Geometric Network Design with Selfish Agents

Martin Hoefer[1] ⋆ and Piotr Krysta[2] ⋆⋆

[1] Department of Computer & Information Science, Konstanz University,
Box D 67, 78457 Konstanz, Germany.
hoefer@inf.uni-konstanz.de
[2] Department of Computer Science, Dortmund University,
Baroper Str. 301, 44221 Dortmund, Germany.
piotr.krysta@cs.uni-dortmund.de

Abstract. We study a geometric version of a simple non-cooperative network creation game introduced in [2], assuming Euclidean edge costs on the plane. The price of anarchy in such geometric games with k players is $\Theta(k)$. Hence, we consider the task of minimizing players incentives to deviate from a payment scheme, purchasing the minimum cost network. In contrast to general games, in small geometric games (2 players and 2 terminals per player), a Nash equilibrium purchasing the optimum network exists. This can be translated into a $(1+\epsilon)$-approximate Nash equilibrium purchasing the optimum network under more practical assumptions, for any $\epsilon > 0$. For more players there are games with 2 terminals per player, such that any Nash equilibrium purchasing the optimum solution is at least $\left(\frac{4}{3} - \epsilon\right)$-approximate. On the algorithmic side, we show that playing small games with best-response strategies yields low-cost Nash equilibria. The distinguishing feature of our paper are new techniques to deal with the geometric setting, fundamentally different from the techniques used in [2].

1 Introduction

The Internet is a powerful and universal artefact in human history and one of the most dynamic driving forces in modern society. An interesting recent research direction is to understand and influence the development of the Internet. A fundamental difference to other networks is that the Internet is built and maintained by a number of independent agents that pursue relatively limited, selfish goals. This motivated a lot of the research in a field now called algorithmic game theory. A major direction in this field is to analyze stable solutions in non-cooperative (networking) games. The most prominent measure is the *price of anarchy* [13], which is the ratio of the worst cost of a Nash equilibrium over the cost of an optimum solution. The price of anarchy has been considered in a variety of fields, such as, load balancing [6, 13], routing [16], and flow control [8]. A slightly different measure – the cost of the best Nash equilibrium instead of the worst was considered in [17]. This is the optimum solution no user has an incentive to defect

⋆ This author is supported by the DFG Graduiertenkolleg 1042 “Explorative Analysis and Visualization of Large Information Spaces” and partially supported by DFG grant Kr 2332/1-1 within Emmy Noether program

⋆⋆ This author is supported by DFG grant Kr 2332/1-1 within Emmy Noether program.

from, hence we will follow [1] and refer to it as the *price of stability*. In this paper we will consider both prices for the geometric version of a network creation game.

Network connection games. Anshelevich et al. [2] proposed a game theoretic model called a *connection game* for building and maintaining the Internet topology, which will be the basis for our paper. Agents are to build a network, and each agent holds a number of terminals at nodes in a graph, which she wants to connect by buying edges of the graph. The cost of edges can be shared among the players. An edge can only be used for connection if fully paid for. However, once it is paid for, any player can use it to connect her terminals. A strategy of a player is a payment function, i.e., her (possibly zero) contribution to paying the cost of each edge. Given strategies of all players form a Nash equilibrium if no player could deviate to a different strategy resulting in a smaller total payment to this player. In this game the problem of finding the cheapest payment strategy for one player is the classic Steiner tree problem. The problem of finding a minimum cost network satisfying all connections and minimizing the sum of all players payments is the Steiner forest problem.

Unfortunately, both the price of anarchy and the price of stability of this game can be in the order of k, the number of players. This is also an upper bound, because if the price of anarchy were more than k, there would be a player that could deviate by purchasing the optimum network all by herself. It is NP-complete to determine that a given game has a Nash equilibrium. Thus, in [2] a different approach was taken, in which a central institution determines a network and payment schemes for players. The goal is twofold: on one hand a cheap network should be purchased, on the other hand each player shall have the least motivation to deviate. As a strict Nash equilibrium might not exist, a payment scheme was presented that determines a 3-approximate Nash equilibrium on the socially optimum network. Finding the minimum cost network, however, is NP-hard. But, approximation algorithms for the Steiner tree (forest) problem [15, 11] can be used to find a $(4.65 + \epsilon)$-approximate Nash equilibrium purchasing a 2-approximate network in polynomial time [2].

A different network creation game was proposed by Fabrikant et al. [9]. Here each player corresponds to a node, and she can only contribute to edges incident to her node. A similar game was also considered by [4, 12] in the context of social networks. Being well-suited in this setting, for the global context of the Internet it is more appropriate to assume that players hold more terminals, can share edge costs and can contribute to costs anywhere in the network.

