On Nash Equilibria for a Network Creation Game

Susanne Albers* Stefan Eilts† Eyal Even-Dar ‡ Yishay Mansour § Liam Roditty ¶

Abstract

We study a network creation game recently proposed by Fabrikant, Luthra, Maneva, Papadimitriou and Shenker. In this game, each player (vertex) can create links (edges) to other players at a cost of α per edge. The player's goal is to minimize the sum consisting of (a) the cost of the links he has created and (b) the sum of the distances to all other players.

Fabrikant et al. [10] conjectured that there exists a constant A such that, for any $\alpha > A$, all non-transient Nash equilibria graphs are trees. In this paper we disprove the tree conjecture. More precisely, we show that for any positive integer n_0, there exists a graph built by $n \geq n_0$ players which contains cycles and forms a non-transient Nash equilibrium, for any α with $1 < \alpha \leq \sqrt{n/2}$. Our construction makes use of some interesting results on finite affine planes. On the other hand we show that for $\alpha \geq 12n \log n$ every Nash equilibrium forms a tree.

The main result of Fabrikant et al. [10] is an upper bound on the price of anarchy of $O(\sqrt{\alpha})$ where $\alpha \in [2, n^2]$. We improve this bound for every α. Specifically, we derive a constant upper bound for $\alpha \leq \sqrt{n}$ and for $\alpha \geq 12n \log n$. For the intermediate values we derive an improved bound of $O(1 + (\min\{\frac{\alpha^2}{n}, \frac{n^2}{\alpha}\})^{1/3})$.

Additionally, we develop characterizations of Nash equilibria and extend our results to a weighted network creation game as well as to scenarios with cost sharing.

*Institut für Informatik, Albert-Ludwigs-Universität Freiburg, Georges-Köhler-Allee 79, 79110 Freiburg, Germany. salbers@informatik.uni-freiburg.de Work supported by the German - Israeli Foundation for Scientific Research & Development, project G-783-61.6/2003.

†Institut für Informatik, Albert-Ludwigs-Universität Freiburg, Georges-Köhler-Allee 79, 79110 Freiburg, Germany. eilts@informatik.uni-freiburg.de

‡School of Computer Science, Tel-Aviv university, email evend@post.tau.ac.il. Supported in part by the IST Programme of the European Community, under the PASCAL Network of Excellence, IST-2002-506778

§School of Computer Science, Tel-Aviv university, email mansour@post.tau.ac.il. Supported in part by the IST Programme of the European Community, under the PASCAL Network of Excellence, IST-2002-506778, by a grant from the Israel Science Foundation and an IBM faculty award.

¶School of Computer Science, Tel-Aviv university, email liamr@post.tau.ac.il

1 Introduction

Network design is a fundamental problem in computer science and operations research. This line of research assumes a central authority that constructs the network and has various optimization criteria to fulfill. In practice, however, many networks are actually formed by selfish players who are motivated by their own interests and their own objective function. For instance, the Internet, networks for exchanging goods and social networks are all formed by many players and not by a single authority. This motivates the research of network creation by multiple selfish players.

In this work we focus on the later model and allow individual users to decide which edges to buy. The appropriate concept for studying such a scenario is that of Nash equilibria [18], where no user has the incentive to deviate from his strategy. We analyze the performance of the resulting network architectures using the *price of anarchy*, introduced by Koutsoupias and Papadimitriou in their seminal paper [17]. Recently, Nash equilibria and their associated price of anarchy have been studied for a wide range of classical computer problems such as job scheduling, routing, facility location and, last but not least, network design and creation, see e.g. [1, 2, 3, 7, 6, 8, 11, 10, 13, 15, 17, 20]. This also includes variants of the price of anarchy, called the price of stability [1, 2, 6].

In this paper we study a network creation game introduced by Fabrikant, Luthra, Maneva, Papadimitriou and Shenker [10]. The game is defined as follows, there are n players, each of which is associated with a separate network vertex. These players have to build a connected, undirected graph. Each player may lay down edges to other players. Once the edges are installed, they are regarded as undirected and may be used in both directions. The resulting network is the set of players (vertices) and the union of all edges laid out. The cost of each player consists of two components. Firstly, a player pays an edge building cost equal to α times the number of edges laid out by him, for some $\alpha > 0$. Secondly, the player incurs a connection cost equal to the sum of the shortest path distances to other players. This game models scenarios in which peers wish to communicate and transfer data. Each peer incurs a hardware cost and pays for the communication delays to other players.

