Strategic Network Creation for Enabling Greedy Routing

Julian Berger [1], Tobias Friedrich [1], Pascal Lenzner[2], Paraskevi Machaira [1], Janosch Ruff [1]

[1]Hasso Plattner Institute, University of Potsdam, Potsdam, Germany
[2]Institute of Computer Science, University of Augsburg, Augsburg, Germany
julian.berger@student.hpi.de, {friedrich, paraskevi.machaira, janosch.ruff}@hpi.de, pascal.lenzner@uni-a.de

Abstract

Today we rely on networks that are created and maintained by smart devices. For such networks, there is no governing central authority but instead the network structure is shaped by the decisions of selfish intelligent agents. A key property of such communication networks is that they should be easy to navigate for routing data. For this, a common approach is greedy routing, where every device simply routes data to a neighbor that is closer to the respective destination.

Networks of intelligent agents can be analyzed via a game-theoretic approach and in the last decades many variants of network creation games have been proposed and analyzed. In this paper we present the first game-theoretic network creation model that incorporates greedy routing, i.e., the strategic agents in our model are embedded in some metric space and strive for creating a network among themselves where all-pairs greedy routing is enabled. Besides this, the agents optimize their connection quality within the created network by aiming for greedy routing paths with low stretch.

For our model, we analyze the existence of (approximate)-equilibria and the computational hardness in different underlying metric spaces. E.g., we characterize the set of equilibria in 1-2-metrics and tree metrics and show that Nash equilibria always exist. For Euclidean space, the setting which is most relevant in practice, we prove that equilibria are not guaranteed to exist but that the well-known Θ-graph construction yields networks having a low stretch that are game-theoretically almost stable. For general metric spaces, we show that approximate equilibria exist where the approximation factor depends on the cost of maintaining any link.

Full version — https://doi.org/10.48550/arXiv.2403.15307

Introduction

Many important real-world networks, like the Internet or peer-to-peer ad-hoc networks among smart devices, are created by the decentralized interaction of various agents that pursue their own goals, like maximizing their centrality, minimizing their latency, or maximizing the network throughput. These agents can thus be considered as acting selfishly and strategic, i.e., they try to optimize their costs for creating and using the network. This observation sparked a whole research area devoted to game-theoretic network formation models (Papadimitriou 2001) and there is an abundance of interesting variants, typically called *network creation games*, that capture how networks emerge in different domains, e.g., communication networks or social networks.

Networks enable communication. To this end, it is commonly assumed that shortest paths are used for routing traffic within the created network. However, to make this work, all agents need to know the global structure of the network. Moreover, in case of dynamic changes to the network, this global view must be updated or otherwise shortest path routing would fail. In case of the Internet, shortest path routing is currently ensured by the usage of extensive routing tables that exactly specify which next-hop neighbor to use for which destination. All these routing tables must be maintained and updated accordingly, even for insignificant structural changes. This can be avoided by using *greedy routing*, where in each network node every incoming packet is simply routed to a neighbor that is closer to the packet's destination. For this, geographic information is needed, e.g., the nodes must be embedded in some underlying metric space and the positions in that space allow for deciding which next hop to use for routing. Exactly this has been proposed for the Internet (Boguná, Papadopoulos, and Krioukov 2010), i.e., to map the nodes of the Internet into a metric space such that greedy routing works. This is called a *greedy embedding* and they can be computed efficiently (Bläsius et al. 2020).

However, centrally computing a greedy embedding of a network does not account for the selfish strategic behavior of the network participants. The agents do not only want to ensure that greedy routing works, but at the same time they want to maximize their connection quality within the created network, e.g., they want to minimize their average stretch. Thus, strategic agents would "rewire" a given greedy embedding, if this reduces their stretch. However, this might endanger that greedy routing is guaranteed to succeed.

In this paper, we set out to investigate this tension between the stability of the network with respect to local structural changes by the agents and maintaining greedy routing.

Model and Preliminaries

We consider n agents that correspond to points $\mathcal{P} = \{p_1, \dots, p_n\}$ in some metric space $\mathcal{M} = (\mathcal{P}, d_\mathcal{M})$, where $d_\mathcal{M}(u, v)$ denotes the distance of $u \in \mathcal{P}$ and $v \in \mathcal{P}$ in $\mathcal{M}$.

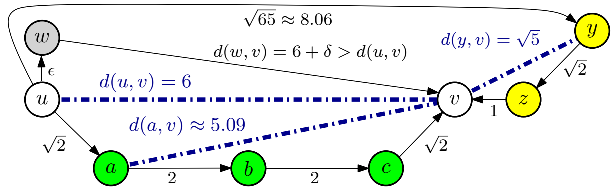

Figure 1: Greedy routing paths in the Euclidean plane. Two such paths from u to v exist: u, a, b, c, v and u, y, z, v. Path u, w, v is not a greedy routing path, since $d(w, v) > d(u, v)$. The shortest greedy routing path is u, a, b, c, v, so $d_G^{\text{greedy}}(u, v) = 2\sqrt{2} + 4 \approx 6.83$ and $\text{stretch}_G(u, v) = d_G^{\text{greedy}}(u, v)/d(u, v) = \frac{2\sqrt{2}+4}{6} \approx 1.14$. For small enough ε and δ, path u, w, v has length less than $d_G^{\text{greedy}}(u, v)$. Thus, the shortest u-v-path may not be a greedy routing path.

Besides arbitrary metric spaces where $d_{\mathcal{M}}$ has to satisfy the triangle inequality, we will also consider *1-2-metrics*, where $d_{\mathcal{M}}(u, v) \in \{1, 2\}$, for any $u, v \in \mathcal{P}$, *tree-metrics*, where the distances are determined by a given weighted spanning tree T, such that $d_{\mathcal{M}}(u, v) = d_T(u, v)$, i.e., the distance between points $u, v \in \mathcal{P}$ is their distance in T, and the *Euclidean-metric*, where the points are located in Euclidean space and the Euclidean distance is used. We omit the reference to $\mathcal{M}$ if it is clear from the context.