In a more recent paper Anshelevich et al. [1] have proposed a slightly different setting for the connection game. Here the focus is put on the Shapley value, a classic cost allocation protocol. Each edge is assumed to be shared equally among the players using it, and an $O(\log k)$ upper bound on the price of stability is shown. They considered bounds on the convergence of best-response dynamics and also studied versions of the game with edge latencies and weighting schemes.

Our contributions and results. In this paper we consider a special case of the connection game, the *geometric* connection game. Geometric edge costs present an interesting special case of the problem, as the connection costs of a lot of large networks can be approximated by the Euclidean distance on the plane [7]. Furthermore, for the geometric versions of combinatorial optimization problems usually improved results can be derived. For example, the geometric Steiner tree problem allows a PTAS [3], which contrasts the inapproximability for the general case [5]. This makes consideration of the geometric connection game attractive, and yields hope for significantly improved properties. In this paper, we present the following results for geometric connection games:

- The price of anarchy for geometric connection games with k players is k, even if we have two terminals per player. This, unfortunately, is the same bound as for general connection games [2].
- For games with 2 players each with 2 terminals, the price of stability is 1. The equilibrium payment scheme assigns payments along an edge according to a continuous function. For cases, in which this is unreasonable, we split an edge into small pieces, and each piece is bought completely by one player. Then a $(1+\epsilon)$-approximate Nash equilibrium can be achieved, for any $\epsilon > 0$.
- One cannot obtain results as above for more complicated games. Namely, for games with three or more players and 2 terminals per player, these results cannot be extended. There is a lower bound of $\left(\frac{4}{3} - \epsilon\right)$, for any $\epsilon > 0$, on approximate Nash equilibria purchasing the optimum network, which is slightly lower than the $\left(\frac{3}{2} - \epsilon\right)$ bound for general connection games in [2]. Thus, our result for geometric games with 2 players and 2 terminals per player is tight.
- If players play the game iteratively with best-response deviations, then in games with 2 players and 2 terminals per player the dynamics arrive at a Nash equilibrium very quickly. Furthermore, the created network is a $\sqrt{2}$-approximation to the cost of an optimum network.

The main difficulty when dealing with these geometric games is due to their inherent continuous nature. Most of our results require specific geometric arguments and new proof techniques that are fundamentally different from the ones previously used by Anshelevich et al. [2] for general connection games. The development of these new techniques is considered as an additional contribution of our paper.

Outline. Section 2 contains a formal definition of the geometric connection game, and Section 3 presents our results on the price of anarchy. Section 4 describes the results on the price of stability (Theorems 2, 3, and 5), and the analysis of the best-response dynamics (Theorem 4). Missing proofs will be given in the full version of the paper.

2 The model and preliminaries

The geometric connection game is defined as follows. Let V be a set of nodes which are points in the Euclidean plane. There are k non-cooperative players, each holding a number of terminals located at a subset of nodes from V. Each player strives to connect all of her terminals into a connected component. To achieve this a player offers money to purchase segments in the plane. The cost of a segment equals its length in the plane. Once the total amount of money offered by all players for a certain segment exceeds its cost, the segment is considered *bought*. Bought segments can be used by *all* players to connect their terminals, even if they contribute nothing to their costs. A strategy of player i is a payment function p_i that specifies how much she contributes to each segment in the plane. A collection of strategies, one for each player, is called

a *payment scheme* $p = (p_1, \ldots, p_k)$. A Nash equilibrium is a payment scheme p, in which no player i can connect her terminals at a lower cost by unilaterally reallocating her payments and switching to another function p'_i. [3] We will denote the social optimum solution, i.e., the minimum cost forest that connects the terminals of each player, by T^*. The subtree of T^* needed by player i to connect her terminals is denoted by T^i.
Constructing a minimum cost network satisfying all connections is the geometric Steiner forest problem. As the components of a Steiner forest are Steiner trees for a subset of players, well-known properties of optimum geometric Steiner trees hold for T^*.

Lemma 1. *[10, 14] Any 2 adjacent edges in an optimal geometric Steiner tree connect with an inner angle of at least 120°.*

Hence, every Steiner point of an optimal geometric Steiner tree has degree 3 and each of the 3 edges meeting at it makes angles of 120° with the other two [10, 14].
Another powerful tool for the analysis of connection games is the notion of a *connection set* that was the key ingredient to the analysis presented in [2].