Formally, we represent the set of players by a vertex set $V = \{1, \ldots, n\}$. A *strategy*, for a player $v \in V$, is a set of vertices $S_v \subseteq V \setminus \{v\}$ such that v creates an edge to every $w \in S_v$. Given a combination of strategies $\vec{S} = (S_1, \ldots, S_n)$, the resulting graph $G(\vec{S}) = (V, E)$ consists of the edge set $E = \bigcup_{v \in V} \bigcup_{w \in S_v} \{v, w\}$. In our analysis it will sometimes be convenient to assume that the edges have a direction. A directed edge (v, w) indicates that the player v built an edge to w. The cost of a player v under $\vec{S}$ is $Cost(v, \vec{S}) = \alpha |S_v| + \sum_{w \in V, w \neq v} \delta(v, w)$, where $\delta(v, w)$ is the length of the shortest path between v and w in $G(\vec{S})$.

A combination of strategies $\vec{S}$ forms a Nash equilibrium if, for any player $v \in V$ and every other combination of strategies $\vec{U}$ that differ from $\vec{S}$ only in v's component, $Cost(v, \vec{S}) \leq Cost(v, \vec{U})$. The induced graph $G(\vec{S})$ is called the equilibrium graph. $\vec{S}$ is a *strong* Nash equilibrium if, for every player v, strict inequality $Cost(v, \vec{S}) < Cost(v, \vec{U})$ holds. Otherwise, it is a *weak* Nash equilibrium. In a weak Nash equilibrium at least one player can change its strategy without affecting its cost. We will also use the notion of *transient* Nash equilibria [10]. A transient Nash equilibrium is a weak equilibrium from which there exists a sequence of single-player strategy changes, which do not change the deviator's cost, leading to a non-equilibrium position.

For a combination of strategies $\vec{S}$, let $Cost(\vec{S}) = \sum_{v \in V} Cost(v, \vec{S})$ be the total cost of all players. Let $Cost(OPT)$ be the cost of the social optimum that achieves the smallest possible value. The price of anarchy is the worst-case ratio $Cost(\vec{S})/Cost(OPT)$, taken over all Nash equilibria $\vec{S}$.

Previous work: Fabrikant et al. [10] main interest was to analyze the price of anarchy of the game. They easily observe that, for $\alpha < 2$ and $\alpha > n^2$, it is constant. Their main contribution is an upper bound of $O(\sqrt{\alpha})$ for $\alpha \in [2, n^2]$. This upper bound can be as large as $O(n)$ when $\alpha = n^2$. Fabrikant et al. pointed out that in their constructions as well as in experiments that they preformed only tree Nash equilibria were found. The only exception was the Petersen graph that represents a transient Nash equilibrium. This fact motivated them to formulate a *tree conjecture* stating that there exists a constant A such that, for any $\alpha > A$, all non-transient Nash equilibria are trees. In other words, every Nash equilibrium that has a cycle in the underlying graph is transient and, in particular, weak. Finally, they proved that the price of anarchy is constant for a tree Nash equilibrium.

In a recent work Corbo and Parkes [5] study the price of anarchy in the model introduced by Fabrikant et al. with a single variation that the edges are not bought by a single player but by both players at the end points of the edge.

There exists a large body of previous work on other network design problems. Anshelevich et al. [1] investigate a network design problem where players, in a given graph, have to connect desired terminal pairs. They analyze the quality of the best Nash equilibrium under Shapley cost sharing. Anshelevich et al. [2] consider connection games where each player has to connect a set of terminals and present algorithms for computing approximate Nash equilibria. Further work on cost sharing in network design includes [12, 15, 19, 16]. Bala and Goyal [3] study a network formation problem in which players incur cost but also benefit from building edges to other players. They trade off the costs of forming links against the potential reward from doing so. Haller and Sarangi [13] build on this work and allow player heterogeneity.

Social and economic networks in which each player is a different vertex in the graph play a major role in the economic literature. For a recent and detailed review of social and economics models see [14].

