Geometric Network Design with Selfish Agents

Martin Hoefer* Piotr Krysta†

Abstract

A simple non-cooperative network creation game has been introduced in [2]. In this paper we study a geometric version of this game, assuming Euclidean metric edge costs on the plane. The price of anarchy in such geometric games with k players is still $\Theta(k)$. Hence, we consider the task of minimizing players incentives to deviate from a payment scheme, purchasing the minimum cost network, which was introduced in [2]. 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 using some more practical assumptions, for any $\epsilon > 0$. For more players, however, 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 is the fact that we needed to develop new techniques to deal with the geometric setting, which are fundamentally different from the techniques used in [2] for general games.

1 Introduction

The network analysis yields a variety of interesting questions, which are important for many areas in research and applications. One of the most dynamic driving forces in modern society is the existence of the Internet–a powerful and universal artefact in human history. An interesting research direction explored recently 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* [16], 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, 7, 16, 20], routing [19, 21, 22], facility location [9, 24], and flow control [10, 23]. A slightly different measure–the cost of the best Nash equilibrium instead of the worst was considered in [22]. This is the optimum solution no user has an incentive to defect 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 recently 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

*Department of Computer and Information Science, Konstanz University, Box D 67, 78457 Konstanz, Germany. hoefer@inf.uni-konstanz.de. This author is partially supported by DFG grant Kr 2332/1-1 within Emmy Noether program and DFG grant GRK 1042/1.

†Department of Computer Science, Dortmund University, Baroper Str. 301, 44221 Dortmund, Germany. piotr.krysta@cs.uni-dortmund.de. This author is supported by DFG grant Kr 2332/1-1 within Emmy Noether program.

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 network design problem. The problem of finding a minimum cost network satisfying all connection needs and minimizing the sum of all players payments is the generalized Steiner tree 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.
Furthermore, it is *N*P-complete to determine, whether a given game has any Nash equilibrium at all. 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 the 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 be possible, a payment scheme was presented that determines a 3-approximate Nash equilibrium on the socially optimum network, i.e., purchases the minimum cost network and allows each player to reduce her costs by at most a factor of 3 by deviating. Finding the minimum cost network, however, is *N*P-hard, and the currently best known approximation algorithms for the (generalized) Steiner tree problem have an approximation factor of (2) 1.55 [13, 18]. Using these algorithms for any game a payment scheme can be found in polynomial time that presents a $(4.65 + \epsilon)$-approximate Nash equilibrium purchasing a 2-approximate network [2].
Connection games are related to the field of network creation. Fabrikant et al. [11] proposed a different network creation game, in which each player corresponds to a node. A player can only contribute to edges that are incident to her node. A similar game was also considered by [4, 14] 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 a classic cost allocation protocol, namely the Shapley value. Each edge is assumed to be shared equally among the players using it. In this setting they could prove an $O(\log k)$ upper bound on the price of stability. They considered bounds on the convergence of best-response dynamics and derived extended results for 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 [8]. Furthermore, for the geometric versions of combinatorial optimization problems usually improved results can be derived by employing the specific Euclidean structure. For example the geometric Steiner tree problem allows a PTAS [3], which contrasts the inapproximability results 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 $\epsilon > 0$. This is a significant improvement over the general case, where games with 2 players and 2 terminals per player exist such that any Nash equilibria purchasing the optimum network is at least $(\frac{6}{5} - \epsilon)$-approximate [15].
- The case of 2 players with 2 terminals per player may seem a very special one, but, it turns out that we 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 corresponding 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, which, e.g., makes the number of possible player's strategies potentially unbounded. Thus, 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 one of the contributions of our paper besides the results listed above. For example, consider how a shared edge is paid for when buying an optimum solution in a game with 2 players and 2 terminals per player. We define an abstract continuous function that specifies the payments on all subintervals of the shared edge. Nash properties are shown to imply some geometric parameterized bounds on this function. We then find such a function as a solution to an appropriate system of functional inequalities, which leads to an exact Nash equilibrium. The continuity assumptions of the payment function, however, may not be realistic. We show how to relax these assumptions by discretizing the payment function, which leads to an $(1+\epsilon)$-approximate Nash equilibrium, for any $\epsilon > 0$.

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 can be found in the appendix.

2 The model and preliminaries

The geometric connection game is defined as follows. Let V be a set of nodes located 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 is determined by 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 for 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* $p = (p_1, \dots, p_k)$. 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.[1] 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.

The problem of constructing a minimum cost network satisfying all connection needs is the geometric Steiner forest problem. As the components of a Steiner forest are Steiner trees for a subset of players, some well-known properties of optimum geometric Steiner trees hold for T^*.

Lemma 1 *[12, 17] Any 2 adjacent edges in an optimal geometric Steiner tree connect with an inner angle of at least 120°.*

Lemma 2 *[12, 17] 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.*

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 somehow 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 $\{j \in \{1, \dots, k\} : e \in T^j\} =$

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

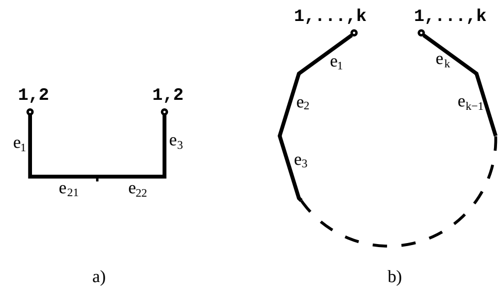

Figure 1: Geometric games with maximum price of anarchy.

$\{j' \in \{1, \ldots, k\} : e' \in T^{j'}\}$. Our proofs will rarely use connection sets. We will rather deploy geometric arguments and use connection sets only to limit the number and types of segments and 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 3 and 4 follow directly from the Triangle Inequality.