Definition 1. *A connection set S of player i is a subset of edges of T^i, such that for each connected component C in $T^* \setminus S$ either $(1°)$ there is a terminal of i in C, or $(2°)$ any player that has a terminal in C has all of its terminals in C.*

Intuitively, after removing a connection set from T^* and reconnecting the terminals of player i the terminals of all players will be connected in the resulting solution. As T^* is the optimal solution, the maximum cost of *any* connection set S for player i is a lower bound for the cost of *any* of her deviations. Connection sets in a game with 2 terminals per player are easy to determine. Each T^i forms a path inside T^*, and two edges e, e' belong to the same connection set for player i iff $i \in \{j \in \{1, \ldots, k\} : e \in T^j\} = \{j' \in \{1, \ldots, k\} : e' \in T^{j'}\}$. We will mainly deploy geometric arguments and use connection sets only to limit the number of cases to be examined.

3 The price of anarchy

Theorem 1. *The price of anarchy for the geometric connection game with k players and 2 terminals per player, is precisely k.*

Proof. We have already argued in the introduction that the price of anarchy is at most k. Let us show now that k is also a lower bound. At first we will somewhat generally consider how a player in the geometric environment is motivated to deviate from a given payment scheme. Suppose we are given a game with 2 terminals per player and a feasible forest T, which satisfies the connection requirement for each player. Furthermore, let p be a payment function, which specifies a payment for each player on each edge. The next two Lemmas 2 and 3 follow directly from the Triangle Inequality.

Lemma 2. *If the deviation for a player i from p includes an edge $e \notin T$, this edge is a straight line segment, with start and end either at a terminal or some other part of T (possibly an interior point of some edge of T). It is located completely inside the Euclidean convex hull of T.*

[3] An α-approximate Nash equilibrium is a payment scheme where each player may reduce her costs by at most a factor of α by deviating.

Using these observations we can specify some properties of Nash equilibria for geometric connection games.

Lemma 3. *In a Nash equilibrium of the geometric connection game for k players, edges e_1, e_2 bought fully by one player are straight segments and meet with other, differently purchased edges with an inner angle of at least 90°. In the case of 2 terminals per player e_1 and e_2 can only meet at a point if they have an inner angle of* $180°$.

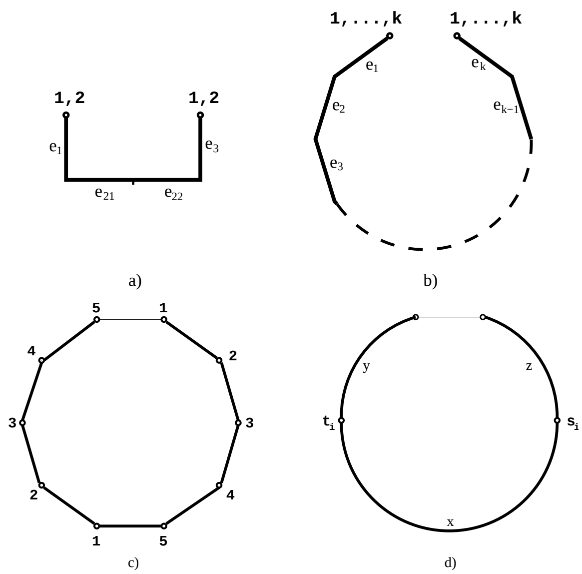

Fig. 1. (a),(b) Geometric games with maximum price of anarchy; (c),(d) Lower bound for approximate Nash purchasing T^*

Consider the game for 2 players and 2 terminals shown in Figure 1a. We have two designated nodes, each containing one terminal of each player. Let $e_2 = e_{21} \cup e_{22}$. The payment scheme purchases T in the following way. Player 1 pays for e_3 and e_{21}. Player 2 pays for e_1 and e_{22}. Let the costs be $e_1 = e_3 = e_{21} = e_{22} = \frac{1}{2}$. e_1 and e_2 as well as e_2 and e_3 are orthogonal. The optimal solution in this network is the direct connection between the terminals. The presented payment scheme, however, forms a Nash equilibrium. Note that the necessary conditions of Lemma 3 are fulfilled. In addition, no player can deviate by simply removing any payment from the network. Lemma 2 restricts the attention to straight segments inside the rectangle, which is the Euclidean hull of T. The argument is given for player 1 – it can be applied symmetrically to player 2. We will consider all meaningful straight segments inside the convex hull of T as deviations. Note that any deviation with both endpoints inside the same edge e_1, e_{21}, etc. (or with endpoints in e_{21} and e_{22}) and any segment between e_1 and e_3 is unprofitable. Now consider a deviation $d = (u, v)$ for player 1 connecting points $u \in e_1$