The goal of the agents is to create a directed network among themselves. For this, each agent strategically decides over its set of outgoing directed edges. The *strategy of agent* u is $S_u \subseteq V \setminus \{u\}$, i.e., agent u can create edges to any subset of the other agents. Let $\mathbf{s} = (S_1, \ldots, S_n)$ denote the *strategy-profile*, which is the vector of strategies of all agents. As shorthand, for any agent $u \in V$ let $\mathbf{s} = (S_u, \mathbf{s}_{-u})$, where $\mathbf{s}_{-u}$ is the vector of strategies of all agents except agent u. Any strategy-profile $\mathbf{s}$ uniquely defines a directed weighted network $G(\mathbf{s}) = (\mathcal{P}, E(\mathbf{s}), \ell)$, where the edge-set $E(\mathbf{s})$ is $\bigcup_{u \in V}\{(u, v) \mid v \in S_u\}$ and the length of any edge $(u, v) \in E(\mathbf{s})$ is equal to the distance of the positions of its endpoints in $\mathcal{M}$, i.e., $\ell(u, v) = d_{\mathcal{M}}(u, v)$.

Given a weighted directed network $G = (\mathcal{P}, E, \ell)$, where the nodes in $\mathcal{P}$ are points in metric space $\mathcal{M}$, a *greedy routing path from* u *to* v *in* G is a path $x_1, x_2, ..., x_j$, with $x_i \in \mathcal{P}$, for $1 \le i \le j$, where $x_1 = u$, $x_j = v$, and $(x_i, x_{i+1}) \in E$, for $1 \le i \le j-1$, such that $d_{\mathcal{M}}(x_i, v) > d_{\mathcal{M}}(x_{i+1}, v)$ holds for all $1 \le i \le j-1$. Thus, such a path is a directed path from u to v in G, where along the path the nodes get strictly closer to the endpoint of the path in terms of their distance in $\mathcal{M}$. For two nodes $u, v \in \mathcal{P}$, we define $d_G^{\text{greedy}}(u, v)$ as the length of the shortest greedy routing path between u and v in G, where the length of a path $x_1, x_2, ..., x_j$ is $\sum_{i=1}^{j-1} \ell(x_i, x_{i+1}) = \sum_{i=1}^{j-1} d_{\mathcal{M}}(x_i, x_{i+1})$. If no greedy routing path exists between u and v in G, then $d_G^{\text{greedy}}(u, v) = \infty$. We will call $d_G^{\text{greedy}}(u, v)$ the *greedy-routing-distance* between u and v in network G. We say that *greedy routing is enabled* in G, if any pair of nodes in G has finite greedy-routing-distance. See Figure 1.

For any two nodes u and v in network G, we will compare their greedy-routing-distance with their distance in $\mathcal{M}$. The ratio of these values is called the *stretch*, i.e., we have

$$\text{stretch}_G(u, v) = \begin{cases} \frac{d_G^{\text{greedy}}(u,v)}{d_{\mathcal{M}}(u,v)} & \text{, if a greedy routing path} \\ & \text{ from } u \text{ to } v \text{ exists in } G, \\ Z & \text{, otherwise,} \end{cases}$$

where Z is some arbitrarily large number that serves as a penalty for not having a greedy routing path. Intuitively, the stretch measures the detour that the best greedy routing path has to take, compared to the shortest possible path, i.e., to having a direct edge to the target node.

Agents choose their strategies to minimize their *cost* within the formed network. The cost of agent u in network $G(\mathbf{s})$ is defined as

$$c_u(\mathbf{s}) = \text{stretchcost}_u(\mathbf{s}) + \text{edgecost}_u(\mathbf{s}),$$

where $\text{stretchcost}_u(\mathbf{s}) = \sum_{v \in \mathcal{P} \setminus \{u\}} \text{stretch}_{G(\mathbf{s})}(u, v)$ and $\text{edgecost}_u(\mathbf{s}) = \alpha |S_u|$, for a given $\alpha > 0$. The latter is a global parameter and allows to adjust the comparison of edge costs versus stretch costs. The *social cost* of a network $G(\mathbf{s})$ is defined as $c(\mathbf{s}) = \sum_{u \in \mathcal{P}} c_u(\mathbf{s})$. For any set of points $\mathcal{P}$, the network $G(\mathbf{s}^*) = (\mathcal{P}, E(\mathbf{s}^*), \ell)$ minimizing the social cost is called the *social optimum network* for $\mathcal{P}$.

An *improving response* of agent u in strategy-profile $\mathbf{s} = (S_u, \mathbf{s}_{-u})$ is a strategy S'_u such that $c_u((S'_u, \mathbf{s}_{-u})) < c_u((S_u, \mathbf{s}_{-u}))$, i.e., by using strategy S'_u, agent u has strictly lower cost compared to strategy S_u. A strategy S^*_u is called a *best response* of agent u in strategy-profile $\mathbf{s} = (S_u, \mathbf{s}_{-u})$, if $c_u(S^*_u, \mathbf{s}_{-u}) \le c_u(S'_u, \mathbf{s}_{-u})$ for any other strategy $S'_u \subseteq V \setminus \{u\}$, i.e., a best response strategy minimizes agent u's cost, given that the other agents' strategies are fixed.

A strategy-profile $\mathbf{s}$ is in *pure Nash Equilibrium* (NE) if no agent has an improving move, i.e., in $\mathbf{s}$ every agent already plays a best response. Since we have a bijection between strategy-profiles $\mathbf{s}$ and the corresponding networks $G(\mathbf{s})$, we will say that network $G(\mathbf{s})$ is in NE, if $\mathbf{s}$ is in NE. A network $G(\mathbf{s})$ is in *Greedy Equilibrium* (GE) (Lenzner 2012) if no agent has an improving response that consists of adding, swapping or deleting a single incident outgoing edge, where a swap is a combination of deleting an incident outgoing edge and adding another one. Notice that any network in NE is also in GE. A network $G(\mathbf{s})$ is in β-approximate NE (β-NE) if no agent u can change its strategy such that its cost decreases below $\frac{1}{\beta} c_u(\mathbf{s})$, i.e., no agent can reduce its cost to a β-fraction by an improving response.

An *improving (best) response path* is a sequence of strategy-profiles $\mathbf{s}_0, \mathbf{s}_1, \mathbf{s}_2, \ldots, \mathbf{s}_k$, such that $\mathbf{s}_i$ results from some agent changing to an improving (best) response in $\mathbf{s}_{i-1}$, for $1 \le i \le k$. An *improving (best) response cycle* (IRC or BRC, respectively) is a cyclic improving response path, i.e., where $\mathbf{s}_0 = \mathbf{s}_k$ holds. The non-existence of improving response cycles, i.e., if every improvement path has finite length, is equivalent to the existence of an ordinal potential function (Monderer and Shapley 1996). The latter implies that NE can be found via natural sequential improvement dynamics. A strategic game is called *weakly acyclic (WA) (weakly acyclic under best response (BR-WA))* if from every strategy vector $\mathbf{s}$ there exists a finite improving (best) response path that starts in $\mathbf{s}$ and ends in a NE.

