

Greedy Routing Reachability Games

Pascal Lenzner
University of Augsburg
Augsburg, Germany
pascal.lenzner@uni-a.de

Paraskevi Machaira
Hasso Plattner Institute
Potsdam, Germany
paraskevi.machaira@hpi.de

ABSTRACT

Today’s networks consist of many autonomous entities that follow their own objectives, i.e., smart devices or parts of large AI systems, that are interconnected. Given the size and complexity of most communication networks, each entity typically only has a local view and thus must rely on a local routing protocol for sending and forwarding packets. A common solution for this is greedy routing, where packets are locally forwarded to a neighbor in the network that is closer to the packet’s destination.

In this paper we investigate a game-theoretic model with autonomous agents that aim at forming a network where greedy routing is enabled. The agents are positioned in a metric space and each agent tries to establish as few links as possible, while maintaining that it can reach every other agent via greedy routing. Thus, this model captures how greedy routing networks are formed without any assumption on the distribution of the agents or the specific employed greedy routing protocol. Hence, it distills the essence that makes greedy routing work.

We study two variants of the model: with directed edges or with undirected edges. For the former, we show that equilibria exist, have optimal total cost, and that in Euclidean metrics they can be found efficiently. However, even for this simple setting computing optimal strategies is NP-hard. For the much more challenging setting with undirected edges, we show for the realistic setting with agents in 2D Euclidean space that the price of anarchy is between 1.75 and 1.8 and for higher dimensions it is less than 2. Also, we show that best response dynamics may cycle, but that in Euclidean space almost optimal approximate equilibria can be computed in polynomial time. Moreover, for 2D Euclidean space, these approximate equilibria outperform the well-known Delaunay triangulation.

KEYWORDS

greedy routing, network creation games, Nash equilibrium, price of anarchy, game theory, geometry

ACM Reference Format:

Pascal Lenzner and Paraskevi Machaira. 2026. Greedy Routing Reachability Games. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 9 pages. <https://doi.org/10.65109/JJND5892>

1 INTRODUCTION

Greedy routing is widely used in many applications, because it enables communication between the nodes of a network without

requiring complete knowledge of the network’s structure. Given that our communication networks today are large, complex, and potentially also dynamically changing, using traditional shortest path based routing, facilitated via fixed routing tables that must be distributed and maintained at all important network nodes, will have to be replaced by local and more adaptive alternatives. One of the most prominent options for this is greedy routing, also known as geographic routing or stateless routing [21, 27]. While common for small ad-hoc peer-to-peer networks of smart devices or sensors, greedy routing was recently proposed for the whole Internet [8].

The key for greedy routing is that every network node has a position in some underlying space, typically a metric space like the Euclidean plane or the hyperbolic disc. These positions allow for deciding locally which network neighbor to use as the next hop in packet routing. If at every step of the greedy routing path always a next hop exists that is strictly closer to the target position than the current node then greedy routing succeeds. If greedy routing succeeds between all pairs of nodes, the network is called *navigable*. Thus, the challenge is to create a navigable network.

Significant research has been conducted on greedy routing but mostly with the goal of finding a suitable greedy embedding, i.e., a mapping of nodes to virtual coordinates in some space to facilitate routing [6, 13, 40]. For this, a fixed network is given and the coordinates from the embedding guide packets on their way to their destination. A different approach and equally important, is the setting where there is no predefined network and the nodes have to construct a network among themselves such that greedy routing is enabled. One prominent application are first responder networks in areas where a natural disaster like an earthquake or tsunami has destroyed the communication infrastructure. The idea is that communication sensors are dropped in the area that then establish a communication network that can be used by rescue teams and disaster relief workers [31, 34]. For such applications it is typically assumed that the process of setting up the network is coordinated in the sense that all devices run the same local algorithm. Another application is database similarity search where constructing navigable networks is mentioned as an open problem [2, 16].

While embeddings work with a fixed network and try to find suitable node positions, shifting the focus away from greedy embeddings highlights two important aspects of the problem. First, and ever more important in the current age of smart devices, it allows the study of a decentralized setting without a central authority dictating the network structure. Instead, following a game-theoretic approach [37] and in particular the well-known paradigm of network creation games [3, 19], the nodes themselves make strategic decisions on which edges to build in order to enable greedy routing, calling for a game-theoretic analysis. Second, removing the constraints imposed by a given network reveals the fundamental mechanisms and structural properties that greedy routing networks



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

necessarily require. This understanding is essential for designing optimal networks that enable greedy routing and that have certain structural properties, such as minimizing the total cost.

In this paper we set out to investigate the structural properties of navigable networks that arise from the interaction of selfish agents. For this, we consider a very basic network formation game, where agents that correspond to nodes with a fixed position in a metric space strategically set up links to other nodes to enable greedy routing paths to all network nodes. This objective aims at greedy routing reachability, i.e., the most basic property needed for all agents. With this, we lay the foundation for more complex models that besides greedy navigability might add more objectives, like robustness or stretch guarantees. We find that almost stable states exist and can be efficiently computed. Moreover, these equilibria have favorable properties since they are guaranteed to be close to centrally optimized networks in terms of cost, i.e., in terms of the number of built edges. Smart agents can make independent decisions. We show that such agents do not need central coordination to create almost optimal navigable networks.

1.1 Model and Preliminaries

We consider n agents corresponding to points $\mathcal{P} = \{p_1, \dots, p_n\}$ in a metric space $\mathcal{M} = (\mathcal{P}, d_{\mathcal{M}})$, where $d_{\mathcal{M}}(u, v)$ denotes the distance between any two points $u, v \in \mathcal{P}$, with $d_{\mathcal{M}}(u, u) = 0$ for all $u \in \mathcal{P}$, and where the triangle inequality holds. If clear from the context, we omit \mathcal{M} . The agents, knowing the positions of all points, aim to create a network that enables greedy routing. For this, each agent decides on a set of incident edges it wants to establish.

The *strategy of agent $u \in \mathcal{P}$* is S_u , a set of edges with endpoints in $\mathcal{P} \setminus u$, i.e., agent u can buy incident edges to any subset of other agents. For each edge in S_u , we say that agent u *owns* the edge. We consider two variants. In the directed case, each edge in S_u is of the form (u, v) , representing that agent u buys a directed edge to v . In the undirected case, each edge in S_u is of the form $\{u, v\}$, denoting that agent u buys an undirected edge to v . Let $\mathbf{s} = (S_1, \dots, S_n)$ be the *strategy-profile*, i.e., the vector of strategies of all agents. Also, for agent $u \in \mathcal{P}$ let $\mathbf{s} = (S_u, \mathbf{s}_{-u})$, where \mathbf{s}_{-u} is the vector of strategies of all agents except u . We sometimes use the notation $S_u(\mathbf{s})$ to explicitly refer to the strategy profile which contains strategy S_u . Any strategy-profile \mathbf{s} defines a network $G(\mathbf{s}) = (\mathcal{P}, E(\mathbf{s}))$, where $E(\mathbf{s}) = \bigcup_{u \in \mathcal{P}} S_u$ in case of directed edges or $E(\mathbf{s}) = \bigcup_{u \in \mathcal{P}} \{\{u, v\} \mid \{u, v\} \in S_u \vee \{u, v\} \in S_v\}$ in case of undirected edges. Vertices of $G(\mathbf{s})$ are called agents, nodes, or points.

