable in terms of being close to a Nash equilibrium state. For this, we focus on studying (β, γ) -networks that are in β -approximate Nash equilibrium and have a social cost of at most γ times the cost of the social optimum. In particular, we considered (β, γ) -networks in the Euclidean version of the Generalized Network Creation Game by Bilò et al. [12], where agents are points in \mathbb{R}^d , and each agent aims to maximize her centrality by creating costly edges. This version has the natural feature that the cost of each edge is proportional to the Euclidean distance between the endpoints. Hence, this model captures many real-world settings for the decentralized creation of communication networks.

Our main contribution is a $O(n^2)$ time algorithm for computing (β, β) -networks with low β . First of all, this result is interesting since it is one of the very few algorithmic results for constructing (approximate) Nash equilibria in the realm of network creation games. Such a centralized algorithm is valuable in a setting with strategic agents since a central designer could propose a network to the strategic agents which then may selfishly deviate from the proposed solution. If this proposed network is (almost) stable, then the agents have no (or only a very low) incentive for deviating. If additionally the proposed network has other beneficial properties like (almost) optimal social cost then this is another compelling reason for accepting the centrally designed proposal.

Our algorithm is simple but non-trivial and relies on techniques from the well-studied t -spanner problem. Moreover, our algorithm creates $(O(1), O(1))$ -networks if $\alpha \leq \sqrt[3]{n}$, i.e., for the case where edges are comparably cheap or where the number of nodes in the network is large. The same holds true for networks on random point sets or on grids, both of which seem to be natural topologies.

In contrast to these positive results, we observed that none of the extreme cases of (β, γ) -networks could guarantee a constant approximation. Namely, a social optimum network can be very unstable and NP-hard to compute, while a NE can have a much higher social cost than the optimal network. Moreover, we have shown that the finite improvement property does not hold for our model. It was shown by Bilò et al. [12] that computing best response strategies is NP-hard. This indicates that there is no efficient way of finding an (almost) stable state in a decentralized way. Hence, computing such a state via a centralized algorithm and then proposing it to the agents could be a way to circumvent this hard problem.

As another important result of the paper, we have shown that the PoA is linear in the parameter α independently of the dimension of the Euclidean space or of edge weights being metric. Although conjectured by Bilò et al. [12], this is surprising because it contradicts the intuition that the PoA should be lower for the metric case, especially for low dimensions.

We focused on three extreme cases of the bicriteria optimization, i.e., when one of the approximation factors is 1 or both approximation factors are equal. Of course, it would be interesting to map the whole Pareto frontier precisely. Another promising direction for future work is to extend our approximation results for the non-metric case. For example, any solution for a k -distributable t -spanner for weighted networks would make our approximation algorithm work even for the Generalized Network Creation Game.