Discussion of the Model Assumptions

Our main goal is to consider a simple but general network creation model that features agents that use geographic information of their neighbors for routing.

Considering Different Metrics: For capturing as many realistic settings as possible, we consider different underlying metrics, starting from the simplest non-trivial metric, 1-2-metrics, that capture settings where pairwise communication can only have low quality (length 2) or high quality (length 1), e.g., due to hardware constraints. Such metrics have often been studied, e.g., for the TSP (Karp 1972; Adamaszek, Mnich, and Paluch 2018). We also consider tree-metrics that are commonly used for approximating other metrics (Fakcharoenphol, Rao, and Talwar 2004) and that naturally emerge if overlay networks with an underlying tree topology are considered. The setting of Euclidean metrics is natural for ad-hoc networks (Ramanathan and Redi 2002) but also for classical wireline networks.

Using Shortest Greedy Routing Paths: Given that an abundance of different greedy routing protocols exist (Mauve, Widmer, and Hartenstein 2001), we present a more general analysis by abstracting away from any concrete greedy routing protocol. To achieve this, our definition of the greedy-routing distance only depends on the network G and the metric space $\mathcal{M}$ but not on any concrete assumption of which neighbor that is closer to the target is used as next hop in routing. This more abstract definition ensures robust bounds, since our distance measure always gives a lower bound on the distance achieved by any particular greedy routing protocol. E.g., always using the neighbor that is closest to the target can yield greedy routing paths that are arbitrarily longer, compared to the shortest greedy routing path. Thus, our greedy routing can be understood as the result of iteratively improving the employed greedy routing paths over time*. Note that the shortest greedy routing path can be found efficiently via Dijkstra's algorithm.

Considering the Stretch: From Moscibroda, Schmid, and Wattenhofer (Moscibroda, Schmid, and Wattenhofer 2006) we adopt measuring the connection quality between agents via their stretch. This assumes that by using the network, agents can infer their network distance to any other node, e.g., by measuring their latency. We note that with slight modifications of the respective proofs all our results carry over if we would directly use the distance in the network, as in (Bilò et al. 2019), without dividing by the respective distance in the underlying metric space.

The Role of the Parameter α: As is common also in other variants of network creation games (Fabrikant et al. 2003; Moscibroda, Schmid, and Wattenhofer 2006; Bilò et al. 2019), the parameter α captures the trade-off between the costs for establishing a link in the network and the costs for communication between two nodes, e.g., the observed latency. We can assume a high alpha for settings where the edge costs dominate the stretch costs and a low alpha for the inverse setting. This allows to model a wide range of applications, e.g., the parameter α could encode the basic set-up costs for establishing a cryptographically secure connection, or for purchasing hardware, e.g., directed antennae or cables. Given these examples and for the sake of simplicity, we assume that the same parameter α applies to all agents.

Related Work

Network Creation Games (NCGs) were first introduced by Fabrikant et al. (2003). In their model, agents that correspond to nodes of a network strategically create incident undirected edges with the goal of minimizing their sum of hop-distances to all other agents. They show that NE always exist, that the social optimum network is either a clique or a star, and that computing a best response is NP-hard. NCGs admit IRCs (Kawald and Lenzner 2013) and any network in GE is in 3-NE (Lenzner 2012). Moreover, many variations of NCGs have been investigated, e.g., a bilateral version (Corbo and Parkes 2005; Friedrich et al. 2023), variants with locality (Bilò et al. 2016; Cord-Landwehr and Lenzner 2015), with robustness (Meirom, Mannor, and Orda 2015; Chauhan et al. 2016; Echzell et al. 2020), with budget constraints (Laoutaris et al. 2008; Ehsani et al. 2011), with non-uniform edge prices (Meirom, Mannor, and Orda 2014; Chauhan et al. 2017; Bilò et al. 2021), with heterogeneous agents (Bullinger, Lenzner, and Melnichenko 2022), or with temporal edges (Bilò et al. 2023).

Closer to our model are NCGs that involve geometry. In the wireline strong connectivity game (Eidenbenz, Kumar, and Zust 2006) agents that correspond to points in the plane create undirected edges to ensure that they can reach all other agents. There, NE exist and can be found efficiently. Geometric spanner games (Abam and Qafari 2019) are similar, but the agents want to ensure a given maximum stretch. Bilò et al. (2019) defined a similar model with undirected edges, where the agents want to minimize the sum of their shortest path distances to all other agents. They show that NE exist in 1-2-metrics and tree-metrics and that every GE is in 3-NE for any metric. Also, even on 1-2-metrics computing a best response is NP-hard. Later, it was shown that $(\alpha + 1)$-NE always exist (Friedemann et al. 2021).

Closest to our model is the work by Moscibroda, Schmid, and Wattenhofer (2006). Their model involves agents that form directed edges and that have the same cost function as in our model. However, classical shortest paths are used for defining the stretch. For Euclidean metrics they provide an instance that does not admit a NE and they show that deciding NE existence is NP-hard. Note that networks in NE might not be in NE in our model an vice versa. Thus, none of their results can be directly transferred to our model.

For background on (Algorithmic) Game Theory and Multiagent Systems, we refer to the textbooks by Nisan et al. (2007) and Shoham and Leyton-Brown (2009).

Our Contribution

We propose and analyze the first network creation game where intelligent selfish agents aim for optimizing their communication quality while at the same time maintaining

*Such a process cannot get stuck in local non-greedy routing since we assume all-pairs communication, i.e., at any time every network node ensures to have a greedy routing path to all other nodes. Thus, a local non-greedy routing would directly urge the corresponding node of a routing dead-end to change its edge-set.