Lemma 3 *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.*

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

Lemma 4 *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 $180°$.*

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 schemes, however, form a Nash equilibrium. Note that the necessary conditions of Lemma 4 are fulfilled. In addition, no player can deviate by simply removing any payment from the network. Lemma 3 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 first that a deviation with both endpoints inside the same edge e_1, e_{21}, etc. (or with endpoints in e_{21} and e_{22}) is not profitable, because the segments are straight. Note further that any segment between e_1 and e_3 is unprofitable, because its length is at least 1. 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 $|e_3^d|, |e_{21}^d| \leq \frac{1}{2}$

$$|d| \geq |e_{21}^d| + \frac{1}{2} \geq |e_3^d| + |e_{21}^d|$$

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 the network with more players assume that each player has one terminal at each of the two designated nodes. The nodes are again 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 4 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 him 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 $90\,°$. 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. This argument is valid for any player i. As the optimum solution is the direct connection of cost 1, the theorem 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 bold networks in Figures 2 and 3 depict the different structures of the optimal solution, denoted T^*, we consider. These are only solutions, in which 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.

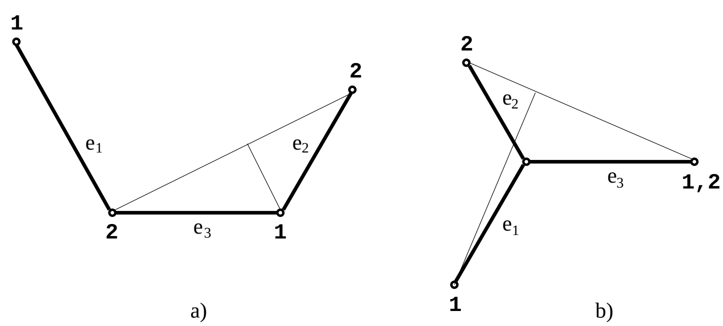

Figure 2: Network types with one edge per connection set

Type 2.1 In this case the network is a path (see Figure 2 a). The following Lemma 5 describes the structure of meaningful deviations. Nash equilibrium requires that everybody pays for what is 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, there are only the cases left as described in the lemma below.

Lemma 5 *Given an optimal network T of Type 2.1 and a payment function p that assigns player i only to pay for edges in T^i, the only deviations player 1 will pick 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 $v \in e_3$.*

Proof. We analyze the payments of player 1. The observation follows symmetrically for player 2. With Lemma 3 and the fact that the segments of T^* are straight we can restrict the possible deviations to 3 possible

types of straight segments $d = (u, v)$: $u \in e_1$, $v \in e_3$; $u \in e_1$, $v \in e_2$ and $u \in e_2$, $v \in e_3$.
The deviations of case 3 allow player 1 to deviate from (parts of) only one connection set, namely $\{e_3\}$. Payments on one connection set, however, provide lower bounds for the deviations, and so they are not profitable. If e_1 and e_2 are located on different sides of the line through e_3, an edge of case 2 always crosses e_3 and therefore decomposes into two edges of the other cases. If e_1 and e_2 are located on the same side (this is depicted in Figure 2a), then by easy observation for each edge of case 2 there is a cheaper edge, which starts at the same point $u \in e_1$ and goes directly to the terminal of player 1 on e_3. However, as 1 does not pay on e_2, the savings on e_1 and e_3 for player 1 are the same for both edges. Hence an edge of case 1 provides a superior way to deviate. This proves the lemma. □

An adjusted version of this lemma will be true for most of the cases we consider in the remaining proof. 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. Lemmas 1 and 2). Hence, we will use an *angle assumption* for the remaining proof, i.e.,
Angle assumption: all the edges connecting in the optimal solution make inner angles of exactly 120 °.

Consider the following payment scheme, which 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 4 are fulfilled. Consider a deviation $d = (u, v)$ from Lemma 5 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 2.2 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). The proof of the next lemma is in the appendix.

Lemma 6 *Given an optimal network T of Type 2.2 and a payment function p that assigns player i only to pay for edges in T^i, the only deviations player 1 will pick 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 always 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.

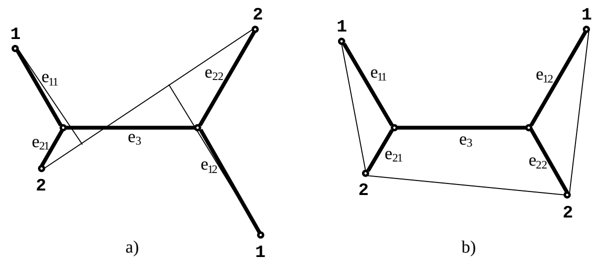

Figure 3: Network types with connections sets e_1 and e_2 having two edges

Type 2.3 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 3a). 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}. The proof of the next lemma can be found in the appendix.

Lemma 7 *Suppose we have Type 2.3 optimum Steiner network. Under the same assumptions as of Lemma 5 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 2.2 above. For player 2, $p_2(e_{21}) = |e_{21}|$, $p_2(e_{22}) = |e_{22}|$ and $p_(e_3') = p_1(e_3')$. 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 2.2.

Type 2.4 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 3b). Here we get some additional deviations that complicate the analysis. The proof of the next lemma is in the appendix.