Given a network $G = (\mathcal{P}, E)$, a *greedy routing path from u to v* in G is a sequence (x_1, x_2, \dots, x_j) , with $x_i \in \mathcal{P}$, for $1 \leq i \leq j$, where $x_1 = u$, $x_j = v$, and $(x_i, x_{i+1}) \in E$ (or $\{x_i, x_{i+1}\} \in E$, for undirected edges), for $1 \leq i \leq j - 1$, such that $d(x_i, v) > d(x_{i+1}, v)$ holds for all $1 \leq i \leq j - 1$. Thus, such a path is a directed (undirected) 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. We say that *greedy routing is enabled* for an agent u in network G , or that u is *greedy connected*, if in G there exists a greedy routing path from u to every other node. If this holds for every node in G , then we say that greedy routing is enabled in G . See Figure 1.




Figure 1: Illustration of greedy routing paths in \mathbb{R}^2 . The pink and blue s - t -paths are greedy routing paths, while the purple s - t -path is not, since $15 = d(s, t) < d(u, t) = \sqrt{12^2 + 15^2}$. Greedy routing is enabled only for node s .

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

$$c_u(\mathbf{s}) = |S_u| + \omega,$$

where ω serves as a penalty term, i.e., it is equal to 0 if greedy routing is enabled for agent u and ∞ otherwise. The *social cost* of network $G(\mathbf{s})$ is $c(\mathbf{s}) = \sum_{u \in \mathcal{P}} c_u(\mathbf{s})$, i.e., the total cost over all agents. For a set of points \mathcal{P} , the network $G(\mathbf{s}^*) = (\mathcal{P}, E(\mathbf{s}^*))$ minimizing the social cost is called the *social optimum network (SO)* for \mathcal{P} .

A strategy S_u^* is called a *best response* of agent u in strategy-profile $\mathbf{s} = (S_u, \mathbf{s}_{-u})$, if $c_u(S_u^*, \mathbf{s}_{-u}) \leq c_u(S'_u, \mathbf{s}_{-u})$ for any other strategy S'_u where in the directed case $S'_u \subseteq \{(u, v) \mid v \in \mathcal{P}, v \neq u\}$, while in the undirected case $S'_u \subseteq \{\{u, v\} \mid v \in \mathcal{P}, v \neq u\}$. That is, a best response strategy minimizes agent u 's cost given that the other agents' strategies are fixed. A strategy-profile \mathbf{s} is a *pure Nash Equilibrium (NE)*, or *stable* for short, if in profile \mathbf{s} every agent already plays a best response. We denote by NE the set of all strategy profiles that are pure Nash equilibria. For the setting with directed edges, we have a bijection between strategy-profiles \mathbf{s} and the corresponding networks $G(\mathbf{s})$. In the case of undirected edges, the same holds if also the edge ownership information is considered. In both cases, we will say that a network $G(\mathbf{s})$ is in NE, if the strategy profile $\mathbf{s} \in \text{NE}^1$. Note that in every NE network greedy routing must be enabled, since every agent could buy direct edges to every other node to avoid the infinite penalty ω .

We also consider weaker versions of stability. A network $G(\mathbf{s})$ is in β -*approximate NE* (β -NE or β -stable) 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 any strategy change. Similarly, $G(\mathbf{s})$ is in γ -*additive NE* ($+\gamma$ -NE or $+\gamma$ -stable) if no agent u can decrease its cost below $c_u(\mathbf{s}) - \gamma$. If greedy routing is enabled, then in a β -stable network, agents buy at most a factor of β too many edges, while in a $+\gamma$ -stable network they buy at most γ extra edges. Thus, 1-stable and $+\infty$ -stable networks are in NE.

The Price of Anarchy (PoA) [29] is defined as the worst-case ratio of the social cost of any NE and the corresponding social optimum.

A *best response path* is a sequence $\mathbf{s}_0, \mathbf{s}_1, \dots, \mathbf{s}_k$ of strategy-profiles, such that \mathbf{s}_i results from some agent switching to a best response in \mathbf{s}_{i-1} , for $1 \leq i \leq k$. A *best response cycle (BRC)* is a cyclic best response path, i.e., $\mathbf{s}_0 = \mathbf{s}_k$, and its existence implies that a game cannot have an ordinal potential function [36]. Hence, it is not a potential games for which the existence of a NE is guaranteed.

¹For undirected edges, in NE no edge can be bought by both endpoints.




Figure 2: (right) The nearest neighbor graph of 20 points \mathcal{P} in \mathbb{R}^2 , (left) The Delaunay triangulation of \mathcal{P} .

For each point $u \in \mathcal{P}$ let the set of *nearest neighbors* of u be $N(u) = \arg \min_{v \in \mathcal{P} \setminus \{u\}} d_{\mathcal{M}}(u, v)$. Thus, the set $N(u)$ consists of points in \mathcal{P} with minimum distance to u . The *Nearest Neighbor Graph (NNG)* on \mathcal{P} , denoted by $G^{NNG}(\mathcal{P}, E)$ is defined as follows. In the directed case, the edge set is $E = \{(u, v) \mid v \in N(u)\}$, i.e., each point has a directed edge to its nearest neighbors, while in the undirected case, $E = \{\{u, v\} \mid v \in N(u) \vee u \in N(v)\}$. See Figure 2 (left) for an example. It is known that the NNG has to be a subgraph of every graph that supports greedy routing [38].

A *Delaunay triangulation (DT)* [15] on a set of points \mathcal{P} in the D -dimensional Euclidean space is a triangulation such that no point in \mathcal{P} is inside the circum-hypersphere of any D -simplex in $DT(\mathcal{P})$ (see Figure 2 right) for an example. An important structural property is that the NNG on \mathcal{P} is a subgraph of $DT(\mathcal{P})$. Moreover, it has been shown in [9] that Delaunay triangulations in two-dimensional Euclidean space support greedy routing. For this case, the Delaunay triangulation is a planar graph, i.e., it has at most $3|\mathcal{P}| - 6$ edges.

The *kissing number* in \mathbb{R}^D , denoted by $K(D)$ is the maximum number of disjoint unit hyperspheres that can simultaneously touch a given unit hypersphere in \mathbb{R}^D . Equivalently [12], $K(D)$ is the maximum number of points that can be placed on the surface of a unit sphere in \mathbb{R}^D so that the angular distance between any two points is at least 60° . We have $K(D) \leq 2^{0.401D}$ but exact values of $K(D)$ are known only for $D \in \{1, 2, 3, 4, 8, 24\}$. See Figure 3.




Figure 3: The kissing number in \mathbb{R}^3 is 12.

For an agent $u \in \mathcal{P}$ embedded in a metric space, we define the *greedy routing set* as a set of edges, such that, after connecting with these edges, agent u has at least one neighbor v satisfying $d(v, w) < d(u, w)$, for all $w \in \mathcal{P} \setminus \{u\}$. A *minimum greedy routing set*, denoted $\Phi(u)$, is the smallest such set for an agent u , and its size $\phi(u) = |\Phi(u)|$ is called the *greedy routing degree* of agent u . In other words, $\phi(u)$ measures the minimum connectivity effort required for agent u to locally enable greedy routing. By $\phi^+(u)$, we denote the number of these edges that belong to the NNG. See Figure 4.