Our contribution: In this paper we first show that the tree conjecture is incorrect. We prove that, for any positive integer n_0, there exists a graph built by $n \geq n_0$ players that contains cycles and forms a strong Nash equilibrium, for any α with $1 < \alpha \leq \sqrt{n/2}$. The graphs we construct are *geodetic*, i.e. the shortest path between any two vertices is unique, and have a diameter of 2. These properties are crucial in showing that the Nash equilibrium is indeed strong. If a player deviates from its original strategy and builds less edges or edges to different players, then — since the original graph was geodetic — the shortest path distance cost increases substantially. If a player decides to build more edges, then — since the graph has diameter 2 — the cost saving is negligible. Our construction resorts to some concepts from graph theory and geometry. In particular, we use results on finite affine planes. To the best of our knowledge, these concepts have never been used in game theoretic investigations and might be helpful when studying other graph oriented games.

We proceed and give improved upper bounds on the price of anarchy. Our main result here is a **constant** upper bound on the price of anarchy for both $\alpha \leq \sqrt{n}$ and $\alpha \geq 12n \log n$. We prove that if $\alpha \geq 12n \log n$, the price of anarchy is in fact not larger than 1.5 and goes to 1 as α increases. Interestingly, the proof shows that if $\alpha \geq 12n \log n$, any Nash equilibrium is indeed a tree. For any α, we prove an upper bound of $O(1+(\min\{\frac{\alpha^2}{n}, \frac{n^2}{\alpha}\})^{1/3})$. Thus, if $\alpha \in O(\sqrt{n})$, the price of anarchy is again constant. For $\alpha \in [\sqrt{n}, n]$ the value increases, reaching a maximum of $O(n^{1/3})$ at $\alpha = n$. For $\alpha > n$, the price of anarchy is decreasing. Hence, we have constant prices of anarchy for large ranges of α and a worst case bound of $O(n^{1/3})$ instead of $O(n)$.

Furthermore, we analyze the structure of Nash equilibria, investigating solutions with short induced cycles. We prove that any Nash equilibrium that forms a chordal graph having induced cycles of length three is indeed transient. We show that such equilibria do exist for all n. Furthermore, we show that if $\alpha < n/2$, then the only tree that forms an equilibrium is the star and that there exists Nash equilibria graphs of n vertices which are not trees.

Additionally, we study a weighted network creation game in which player v wishes to send a certain amount of traffic to player u, for any v and u. In the cost of player v, the shortest path distance to u is multiplied by this traffic amount. We also provide an upper bound on the price of anarchy. For a uniform traffic matrix, we obtain for the weighted game the same bounds as our bounds for the unweighted game.

Finally, we consider settings with cost sharing where players can pay for a fraction of an edge. The edge exists if the total contribution by all players is at least α. We show that in both the unweighted and weighted games part of our upper bounds on the price of anarchy carry over. We also prove that there exist strong

Nash equilibria with cycles in which the cost is split evenly among players.

2 Disproving the tree conjecture

We will present a family of graphs that form strong Nash equilibria and have induced cycles of length three and five. To construct these graphs, we have to define affine planes, see e.g. Mac Williams and Sloane [21].

Definition 1 *An affine plane is a pair $(A, \mathcal{L})$, where A is a set (of points) and $\mathcal{L}$ is a family of subsets of A (of lines) satisfying the following four conditions.*

- *For any two points,there is a unique line containing these points.*
- *Each line contains at least two points.*
- *Given a point x and a line L that does not contain x, there is a unique line L' that contains x and is disjoint from L.*
- *There exists a triangle, i.e. there are three distinct points which do not lie on a line.*

If A is finite, then the affine plane is called finite.

Two lines are *parallel*, in signs $\|$, if the lines are disjoint or if they are equal. Given a point x and a line L, we denote by $(x\|L)$ the unique line that is parallel to L and contains x. Parallelism defines an equivalence relation on the lines, and the equivalence class of L is denoted by $[L]$.

If q is a prime power, then for the field $F = GF(q)$ the sets $A = F^2$ and $\mathcal{L} = \{a+bF \mid a, b \in A, b \neq 0\}$ are an affine plane of order q, denoted by $AG(2, q)$. The plane contains q^2 points and $\binom{q^2}{2}/\binom{q}{2} = q(q+1)$ lines. There are $q+1$ equivalence classes ($q-1$ real slopes, horizontal and vertical lines). Each class has q lines and each such line contains q points.

We are now ready to describe the graphs representing strong Nash equilibria. The graphs were also constructed by Blokhuis and Brouwer [4] as instances of geodetic graphs. For an affine plane $AG(2, q)$ we define a graph $G = (V, E)$ with $V = A \cup \mathcal{L}$. In the following, when we refer to a point or a line, we often mean the corresponding vertex or player. The edge set E is specified as follows.