Lemma 8 *Given a network T of Type 2.4 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 $f(x, y) \in [0, 1], 0 \le x \le y \le 1$ 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 following bounds:

$$|e_{11}| + yf(0, y) \le \sqrt{|e_{11}|^2 + y^2 + |e_{11}|y}, \quad (1)$$
$$|e_{12}| + yf(1 - y, 1) \le \sqrt{|e_{12}|^2 + y^2 + |e_{12}|y} \quad (2)$$

For player 2 the symmetric requirements lead to

$$|e_{21}| + y(1 - f(0, y)) \le \sqrt{|e_{21}|^2 + y^2 + |e_{21}|y}, \quad (3)$$
$$|e_{22}| + y(1 - f(1 - y, 1)) \le \sqrt{|e_{22}|^2 + y^2 + |e_{22}|y}. \quad (4)$$

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

$$|d_1| + |d_2| \geq |e_{11}| + |e_{12}| + |e_{21}| + |e_{22}| + 1. \tag{5}$$

As we strive for a Nash payment scheme, d_1 and d_2 are not cheaper than the contribution of the players:

$$\begin{aligned} |d_1| &\geq |e_{11}| + |e_{12}| + f(0,1), & (6)\\ |d_2| &\geq |e_{21}| + |e_{22}| + 1 - f(0,1). & (7) \end{aligned}$$

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 we get

$$\begin{aligned} |e_{21}| + |e_{22}| + 1 - |d_2| \quad &\leq \quad f(0,1) \quad \leq \quad |d_1| - |e_{11}| - |e_{12}|, & (8)\\ 1 - h\left(\frac{|e_{21}|}{y}\right) \quad &\leq \quad f(0,y) \quad \leq \quad h\left(\frac{|e_{11}|}{y}\right), & (9)\\ 1 - h\left(\frac{|e_{22}|}{y}\right) \quad &\leq \quad f(1-y,1) \quad \leq \quad h\left(\frac{|e_{12}|}{y}\right). & (10) \end{aligned}$$

Now consider the behavior of $h(x)$ in (9) and (10) when altering the constants $|e_{11}|$ and $|e_{12}|$. We observe that for the derivative of $h(x)$

$$h'(x) = \frac{2x + 1 - 2\sqrt{x^2+x+1}}{2\sqrt{x^2+x+1}} < 0 \tag{11}$$

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.

In addition to these bounds 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 $f(1 - y, 1)$:

$$f(x,y) = \frac{yf(0,y) - xf(0,x)}{y-x}, \quad f(1-y,1-x) = \frac{yf(1-y,1) - xf(1-x,1)}{y-x}, \quad 0 \leq x \leq y \leq 1. \tag{12}$$

In particular, we will focus on symmetric payment functions, i.e., we will assume that $f(0, y) = f(1 - y, 1)$ for any $y \in [0, 1]$. Observe, that 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 strive to provide a feasible function $f(0, y)$, which obviously must obey all bounds (8)-(10). First, we pay some attention to the feasibility of the bounds.

Lemma 9 *The bounds (8)-(10) do not imply a contradiction. In particular the interior bounds (9), (10) 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 $x, x' > 0$

$$\begin{aligned} \lim_{x\to\infty}(1 - h(x)) \quad &= \quad \tfrac{1}{2} \quad = \quad \lim_{x\to\infty} h(x), & (13)\\ 1 - h(x) \quad &\leq \quad \tfrac{1}{2} \quad \leq \quad h(x'). & (14) \end{aligned}$$

This proves the second part of the lemma.
With (5) it is obvious that (8) is no contradiction. Assume that the bounds in (8) and (9) form a contradiction. Then at least one of

$$\begin{aligned} 1 - h(|e_{21}|) \quad &> \quad |d_1| - |e_{11}| - |e_{12}|, & (15)\\ |e_{21}| + |e_{22}| + 1 - |d_2| \quad &> \quad h(|e_{11}|) & (16) \end{aligned}$$

must hold. However, with (5) we have

$$|e_{22}| + \sqrt{|e_{21}|^2 + 1 + |e_{21}|} \quad < \quad |d_2|, \tag{17}$$

$$|e_{12}| + \sqrt{|e_{11}|^2 + 1 + |e_{11}|} \quad < \quad |d_1| \tag{18}$$

and it is easy to see that both bounds give a contradiction with the Triangle Inequality. Thus the bounds above are feasible. The first part of the lemma follows with a similar argument for the bounds of (8) and (10). □

This result supports our proofs for the previous network Types 2.3 and 2.2. The function used is the linear function $f(x, y) = \frac{1}{2}$ and satisfies the bounds (9) and (10), which are the only ones present. For network Type 2.4 a solution is possible as well. In the easiest case if

$$|e_{21}| + |e_{22}| + 1 - |d_2| \le \frac{1}{2} \le |d_1| - |e_{11}| - |e_{12}| \tag{19}$$

holds, $f(0, y) = f(x, y) = \frac{1}{2}$ again is a solution. The remaining analysis for this case will then be the same as for the previous network Type 2.3 and the same results follow. Hence, for the remainder of the proof we will assume (19) is not valid. A solution for this more complicated situation is presented in the next lemma.

Lemma 10 *There is a constant t and one of the two functions*

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

that allow us to construct a payment scheme forming a Nash equilibrium in network of Type 2.4.

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 $y = 1$:

$$f_1(0, y) = h\left(\frac{t}{y}\right) = \frac{\sqrt{t^2 + ty + y^2} - t}{y}, \quad t = \frac{1 - r^2}{2r - 1}.$$

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 $y = 1$:

$$f_2(0, y) = 1 - h\left(\frac{t}{y}\right) = 1 - \frac{\sqrt{t^2 + ty + y^2} - t}{y}, \quad t = \frac{1 - r^2}{2r - 1} - 1.$$

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 9 and the monotonicity of $h(x)$ we see that the functions f_1, f_2 obey the bounds (8)-(10) for any $y \in [0, 1]$. Namely, Lemma 9 says that the upper bound functions of (9) and (10) 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 also feasible, because the constants t involved in f_1 are smaller than appropriate constants in the upper bounds and we use the monotonicity of $h(\cdot)$. 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 proof of Theorem 2. □

The proof of the theorem relies on the fact that an edge e_3 can be purchased such that the payments of players on the intervals of e_3 follow a continuous differentiable function. This seems a rather 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. The adjustment slightly increases the incentives to deviate. Another way to overcome this is to let players play the game according to best-response strategies. This will lead into a strict low-cost Nash equilibrium. A *divisible* payment scheme $p = (p_1, \ldots, p_k)$ for a geometric connection game is a payment scheme such that there exists a partition $\mathcal{P}$ of the plane into segments such that $p_i(e) = 0$ or $p_i(e) = |e|$ for all $i = 1, \ldots, k$ and all $e \in \mathcal{P}$. For proofs of the next two theorems see the appendix.