and $v \in e_{22}$, which are the two segments paid for by player 2. Suppose $d \neq e_{21}$ then $|d| > \frac{1}{2}$. Using d, however, player 1 can save only a cost of $\frac{1}{2}$ by dropping e_{21}. If $u \in e_3$ and $v \in e_{21}$, then d connects segments purchased by player 1. Suppose she defects to such an edge. Let e_3^d be the part of e_3 inside the cycle introduced by d in T (e_{21}^d accordingly). Then with the Phythagorean Theorem and $|e_3^d|, |e_{21}^d| \leq \frac{1}{2}$ the lower bound $|d| \geq |e_{21}^d| + \frac{1}{2} \geq |e_3^d| + |e_{21}^d|$ holds, so d is not profitable for player 1. Hence, all edges player 1 would consider for a deviation are unprofitable. With the symmetric argument for player 2 it follows that the payment scheme represents a Nash equilibrium. Since the optimum solution is half of the cost of T, the theorem follows for games with 2 players and 2 terminals per player.
In a network with more players assume that each player has one terminal at each of the two designated nodes. The nodes are separated by a distance of 1. Construct a path between the nodes, which approximates a cycle with k straight edges of cost 1 each (see Figure 1b). Each player i is assigned to pay for one edge e_i of cost 1. Observe that the necessary conditions of Lemma 3 are fulfilled. Now consider the deviations for a player i. She will neither consider segments that cost more than 1 nor segments that do not allow her to save on e_i. Of the remaining deviations none will yield any profit, because the cyclic structure makes the interior angles between the edges amount to at least $90\,^\circ$. Any deviation $d = (u, v)$ from a point $u \in e_i$ to any other point v will be longer than the corresponding part e_i^d that it allows to save. As the optimum solution is the direct connection of cost 1, Theorem 1 follows. □

This result is contrasted with a result on the price of stability, i.e., the cost of the best Nash equilibrium over the cost of the optimum network.

4 The price of stability

Theorem 2. *The price of stability for geometric connection games with 2 players and 2 terminals per player is 1.*

Proof. We will consider all different games classifying them by the structure of their optimum solution network. The networks in Figure 2 depict the different structures of T^* we consider. In each of them there is an edge $e_3 \in T^1$ and $e_3 \in T^2$. If there is no such edge, the solution is composed of one connection set per player, and a Nash equilibrium can be derived by assigning each player to purchase her subtree T^i.

Type 21 In this case the network is a path (see Figure 2a). The following Lemma 4 describes the structure of meaningful deviations. Nash equilibrium requires that everybody contributes only inside her subtree. That means e_1 is paid by player 1 and e_2 by player 2. The length of every segment lower bounds the deviation costs (by the connection set property)–so no deviation between points of the same segment is meaningful. Furthermore, from the above we know that only straight segments inside the convex hull need to be considered. Hence, the only cases left are described in the lemma below.

Lemma 4. *Given an optimal network T of Type 21 and a payment function p that assigns player i only to pay for edges in T^i, the only deviations for player 1 are straight segments from a point $u \in e_1$ to a point $v \in e_3$, player 2 only from $u \in e_2$ to $v \in e_3$.*

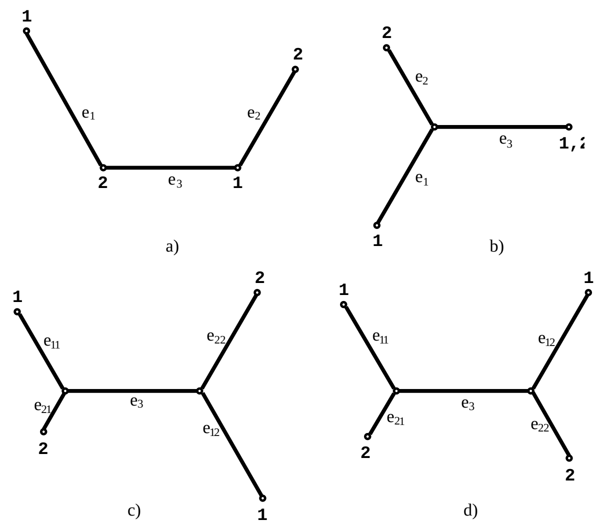

Fig. 2. Network types for T^*

Proof. (sketch) We analyze the payments of player 1. The claim follows for player 2 by symmetry. By Lemma 2 and the fact that the segments of T^* are straight we can restrict the possible deviations to 3 possible cases of straight segments $d = (u, v)$: (1) $u \in e_1$, $v \in e_3$; (2) $u \in e_1$, $v \in e_2$; (3) $u \in e_2$, $v \in e_3$. Cases 2 and 3 can both be disregarded, because either they allow a player to deviate from (parts of) only one connection set, or they can be decomposed into or bounded by deviations from Case 1. □

An adjusted version of Lemma 4 will be true for most of the cases we consider in the remaining proof (cf. Lemmas 5, 6 and 7). Proofs of these lemmas will be omitted.
By the Cosine Theorem the deviation lengths between two adjacent segments are minimized if the angle between segments is minimized, i.e., amounts to 120 ° (cf. Lemma 1). Hence, for the remaining proof we will use an
Angle assumption: all the edges connecting in the optimal solution make inner angles of exactly 120 °.