- A point and a line are connected by an edge if and only if the line contains the point.
- Two lines are connected by an edge if and only if they are parallel.
- No two points are connected by an edge.

There are no self-loops or multiple copies of an edge. We have to give orientations to these edges. Every equivalence class of a line L defines a complete subgraph K_q of G. Let $r(L)$ and $s(L)$ denote the indegree and outdegree of L in K_q, respectively. One can easily show by induction that there exists an orientation of the edges of K_q such that, for every line L in K_q, $|r(L) - s(L)| = 0$ if q is odd and $|r(L) - s(L)| = 1$ if r is even. In order to define an orientation for the edges between points and lines, we choose a representative line L^i, $0 \leq i \leq q$, for each of the $q+1$ equivalence classes. The lines of $[L^q] = \{L_0^q, \ldots, L_{q-1}^q\}$ do not build edges to their points; rather the existing edges are built by the points. As for the other equivalence classes, a line $L \in [L^i]$, $0 \leq i \leq q-1$, builds edges to the two points $L \cap L_i^q$ and $L \cap L_{i+1 (\mathrm{mod}\ q)}^q$. All the other edges are built by the points. Every point x is contained in a line $(x \parallel L^q) =: L_j^q$ and has exactly two incoming edges from the lines $(x \parallel L^j)$ and $(x \parallel L^{j-1 (\mathrm{mod}\ q)})$. For $q = 2$, we obtain the Petersen graph.

Figure 1 shows the graph structure relative to a line $L \notin [L^q]$. Let $x_1, \ldots, x_q$ be the q points contained in L. We number these points such that L builds edges to x_1 and x_2. Let $L_1, \ldots, L_{q-1}$ be the $q-1$ lines parallel to L. We number these lines such that the first $r = r(L)$ lines build edges to L while L builds edges to the remaining $q-1-r$ lines. For any point x_i, $1 \leq i \leq q$, we denote by $L_1^{x_i}, \ldots, L_q^{x_i}$ the other q lines that contain x_i. These sets of q lines are disjoint for different x_i since for every pair of points there is a unique line containing this pair. Furthermore these lines are different from $L_1, \ldots, L_{q-1}$. For any line L_i,

$1 \le i \le q-1$, let $x_1^i, \dots x_q^i$ be the q points contained in L_i. Again these point sets are disjoint for different L_i and are also different from $x_1, \dots, x_q$ since the lines L and $L_1, \dots, L_{q-1}$ are parallel. If $L \in [L^q]$, then the structure of the graph is the same except that the edges between L and its points are all built by the points. If $L \notin [L^q]$ then the cost of the player representing L is $(2+s)\alpha+(2q-1)+2(2q-1)q = (s+2)\alpha+4q^2-1$, where $s = s(L) = q-1-r$. If $L \in [L^q]$, then the cost is $s\alpha + 4q^2 - 1$.

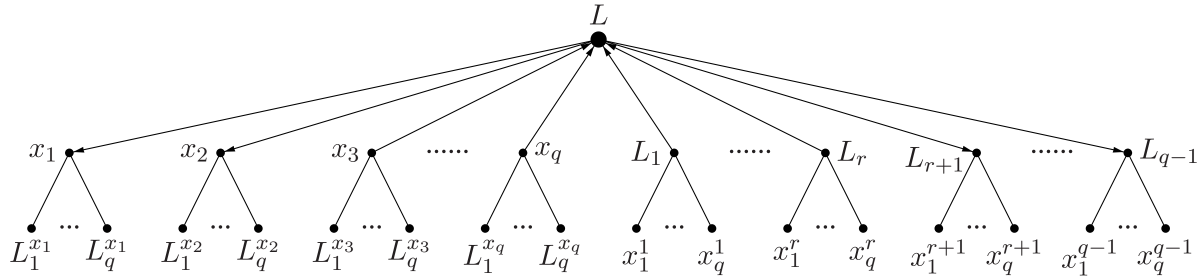

Figure 1: The distances with respect to a line L.

Figure 2 depicts the graph structure relative to a point x. Lines $L_1^x, \dots, L_{q+1}^x$ are the $q+1$ lines containing x. For a line L_i^x, $1 \le i \le q+1$, let $x_1^i, \dots, x_{q-1}^i$ be the other $q-1$ points of L_i^x and let $L_1^i, \dots, L_{q-1}^i$ be the $q-1$ lines parallel to L_i^x. These sets of $q-1$ points and lines are disjoint for different i. Thus the cost of the player representing x is $(q-1)\alpha + (q+1) + 2(q+1)(2(q-1)) = (q-1)\alpha + 4q^2 + q - 3$.