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 best-response dynamics, which is a $\sqrt{2}$-approximation to the optimum solution.*

Unfortunately, the nice results for small games cannot be generalized to games with more players. For games with 3 or more players there exists a constant lower bound on approximate Nash equilibria, hence we cannot achieve deviation factors arbitrarily close to 1 when purchasing an optimum solution network.

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. The proof is similar to the one for the general case given in [2]. 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. Actually, their cost exceeds 2 once the number of players is more than 4, which then is more than the asymptotical payment of each player in the best payment scheme. So no player will consider them as a reasonable alternative. Consider the game in Figure 4, in which every edge of T^* has

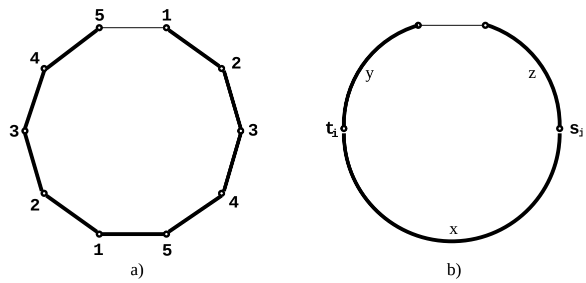

Figure 4: Lower bound for approximate Nash purchasing the optimal network.

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 can be simplified to the view of Figure 4b. 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 is assigned to pay 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 $\left(\frac{4k-2}{3k-1} - \epsilon\right)$. □

Observe that this lower bound proof applies exclusively to games with $k \geq 3$ players. For 2 players the indicated T^* would not be the optimum solution, as it would involve inner angles of less than 120 °.

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] A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibira. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 413–420, 2002.

[8] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry - Algorithms and Applications*. Springer Verlag, 1997.

[9] N. Devanur, N. Garg, R. Khandekar, V. Pandit, A. Saberi, and V. Vazirani. Price of anarchy, locality gap, and a network service provider game. Unpublished manuscript, 2004.

[10] 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.

[11] 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 Processing (PODC)*, pages 347–351, 2003.

[12] E. Gilbert and H. Pollak. Steiner Minimal Trees. *SIAM Journal on Applied Mathematics*, 16:1–29, 1968.

[13] M. Goemams and D. Williamson. A general approximation technique for constrained forest problems. *SIAM Journal on Computing*, 24(2):296–317, 1995.

[14] H. Heller and S. Sarangi. Nash networks with heterogeneous agents. Technical Report Working Paper Series, E-2001-1, Virginia Tech, 2001.

[15] M. Hoefer. Network connnection games. Master's thesis, Technical University of Clausthal, September 2004.

[16] 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.

[17] Z. Melzak. On the problem of Steiner. *Canadian Mathematical Bulletin*, 4:143–148, 1961.

[18] 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.

[19] T. Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.

[20] T. Roughgarden. Stackelberg scheduling strategies. *SIAM Journal on Computing*, 33(2):332–350, 2004.

[21] T. Roughgarden and É.Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.

[22] A. Schulz and N. Stier Moses. Selfish routing in capacitated networks. *Mathematics of Operations Research*, 29(4):961–976, 2004.

[23] S. Shenker. Making greed work in networks: A game-theoretic analysis of switch service disciplines. *IEEE/ACM Transactions on Networking*, 3(6):819–831, 1995.

[24] A. Vetta. Nash equilibria in competitive societies with application to facility location, traffic routing and auctions. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, page 416, 2002.

A Appendix

A.1 Proofs of three lemmas

Proof. (of Lemma 6) This is straightforward because every other deviation allows the corresponding player to save on (parts of) only one connection set, but these parts provide lower bounds for the deviations. □

Proof. (of Lemma 7) We apply the corresponding arguments from the previous network types to exclude nearly all other deviations. The final argument is provided by noting that an edge between $u \in e_{11}$ and $v \in e_{12}$ crosses e_3 and therefore decomposes to two edges of the specified form. The statement for player 2 follows by symmetry. □

Proof. (of Lemma 8) The previous arguments help to exclude all other deviations. Unlike Type 2.3, a direct connection between e_{11} and e_{12} does not cross e_3 and the case does not decompose. If, however, player 1 picks such a deviation, by the Triangle Inequality she will be most profitable to pick the direct connection between her terminals. □

A.2 Discretization

Proof of Theorem 3. In this situation we assign certain parts of the network to be paid exclusively by one player. This can be achieved with a technical adjustment: We split the network in small pieces, which are assigned to be bought by one of the players. In this way we can generate divisible payment schemes with a low incentive to deviate depending on the grain of the splitting intervals. Let us again consider the different types of games leaving aside Type 2.1, where a divisible payment scheme was already constructed.