1.2 Related Work

In greedy routing [21], each node forwards a packet to the neighbor that is geographically closest to the destination, using only local information. However, a packet may get stuck at a *local minimum* in the network. Combining greedy forwarding with *perimeter routing*




Figure 4: (left) A greedy routing set of agent u , (right) the minimum greedy routing set with $\phi(u) = 3$ and $\phi^+(u) = 1$.

to recover from such local minima is introduced independently by Karp and Kung [27] and Bose et al. [10]. Subsequent studies extended this approach to handle routing voids and to improve delivery guarantees [14, 20, 30, 33]. Another method for overcoming voids is assigning virtual coordinates to the nodes, constructing a greedy embedding [6, 18, 28]. While such embeddings can be computed efficiently, they introduce limitations and do not express real distances. Despite the extensive research on greedy routing, not much is known for the problem of constructing navigable networks.

The Delaunay triangulation (DT) [15] plays a central role in navigable networks as it enables greedy routing in \mathbb{R}^2 [9] and \mathbb{R}^3 [39]. Also, if the destination can be any point in the continuous space rather than a discrete set of nodes, the DT without degenerate edges is a subgraph of every graph that supports greedy routing [24].

Navigable networks in higher dimensional spaces have attracted interest because of their application in distributed similarity search and high-dimensional data indexing [16, 23]. Structures like hierarchical navigable small-world networks [35] use greedy routing principles to perform efficient approximate nearest-neighbor search. Most of the work on this is focused on empirical studies.

Since navigable networks share many structural properties with real-world complex networks [7], game theory provides a useful framework for understanding how such networks form and evolve over time. Bala and Goyal [3] introduced a non-cooperative network formation game where agents correspond to network nodes and buy incident edges to reach every other agent. In contrast to our model, any path connecting two nodes establishes reachability. They show that NE exist and can be easily found via suitable game dynamics. More involved are the network creation games (NCGs) by Fabrikant et al. [19], where the agents buy incident undirected edges to minimize the sum of hop distances to all other agents. They show that NE exist and have a low PoA, but computing a best response is NP-hard. NCG variants that involve geometry are proposed by Eidenbenz et al. [17], where agents corresponding to points in the plane establish undirected edges to ensure reachability, by Abam et al. [1], where the agents aim for a low maximum stretch (defined as the ratio of graph distance to metric space distance), and by Bilò et al. [5] and Friedemann et al. [22], where agents aim to minimize their sum of shortest path distances to the other agents.

The works most closely to ours, where greedy routing is studied within the network creation framework, are those of Gulyás et al. [25] and Berger et al. [4]. In [25] agents aim to reach other agents via greedy routing in hyperbolic space. They study the directed case where agents do not influence each other and prove that computing a minimum greedy routing set is NP-hard. In [4], each connection has a cost $\alpha > 0$, and agents aim to create a directed network in which greedy routing works, while minimizing connection costs and the sum of stretches to the other agents. Most results are limited to 1-2 metrics (where $d_{\mathcal{M}}(u, v) \in \{1, 2\}$ for any $u, v \in \mathcal{P}$) and tree

metrics, although they prove the existence of a 5-approximate NE in \mathbb{R}^2 , and an $(\alpha + 1)$ -NE in general metric spaces.

1.3 Our Contribution

Our work establishes the theoretical foundations that bridge game theory and computational geometry, introducing tools for developing more complex models for network creation. While most of the literature focuses on greedy embeddings or on greedy routing in families of pre-constructed networks with directed edges, we study the process of creating networks, directed or undirected, that enable greedy routing. In contrast to existing work, our model captures how greedy routing networks are formed without assumptions on the distribution of the agents or the specific greedy routing protocol. Hence, it distills the essence that makes greedy routing work.

In particular, for agents in 2D Euclidean metrics we give an algorithm that efficiently constructs a navigable network that is almost stable from a game-theoretic point of view, while having at most 80% more edges than the sparsest possible navigable network. This improves on the approximation factor of 3 achieved by the well-known Delaunay triangulation. Moreover, this can be seen as the basic step needed for designing efficient algorithms for creating networks that incorporate additional aspects like stretch or robustness guarantees. The created networks are useful for applications such as sensors that are randomly dropped over an area, IoT networks, and other situations where a communication network has to be formed among smart devices. Also, our approach adapts well to dynamic network topologies, is independent of the employed routing protocol, and it is more robust, since it only relies on local communication between immediate neighbors.

While focusing on 2D Euclidean space, we show that our approach also works in general metrics and higher dimensions, yielding navigable networks with at most twice as many edges as the optimum. This increase is due to agent selfishness, and we study this by giving almost tight bounds on the price of anarchy.

All omitted details can be found in the full version [32].

2 PROPERTIES FOR GREEDY ROUTING

We establish some structural properties that will be used in our analysis of both directed and undirected variants. These properties are essential for enabling greedy routing.

General Metrics: We start with properties that hold for any metric space. Thus, these properties are the foundation for greedy routing in any network.

Lemma 1. *Consider two agents $u, v \in \mathcal{P}$ connected by an edge (u, v) (or an undirected edge $\{u, v\}$). If greedy routing is enabled for agent v , then the agents that can be reached by a greedy routing path from u via v are independent of agent v 's strategy.*

Next, we have another rather obvious property.

Lemma 2. *In a NE, an agent u builds at most $\phi(u)$ many edges.*

Remark 1. *In a NE network, for any agent $u \in \mathcal{P}$ the value of $\phi^+(u)$ is independent of the strategy profile, since all edges of the NNG have to be present in any NE network.*

Euclidean Metric: Now we move to properties that hold for Euclidean metrics. The following two lemmas were proven by Clarkson [11]. For completeness, we provide the proofs in our context.

Lemma 3. *In Euclidean metrics, consider nodes $s, u, v \in \mathcal{P}$, where $d(s, u) \leq d(s, v)$, and let θ denote the angle $\angle usv$. If $\theta < \pi/3$ then the sequence $P = (s, u, v)$ is a greedy routing path from node s to node v .*

Lemma 4. *For 2D Euclidean metrics, the greedy routing degree of any agent can be computed in polynomial time and is at most 6.*

Remark 2. *For Euclidean metrics in dimension $D > 0$, the greedy routing degree of any agent can be computed in polynomial time and it is at most the kissing number $K(D)$.*

The above follows by the definition of the kissing number by Conway and Sloane [12] and by Lemma 3.

3 GREEDY ROUTING: DIRECTED EDGES

In this section we explore the directed variant of our model beginning with some results that hold in every metric space.

3.1 Nash Equilibria

General Metrics. We start by proving that Nash equilibria always exist for arbitrary metric spaces.

THEOREM 1. *Every NE is a SO and every SO is a NE.*

This directly implies the following positive result.

Corollary 1. *The price of anarchy in the directed version is 1.*

Next, we show that best response cycles cannot exist. This yields that stable states can be found via natural game dynamics.

Lemma 5. *Best response cycles do not exist.*

3.2 Computational Complexity

General Metrics. Here, we investigate the computational complexity of computing a best response in general metric spaces, and computing a NE in Euclidean metrics.

Although computing the greedy routing degree can be done in polynomial time, computing a best response is a hard problem, even in one-dimensional metric spaces.

THEOREM 2. *For Euclidean metrics (hence also for general metrics) computing a best response of an agent is NP-hard.*

Euclidean Metrics. While computing a best response is hard for Euclidean metrics, we now prove the contrasting result that in Euclidean metrics, computing a NE, a social optimum, and deciding if a given strategy profile is a NE is efficiently computable.