	1-2 Metrics	Tree Metrics	Euclidean Metrics	General Metrics
NE Existence	NE always exist NE characterization	NE always exist GE & NE unique	no existence	no existence
Dynamics	IRCs exist, no BRCs	BR-WA	IRCs+BRCs	IRCs+BRCs
Complexity	$\alpha > 1/2$: BR & NE-ver NP-hard $\alpha \leq 1/2$: BR & NE-ver in P	BR NP-hard NE-dec in P	BR NP-hard	BR NP-hard {NE,GE}-dec NP-hard
Approx-NE	$\mathcal{O}(\log n)$	1	5	$\alpha + 1$

Table 1: Result overview. Abbreviations used: IRC, BR(C) (Improving/Best Response (Cycle)), BR-WA (Weakly Acyclic under Best Response), NE-ver (NE-verification), {NE,GE}-dec ({NE,GE}-decision).

that all-pairs greedy routing is enabled. Given the favorable properties of greedy routing for dynamically changing networks among smart devices, this is highly relevant for current and future multi-agent network formation scenarios.

As our main technical contribution, we uncover the influence on the underlying metric space on the existence of equilibria, the game dynamics, and the computational complexity of computing best response strategies and on deciding equilibrium existence. See Table 1 for an overview.

We show that NE always exist on 1-2 metrics and tree metrics but on Euclidean metrics there are instances that do not admit NE. However, we show that in all considered metrics approximate-NE with low approximation factor do exist and can be constructed efficiently. Regarding the dynamics, we show that improving response cycles exist for almost all variants but that on 1-2 metrics and tree metrics convergence to an NE is guaranteed if only best responses are played. On the complexity side, we show that computing such best responses is NP-hard in most cases. Also deciding NE existence is NP-hard for some cases.

In comparison to (Moscibroda, Schmid, and Wattenhofer 2006), where shortest path routing instead of greedy routing is used, our results on Euclidean metrics are similar but this is not obvious at all. Adding the constraint of greedy routing forces the agents to create certain edges which yields entirely different behavior[†]. Thus, their results do not carry over to our model. Moreover, we show the even stronger result that not even GE are guaranteed to exist.

We emphasize, that the efficient construction of approximate equilibria for Euclidean and general metrics directly carries over to the model from Moscibroda, Schmid, and Wattenhofer (2006). This is true, since our agents use greedy routing paths instead of shortest paths, so the path choice is more restricted in our model. Thus, the stretch can only decrease when considering shortest paths. Hence, we present the first positive results on equilibrium existence for the model in (Moscibroda, Schmid, and Wattenhofer 2006). Also, our constant approximate NE for Euclidean metrics via Θ-graphs are highly relevant in practice. Θ-graphs are well-known geometric spanners with many applications.

[†]For example, in a metric space where all the nodes have the same distance to each other and $\alpha = 1$, in the model of (Moscibroda, Schmid, and Wattenhofer 2006) a star is a NE while in our model the only NE is the complete graph.

Due to space constraints, we focus our presentation on the results for the practically most relevant settings, i.e., on tree metrics and Euclidean metrics. All omitted details, in particular the technically interesting analysis of 1-2 metrics, can be found in the full version (Berger et al. 2024).

1-2 Metrics

As a warm-up we consider the simplest non-trivial setting where the edge lengths in the underlying space fulfill the triangle inequality. For 1-2 metrics, where all edges have either length 1 or length 2 we present a complete picture since we are able to completely characterize the set of NE. This is rare in the realm of network creation games. Moreover, we completely map the boundary of computational hardness for best response computation and deciding equilibrium existence.

The results are shown in Table 1.

Tree Metrics

As the next step, we examine networks that are created with an underlying tree metric. In a tree metric, a positively weighted undirected spanning tree T is given, such that for all nodes $u, v \in \mathcal{P}$, we have $d(u, v) = d_T(u, v)$. In the following, we always use T to denote the given tree.

Let $\mathbf{s}^T$ denote the strategy-profile, where every edge of T is created in both directions and let $G^T = G(\mathbf{s}^T)$ be the corresponding network. For our analysis, we will consider the network $G^T = (\mathcal{P}, E^T, \ell)$ *rooted at a node* $r \in \mathcal{P}$, denoted as G^T_r. This is defined analogous to rooting the tree T at node r: node v *is the child of node* u *in* G^T_r, if $(u, v) \in E^T$ and if $d(u, r) < d(v, r)$, i.e., if u is closer to r than v. Moreover, in G^T_r node w is a *descendant* of node u, if a path $u = x_1, x_2, \ldots, x_k = w$ exists, such that x_{i+1} is the child of x_1, for $1 \leq i \leq k - 1$. Also, u is a descendant of itself.

Since for all our purposes the network G^T_r behaves exactly like the tree T rooted at r, we will from now on use the terminology from trees, when working with G^T or G^T_r. For example, for G^T_r we let $subtree(u)$ denote the subtree of G^T_r rooted at node u that includes all descendants of u (including u). Furthermore, let $below(u)$ refer to the set of subtrees $\{subtree(v) \mid v \text{ is a child of } u\}$. Using the above definitions, we get the following useful statements:

Lemma 1. *In a tree metric, a greedy routing path from node* u *to a node* v *can only consist of nodes that are in the same subtree from* $below(u)$ *in* T *rooted at* u *containing* v.

Lemma 2. *In any strategy profile* $\mathbf{s}$ *that enables greedy routing with a tree metric, for any node* u*, in* $G(\mathbf{s})$ *agent* u *needs to have an edge to some node in every subtree in* $below(u)$ *of* T *rooted at* u*.*

Tree Metrics: Equilibrium Existence Here, we show the existence of equilibria and we completely characterize them.

Theorem 3. *In a tree metric, the network* G^T *is always a NE and a social optimum.*

Next, we show that GE are unique. This completely characterizes all GE and NE.

Theorem 4. *In a tree metric, the network* G^T *is the only GE.*

Proof. We prove the statement in two steps while assuming, for the sake of contradiction, that there is a GE network G that differs from G^T. First, we show that in this case, network G^T cannot be a proper subgraph of G, i.e., network G contains all the edges of G^T and at least one other edge. Then, in a second step, we show that if G^T is not a subgraph of G, i.e., if G does not contain all the edges of G^T, then G is not in GE. Hence, any GE network is identical to G^T.

To show that $G^T = (\mathcal{P}, E^T, \ell)$ cannot be a proper subgraph of $G = (\mathcal{P}, E, \ell)$, we observe that if this was the case, then $E^T \subset E$ holds, i.e., $E \setminus E^T \neq \emptyset$. Thus, there exists at least one edge $e \in E \setminus E^T$. However, by definition of G^T, the total strechcosts and the distances are minimized in the network G^T. Hence, removing the edge $e \in E \setminus E^T$ would be an improving move. A contradiction.