Type 2.2 We assume the segment e_3 is discretized into $\frac{|e_3|}{t}$ segments $e_{3,j}, j = 1, \ldots, \frac{|e_3|}{t}$ of length t. Let the enumeration start at the Steiner vertex–see Fig. 2b. For the ease of presentation we define the connection point between $e_{3,j}$ and $e_{3,j+1}$ to be the *right end* of $e_{3,j}$ and the *left end* of $e_{3,j+1}$. One of the two players must buy the piece $e_{3,1}$ of length t, whose left end is the Steiner point. This player then buys two consecutive straight lines and violates the conditions from Lemma 4. She will be able to reduce her payments by connecting directly between the endpoints.
Consider the following payment function: The segments $e_{3,j}$ are assigned alternating such that every two neighboring segments are paid by different players. Player 1 buys the first and the other 'odd' segments. d is a deviation from $u \in e_1$ to $v \in e_3$. Note that once v is in a segment $e_{3,j}$ bought by player 2, there is a deviation to the left end of $e_{3,j}$ that is cheaper and allows player 1 the same savings. She will only deviate to points located on segments $e_{3,j}$, which she buys. Furthermore, we show below that player 1 can achieve the best cost reduction by deviating to $e_{3,1}$ (see Figure 5).

Lemma 11 *For every deviation d from $u \in e_1$ to a point in segment $e_{3,j}$, $j \neq 1$ there is a deviation d' from u to a point in $e_{3,1}$, which yields a higher cost reduction for player 1.*

Proof. Assume $d = (u, v)$ with $u \in e_1, v \in e_{3,j} \neq e_{3,1}$ and e_1^d, e_3^d being the corresponding segments in the cycle. $e_{3,j}$ is paid by player 1. Let r be the number of full segments player 1 buys on e_3^d. Let e_3' be the segment between the left end of $e_{3,j}$ and v. Thus $p_1(e_3^d) = rt + e_3'$ and $|e_3^d| = 2rt + |e_3'|$. v' is to be the point on $e_{3,1}$ that has distance $|e_3'|$ to the left end of $e_{3,1}$ and $d' = (u, v')$ the deviation to this point. The lemma states that

$$\frac{p_1(e_1^d) + p_1(e_3)}{|d'| + p_1(e_3) - |e_3'|} \geq \frac{p_1(e_1^d) + p_1(e_3)}{|d| + p_1(e_3) - rt - |e_3'|}$$

Note that we consider a fixed point of e_1, however, two different points on e_3. We have to consider that the deviation to v is more costly than d', however, also allows player 1 to decrease the contribution to the cost of e_3 by an additional amount of rt. It follows that

$$|d| \geq |d'| + rt$$

As all terms are nonnegative, the inequality allows squaring on both sides:

$$|d|^2 \geq |d'|^2 + 2rt|d'| + (rt)^2.$$

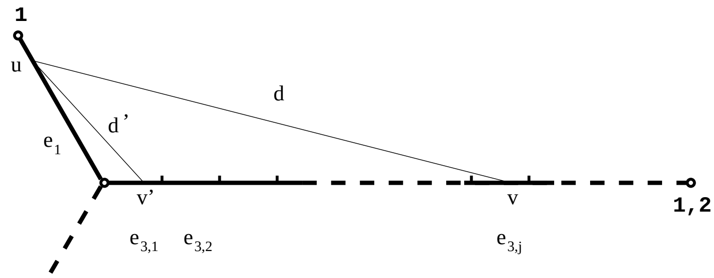

Figure 5: d' is more profitable than d

Then with

$$\begin{aligned} |d| &= \sqrt{|e_1^d|^2 + (2rt + |e_3'|)^2 + (2rt + |e_3'|)|e_1^d|} \\ |d'| &= \sqrt{|e_1^d|^2 + |e_3'|^2 + |e_1^d||e_3'|} \end{aligned}$$

this reduces to

$$\frac{3}{2}(rt)^2 + 2|e_3'| + |e_1^d| \geq |e_3'| + |e_1^d| \geq |d'|,$$

which holds by the Triangle Inequality. □

Having established this result, we continue by upper bounding the cost reduction player 1 can achieve. Player 1 will always deviate from her terminal to the right end of $e_{3,1}$. The relative reduction will be maximized, if she has no additional payments than $|e_1^d| + |e_{3,1}|$. Then with

$$\epsilon_1(x,t) = \frac{x+t}{\sqrt{x^2+t^2+xt}}$$

the deviation factor becomes

$$1 + \epsilon(t) \leq \epsilon_1(|e_1|, t).$$

Observe, that $\lim_{t\to 0} \epsilon_1(|e_1|, t) = 1$. Finally, it is obvious that player 2 pays at most half of e_3 for any deviation from e_2 to e_3. She therefore has a strict equilibrium. It is only the player buying $e_{3,1}$ and the odd segments who can improve her payments. It is possible to pick the player in the beginning such that

$$1 + \epsilon(t) = \min\{\epsilon_1(|e_1|, t), \epsilon_1(|e_2|, t)\}$$

Type 2.3 With the discretization of e_3 described for Type 2.2, we lose a factor of $(1 + \epsilon(t))$ for this type, too. Now both players might deviate from both sides of e_3, however, the most profitable deviations again are the ones to the segments $e_{3,1}$ and $e_{3,|e_3|/t}$, which connect to the Steiner points. To see this, consider two deviation edges $d_a = (u_a, v_a)$ and $d_b = (u_b, v_b)$ with $u_a \in e_{11}$, $u_b \in e_{12}$, $v_a, v_b \in e_3$. If the cycles introduced by d_a and d_b intersect on e_3, changing the endpoints, i.e., edges $d = (u_a, v_b)$ and $d' = (u_b, v_a)$ lead to a superior, non-intersecting deviation. However, once the cycles do not intersect on e_3, we can decompose the situation and consider each side with e_{11} and e_{12} separately (see Figure 6). The rest of the proof then follows analogously to the previous type. Note that our bound for $\epsilon(t)$ can be adjusted to hold

$$1 + \epsilon(t) \leq \max\{\min\{\epsilon_1(|e_{11}|, t), \epsilon_1(|e_{21}|, t)\}, \min\{\epsilon_1(|e_{12}|, t), \epsilon_1(|e_{22}|, t)\}\}$$

if we assign the $e_{3,j}$ alternating and the segments $e_{3,1}, e_{3,|e_3|/t}$ to the player, which yields the smaller factor, respectively. Furthermore the proof of this type includes all previously considered network types, especially all types for which any of $|e_{11}|, |e_{12}|, |e_{21}|, |e_{22}| = 0$. Note that in these networks we might not have Steiner points and the inner angle between the segments could be greater than 120 °. Greater angles, however, make it less favorable to deviate from a payment scheme, which was already observed in the proof of Theorem 2 (Angle Assumption).