Corollary 2. *In Euclidean metrics of dimension $D > 0$, computing a social optimum, and thus also a NE, is polynomial time computable.*

Since in a NE every agent must buy the edges in its minimum greedy routing set, and because this can be computed efficiently, we get the following.

Corollary 3. *In Euclidean metrics of dimension $D > 0$, deciding if a strategy profile is a NE is polynomial time computable.*

4 GREEDY ROUTING: UNDIRECTED EDGES

In this section we explore the much more challenging undirected variant of our model. We start with results on two-dimensional Euclidean metrics. Later, we will generalize some of our findings to higher dimensions and general metric spaces.

4.1 Price of Anarchy

2D Euclidean Metric. We start with a technical lemma that will be essential for deriving upper bounds on the price of anarchy.

Lemma 6. *Let $G(\mathbf{s})$ be a NE network over a set of points \mathcal{P} and let $G_{NNG}(\mathcal{P})$ be the corresponding NNG. Let $C \subseteq \mathcal{P}$ such that all nodes in C belong to the same connected component of $G_{NNG}(\mathcal{P})$. Then, the total number of edges built by the agents in C in $G(\mathbf{s})$ that do not appear in G_{NNG} is at most $2 + 3|C|$. (See Figure 5.)*




Figure 5: A NNG component of two nodes u, v . The maximum number of edges that have to be built in order for agents u and v to be greedy routing connected is 8.

We will now use Lemma 6 to prove one of our main results: That NE networks are close to optimal in terms of density. In the worst case, selfish behavior only increases the number of edges by 80%.

THEOREM 3. *In 2D Euclidean metrics, the PoA is at most 1.8.*




Figure 6: The (a) SO, (b) NNG, and (c) a NE of a network with four nodes. Directed arcs show the edge-ownership.

PROOF. As stated above, any network that supports greedy routing must contain the NNG as a subgraph. Thus, every social optimum and every NE network also contains the NNG as a subgraph. See Figure 6 for an example.

The NNG consists of connected components, each containing at least two nodes. Consider a NE network $G(\mathbf{s})$ defined over a set of points \mathcal{P} , and let $G_{NNG}(\mathcal{P})$ be the corresponding NNG. Partition \mathcal{P} into k disjoint sets C_1, C_2, \dots, C_k , where k is the number of connected components of $G_{NNG}(\mathcal{P})$. Each set C_i contains the nodes of the i -th connected component, meaning that any two nodes $u, v \in C_i$ are in the same set C_i if and only if they belong to

same connected component of $G_{NNG}(\mathcal{P})$. Then, the social cost is equal to:

$$c(\mathbf{s}) = \sum_{i=1}^k \sum_{u \in C_i} |S_u(\mathbf{s})| = \sum_{i=1}^k \sum_{u \in C_i} (|S_u^-(\mathbf{s})| + |S_u^+(\mathbf{s})|),$$

where $S_u^+(\mathbf{s})$ denotes the set of edges built by agent u in $G(\mathbf{s})$ that also appear in $G_{NNG}(\mathcal{P})$, and $S_u^-(\mathbf{s})$ denotes the edges built by agent u that do not appear in $G_{NNG}(\mathcal{P})$. Equivalently, we define $deg_{G(\mathbf{s})}^+(u)$ to be the number of edges incident to u in $G(\mathbf{s})$ that also appear in $G_{NNG}(\mathcal{P})$ and $deg_{G(\mathbf{s})}^-(u)$ the number of edges incident to u in $G(\mathbf{s})$ that do not appear in $G_{NNG}(\mathcal{P})$.

For every agent u there is a set $W_u \subseteq \mathcal{P}$ of agents that agent u can reach via greedy routing paths starting with an edge from $G_{NNG}(\mathcal{P})$ incident to agent u . Let $\phi'(u)$ denote the minimum number of edges that u needs to be incident to in order to be able to reach all other agents, i.e., the set $\mathcal{P} \setminus W_u$, via greedy routing paths.

In the worst case, agent u builds all $\phi'(u)$ many edges by itself. Furthermore, by Lemma 2 and Remark 1, we know that agent u will not build more than $\phi'(u)$ non-NNG edges. Therefore,

$$|S_u^-(\mathbf{s})| \leq \phi'(u). \quad (1)$$

Let \mathbf{s}^* be a strategy profile that achieves the social optimum, and let $G(\mathbf{s}^*)$ be the corresponding network. Not counting the edges from $G_{NNG}(\mathcal{P})$, every agent u in $G(\mathbf{s}^*)$ must be incident to at least $\phi'(u)$ edges in order to be greedy routing connected. Thus,

$$\sum_{u \in \mathcal{P}} \phi'(u) \leq \sum_{u \in \mathcal{P}} deg_{G(\mathbf{s}^*)}^-(u) = 2 \sum_{u \in \mathcal{P}} |S_u^-(\mathbf{s}^*)|. \quad (2)$$

Combining (1) and (2) we get:

$$\begin{aligned} \sum_{u \in \mathcal{P}} \phi'(u) &\leq 2 \sum_{u \in \mathcal{P}} |S_u^-(\mathbf{s}^*)| \\ \frac{1}{2} \sum_{u \in \mathcal{P}} |S_u^-(\mathbf{s})| &\leq \sum_{u \in \mathcal{P}} |S_u^-(\mathbf{s}^*)| \\ \frac{1}{2} \sum_{u \in \mathcal{P}} |S_u^-(\mathbf{s})| + \sum_{u \in \mathcal{P}} |S_u^+(\mathbf{s})| &\leq \sum_{u \in \mathcal{P}} |S_u^-(\mathbf{s}^*)| \\ &\quad + \sum_{u \in \mathcal{P}} |S_u^+(\mathbf{s}^*)|, \end{aligned} \quad (3)$$

where the last inequality holds since $\sum_{u \in \mathcal{P}} |S_u^-(\mathbf{s}^*)| = \sum_{u \in \mathcal{P}} |S_u^+(\mathbf{s}^*)|$, as in both strategy profiles the same number of NNG edges must be bought. Therefore, for the PoA, we have:

$$PoA \leq \frac{c(\mathbf{s})}{c(\mathbf{s}^*)} = \frac{\sum_{u \in \mathcal{P}} |S_u(\mathbf{s})|}{\sum_{u \in \mathcal{P}} |S_u(\mathbf{s}^*)|}.$$

Using (3) and summing over the connected components we get:

$$PoA \leq \frac{\sum_{u \in \mathcal{P}} |S_u^-(\mathbf{s})| + \sum_{u \in \mathcal{P}} |S_u^+(\mathbf{s})|}{\frac{1}{2} \sum_{u \in \mathcal{P}} |S_u^-(\mathbf{s})| + \sum_{u \in \mathcal{P}} |S_u^+(\mathbf{s})|} \quad (4)$$

$$= \frac{\sum_{i=1}^k \sum_{u \in C_i} |S_u^-(\mathbf{s})| + \sum_{i=1}^k \sum_{u \in C_i} |S_u^+(\mathbf{s})|}{\frac{1}{2} \sum_{i=1}^k \sum_{u \in C_i} |S_u^-(\mathbf{s})| + \sum_{i=1}^k \sum_{u \in C_i} |S_u^+(\mathbf{s})|}. \quad (5)$$