Let us now consider the case that G^T is not a subgraph of G, while G is a GE. We define $f_a(b)$ to be the number of descendants of node b in G_a^T. We consider a specific edge, for this, let $(u, v) = \arg\min_{(a,b)\in E^T \setminus E} f_a(b)$, i.e., the tuple (u, v) that is an edge in G^T but not in G, that minimizes the number of descendants of node v in G_u^T. Notice that such an edge always exists, since we assume that G^T is not a subgraph of G and thus, $E^T \setminus E \neq \emptyset$.

We now examine the greedy routing path from node u to node v in network G. Since $(u, v) \notin E$, in conjunction with the assumption that G is a GE, there must exist another node $x \neq v$, such that $(u, x) \in E$, that enables a greedy routing path from u to v in G, i.e., the edge (u, x) is the first edge on this greedy routing path. By Lemma 1, node x must belong to the same subtree T' as v in $below(u)$ in T rooted at u. Since v is a child of u in T rooted at u, we have that $T' = subtree(v)$. It follows that $(u, x) \notin E^T$, since in G_r^T node u only has v as child in T'.

Next, recall that we assume that $G = G(\mathbf{s}) = G(S_u, \mathbf{s}_{-u})$ is a GE. It follows that for agent u it is not an improving move to swap the edge $(u, x) \in E$ for the edge $(u, v) \notin E$. Now, let $S'_u = S_u \setminus \{x\}$ and let $\mathbf{s}^v = (S'_u \cup \{v\}, \mathbf{s}_{-u})$. For convenience, let $\mathbf{s}^x = \mathbf{s} = (S'_u \cup \{x\}, \mathbf{s}_{-u})$.

We first show, that after the swap, i.e., in $G(\mathbf{s}^v)$, agent u still has a greedy routing path to all nodes. Since in $G(\mathbf{s}^v)$ nothing changes for greedy routing paths that do not use x as first hop, it suffices to focus on such paths. Thus, consider a greedy routing path P from u to $z \in \mathcal{P}$ in $G(\mathbf{s}^x)$ that uses x as first hop. Since G is a GE, agent v has a greedy routing path Q to x in G. Moreover, since x is in $subtree(v)$, it follows that $d_T(v, x) < d_T(u, x)$. Thus, in $G(\mathbf{s}^v)$ agent u has a greedy routing path to x via node v and Q. This implies that also in $G(\mathbf{s}^v)$ agent u has a greedy routing path to z, since the path via v, Q and P is a greedy routing path.

The swap from (u, x) to (u, v) decreases agent u's stretch to v but it does not change u's edgecost. Since the swap does not reduce agent u's cost, there must be a node $w \in \mathcal{P}$ to which agent u's stretch increases after the swap, i.e., $\text{stretch}_{G(\mathbf{s}^x)}(u, w) < \text{stretch}_{G(\mathbf{s}^v)}(u, w)$. It follows, that in $G = G(\mathbf{s}^x)$, agent u has a greedy routing path via x to w.

By Lemma 1, it follows that both nodes x and w must belong to the same subtree T' in $below(u)$ in T rooted at u, which must be $T' = subtree(v)$.

Now, if $\text{stretch}_G(v, w) = 1$, then agent u's stretch to w must be optimal in $G(\mathbf{s}^v)$. Hence, the stretch cannot have increased compared to $G(\mathbf{s}^x)$. Thus, $\text{stretch}_G(v, w) > 1$ can be assumed. Then there there must be a vertex p along the path from v to w in T that in G does not create the edge to the next node q on that path. Thus, the edge (p, q) is in $E^T \setminus E$. Note that $f_p(q) < f_u(v)$, since q's subtree in G_p^T is completely contained in v's subtree in G_u^T. This is a contradiction to the choice of edge (u, v), i.e., to (u, v) having the minimum f-value. Thus, every edge of G^T must be contained in G. Hence, G^T is the only GE. □

Tree Metrics: Dynamic Properties We investigate if NE networks in tree metrics can be found by iteratively selecting best responses. The answer is affirmative.

Theorem 5. *In a tree metric our game is weakly acyclic under best responses.*

Proof. Fix any $u \in \mathcal{P}$ and root T in u. Consider any subtree $T' \in below(u)$. Recall that, by Lemma 2, agent u needs to build an edge (u, v), where $v \in T'$, in order to enable greedy routing from u to the set of nodes in T'. Note that if in a strategy-profile $\mathbf{s}$ all edges of T' are created in both directions, then the best response of agent u in $\mathbf{s}$ is to create the edge (u, u'), where u' is the root of subtree T', since $u' = \arg\min_{v\in V(T')} d_T(u, v)$. We use this to get from any strategy-profile $\mathbf{s}$ to the strategy-profile $\mathbf{s}^T$ by a finite sequence of best responses. For this, we root T in u and activate the agents of T in a bottom-up fashion, i.e., starting with the leaves and then moving upwards. This ensures that all edges of G^T are eventually created.

After all edges of G^T are created, we activate all agents that still have edges that do not belong to G^T. Any such edge will be removed since the edges of G^T already suffice to achieve optimal stretchcosts. Thus, since G^T is a NE, our game is weakly acyclic under best responses. □

Tree Metrics: Computational Complexity Here, we investigate the computational complexity of computing a best response in tree metrics and show that this is NP-hard. However, since by Theorem 4 G^T is the unique GE, deciding NE existence and computing a NE is tractable.

Theorem 6. *In tree metrics, computing a best response strategy is NP-hard.*

Euclidean Metrics

We study Euclidean metrics, which are metrics where there is a function that maps agents to points in a d-dimensional Euclidean space, such that the distances in the metric between agents correspond to the distances of their points in the Euclidean space, measured via the 2-norm. We focus on 2D-Euclidean metrics but all results regarding the existence of equilibria and computational complexity directly apply to higher dimensional spaces as well. Given three nodes $u, v, w \in \mathcal{P}$ in a Euclidean metric, let $\angle vuw$ be the angle of u formed by rays $\overrightarrow{uv}$ and $\overrightarrow{uw}$. We always consider the positive angle, which is at most π.

Euclidean Metrics: Equilibrium Existence First, we show the negative result that GE may not exist. We modify a proof from Moscibroda, Schmid, and Wattenhofer (2006).