The following payment scheme forms a Nash equilibrium. Let $e_{3,1}$ be a half subsegment of e_3 connecting the center of e_3 with the terminal of player 1. Similarly, $e_{3,2}$ is the other half subsegment of e_3, connecting the center of e_3 with the terminal of player 2. Then, for player 1, $p_1(e_1) = |e_1|$, $p_1(e_{3,1}) = |e_{3,1}|$, and $p_1 = 0$ elsewhere. For player 2, $p_2(e_2) = |e_2|$, $p_2(e_{3,2}) = |e_{3,2}|$, and $p_2 = 0$ elsewhere.

Note first that the necessary conditions from Lemma 3 are fulfilled. Consider a deviation $d = (u, v)$ in Lemma 4 for player 1 with $u \in e_1$ and $v \in e_3$. As the angle between e_1 and e_3 is exactly 120 °, the length (and cost) of this segment by the Cosine Theorem is

$$|d| = \sqrt{|e_1^d|^2 + |e_3^d|^2 + |e_1^d||e_3^d|},$$

where e_1^d and e_3^d are the segments of e_1 and e_3 in the cycle in $T + d$. The payment of player 1 that can be removed when buying d is $p_1(e_1^d) + p_1(e_3^d) = |e_1^d| + \max(|e_3^d| - \frac{|e_3|}{2}, 0)$. Once v lies in $e_{3,2}$ paid by player 2, $|e_3^d| < \frac{|e_3|}{2}$ and the deviation cannot be cheaper than $|e_1^d|$. Otherwise when $|e_3^d| \geq \frac{|e_3|}{2}$ we can see that $|e_1^d||e_3| + |e_3^d||e_3| - |e_1^d||e_3^d| \geq \frac{|e_3|^2}{4}$. Then it follows that

$$|e_1^d|^2 + |e_3^d|^2 + |e_1^d||e_3^d| \geq |e_1^d|^2 + |e_3^d|^2 + \frac{|e_3|^2}{4} - |e_1^d||e_3| - |e_3^d||e_3| + 2|e_1^d||e_3^d|.$$

Finally we get $|d| \geq |e_1^d| + |e_3^d| - \frac{|e_3|}{2} = p_1(e_1^d) + p_1(e_3^d)$ and see that player 1 has no way of improving her payments. By symmetry the same is true for player 2 and the proof for this network type is completed. □

Type 22 This network type consists of a star, which has a Steiner vertex in the middle and three leaves containing the terminals of the players (see Figure 2b).

Lemma 5. *Given an optimal network T of Type 22 and a payment function p that assigns player i only to pay for edges in T^i, the only deviations for player 1 are straight segments from a point $u \in e_1$ to a point $v \in e_3$, player 2 only from $u \in e_2$ to $v \in e_3$.*

The following is a Nash equilibrium payment scheme. Let $e_3' = (u, v)$, with $u, v \in e_3$, be any subsegment of e_3, where u, v are two interior points on e_3. Then, in the strategy for player 1, $p_1(e_1) = |e_1|$ and $p_1(e_3') = \frac{|e_3'|}{2}$ for any such subsegment e_3' of e_3. For player 2, $p_2(e_2) = |e_2|$ and $p_2(e_3') = p_1(e_3')$ for any subsegment e_3' of e_3. For any other segments in the plane, $p_1 = 0$ and $p_2 = 0$. Consider a deviation $d = (u, v)$ for player 1 with $u \in e_1$ and $v \in e_3$. The amount of payment player 1 can save with this edge is

$$|e_1^d| + \frac{|e_3^d|}{2} = \sqrt{|e_1^d|^2 + |e_1^d||e_3^d| + \frac{|e_3^d|^2}{4}} = \sqrt{|d|^2 - \frac{3|e_3^d|^2}{4}} < |d|.$$

Hence, the deviation is more costly than the possible cost saving for player 1. The proof of a strict Nash for this type follows from the symmetric argument for player 2. □

Type 23 In this network type we have two Steiner points and the terminals of a player are located on different sides of the line through e_3 (see Figure 2c). The connection set for player i formed by edges that are only in T^i now consist of two edges e_{i1} and e_{i2}.