The PoA is maximized when $\sum_{i=1}^k \sum_{u \in C_i} |S_u^-(\mathbf{s})|$ is as large as possible and $\sum_{i=1}^k \sum_{u \in C_i} |S_u^+(\mathbf{s})|$ is as small as possible. Therefore, using Lemma 6, we have

$$\sum_{i=1}^k \sum_{u \in C_i} |S_u^-(\mathbf{s})| \leq \sum_{i=1}^k \sum_{u \in C_i} (3|C_i| + 2),$$

while

$$\sum_{i=1}^k \sum_{u \in C_i} |S_u^+(s)| \geq \sum_{i=1}^k \sum_{u \in C_i} (|C_i| - 1),$$

since C_i is connected. Therefore, we have:

$$\begin{aligned} PoA &\leq \frac{\sum_{i=1}^k (3|C_i| + 2) + \sum_{i=1}^k (|C_i| - 1)}{\frac{1}{2} \sum_{i=1}^k (3|C_i| + 2) + \sum_{i=1}^k (|C_i| - 1)} \\ &= \frac{\sum_{i=1}^k (4|C_i| + 1)}{\frac{1}{2} \sum_{i=1}^k (3|C_i| + 2) + \frac{2}{2} \sum_{i=1}^k (|C_i| - 1)} \\ &= \frac{2 \sum_{i=1}^k (4|C_i| + 1)}{\sum_{i=1}^k (3|C_i| + 2) + \sum_{i=1}^k (2|C_i| - 2)} \\ &= \frac{8|\mathcal{P}| + 2k}{5|\mathcal{P}|}. \end{aligned} \quad (6)$$

The maximum number of connected components in the NNG is $\frac{|\mathcal{P}|}{2}$ since every connected component contains at least two nodes. Therefore,

$$PoA \leq \frac{(1 + 8)|\mathcal{P}|}{5|\mathcal{P}|} = 1.8. \quad \square$$

Now we proceed to show that our PoA upper bound is almost tight. For this, we give an almost matching lower bound, showing that selfish behavior can increase the number of built edges by 75%.

Lemma 7. *In 2D Euclidean metrics the PoA is at least 1.75.*

PROFSKETCH. Consider the network in Figure 7(a). Let C denote the set of connected components of the corresponding NNG of the instance. Each NNG component $C \in C$ consists of two nodes (connected with a red edge) and is connected to nodes outside of C via 6 edges, the minimum number of required edges for greedy connectivity in this instance. This is a SO, since every non-NNG edge originating from an agent u has equal length and contributes to the greedy routing connectivity of both endpoints. Thus, the total number of edges is $|C|(\frac{6}{2} + 1) = 4|C|$. In contrast, the network in Figure 7(b), is a NE, where each non-NNG edge contributes to the greedy connectivity of only one endpoint. Hence, the number of the non-NNG edges is doubled compared to the SO, i.e., it is $|C|(6 + 1) = 7|C|$. Therefore, the social cost ratio is $\frac{7}{4} = 1.75$. \square

Higher Dimensions and General Metrics. We extend our analysis to obtain upper bounds on the PoA for higher dimensional Euclidean spaces and for general metrics.

Corollary 4. *In Euclidean metrics of dimension D , the PoA is at most $2 - \frac{1}{K(D)}$, while in general metric spaces it is less than 2.*

4.2 Approximate Nash Equilibria

General Metrics. We first study the simplest way for finding a Nash equilibrium: best response dynamics. Unfortunately, these are not guaranteed to converge.

THEOREM 4. *Best response cycles exist.*

PROFSKETCH. A best response cycle is shown in Figure 8. \square

Next, we introduce some definitions and results that are essential for proving the existence of approximate NE.

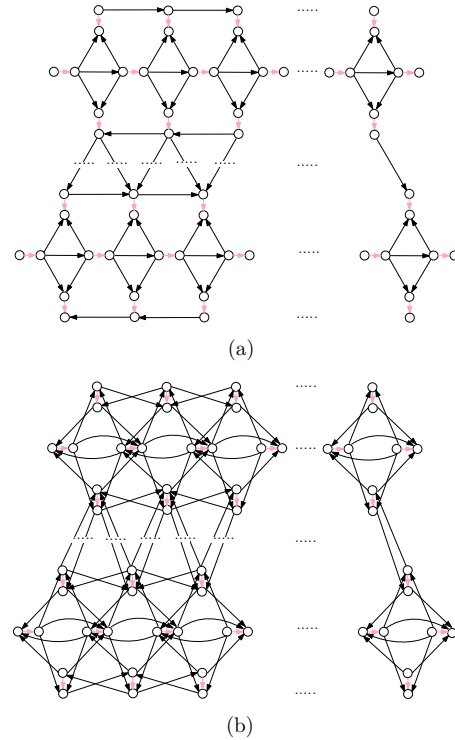


Figure 7: (a) social optimum (SO), (b) a NE with social cost 1.75 times the cost of SO. Red edges belong to the NNG. All edges are undirected, the arcs depict edge-ownership.




Figure 8: The best response cycle for Theorem 4.

Definition 1. *Consider graph $G(\mathcal{P}, E)$ that supports greedy routing. For an agent $u \in \mathcal{P}$ with incident edges $E_u \subseteq E$, the critical incident set $H_u \subseteq E_u$ is a maximal set of incident edges such that removing any single edge from H_u breaks agent u 's greedy connectivity.*

Definition 2. *Let $G(\mathcal{P}, E)$ be a graph that supports greedy routing and let H_u be a critical incident set of agent u . Consider the modified network $G' = G(\mathcal{P}, E \setminus H_u)$, where the edge set H_u is removed. The critical best response strategy is defined as agent u 's best response strategy S_u^{best} in G' . If multiple best responses exist, we select the one that maximizes the overlap with the critical incident set, i.e., the one for which the quantity $|S_u^{best} \cap H_u|$ is maximized.*

Next, we define the function $\alpha(u)$, for every agent $u \in \mathcal{P}$, which reflects how the network’s structure influences an agent’s strategy.

Definition 3. Let $G(\mathcal{P}, E)$ be a graph that supports greedy routing. Suppose that exactly $k \geq 0$ edges in S_u^{best} are new (not incident to node u in $G(\mathcal{P}, E)$). Then $\alpha(u) := k$ is defined as the number of those edges, i.e., $\alpha(u) := |S_u^{best} \setminus H_u|$. (See Example 5.)




Figure 9: Illustration of the function $\alpha(u)$.

Example 5. We want to determine the owner of each edge in Figure 9 assuming that the shown network is part of a NE network.

The sets H_u, S_u^{best} and parameter $\alpha(u)$ are used as tools to achieve this goal. In Figure 9 (left), assume that agent u is greedy connected under greedy routing set A, while $H_u = \{\{u, v\}, \{u, b\}, \{u, x\}\}$. Therefore, agent u has no incentive to buy the edges $\{u, z\}, \{u, a\}$, as they do not contribute to its greedy connectivity. Thus, if this was a NE, these edges must therefore be bought by agents z and a , respectively. Moreover, assume that agent u would be greedy connected under greedy routing set B if there were an edge to node p . Since, the critical best response S_u^{best} is defined as the best response in the network without the H_u edges, and the edges $\{z, u\}, \{a, u\}$ are bought by z and a , respectively, we have that $S_u^{best} = \{\{u, v\}, \{u, p\}\}$. As node v is already incident to u , we have $\alpha(u) = 1$. We can determine which edges agent u is willing to buy in this instance. First, agent u has an incentive to buy the edge $\{u, v\}$ because it belongs to both greedy routing sets, i.e., to $S_u^{best} \cap H_u$. Additionally, agent u can buy at most one edge from $H_u \setminus S_u^{best}$, because $\alpha(u) = 1$. Otherwise, agent u could drop the edges $\{u, b\}$ and $\{u, x\}$ and instead buy an edge to node p , reducing its overall cost, since the edges $\{u, z\}$ and $\{u, a\}$ are bought by z and a , respectively. Note that while the edges $\{u, a\}$ and $\{u, z\}$ do not contribute to the greedy connectivity of agent u , they still affect the value of $\alpha(u)$.