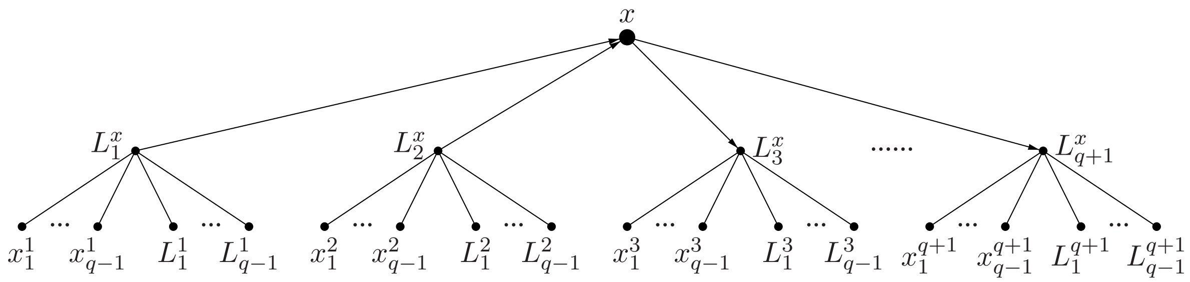

Figure 2: The distances with respect to a point x.

Lemma 1 *Let $q > 10$. For α in the range $1 < \alpha < q+1$, no player associated with a line L has a different strategy that achieves a cost equal to or smaller than that of L's original one. For α in the range $1 \le \alpha \le q+1$, L has no strategy with a smaller cost.*

Proof. We prove the lemma for a line $L \notin [L^q]$, which builds two edges to points. This implies that the lemma also holds for lines $L' \in [L^q]$ which do not build edges to points. For, if a line $L' \in [L^q]$ had a different strategy with the same or a smaller cost, then any line $L \notin [L^q]$ could adopt the same strategy change while maintaining the two edges built to points. This would result in the same or a smaller cost, respectively. As we will show in the following, this is impossible.

Fix a line $L \notin [L^q]$. We consider all possible strategy changes. First, if L builds $l > s+2$ edges, then at best there are $l - s - 2 + 2q - 1$ vertices at distance 1 while the other vertices are at distance 2 from L. In L's original strategy there are $2q-1$ vertices at distance 1 while all other vertices are at distance 2. Thus, L's original strategy has a cost which is at least $\alpha(l-s-2) - (l-s-2)$ smaller than that of S, and this expression is strictly positive for $\alpha > 1$. Thus buying more than $s+2$ edges does not pay off.

In the remainder of this proof we study the case that L builds at most $s+2$ edges and start with the strategy S_0 in which L does not build any edges at all. The resulting shortest path tree of L is given

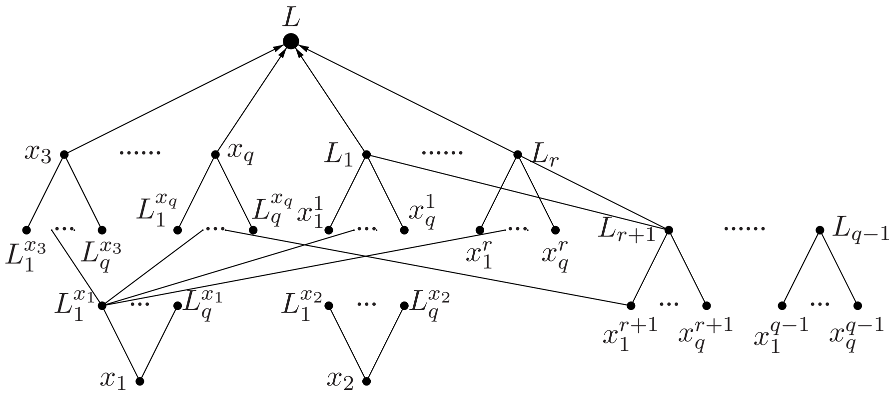

Figure 3: Strategy change S_0.