Lemma 6. *Suppose we have Type 23 optimum Steiner network. Under the same assumptions as of Lemma 4 the only deviations player 1 will consider in this game are straight edges from a point $u \in e_{11}$ or a point $u \in e_{12}$ to a point $v \in e_3$, player 2 only from $u \in e_{21}$ or $u \in e_{22}$ to $v \in e_3$.*

We construct an equilibrium payment as follows. For player 1, $p_1(e_{11}) = |e_{11}|$, $p_1(e_{12}) = |e_{12}|$ and $p_1(e_3') = \frac{|e_3'|}{2}$ with e_3' being any subsegment of e_3, as for the scheme of Type 22 above. For player 2, $p_2(e_{21}) = |e_{21}|$, $p_2(e_{22}) = |e_{22}|$ and $p_2(e_3') = p_1(e_3')$. Otherwise, $p_1 = 0$ and $p_2 = 0$. For all possible deviations, the cost is greater than the contribution to T^* a player could save. This follows with the proof of Type 22. □

Type 24 The last network type considered is the one including two Steiner points where the terminals of a player are located on the same side of the line through e_3 (see Figure 2d). Here we get some additional deviations that complicate the analysis.

Lemma 7. *Given a network T of Type 24 and a payment function that assigns payments to player i only in her subtree T^i. Then the only deviations player 1 considers are straight edges between $u \in e_{11}$ or $u \in e_{12}$ and $v \in e_3$ as well as the direct connection between her terminals. For player 2 the symmetric claim holds.*

To present the payment function, we scale our game such that e_3 has length 1. We now treat e_3 as an interval $[0, 1]$ and introduce a function $f(x, y) \in [0, 1], 0 \leq x \leq y \leq 1$ that specifies the fraction of the cost player 1 pays in the interval $[x, y]$ of e_3, i.e., the payment of player 1 on $[x, y]$ is $(y - x)f(x, y)$. Let, w.l.o.g., the Steiner point of e_{11} be point 0 of e_3 and the other Steiner point be point 1. We now have to ensure that for every deviation from e_{11} or e_{12} to a point $y, 1 - y \in e_3$ the savings on the segments do not exceed the cost of the deviation. This results in the bounds $|e_{11}| + yf(0, y) \leq \sqrt{|e_{11}|^2 + y^2 + |e_{11}|y}$ and $|e_{12}| + yf(1 - y, 1) \leq \sqrt{|e_{12}|^2 + y^2 + |e_{12}|y}$. For player 2 the symmetric requirements lead to similar bounds with $|e_{21}|$ and $|e_{22}|$. Furthermore, we can derive bounds from the direct connections between the terminals. They will be denoted as d_1 and d_2 for players 1 and 2, respectively. With the optimality of our network and $|e_3| = 1$ we have

$$|d_1| + |d_2| \geq |e_{11}| + |e_{12}| + |e_{21}| + |e_{22}| + 1. \tag{1}$$

As we strive for a Nash payment scheme, d_1 and d_2 are not cheaper than the contribution of the players, hence $|d_1| \geq |e_{11}| + |e_{12}| + f(0, 1)$ and $|d_2| \geq |e_{21}| + |e_{22}| + 1 - f(0, 1)$. The nature of these edges implies that their bounds only apply to the payment on the whole segment e_3, i.e., they do not restrict the partition of the payment inside the segment. Using a function $h(x) = \sqrt{x^2 + x + 1} - x$ and solving for f the previous bounds can be turned into

$$\begin{aligned} |e_{21}| + |e_{22}| + 1 - |d_2| \quad &\leq \quad f(0,1) \quad \leq \quad |d_1| - |e_{11}| - |e_{12}|, && (2) \\ 1 - h\left(\frac{|e_{21}|}{y}\right) \quad &\leq \quad f(0,y) \quad \leq \quad h\left(\frac{|e_{11}|}{y}\right), && (3) \\ 1 - h\left(\frac{|e_{22}|}{y}\right) \quad &\leq f(1-y,1) \leq \quad h\left(\frac{|e_{12}|}{y}\right). && (4) \end{aligned}$$

Now consider the behavior of $h(x)$ in (3) and (4) when altering the constants $|e_{11}|$ and $|e_{12}|$. We observe that for the derivative $h'(x) < 0$ holds. The function is monotone decreasing in x, and increasing $|e_{11}|, |e_{12}|, |e_{21}|, |e_{22}|$ tightens lower and upper bounds. So we will only consider deviations from terminals to e_3, as this results in the strongest bounds for the Nash payments.

We also require that payments can be feasibly split to subintervals. The payment of player 1 on an interval $[x, y]$ has to be the sum of the payments on the two subintervals $[x, v]$ and $[v, y]$ for any $v \in [x, y]$. Using this property, we can define $f(x, y)$ by using the functions $f(0, y)$ and $f(1 - y, 1)$. Namely, for $0 \leq x \leq y \leq 1$, we have:

$$f(x, y) = \frac{yf(0, y) - xf(0, x)}{y - x}, \quad f(1 - y, 1 - x) = \frac{yf(1 - y, 1) - xf(1 - x, 1)}{y - x}. \tag{5}$$

In particular, we will focus on symmetric payment functions, i.e., we will assume that $f(0, y) = f(1-y, 1)$ for any $y \in [0, 1]$. This also implies that $f(x, y) = f(1-y, 1-x)$, where $0 \leq x \leq y \leq 1$. For the rest of the proof we will provide a feasible function $f(0, y)$, which obviously must obey all bounds (2)-(4). First, we pay some attention to the feasibility of the bounds.