In Figure 9 (right) the greedy degree of agent u is 3 and agent u is simultaneously greedy connected by two different greedy routing sets, while $H_u = \{\{u, v\}\}$. Thus, $S_u^{best} = \{\{u, v\}\}$, i.e., $\alpha(u) = 0$. Therefore, agent u has an incentive to buy only the edge $\{u, v\}$. \blacktriangleleft

Lemma 8. Let $G = (\mathcal{P}, E)$ be a graph that supports greedy routing, and let $u \in \mathcal{P}$ be an agent. In any Nash equilibrium of G , agent u is willing to buy only the edges in $S_u^{best} \cap H_u$, along with exactly $\alpha(u)$ many edges from $H_u \setminus S_u^{best}$.

PROOF. Given a graph $G(\mathcal{P}, E)$ that supports greedy routing, suppose there exists an ownership assignment such that G is in a NE. Our goal is to determine which incident edges an agent is willing to buy in order for G to remain in a NE.

In a NE every agent u buys a subset $S_u(\mathbf{s}) \subseteq E_u$ of its incident edges, while the other edges $E_u \setminus S_u$ must be bought by its neighbors.

First, we show that in a NE, an agent u does not buy any edges in $E_u \setminus H_u$. Suppose that there is no set $W \subseteq (E_u \setminus H_u)$ such that while the removal of any single edge of them does not affect the greedy connectivity of u , the removal of all the edges in W does. In this case, agent u has no incentive to buy any edge in $E_u \setminus H_u$ since they don’t contribute to its greedy connectivity. In the case that such a set W exists, agent u must be greedy routing connected in G by more than one strategy. Otherwise, the removal of any single edge in W would affect the greedy connectivity of agent u , which would imply $W_u \subseteq H_u$ (see Example 5). Suppose that agent u buys a set of edges in $E_u \setminus H_u$ that belong only to one strategy that enables greedy routing for agent u . Since the rest of the edges are bought by its neighbors, along with them will be the edges that belong to another strategy that enables greedy routing for agent u . Thus, agent u would be greedy connected by more than one strategy and could reduce its cost by dropping all of its edges. Alternatively, suppose that agent u buys a set of edges in $E_u \setminus H_u$ that belong to more than one greedy routing strategy. Since all the incident edges are bought by some agent, agent u could reduce its cost by keeping only the edges of one strategy and dropping the rest. Thus, in every case agent u will not buy any edge in $E_u \setminus H_u$ when G is in NE.

On the other hand, agent u would be willing to buy some of the edges in H_u in order to remain greedy routing connected. In any NE, agent u is willing to buy all the edges from $S_u^{best} \cap H_u$, since these edges are part of the best response. In addition, agent u is willing to buy up to $\alpha(u)$ many edges from the remaining edges in H_u , i.e., from $H_u \setminus S_u^{best}$. Otherwise, agent u could reduce its cost by removing the edges in H_u that do not appear in $S_u^{best} \cap H_u$ and instead buying these k edges. This completes the proof. \square

In a NE graph $G(\mathbf{s})$ every edge contributes to the greedy connectivity of at least of one of its endpoints. We distinguish between two edge types, *single-edges* and *double-edges* defined as follows:

Definition 4. In a NE graph $G(\mathbf{s})$ the strategy of an agent $u \in \mathcal{P}$ can be partitioned into two sets: the set of single-edges S_u^s and the set of double-edges S_u^d . An edge $\{u, v\} \in S_u^s$ is called single-edge if it contributes to the greedy connectivity of only agent u , while an edge $\{u, v\} \in S_u^d$ is called double if it contributes to the greedy connectivity of both endpoints u and v .

Note that although a single-edge $\{u, v\}$ contributes to the greedy connectivity of only u it may affect the value of $\alpha(v)$ of agent v . This can happen when a single-edge $\{u, v\}$ bought by agent u is required to be present in a best response of v . (See Example 5.)

Another tool that will be useful is Hakimi’s theorem [26]:

THEOREM 6. (Hakimi’s Theorem) Let $G(V, E)$ be an undirected graph and let δ be an integer-valued function on V . Then G admits an orientation such that the in-degree of each vertex $u \in V$ is exactly $\delta(u)$, if and only if for every subset $U \subseteq V$ holds $\sum_{u \in U} \delta(u) \geq |E(U)|$, where $E(U)$ denotes the set of edges in G with both endpoints in U .

We proceed by proving that every social optimum is a 2-NE.

THEOREM 7. In general metric spaces, every SO is a 2-NE.

PROOF. The proof is structured as follows. We begin with a SO graph $G = (\mathcal{P}, E)$, where no ownership is assigned to its edges. Our

goal is to prove that there is a valid ownership assignment such that G is an approximate NE.

To achieve this, we first construct the auxiliary graph $\tilde{G}(\mathcal{P}, \tilde{E})$. For each agent $u \in \mathcal{P}$, the graph $\tilde{G}(\mathcal{P}, \tilde{E})$ includes all the single-edges of agent u as well as the edges that belong to its critical best response S_u^{best} , i.e. $\tilde{E} = \bigcup_{u \in \mathcal{P}} (S_u^s \cup S_u^{best})$. By the definition of a critical best response, and since all the single-edges are present, this graph supports greedy routing.

Since G is a SO, \tilde{G} has to have at least as many edges as G . Therefore we have:

$$\begin{aligned} |\tilde{E}| &\geq |E| \\ \sum_{u \in \mathcal{P}} (|S_u^{best} \cup S_u^s|) &\geq |E| \\ \sum_{u \in \mathcal{P}} (|S_u^{best} \setminus H_u| + |S_u^{best} \cap H_u| + |S_u^{s-}|) &\geq |E|, \end{aligned}$$

where S_u^{s-} denotes the set of single-edges associated with agent u that are not in $(S_u^{best} \cap H_u)$. Since $\alpha(u) = |S_u^{best} \setminus H_u|$, we have:

$$\sum_{u \in \mathcal{P}} (\alpha(u) + |S_u^{best} \cap H_u| + |S_u^{s-}|) \geq |E|. \quad (7)$$

Now that we obtained this inequality, we go back to the SO graph G . By Lemma 8, in any NE of \mathcal{P} agent u is willing to buy all the edges in $S_u^{best} \cap H_u$ and exactly $\alpha(u)$ many edges from $H_u \setminus S_u^{best}$.

First, we assign all the edges in $S_u^{best} \cap H_u$ to agent u . If an edge belongs to the critical best response of both endpoints we choose one arbitrarily. Then, observe that the number of the remaining single-edges contributing to the greedy connectivity of an agent $u \in \mathcal{P}$ cannot be more than $\alpha(u)$, because otherwise we could replace the edges in S_u^{s-} in $G(\mathcal{P}, E)$ with $\alpha(u)$ many edges from $S_u^{best} \setminus H_u$ without affecting the greedy connectivity of the other agents, creating a graph with less edges than the social optimum, a contradiction. Next, we assign the S_u^{s-} edges (which are at most $\alpha(u)$ many) to agent u , for every $u \in \mathcal{P}$.