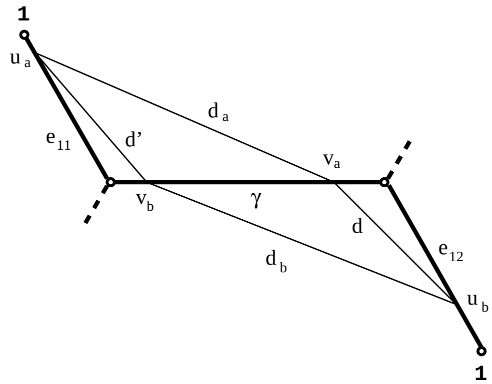

Figure 6: d and d' are better than d_a and d_b

Type 2.4 For this case we will use a slightly different discretization. We assume that e_3 is split into segments $e_{3,j}$ of length t. Now we do not assign the segments to a single player. We rather split each $e_{3,j}$ into two subsegments (possibly with different lengths) and assign each player to purchase one subsegment. Of course, one can further refine this discretization with a smaller t to obey the previously used assumption that each of the equally sized segments $e_{3,j}$ has to be paid completely by one player. This, however, would lead to negligibly perturbed results but involve much more technical details, therefore we refrained from doing so. Let the segments $e_{3,j}$ be defined as for the previous types. We start our enumeration at an arbitrary Steiner point and split each segment $e_{3,j}$ into two subsegments. The subsegment connected to the left (right) end of $e_{3,j}$ is called the left (right) share of $e_{3,j}$. Assume that player 1 buys every left share and player 2 every right share. The length of the shares is possibly different in each $e_{3,j}$ and specified by the value of the function $f(0,x)$, which was given in the proof of Theorem 2 for Nash equilibrium payments. With a similar calculation we can verify the same statement as in Lemma 11, however, now the corresponding point on segment $e_{3,1}$ might be in the part paid by player 2. Thus, the bound might be slightly lower, because player 1 is only willing to deviate to the end of the part of the first segment paid by him. We can also exclude intersecting devitaions shown in Figure 6. With the worst-case assumption that each player pays only for one segment on e_3, we get the following pessimistic upper bound for $\epsilon(t)$:

$$1 + \epsilon(t) \leq \max\{\epsilon_1(a, tb) \mid a \in \{|e_{11}|, |e_{12}|, |e_{21}|, |e_{22}|\},\ b \in \{1, f(0,t), 1 - f(0,t)\}\} - 1$$

This proves Theorem 3. □

A.3 Best response dynamics

Proof of Theorem 4. The theorem will be proven using best-response dynamics. The players repeatedly apply a local improvement step and switch to the cheapest possibility to connect their terminals. To prove that these dynamics arrive at a Nash equilibrium, we need the following lemma.

Lemma 12 *Suppose player 1 purchases the direct connection d_1 between her terminals and player 2 purchases straight segments s_i, $i \in \mathbb{N}$ that meet with d_1 with inner angles of more than 90°. Player 1 will not be able to decrease her payments below $|d_1|$ using any portion of any segment s_i.*

Proof. Directly from the Triangle Inequality (see Figure 7). Divide the plane into three regions based on d_1 (see Figure 7). The dashed lines are orthogonal to segment d_1. Assume for the moment that the lemma holds for segments that hit the middle region A_2. If a straight segment s of player 2 meets d_1 at a terminal, it must never cross to the A_2, because otherwise the inner angle would be less than 90°. As we assume the lemma holds for A_2, it suffices to consider only deviations that start at a point $u \in d_1$. However, from u player 1 needs less cost to reach her terminal than any other point on s. Hence, she will never be motivated to use

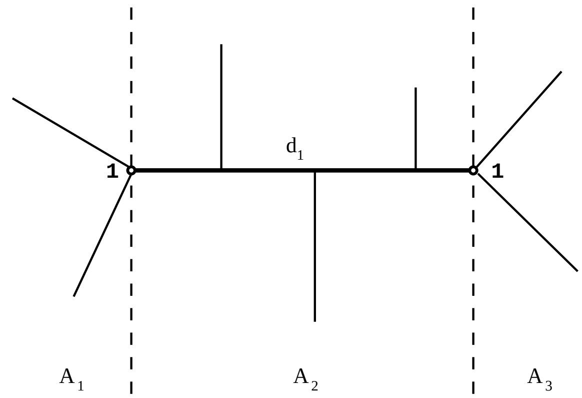

Figure 7: d_1 stays the cheapest connection.

some portion of s.
Turning to the middle region A_2 consider a segment s that meets at an inner point of d_1. It is obvious that s must be orthogonal to d_1 to provide an inner angle of 90 °. Suppose at first that there is only one segment s, and that player 1 uses some portion of it in a deviation. Then she has to pay the connection from two points on d_1 to two points on s. This is obviously more costly than the part on d_1 she can save. Assume now there are more segments s_i, $i \in \mathbb{N}$ connecting to d_1. Again the cheapest connections between any s_i and s_j as well as the cheapest connections between the terminals and any segment s_i are parallel to d_1. Therefore, d_1 provides the cheapest connection between the terminals of player 1 and the lemma is proven. □

In the first step player 1 chooses to pay for her direct connection d_1. For the second step let us at first consider the case that player 2 connects to d_1 and uses some portion of it. She will connect in the cheapest way, obeying the properties described in Lemma 3. Her connections will be straight segments connecting to d_1 with at least 90 ° inner angles when connecting to a terminal and exactly 90 ° angles when connecting to an interior point of d_1. Thus, Lemma 12 can be used to argue that a Nash equilibrium is achieved.
Otherwise, it could happen that player 2 picks her direct connection d_2 as well. Then player 1 could possibly improve her payments using the payments of player 2. If she can, her new connections will allow to argue that Lemma 12 holds for player 2 and an equilibrium is achieved. Otherwise we obviously have a Nash equilibrium, because no player can decrease her payments anymore.