Theorem 7. *In a 2D Euclidean metric there are instances of our game that do not admit GE.*

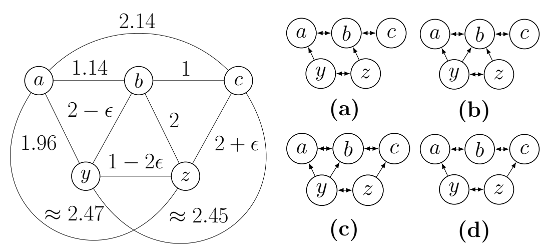

Figure 2: Left: 2D Euclidean instance that does not admit a GE. Right: (a)–(d): Equilibrium strategies for y and z.

Proofsketch. Let $\epsilon > 0$ be an arbitrarily small constant. We show that there are no GE in the instance in Figure 2 (left) for $\alpha = 0.6$. To that end, we first establish that only the strategies for agents y and z shown in Figure 2 (right) could possibly be part of any GE strategy-profile. Then, we show that none of these strategies can be part of any GE. □

Euclidean Metrics: Computational Complexity Next, we show that finding a best response strategy is hard.

Theorem 8. *In a 2D Euclidean metric, computing a best response is NP-hard.*

Euclidean Metrics: Approximate Equilibria Since GE do not always exist and finding a best response is computationally hard, approximate equilibria are the only remaining option to construct almost stable networks for the practically important setting of Euclidean metrics in polynomial time.

To that end, we employ Θ_k-graphs for 2D-Euclidean metrics. First introduced independently by Clarkson (Clarkson 1987) and Keil (Keil 1988), Θ_k-graphs are constructed as follows: Each node u partitions the plane into k disjoint cones with itself as the apex, each having an aperture of $\frac{2\pi}{k}$. Then, for each cone, node u adds an edge to the node v whose projection onto the bisector of the cone is closest to u.

Θ-routing is a way of selecting paths in a Θ_k-graph. With this, to get from a node u to a node v, the edge is used that u created into the cone that contains node v. This procedure is repeated for each hop until node v is reached. First, we show that if k is too small, Θ_k-graphs are not suited as approximate NEs. This result is well-known (Narasimhan and Smid 2007), but we prove it for the sake of completeness.

Theorem 9. *For every $k \leq 5$, there exist 2D-Euclidean metrics where the Θ_k graph does not enable greedy routing.*

With this limitation in mind, we give a general upper bound on the approximation ratio. Let $f(k) = \frac{1}{1-2sin(\frac{\pi}{k})}$ be the maximum stretch of the Θ-routing in a Θ_k-graph with $6 < k < n$ (Ruppert and Seidel 1991).

Theorem 10. *If $6 < k < n$, any 2D-Euclidean instance of our game has a $f(k) + \alpha\frac{k}{n-1}$-NE.*

While this does not yield a constant approximation ratio, we note that, as $n \to \infty$, the ratio goes to $f(k)$, which can be arbitrarily close to 1; for $k \geq 15$, it is below 2.

Now, we establish our main result: The existence of a constant factor approximation that can be efficiently computed.

Theorem 11. *Every 2D-Euclidean instance of the game has a 5-approximate NE.*

Proof. Consider the Θ_8-graph on $\mathcal{P}$. By Bose, De Carufel, Morin, van Renssen and Verdonschot (Bose et al. 2016), Θ-routing in any Θ_{4k+4}-graph gives stretches of at most $1 + \frac{2\sin(\frac{\pi}{4k+4})}{\cos(\frac{\pi}{4k+4})-\sin(\frac{\pi}{4k+4})}$. Thus, the maximum stretch of Θ-routing in the Θ_8-graph is $1 + \frac{2\sin(\frac{\pi}{8})}{\cos(\frac{\pi}{8})-\sin(\frac{\pi}{8})} = 1 + \sqrt{2}$, which yields a greedy routing path. In the best response of an agent, its stretch to any other agent is at least 1.

Every agent u for whom the other agents are not within a cone with angle at most π must build at least two edges in its best response. This is because an edge to any node v can only be part of a greedy routing path to a node i with $\angle vui \leq \frac{\pi}{2}$, because otherwise, by the law of cosines,

$$d(v,i) = \sqrt{d(u,i)^2 + d(u,v)^2 - 2d(u,i)d(u,v)\cos(\angle vui)} > \sqrt{d(u,i)^2} = d(u,i),$$

and as such the path would not be a greedy routing path.

For every agent u for whom the other agents are in a cone with angle at most π, at least three cones of the Θ_8-graph are empty. Thus, agent u only builds at most five edges.

Finally, we consider our stretchcost function. Edgecosts are at least α or 2α in the best response, and at most 5α or 8α in the Θ_8-graph, depending on whether all other agents are within a cone with angle at most π or not.

Let $u \in \mathcal{P}$. Also, let $\mathbf{s}_\Theta$ be the strategy profiles of the Θ_8-graph in a metric, where for agent u not all other agents are in a cone with angle at most π. Also, let $\mathbf{s}_\Theta^\pi$ be the strategy profiles of the Θ_8-graph where all other agents are within a cone with angle at most π. Let $\mathbf{s}_{br}$ and $\mathbf{s}_{br}^\pi$ be the corresponding strategy profiles, where agent u's strategy is changed to its best response, while all other strategies stay unchanged.

For our stretchcost function, we get that the Θ_8-graph is a

$$\max\left(\frac{c_u(s_\Theta)}{c_u(s_{br})}, \frac{c_u(s_\Theta^\pi)}{c_u(s_{br}^\pi)}\right) =$$

$$\max\left(\frac{8\alpha + (1+\sqrt{2})(n-1)}{2\alpha + (n-1)}, \frac{5\alpha + (1+\sqrt{2})(n-1)}{\alpha + (n-1)}\right)$$

-approximate NE, because the stretch to all $n-1$ other agents is at least 1 in a best response and at most $1+\sqrt{2}$ in the Θ_8-graph. This approximation factor is at most 5. □

Note that Theorem 11 does not imply a constant bound for 1-2 metrics, since there are instances in 1-2 metrics that cannot be embedded into the Euclidean plane.

The following theorem gives a lower bound on the approximation ratio of Θ_k-graphs. For the Θ_8-graphs we used in the last theorem, this bound is tight.