Let $G'(\mathcal{P}, E')$ be the subgraph of G containing all remaining unassigned edges $E' = E \setminus ((S_u^{best} \cap H_u) \cup S_u^{s-})$. Therefore we have:

$$\begin{aligned} \sum_{u \in \mathcal{P}} (\alpha(u) + |S_u^{best} \cap H_u| + |S_u^{s-}|) &\geq |E| \\ \sum_{u \in \mathcal{P}} \alpha(u) &\geq |E| - \sum_{u \in \mathcal{P}} (|S_u^{best} \cap H_u| + |S_u^{s-}|) \\ \sum_{u \in \mathcal{P}} \alpha(u) &\geq |E'|. \end{aligned}$$

Every edge in $(u, v) \in E'$ contributes to the greedy connectivity of both agents u and v , so both of them have an incentive to buy it as soon as their remaining α -value permits.

Suppose towards a contradiction, that there exists a subset $U \subseteq \mathcal{P}$ such that $\sum_{u \in U} \alpha(u) < |E'(U)|$, where $E'(U)$ denotes the set of edges in G' with both endpoints in U . In this case, we could replace the edges in $E'(U)$ with the corresponding $\alpha(u)$ many edges in $S_u^{best} \setminus H_u$ for each agent u , and then add the assigned edges in $E \setminus E'$ having as a result a graph that supports greedy routing and has fewer edges than the social optimum, a contradiction.

Hence, Hakimi's theorem (Theorem 6) applies to $G'(\mathcal{P}, E')$. Let $indeg_{G'}(u)$ denote the in-degree of a node $u \in \mathcal{P}$ in graph G' . By Hakimi's theorem, the graph G' can be oriented such that for each

$u \in \mathcal{P}$, we have $indeg_{G'}(u) \leq \alpha(u)$. Interpreting the orientation as edge ownership, there exists an ownership assignment in G' such that agent u buys at most $\alpha(u)$ many edges. Using this assignment in the edges $E \cap E'$ in G , we end up with an ownership assignment where every agent u builds $|S_u^{best} \cap H_u| + 2\alpha(u)$ many edges, while in a NE it would have built $|S_u^{best} \cap H_u| + \alpha(u)$ many. Therefore, the social optimum network G is in 2-NE. \square

2D Euclidean Metrics. We now focus on 2D Euclidean metrics and compared to Theorem 7, we prove the stronger result that every social optimum is a +2-NE.

THEOREM 8. *In 2D Euclidean metrics, every SO is a +2-NE.*

4.3 Computational Complexity

General Metrics. We now show that computing a best response is NP-hard, while an approximate NE can be computed in polynomial time in Euclidean metrics.

THEOREM 9. *Computing a best response is NP-hard.*

Euclidean Metrics. Here, we prove that an approximate NE with at most 1.8 times the optimal number of edges in 2D Euclidean metrics, and $2 - 1/K(D)$ times the optimal number of edges in D -dimensional Euclidean metrics can be computed in polynomial time. Note that this bound improves on the approximation ratio of 3 achieved by the Delaunay triangulation in 2D Euclidean metrics.

THEOREM 10. *A 2-NE in Euclidean metrics of dimension $D > 0$, and a +2-NE in 2D Euclidean metric that has at most 1.8 times the optimal number of edges in 2D Euclidean metrics, and $2 - 1/K(D)$ times the optimal number of edges in D -dimensional Euclidean metrics, can be computed in polynomial time.*

Corollary 5. *In Euclidean metric of dimension $D > 0$, deciding if a strategy profile is a NE is polynomial time computable.*

PROOF. Since, as shown above, the set of the edges that are bought by each agent u in any NE can be computed in polynomial time, we check that the corresponding edges are assigned to each agent u . For the other edges, we check if the remaining α -value is at most the number of edges assigned to each agent u . \square

5 CONCLUSION

We study decentralized network creation to enable greedy routing. We uncover that for the case with directed edges equilibria exist and are optimal, while the setting with undirected edges is much more challenging. For the latter, we give an algorithm that constructs approximate equilibria with only slightly more edges compared to the social optimum. Moreover, we give an almost tight bound on the price of anarchy, which shows that this increase in the number of edges due to the selfishness of the agents cannot be avoided. However, we emphasize that our constructed networks still significantly outperform the well-known Delaunay triangulation.

Future work can build on our basis that combines network creation games with computational geometry. In particular, enhancing the construction so that additional features like stretch or robustness guarantees can be achieved. Moreover, proving that exact equilibria exist is an exciting open problem. We believe that our construction might already yield Nash equilibria, but proving this is challenging and requires additional structural insights.