Lemma 13 *The Nash equilibrium network T_E designed with the local improvement steps is at most a $\sqrt{2}$ approximation of the optimal centralized network T^*.*

Proof. Let the equilibrium network be denoted by T_E. The lemma states that for the ratio

$$r = \frac{|T_E|}{|T^*|} \leq \sqrt{2} \tag{20}$$

holds. As we have seen before, each T_E always contains at least one direct connection between terminals of a player. For the rest of the proof assume that the connection is present for player 2. Observe that $|T_E| \leq |d_1| + |d_2|$ holds, i.e., the cost of the direct deviations is always an upper bound for the cost of the Nash equilibrium. We will again consider all possible geometric connection games for 2 players and 2 terminals per player. If T^* consists of two components or two edges, which are exactly d_1 and d_2, the method finds an exact Nash equilibrium purchasing it. So suppose T^* is a connected tree, which has an edge e_3 present in both subtrees T^1 and T^2. We will consider the games as in the proof of Theorem 2 according to their optimum solution network as shown in Figures 2 and 3. In the following transformations of T^* and T_E are described that construct different networks. Each transformation will increase the bound on r. In the figures thin lines indicate T_E and thick lines T^*.

Type 5.1 Consider a game, in which the optimal network is of the type depicted in Figure 2a. It has no Steiner point, thus the angles between the segments are possibly greater than $120\,°$. T_E contains d_2. In the network in Figure 2a player 1 uses e_1 in T_E to connect her terminal. e_1 is also part of T^*. (Observe, that if e_1 is not part of T_E, then this transformation decreases $|T_E|$ by less than it decreases $|T^*|$. Thus, ratio r again can only increase.) Thus, the ratio can be increased by dropping e_1 from both T^* and T_E. The result of this is depicted in Figure 8. Remember Figure 8 does not depict the original game or network but a reduced

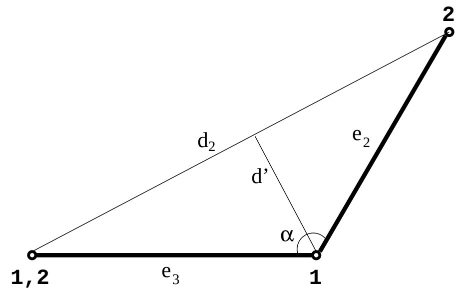

Figure 8: T^* and T_E after removing e_1.

construction to bound the ratio in a game with an optimal network of Figure 2a. Notice that with the Cosine Theorem

$$|d_2| = \sqrt{|e_2|^2 + |e_3|^2 - 2|e_2||e_3|\cos\alpha},$$

where $\alpha \in \left[\frac{2}{3}\pi, \pi\right]$ is the inner angle between e_2 and e_3. Let d' be a segment orthogonal to d_2 and connecting terminal 1 to d_2, see Figure 8. Using the Sine Theorem, the length of the small segment d' is given as

$$|d'| = \frac{|e_2||e_3|\sin\alpha}{\sqrt{|e_2|^2 + |e_3|^2 - 2|e_2||e_3|\cos\alpha}}.$$

Thus, the ratio is bounded by

$$r \leq \frac{|d_2| + |d'|}{|e_2| + |e_3|}.$$

How large can this bound be ? We assume the cost of T^* to be fixed to $|T^*| = |e_2| + |e_3| = 1$ and find a game, which distributes this cost to e_2 and e_3 as badly as possible. Let $x = |e_2|$ and $1 - x = |e_3|$, then

$$r = r(x, \alpha) \leq \frac{(\sin\alpha - 2\cos\alpha - 2)x(1-x) + 1}{\sqrt{2x(1+\cos\alpha)(x-1) + 1}}.$$

Maximizing this expression with the derivative for x results in

$$\frac{(\sin\alpha - 2\cos\alpha - 2)(2x(1+\cos\alpha)(x-1)+1)}{((\sin\alpha - 2\cos\alpha - 2)x(1-x)+1)(1+\cos\alpha)}(1-2x) = 2x - 1.$$

The square root has a real value only for $x = \frac{1}{2}$, which is a maximum point. For any given value of α the ratio is maximized by setting $|e_2| = |e_3|$. Then

$$r(0.5, \alpha) \leq \frac{\sin\alpha}{2\sqrt{2 - 2\cos\alpha}} + \frac{\sqrt{2 - 2\cos\alpha}}{2}.$$

In the interesting range for $\alpha \in [\frac{2}{3}\pi, \pi]$ the derivative of $r(0.5, \alpha)$ has only one extreme point at $\alpha = \arccos\left(-\frac{3}{5}\right) \approx 2.2143$, which corresponds to an angle of $126.87\,°$. The plot in Figure 9 shows that it is a maximum point. In any case the value of r is less than $1.12 < \sqrt{2}$. These steps to upper bound the ratio can be applied to the majority of games. The crucial point in the application is that both direct deviations d_1, d_2