Theorem 12. *There are 2D-Euclidean instances of our game, where the Θ_k-graph is not a $(\lceil \frac{k}{2} \rceil + 1 - \epsilon)$-NE.*

Proof. We equip each node with polar coordinates in the Euclidean plane. For now, we assume that the reference direction of the coordinate system is equal to either the bisector of some cone used in the construction of the Θ_k-graph if $\lceil \frac{k}{2} \rceil + 1$ is odd or the edge of a cone otherwise. Let u be located at the origin and let v be at distance 1 and angle 0. Let the nodes $w_1, ..., w_{\lceil \frac{k}{2} \rceil + 1}$ be at angles evenly spaced out over the interval $[-\frac{\pi}{2} + \frac{2\pi}{5k}; \frac{\pi}{2} - \frac{2\pi}{5k}]$ and at distances $\frac{1}{\cos(\angle vuw_i)}$ (this is possible, as $\angle vuw_i < \frac{\pi}{2}$ and thus $\cos(\angle vuw_i) > 0$). We call this construction C. Agent u builds edges to $\lceil \frac{k}{2} \rceil + 1$ nodes in the Θ_k-graph because each w_i is in a different cone. For all i, we have, by the law of cosines, that

$$\begin{aligned} d(v, w_i) &= \sqrt{1^2 + d(u, w_i)^2 - 2 \cdot 1 \cdot d(u, w_i) \cos(\angle vuw_i)} \\ &= \sqrt{d(u, w_i)^2 - 1} < d(u, w_i). \end{aligned}$$

As such, agent v cannot use u on greedy routing paths to any w_i, whereas u can use v for all of them. If v would not have a greedy routing path to some w_i, this could not be a $(\lceil \frac{k}{2} \rceil + 1 - \epsilon)$-approximate NE because adding edges to all w_i without a greedy routing path would improve v's costs from at least Z to less than Z, which by definition of Z is a large improvement, also by more than a factor of $(\lceil \frac{k}{2} \rceil + 1 - \epsilon)$. Thus, agent u could replace all of its edges with an edge to v and still retain greedy routing paths to all nodes. Let sc and sc' be agent u's stretchcosts before and after this move. Agent u's total costs would improve by a factor of $\frac{(\lceil \frac{k}{2} \rceil + 1)\alpha + sc'}{\alpha + sc}$ which is at least $\lceil \frac{k}{2} \rceil + 1 - \epsilon$ for $\alpha \geq \frac{(\lceil \frac{k}{2} \rceil + 1)sc - sc' - \epsilon}{\epsilon}$.

In a slightly modified construction, even a global rotation of the cones (i.e. the reference direction of the coordinate system not lining up with the bisector/edge of any cone) cannot achieve better results: first, we note that with cone-rotations of less than $\frac{4\pi}{5k}$ in either direction, the same cones stay occupied by nodes and as such, this construction still holds. Also, a rotation by β is equal to a rotation by $\beta \mod \frac{2\pi}{k}$. Thus, we place two copies of C, which we call C_1 and C_2. The copy C_2 is rotated by $i\frac{2\pi}{k} + \frac{\pi}{k}$, for some $i \in \mathbb{N}$, compared to C_1, and its center is displaced. With this, in any rotation, one of u_1 or u_2 still has to build $\lceil \frac{k}{2} \rceil + 1$ many edges. Additionally, by placing the C_1 and C_2 far enough apart and choosing i for the rotation to ensure that they are each contained in cones of u of the other that are occupied by some w_i, these do not impact the number of edges agent u builds in the Θ_k-graph or needs to build in its best response. □

Thus, we cannot get a better bound on the approximation ratio of Θ_8-graphs. Also, the approximation factor does not improve by choosing a smaller k: With $k = 7$ being an odd number, Theorem 12 still gives the same lower bound of $5 - \epsilon$. With $k = 6$ the known bound on the stretch goes up to at least 7 (and that is not along necessarily greedy routing paths; to get greedy routing paths the stretch might be as large as the bound on Θ-routing of $12\sqrt{3}$) (Akitaya, Biniaz, and Bose 2022). For smaller values of k, by Theorem 9, there might be no greedy routing paths between some nodes.

General Metrics

Now we consider general metric spaces. Naturally, all negative results from the preceding sections, like hardness and non-convergence results, immediately carry over. Specifically, the results that are not implicit in the proofs in this section is that computing best responses is NP-hard, that GE are not $\Omega(\frac{\alpha n}{\alpha + n})$ -approximate NE and that GE may not exist.

General Metrics: Hardness of Equilibrium Existence Here, we show that deciding whether GE and NE exist in a given instance with an arbitrary metric is NP-hard.

Theorem 13. *In general metric spaces, it is both NP-hard to decide, if an instance admits a NE or a GE.*

Approximate Equilibria We show the existence of very general approximate equilibria, where the approximation factor does not depend linearly on n or on the distances given by the metric but on the edge price parameter α. Depending on the setting, this can be seen as a positive result.

Theorem 14. *Every instance of our game has a $(\alpha + 1)$-approximate NE.*

Conclusion

We give the first game-theoretic network formation model that focuses on the creation of networks where all-pairs greedy routing is enabled and where the agents optimize their connection quality, i.e., their stretches.

We believe that this is only the first step to further models that guarantee even more favorable properties to hold in the created networks. For example, a guaranteed maximum stretch and robustness aspects like coping with edge or node failures. Another avenue for further research is to consider different edge price functions. Our goal was to first provide insights for the simplest setting, i.e., where every edge has a uniform cost of α, that will then be useful for investigating more complex variants. And this actually works: Based on the presented results we have preliminary results that hold for the version where edge costs are proportional to the edge length, as considered in (Bilò et al. 2019). Finally, it is interesting to explore other techniques for constructing approximate Nash equilibria in the Euclidean setting.

Acknowledgments

This work was supported by DFG project GEONET under grant DFG 442003138.

References

Abam, M.; and Qafari, M. 2019. Geometric Spanner Games. *Theoretical Computer Science*, 795.

Adamaszek, A.; Mnich, M.; and Paluch, K. 2018. New Approximation Algorithms for (1,2)-TSP. In *ICALP'18*, 9:1–9:14.

Akitaya, H. A.; Biniaz, A.; and Bose, P. 2022. On the Spanning and Routing Ratios of the Directed Θ6-graph. *Computational Geometry*, 105–106: 101881.