Lemma 8. *The bounds (2)–(4) do not imply a contradiction. In particular the interior bounds (3), (4) can be fulfilled by* $f(0, y) = \frac{1}{2}$.

Proof. We already know that the upper bound function $h(x)$ is monotone decreasing in x. We observe that for any $x, x' > 0$

$$\lim_{x \to \infty} (1 - h(x)) \;=\; \frac{1}{2} \;=\; \lim_{x \to \infty} h(x), \quad \text{and} \quad 1 - h(x) \;\leq\; \frac{1}{2} \;\leq\; h(x').$$

This proves the second part of the lemma. Regarding the first part, a negation of the bounds leads to a contradiction with the Triangle Inequality with d_1 and d_2. □

Lemma 8 supports our proofs for the previous network Types 23 and 22. The function used is the linear function $f(x, y) = \frac{1}{2}$ and satisfies bounds (3), (4), which are the only ones present. For Type 24 network a solution is possible as well. In the easiest case if

$$|e_{21}| + |e_{22}| + 1 - |d_2| \leq \frac{1}{2} \leq |d_1| - |e_{11}| - |e_{12}| \tag{6}$$

holds, $f(0, y) = f(x, y) = \frac{1}{2}$ again gives a Nash equilibrium payment function. Hence, for the remainder of the proof we will assume (6) is not valid. A solution for this more complicated situation is presented in the next lemma.

Lemma 9. *There is a constant* t *and one of the two functions*

$$f_1(0, y) = h\left(t/y\right) \quad \textit{or} \quad f_2(0, y) = 1 - h\left(t/y\right)$$

that allow us to construct a payment scheme forming a Nash equilibrium in a network of Type 24.

Proof. In the first case we assume that $r = |e_{21}| + |e_{22}| + 1 - |d_2| > \frac{1}{2}$. Then $f(0, y)$ must behave like the upper bounds and achieve a value of r for $y = 1$, so $f_1(0, y) = h(t/y)$ with $t = \frac{1-r^2}{2r-1}$.
In the second case we assume $r = |d_1| - |e_{11}| - |e_{12}| < \frac{1}{2}$. Then $f(0, y)$ must behave like the lower bounds and achieves a value of r for $y = 1$, so $f_2(0, y) = 1 - h(t/y)$ with $t = \frac{1-r^2}{2r-1} - 1$. To achieve a consistent definition of $f(x, y)$ we define $f(1, 1) = f(0, 0) := \lim_{y \to 0} f(0, y) = \frac{1}{2}$.
Then, with Lemma 8 and the monotonicity of $h(x)$ we see that the functions f_1, f_2 obey bounds (2)–(4) for any $y \in [0, 1]$. Lemma 8 says that the upper bound functions of (3)

and (4) map only to $[0.5, 1]$, and the lower bound functions map only to $[0, 0.5]$. Since f_1 maps only to $[0.5, 1]$, all lower bounds are feasible. The upper bounds for f_1 are feasible by the monotonicity of $h(\cdot)$ – the constants t in f_1 are smaller than appropriate constants in the upper bounds. Similarly for f_2. Functions f_1 and f_2 allow to construct a Nash equilibrium payment function. If the payment of player 1 is given by f_1 (f_2), then the payment of player 2 is given by a function f_2 (f_1) with the same constant r as for player 1. This concludes the proofs of Lemma 9 and Theorem 2. □

The proof of the theorem requires that an edge e_3 is purchased such that the payments of players on the intervals follow a continuous differentiable function. This is a quite strong and very unrealistic property. We present two possible alternatives to avoid this. First a discretization of the payment scheme on e_3 is considered such that subsegments of the network are assigned to be purchased completely by single players. It slightly increases the incentives to deviate. Second we let players play the game according to best-response strategies. This will lead into a low-cost strict Nash equilibrium. Full proofs are omitted due to space constraints. A *divisible* payment scheme $p = (p_1, \ldots, p_k)$ for a geometric connection game is a payment scheme such that there is a partition $\mathcal{P}$ of the plane into segments with $p_i(e) = 0$ or $p_i(e) = |e|$ for all $i = 1, \ldots, k$ and all $e \in \mathcal{P}$.