REFERENCES

- [1] Mohammad Ali Abam and Mahnaz Sadat Qafari. 2019. Geometric spanner games. *Theoretical Computer Science* 795 (2019), 398–407. <https://doi.org/10.1016/J.TCS.2019.07.020>
- [2] Yousef Al-Jazzazi, Haya Diwan, Jinrui Gou, Cameron Musco, Christopher Musco, and Torsten Suel. 2025. Distance Adaptive Beam Search for Provably Accurate Graph-Based Nearest Neighbor Search. *arXiv preprint arXiv:2505.15636* (2025).
- [3] Venkatesh Bala and Sanjeev Goyal. 2000. A noncooperative model of network formation. *Econometrica* 68, 5 (2000), 1181–1229. <https://doi.org/10.1111/1468-0262.00155>
- [4] Julian Berger, Tobias Friedrich, Pascal Lenzen, Paraskevi Machaira, and Janosch Ruff. 2025. Strategic Network Creation for Enabling Greedy Routing. In *AAAI'25*, Vol. 39. 13622–13630. <https://doi.org/10.1609/AAAI.V39I13.33488>
- [5] Davide Bilò, Tobias Friedrich, Pascal Lenzen, and Anna Melnichenko. 2024. Geometric Network Creation Games. *SIAM J. Discret. Math.* 38, 1 (2024), 277–315. <https://doi.org/10.1137/20M1376662>
- [6] Thomas Bläsius, Tobias Friedrich, Maximilian Katzmann, and Anton Krohmer. 2020. Hyperbolic Embeddings for Near-Optimal Greedy Routing. *ACM Journal on Experimental Algorithmics* 25 (2020), 1–18. <https://doi.org/10.1145/3381751>
- [7] Marián Boguñá, Dmitri V. Krioukov, and Kimberly C. Claffy. 2009. Navigability of complex networks. *Nature Physics* 5, 1 (2009), 74–80. <https://doi.org/10.1038/nphys1130>
- [8] Marián Boguñá, Fragkiskos Papadopoulos, and Dmitri Krioukov. 2010. Sustaining the internet with hyperbolic mapping. *Nature communications* 1, 1 (2010), 62. <https://doi.org/10.1038/ncomms1063>
- [9] Prosenjit Bose and Pat Morin. 2004. Online Routing in Triangulations. *SIAM J. Comput.* 33, 4 (Jan. 2004), 937–951. <https://doi.org/10.1137/S0097539700369387>
- [10] Prosenjit Bose, Pat Morin, Ivan Stojmenovic, and Jorge Urrutia. 2001. Routing with Guaranteed Delivery in Ad Hoc Wireless Networks. *Wirel. Networks* 7, 6 (2001), 609–616. <https://doi.org/10.1023/A:1012319418150>
- [11] Kenneth L. Clarkson. 1987. Approximation Algorithms for Shortest Path Motion Planning (Extended Abstract). In *STOC'87*, Alfred V. Aho (Ed.). ACM, 56–65. <https://doi.org/10.1145/28395.28402>
- [12] John Horton Conway and Neil James Alexander Sloane. 2013. *Sphere packings, lattices and groups*. Vol. 290. Springer Science & Business Media.
- [13] Andrej Cvetkovski and Mark Crovella. 2009. Hyperbolic Embedding and Routing for Dynamic Graphs. In *INFOCOM'09*. IEEE, 1647–1655. <https://doi.org/10.1109/INFCOM.2009.5062083>
- [14] Susanta Datta, Ivan Stojmenovic, and Jie Wu. 2002. Internal node and shortcut based routing with guaranteed delivery in wireless networks. *Cluster computing* 5, 2 (2002), 169–178. <https://doi.org/10.1023/A:1013985610753>
- [15] Mark de Berg, Otfried Cheong, Marc J. van Kreveld, and Mark H. Overmars. 2008. *Computational geometry: algorithms and applications, 3rd Edition*. Springer. <https://doi.org/10.1007/978-3-540-77974-2>
- [16] Haya Diwan, Jinrui Gou, Cameron Musco, Christopher Musco, and Torsten Suel. 2024. Navigable Graphs for High-Dimensional Nearest Neighbor Search: Constructions and Limits. In *NeurIPS'24*. http://papers.nips.cc/paper_files/paper/2024/hash/6dc63b4063c978cf195bc15178e8152a-Abstract-Conference.html
- [17] Stephan J. Eidenbenz, V. S. Anil Kumar, and Sibylle Züst. 2006. Equilibria in Topology Control Games for Ad Hoc Networks. *Mobile Networks and Applications* 11, 2 (2006), 143–159. <https://doi.org/10.1007/s11036-005-4468-y>
- [18] David Eppstein and Michael T Goodrich. 2008. Succinct greedy graph drawing in the hyperbolic plane. In *GD'08*. Springer, 14–25. https://doi.org/10.1007/978-3-642-00219-9_3
- [19] Alex Fabrikant, Ankur Luthra, Elitza Maneva, Christos Papadimitriou, and Scott Shenker. 2003. On a Network Creation Game. In *PODC'03*. 347–351. <https://doi.org/10.1145/872035.872088>
- [20] Qing Fang, Jie Gao, and Leonidas J Guibas. 2006. Locating and bypassing holes in sensor networks. *Mobile networks and Applications* 11, 2 (2006), 187–200. <https://doi.org/10.1007/s11036-006-4471-y>
- [21] Gregory G Finn. 1987. *Routing and addressing problems in large metropolitan-scale internetworks*. Technical Report.
- [22] 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 *SPAA '21*. 232–242. <https://doi.org/10.1145/3409964.3461807>
- [23] Cong Fu, Chao Xiang, Changxu Wang, and Deng Cai. 2017. Fast approximate nearest neighbor search with the navigating spreading-out graph. *arXiv preprint arXiv:1707.00143* (2017).
- [24] Mohsen Ghaffari, Behnoosh Hariri, and Shervin Shirmohammadi. 2010. On the necessity of using Delaunay Triangulation substrate in greedy routing based networks. *IEEE Communications Letters* 14, 3 (March 2010), 266–268. <https://doi.org/10.1109/LCOMM.2010.03.092045>
- [25] András Gulyás, József J. Biró, Attila Kőrösi, Gábor Rétvári, and Dmitri Krioukov. 2015. Navigable Networks as Nash Equilibria of Navigation Games. *Nature Communications* 6, 1 (2015), 7651. <https://doi.org/10.1038/ncomms8651>
- [26] S Louis Hakimi. 1965. On the degrees of the vertices of a directed graph. *Journal of the Franklin Institute* 279, 4 (1965), 290–308.
- [27] Brad Karp and Hsiang-Tsung Kung. 2000. GPSR: Greedy perimeter stateless routing for wireless networks. In *MOBICOM'00*. 243–254. <https://doi.org/10.1145/345910.345953>
- [28] Robert Kleinberg. 2007. Geographic routing using hyperbolic space. In *INFOCOM'07*. IEEE, 1902–1909. <https://doi.org/10.1109/INFCOM.2007.221>
- [29] Elias Koutsoupias and Christos H. Papadimitriou. 2009. Worst-case equilibria. *Comput. Sci. Rev.* 3, 2 (2009), 65–69. <https://doi.org/10.1016/J.COSREV.2009.04.003>
- [30] Fabian Kuhn, Roger Wattenhofer, and Aaron Zollinger. 2003. Worst-case optimal and average-case efficient geometric ad-hoc routing. In *MOBICOM'03*. 267–278. <https://doi.org/10.1145/778415.778447>
- [31] Vijay Kumar, Daniela Rus, and Sanjiv Singh. 2004. Robot and Sensor Networks for First Responders. *IEEE Pervasive Comput.* 3, 4 (2004), 24–33. <https://doi.org/10.1109/MPRV.2004.17>
- [32] Pascal Lenzen and Paraskevi Machaira. 2026. Greedy Routing Reachability Games. *arXiv preprint arXiv:2601.23126* (2026). [arXiv:2601.23126 \[cs.GT\]](https://arxiv.org/abs/2601.23126) <https://doi.org/10.1109/abs/2601.23126>
- [33] Wen-Jiunn Liu and Kai-Ten Feng. 2008. Greedy routing with anti-void traversal for wireless sensor networks. *IEEE transactions on mobile computing* 8, 7 (2008), 910–922. <https://doi.org/10.1109/TMC.2008.162>
- [34] Konrad Lorincz, David J. Malan, Thaddeus R. F. Fulford Jones, Alan Nawoj, Antony Clavel, Victor Shnayder, Geoffrey Mainland, Matt Welsh, and Steve Moulton. 2004. Sensor Networks for Emergency Response: Challenges and Opportunities. *IEEE Pervasive Comput.* 3, 4 (2004), 16–23. <https://doi.org/10.1109/MPRV.2004.18>
- [35] Yu A Malkov and Dmitry A Yashunin. 2018. Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. *IEEE transactions on pattern analysis and machine intelligence* 42, 4 (2018), 824–836. <https://doi.org/10.1109/TPAMI.2018.2889473>
- [36] Dov Monderer and Lloyd S Shapley. 1996. Potential games. *Games and economic behavior* 14, 1 (1996), 124–143. <https://doi.org/10.1006/game.1996.0044>
- [37] Christos H. Papadimitriou. 2001. Algorithms, games, and the internet. In *STOC'01*. 749–753. <https://doi.org/10.1145/380752.380883>
- [38] Christos H. Papadimitriou and David Ratajczak. 2005. On a conjecture related to geometric routing. *Theor. Comput. Sci.* 344, 1 (2005), 3–14. <https://doi.org/10.1016/J.TCS.2005.06.022>
- [39] Yu Wang, Chih-Wei Yi, Minsu Huang, and Fan Li. 2013. Three-dimensional greedy routing in large-scale random wireless sensor networks. *Ad hoc networks* 11, 4 (2013), 1331–1344. <https://doi.org/10.1016/J.ADHOC.2010.10.003>
- [40] Cédric Westphal and Guanhong Pei. 2009. Scalable Routing Via Greedy Embedding. In *INFOCOM'09*. IEEE, 2826–2830. <https://doi.org/10.1109/INFCOM.2009.5062240>