Berger, J.; Friedrich, T.; Lenzner, P.; Machaira, P.; and Ruff, J. 2024. Strategic Network Creation for Enabling Greedy Routing. *CoRR*, abs/2403.15307.

Bilò, D.; Cohen, S.; Friedrich, T.; Gawendowicz, H.; Klodt, N.; Lenzner, P.; and Skretas, G. 2023. Temporal Network Creation Games. In *IJCAI 2023*, 2511–2519.

Bilò, D.; Friedrich, T.; Lenzner, P.; Lowski, S.; and Melnichenko, A. 2021. Selfish Creation of Social Networks. In *AAAI'21*, 5185–5193.

Bilò, D.; Friedrich, T.; Lenzner, P.; and Melnichenko, A. 2019. Geometric Network Creation Games. In *SPAA '19*, 323–332.

Bilò, D.; Gualà, L.; Leucci, S.; and Proietti, G. 2016. Locality-Based Network Creation Games. *ACM Trans. Parallel Comput.*, 3(1): 6:1–6:26.

Bläsius, T.; Friedrich, T.; Katzmann, M.; and Krohmer, A. 2020. Hyperbolic Embeddings for Near-Optimal Greedy Routing. *ACM Journal on Experimental Algorithmics*, 25: 1–18.

Boguná, M.; Papadopoulos, F.; and Krioukov, D. 2010. Sustaining the internet with hyperbolic mapping. *Nature communications*, 1(1): 62.

Bose, P.; De Carufel, J.-L.; Morin, P.; van Renssen, A.; and Verdonschot, S. 2016. Towards Tight Bounds on Theta-Graphs: More Is Not Always Better. *Theoretical Computer Science*, 616: 70–93.

Bullinger, M.; Lenzner, P.; and Melnichenko, A. 2022. Network Creation with Homophilic Agents. In *IJCAI 2022*, 151–157.

Chauhan, A.; Lenzner, P.; Melnichenko, A.; and Molitor, L. 2017. Selfish Network Creation with Non-uniform Edge Cost. In *SAGT'17*, 160–172.

Chauhan, A.; Lenzner, P.; Melnichenko, A.; and Münn, M. 2016. On Selfish Creation of Robust Networks. In *SAGT'16*, 141–152.

Clarkson, K. 1987. Approximation Algorithms for Shortest Path Motion Planning. In *STOC '87*, 56–65.

Corbo, J.; and Parkes, D. C. 2005. The price of selfish behavior in bilateral network formation. In *PODC'05*, 99–107.

Cord-Landwehr, A.; and Lenzner, P. 2015. Network Creation Games: Think Global - Act Local. In *MFCS'15*, 248–260.

Echzell, H.; Friedrich, T.; Lenzner, P.; and Melnichenko, A. 2020. Flow-Based Network Creation Games. In *IJCAI'20*, 139–145.

Ehsani, S.; Fazli, M.; Mehrabian, A.; Sadeghabad, S. S.; Safari, M.; Saghafian, M.; and ShokatFadaee, S. 2011. On a bounded budget network creation game. In *SPAA'11*, 207–214.

Eidenbenz, S. J.; Kumar, V. S. A.; and Zust, S. 2006. Equilibria in Topology Control Games for Ad Hoc Networks. *Mobile Networks and Applications*, 11(2): 143–159.

Fabrikant, A.; Luthra, A.; Maneva, E.; Papadimitriou, C.; and Shenker, S. 2003. On a Network Creation Game. In *PODC'03*, 347–351.

Fakcharoenphol, J.; Rao, S.; and Talwar, K. 2004. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. Syst. Sci.*, 69(3): 485–497.

Friedemann, W.; Friedrich, T.; Gawendowicz, H.; Lenzner, P.; Melnichenko, A.; Peters, J.; Stephan, D.; and Vaichenker, M. 2021. Efficiency and Stability in Euclidean Network Design. In *SPAA '21*, 232–242.

Friedrich, T.; Gawendowicz, H.; Lenzner, P.; and Zahn, A. 2023. The Impact of Cooperation in Bilateral Network Creation. In *PODC'23*, 321–331.

Karp, R. M. 1972. *Reducibility among Combinatorial Problems*, 85–103.

Kawald, B.; and Lenzner, P. 2013. On dynamics in selfish network creation. In *SPAA '13*, 83–92.

Keil, J. M. 1988. Approximating the Complete Euclidean Graph. In *SWAT 88*, 208–213.

Laoutaris, N.; Poplawski, L. J.; Rajaraman, R.; Sundaram, R.; and Teng, S. 2008. Bounded budget connection (BBC) games or how to make friends and influence people, on a budget. In *PODC'08*, 165–174.

Lenzner, P. 2012. Greedy Selfish Network Creation. In *WINE'12*, 142–155.

Mauve, M.; Widmer, J.; and Hartenstein, H. 2001. A survey on position-based routing in mobile ad hoc networks. *IEEE Netw.*, 15(6): 30–39.

Meirom, E. A.; Mannor, S.; and Orda, A. 2014. Network formation games with heterogeneous players and the internet structure. In *EC'14*, 735–752.

Meirom, E. A.; Mannor, S.; and Orda, A. 2015. Formation games of reliable networks. In *INFOCOM'15*, 1760–1768.

Monderer, D.; and Shapley, L. S. 1996. Potential games. *Games and economic behavior*, 14(1): 124–143.

Moscibroda, T.; Schmid, S.; and Wattenhofer, R. 2006. On the Topologies Formed by Selfish Peers. In *PODC '06*, 133–142.

Narasimhan, G.; and Smid, M. 2007. *Geometric Spanner Networks*. Cambridge University Press.

Nisan, N.; Roughgarden, T.; Tardos, É.; and Vazirani, V. V., eds. 2007. *Algorithmic Game Theory*. Cambridge University Press. ISBN 9780511800481.

Papadimitriou, C. H. 2001. Algorithms, games, and the internet. In *STOC'01*, 749–753.

Ramanathan, R.; and Redi, J. 2002. A brief overview of ad hoc networks: challenges and directions. *IEEE Commun. Mag.*, 40(5): 20–22.

Ruppert, J.; and Seidel, R. 1991. Approximating the D-Dimensional Complete Euclidean Graph. In *3rd Canadian Conference on Computational Geometry*.

Shoham, Y.; and Leyton-Brown, K. 2009. *Multiagent Systems - Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press. ISBN 978-0-521-89943-7.