REFERENCES

- [1] S. Albers, S. Eilts, E. Even-Dar, Y. Mansour, and L. Roditty. On Nash equilibria for a network creation game. In *SODA '06*, pages 89–98, 2006.
- [2] S. Albers and P. Lenzner. On approximate Nash equilibria in network design. *Internet Mathematics*, 9(4):384–405, 2013.
- [3] N. Alon, E. D. Demaine, M. T. Hajiaghayi, and T. Leighton. Basic network creation games. *SIAM Journal on Discrete Mathematics*, 27(2):656–668, 2013.
- [4] C. Álvarez and A. Messegué. Network creation games: Structure vs anarchy. *arXiv:1706.09132*, 2017.
- [5] C. Álvarez and A. Messegué. On the constant price of anarchy conjecture. *arXiv:1809.08027*, 2018.
- [6] C. Álvarez and A. Messegué. On the price of anarchy for high-price links. In *WINE'19*, pages 316–329, 2019.
- [7] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.
- [8] E. Anshelevich, A. Dasgupta, É. Tardos, and T. Wexler. Near-optimal network design with selfish agents. *Theory of Computing*, 4(1):77–109, 2008.
- [9] V. Bala and S. Goyal. A noncooperative model of network formation. *Econometrica*, 68(5):1181–1229, 2000.
- [10] D. Bilò and P. Lenzner. On the tree conjecture for the network creation game. In *STACS'18*, pages 14:1–14:15, 2018.
- [11] D. Bilò, T. Friedrich, P. Lenzner, S. Lowski, and A. Melnichenko. Selfish creation of social networks. *CoRR*, abs/2012.06203, 2020.
- [12] D. Bilò, T. Friedrich, P. Lenzner, and A. Melnichenko. Geometric network creation games. In *SPAA'19*, pages 323–332, 2019.
- [13] D. Bilò, T. Friedrich, P. Lenzner, A. Melnichenko, and L. Molitor. Fair tree connection games with topology-dependent edge cost. In *FSTTCS'20*, pages 15:1–15:15, 2020.
- [14] M. Brautbar and M. J. Kearns. A clustering coefficient network formation game. In *SAGT'11*, pages 224–235, 2011.
- [15] P. Carmi and L. Chaitman-Yerushalmi. Minimum weight euclidean t-spanner is np-hard. *Journal of Discrete Algorithms*, 22:30–42, 2013.
- [16] A. Chauhan, P. Lenzner, A. Melnichenko, and L. Molitor. Selfish network creation with non-uniform edge cost. In *SAGT'17*, pages 160–172, 2017.
- [17] A. Chauhan, P. Lenzner, A. Melnichenko, and M. Münn. On selfish creation of robust networks. In *SAGT'16*, pages 141–152, 2016.
- [18] S. Chechik and C. Wulff-Nilsen. Near-optimal light spanners. *ACM Trans. Algorithms*, 14(3):33:1–33:15, 2018.
- [19] H. Chen and T. Roughgarden. Network design with weighted players. *Theory Comput. Syst.*, 45(2):302–324, 2009.
- [20] H. Chen, T. Roughgarden, and G. Valiant. Designing networks with good equilibria. In *SODA'08*, pages 854–863, 2008.
- [21] A. Cord-Landwehr and P. Lenzner. Network creation games: Think global - act local. In *MFCS'15*, pages 248–260, 2015.
- [22] J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. Selfish routing in capacitated networks. *Mathematics of Operations Research*, 29(4):961–976, 2004.
- [23] E. D. Demaine, M. T. Hajiaghayi, H. Mahimi, and M. Zadimoghaddam. The price of anarchy in network creation games. *ACM Trans. on Algorithms*, 8(2):13, 2012.
- [24] H. Ehzell, T. Friedrich, P. Lenzner, and A. Melnichenko. Flow-based network creation games. In *IJCAI'20*, pages 139–145, 2020.
- [25] S. Eidenbenz, V. A. Kumar, and S. Züst. Equilibria in topology control games for ad hoc networks. *Mobile Networks and Applications*, 11(2):143–159, 2006.
- [26] M. Elkin, O. Neiman, and S. Solomon. Light spanners. *SIAM J. Discret. Math.*, 29(3):1312–1321, 2015.
- [27] A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. Shenker. On a network creation game. In *PODC'03*, pages 347–351, 2003.
- [28] A. Filtser and S. Solomon. The greedy spanner is existentially optimal. *SIAM J. Comput.*, 49(2):429–447, 2020.
- [29] W. Friedemann, T. Friedrich, H. Gawendowicz, P. Lenzner, A. Melnichenko, J. Peters, D. Stephan, and M. Vaichenker. Efficiency and stability in euclidean network design. *arXiv:2104.11618*, 2021.
- [30] T. Friedrich, S. Ihde, C. Keßler, P. Lenzner, S. Neubert, and D. Schumann. Efficient best response computation for strategic network formation under attack. In *SAGT'17*, pages 199–211, 2017.
- [31] M. Gairing, T. Harks, and M. Klimm. Complexity and approximation of the continuous network design problem. In *APPROX/RANDOM'14*, pages 226–241, 2014.
- [32] M. R. Garey and D. S. Johnson. *Computers and intractability*, volume 29. wh freeman New York, 2002.
- [33] A. Gulyás, J. J. Biró, A. Kőrösi, G. Rétvári, and D. Krioukov. Navigable networks as Nash equilibria of navigation games. *Nature communications*, 6:7651, 2015.
- [34] A. Gupta and J. Könemann. Approximation algorithms for network design: A survey. *Surveys in Operations Research and Management Science*, 16(1):3–20, 2011.
- [35] M. Hoefer. Non-cooperative tree creation. *Algorithmica*, 53(1):104–131, 2009.
- [36] M. Hoefer and P. Krysta. Geometric network design with selfish agents. In *COCOON'05*, pages 167–178, 2005.
- [37] D. S. Johnson, J. K. Lenstra, and A. R. Kan. The complexity of the network design problem. *Networks*, 8(4):279–285, 1978.
- [38] R. M. Karp. Reducibility among combinatorial problems. In *Complexity of computer computations*, pages 85–103. Springer, 1972.
- [39] B. Kawald and P. Lenzner. On dynamics in selfish network creation. In *SPAA'13*, pages 83–92, 2013.
- [40] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *STACS'99*, pages 404–413, 1999.
- [41] H. Le and S. Solomon. Truly optimal euclidean spanners. In *FOCS'19*, pages 1078–1100, 2019.
- [42] P. Lenzner. On dynamics in basic network creation games. In *SAGT'11*, pages 254–265, 2011.
- [43] P. Lenzner. Greedy selfish network creation. In *WINE'12*, pages 142–155, 2012.
- [44] T. L. Magnanti and R. T. Wong. Network design and transportation planning: Models and algorithms. *Transportation science*, 18(1):1–55, 1984.
- [45] A. Mamageishvili, M. Mihalák, and D. Müller. Tree Nash equilibria in the network creation game. *Internet Mathematics*, 11(4-5):472–486, 2015.
- [46] M. Mihalák and J. C. Schlegel. Asymmetric swap-equilibrium: A unifying equilibrium concept for network creation games. In *MFCS'12*, pages 693–704, 2012.
- [47] M. Mihalák and J. C. Schlegel. The price of anarchy in network creation games is (mostly) constant. *Theory Comput. Syst.*, 53(1):53–72, 2013.
- [48] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124 – 143, 1996.
- [49] T. Moscibroda, S. Schmid, and R. Wattenhofer. Topological implications of selfish neighbor selection in unstructured peer-to-peer networks. *Algorithmica*, 61(2):419–446, 2011.
- [50] G. Narasimhan and M. Smid. *Geometric spanner networks*. Cambridge University Press, 2007.
- [51] M. Newman. *Networks: an introduction*. Oxford university press, 2010.