Theorem 3. *Given any $\epsilon > 0$ and any geometric connection game with 2 players and 2 terminals per player, there exists a divisible payment scheme, which is a $(1 + \epsilon)$-approximate Nash equilibrium as cheap as the optimum solution.*

Theorem 4. *In any geometric connection game with 2 players and 2 terminals per player, there exists a Nash equilibrium generated by at most 3 steps of the best-response dynamics, which is a $\sqrt{2}$-approximation to the optimum solution.*

The results cannot be generalized to games with more players. For games with 3 or more players there is a constant lower bound on approximate Nash equilibria.

Theorem 5. *For any $k \geq 3, \epsilon > 0$ there exists a game with k players and 2 terminals per player, for which every optimum solution is at least a $(\frac{4k-2}{3k-1} - \epsilon)$-approximate Nash equilibrium.*

Proof. In the class of games delivering the bound there is a circle of terminals with unit distance, and the optimal solution is a minimum spanning tree of cost $2k - 1$. In the geometric environment edges crossing the interior of the circle are not of interest, because their cost is always larger than 1. So no player will consider them as a reasonable alternative. Consider the game in Figure 1c, in which every edge of T^* has cost 1. T^* is depicted with an additional edge of cost 1, which will be the only deviation edge considered. The situation for a player i can be simplified to the view of Figure 1d. Note that for players 1 and k, $z = 0$ and $y = 0$, respectively. For every player there are at least two ways to deviate, either she just contributes to one half of the cycle by paying part x of this half, or she completes the other half of the cycle by paying $y + z + 1$, where y, z are the parts she pays on the depicted portions of the cycle. Thus her deviation factor will be at least $\max\left\{\frac{x+y+z}{x}, \frac{x+y+z}{y+z+1}\right\}$. Minimizing this expression with $x = y + z + 1$ there is at least one player, who pays for $x + y + z = 2x - 1 \geq \frac{2k-1}{k}$. Solving for x and combining with $x = y + z + 1$ results in: $\frac{x+y+z}{x} = \frac{2x-1}{x} \geq \frac{4k-2}{3k-1}$. Now move the terminals s_1 and t_k a little further to the outside keeping the lengths of the edges (s_1, s_2) and (t_{k-1}, t_k) to 1, but increasing the length of (s_1, t_k) to length $(1 + \epsilon)$. T^* will then be the unique optimal solution, and the factor becomes at least $\left(\frac{4k-2}{3k-1} - \epsilon\right)$. □

References

1. E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 295–304, 2004.
2. E. Anshelevich, A. Dasgupta, É. Tardos, and T. Wexler. Near-optimal network design with selfish agents. In *Proceedings of the 35th Annual Symposium on Theory of Computing (STOC)*, pages 511–520, 2003.
3. S. Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *Journal of the ACM*, 45(5):753–782, 1998.
4. V. Bala and S. Goyal. A non-cooperative model of network formation. *Econometrica*, 68:1181–1229, 2000.
5. M. Chlebík and J. Chlebíková. Approximation hardness of the Steiner tree problem in graphs. In *Proceedings of the 8th Scandinavian Workshop on Algorithm Theory (SWAT)*, pages 170–179, 2002.
6. A. Czumaj, P. Krysta, and B. Vöcking. Selfish traffic allocation for server farms. In *Proceedings of the 34th Annual ACM Symposium on the Theory of Computing (STOC)*, pages 287–296, 2002.
7. M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry - Algorithms and Applications*. Springer Verlag, 1997.
8. D. Dutta, A. Goel, and J. Heidemann. Oblivious AQM and Nash equilibrium. In *Proceedings of the 22nd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2003.
9. A. Fabrikant, A. Luthera, E. Maneva, C. Papadimitriou, and S. Shenker. On a network creation game. In *Proceedings of the 22nd Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 347–351, 2003.
10. E. Gilbert and H. Pollak. Steiner Minimal Trees. *SIAM Journal on Applied Mathematics*, 16:1–29, 1968.
11. M. Goemams and D. Williamson. A general approximation technique for constrained forest problems. *SIAM Journal on Computing*, 24(2):296–317, 1995.
12. H. Heller and S. Sarangi. Nash networks with heterogeneous agents. Technical Report Working Paper Series, E-2001-1, Virginia Tech, 2001.
13. E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 404–413, 1999.
14. Z. Melzak. On the problem of Steiner. *Canadian Mathematical Bulletin*, 4:143–148, 1961.
15. G. Robins and A. Zelikovsky. Improved Steiner tree approximation in graphs. In *Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 770–779, 2000.
16. T. Roughgarden and É.Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.
17. A. Schulz and N. Stier Moses. Selfish routing in capacitated networks. *Mathematics of Operations Research*, 29(4):961–976, 2004.