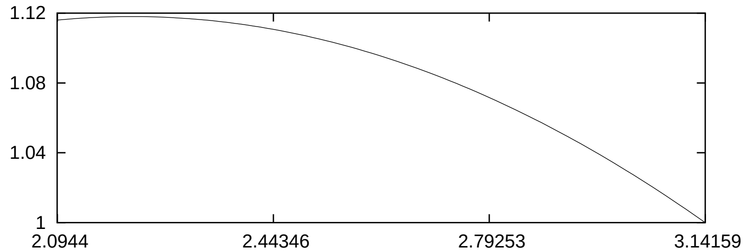

Figure 9: Behavior of the ratio $r(0.5, \alpha)$ for angles between $120\,^\circ$ and $180\,^\circ$.

must share a point with edge e_3. Then d_2, e_3 and a third edge, which is only in the subtree of player 2, will always form at least one such triangle depicted in Figure 8. Then some parts of the payments of player 1 can be upper bounded by the cost of the edges, which are only present in her subtree. This will always enable to transform the networks to arrive at a constellation depicted in Figure 8 and to upper bound r by 1.12.

To illustrate this fact and the possible transformations two additional types of games are analyzed. Notice that both types have an optimal solution T^* such that both d_1 and d_2 share a point with e_3.

Type 5.2 In this type games have optimal networks T^* including exactly one Steiner point and three edges (see Figure 2b). T_E consists of d_2 and the segment connecting the terminal of player 1. This segment can be upper bounded by e_1 and the shortest connection between d_2 and the Steiner point. Then using this upper bound for T_E the ratio can be further increased by dropping e_1. This leads directly to a network construction of Figure 8, thus in this case ratio r again is bounded by 1.12.

Type 5.3 Consider a game, which has an optimal solution of the type depicted in Figure 3a. It has two Steiner points and the terminals of the players are located on different sides of the line through e_3. Then player 1 connects directly to d_2 from her terminals. Suppose she buys e_{11} and e_{12} instead and connects to d_2 from the Steiner points. This will be more costly than the connections she actually chooses in T_E, hence one can upper bound T_E with this assumption. However, now in the cost of T^* and in the upper bound of $|T_E|$ the costs of the edges e_{11} and e_{12} are present. Hence, we can increase the upper bound by removing these costs from both expressions. This is equivalent to adjusting the network as depicted in Figure 10. This new network construction, however, can easily be identified to consist of two networks that were already considered in the previous types. Thus, the ratio r again is bounded by 1.12.

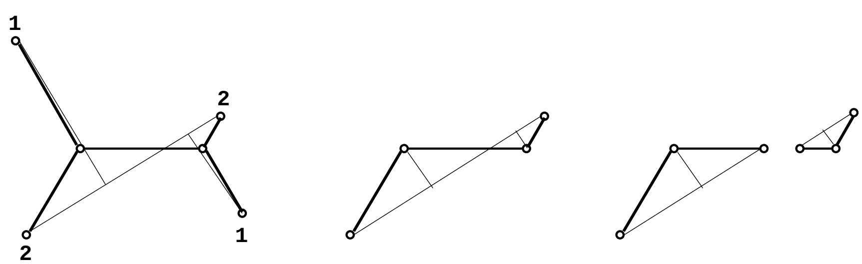

Figure 10: Transformations increasing the upper bound for r.

Type 5.4 Finally, consider the type of games, in which at least one of the deviations does not meet or cross e_3. Assume that this is deviation d_2–otherwise the previous reductions can be applied. The connection set of T^* that belongs only to T^2 consists of two edges e_{21} and e_{22}. Otherwise at least one terminal of player 2 would be located on e_3 and d_2 would hit e_3. Let T_E be such that player 1 chooses to connect to d_2. This of course implies that at least one terminal of player 1 is not located on e_3. Otherwise again the transformations described above can be used to bound the ratio by 1.12. So consider the case in which at least one terminal of player 1 is not located on e_3 (see Figure 3b for an example). Then these terminals connect to a Steiner point in

T^*. Here it is possible to transform similarly to the previous cases and bound the length of a direct connection from the terminal of 1 to d_2. These direct connections are in length upper bounded by the lengths of edges e_{11} or e_{12} between the terminal and the Steiner point in addition to the lengths of the direct connections d' and d'' between the Steiner point and d_2. Hence, it is again possible to remove e_{11} or e_{12} from T^* and the upper bound of T_E while to increase the upper bound for r. Apparently the ratio r in any game with a tree T_E of the described form (a direct deviation that does not cross e_3) can be upper bounded by a network construction depicted in Figure 11. The middle part of d_2 between d' and d'' is always shorter than e_3. Thus, removing

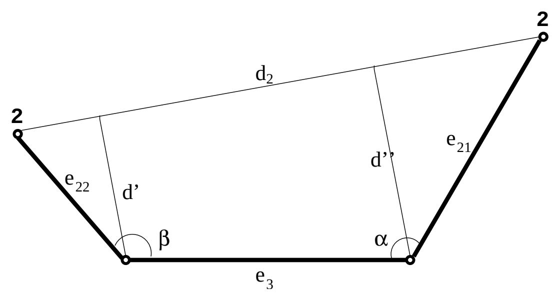

Figure 11: Network structure to bound the ratio when d_2 does not cross e_3.

these parts from consideration further increases the upper bound on r. Finally, two right-angled triangles are left. Hence, our ratio r is upper bounded by the sum of the lengths of the two legs over the length of the hypotenuse in a right-angled triangle, hence $r \leq \sqrt{2}$.

This concludes the proof of Lemma 13. □

With the proof of Lemma 13 the proof of Theorem 4 is also completed. □

Note that the analysis is far from being tight. Actually, $\sqrt{2}$ is only derived for the last type of games. Deriving a tight analysis remains as an open problem.