in Figure 3. Lines $L_{r+1}, \ldots, L_{q-1}$ are a distance of 2 away from L since these lines are connected to $L_1, \ldots, L_r$. Lines $L_i^{x_1}$ and $L_i^{x_2}$, $1 \leq i \leq q$, are a distance of 3 away from L because they do not contain $x_3, \ldots, x_q$ and are not parallel to $L_1, \ldots, L_r$ but are connected to one line from $L_1^{x_j}, \ldots, L_q^{x_j}$, for any j with $3 \leq j \leq q$, and are also connected to one point from $x_1^j, \ldots x_q^j$, for any j with $1 \leq j \leq r$. Points $x_1^i, \ldots, x_r^i$, with $r+1 \leq i \leq q-1$, are a distance of 3 away because they are not contained in $L_1, \ldots, L_r$ but are connected to one line from $L_1^{x_j}, \ldots, L_q^{x_j}$, for any $3 \leq j \leq q$. Finally points x_1 and x_2 are a distance of 4 away from L because these points are only contained in lines $L_1^{x_1}, \ldots, L_q^{x_1}$ and $L_1^{x_2}, \ldots, L_q^{x_2}$, respectively, at distance 3. The cost difference between S_0 and L's original strategy is $-(s+2)\alpha + s(q+1) + 2q + 6 = (q+1-\alpha)(s+2) + 4 > 0$ and hence S_0 is a worse strategy.

Next suppose that L does build edges. The edges can be of six different types: L builds an edge to (a) a line $L_j^{x_i}$ for some $3 \leq i \leq q$ and $1 \leq j \leq q$; (b) a point x_j^i, for some $1 \leq i \leq r$ and $1 \leq j \leq q$; (c) an edge $L_j^{x_1}$ or $L_j^{x_2}$, for some $1 \leq j \leq q$; (d) a point x_j^i, for some $r+1 \leq i \leq q-1$ and $1 \leq j \leq q$; (e) a line L_i, for some $r+1 \leq i \leq q-1$; (f) a point x_1 or x_2. In the following we investigate all of these cases, which are also depicted in Figure 4.

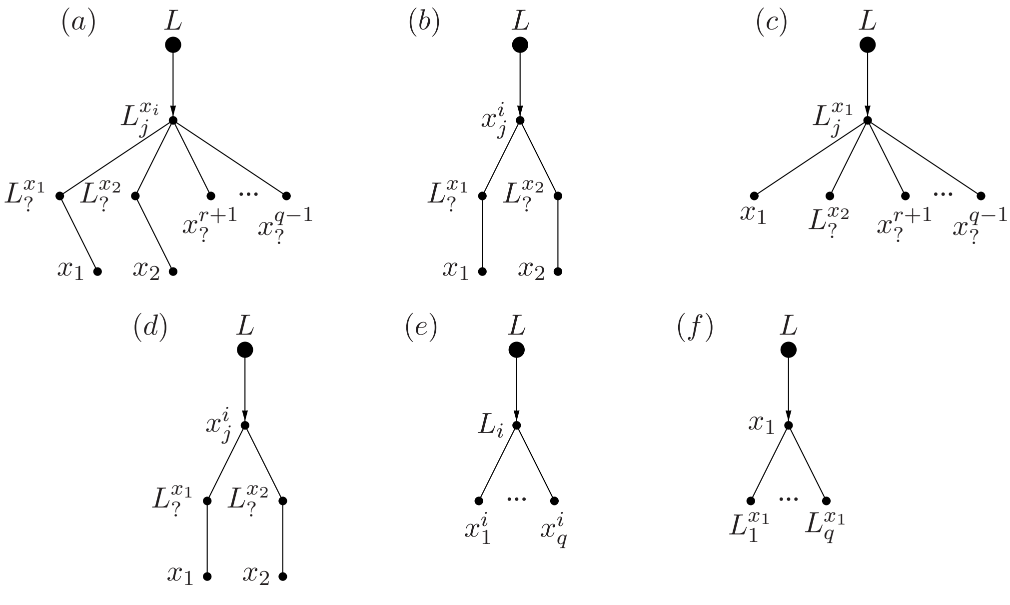

Figure 4: The effect of edges of types (a – f).

Case (a): The line $L_j^{x_i}$ is connected to one line from $L_1^{x_1}, \ldots, L_q^{x_1}$, which is linked to x_1, and to one line from $L_1^{x_2}, \ldots, L_q^{x_2}$, which is linked to x_2. Additionally $L_j^{x_i}$ is connected to one point from $x_1^k, \ldots, x_q^k$, for any $r+1 \leq k \leq q-1$. Thus, setting a link to $L_j^{x_i}$, line L can save a cost of at most $s+5$ relative to S_0. Hence L can save a cost of at most $s+5$ no matter how other links are laid out by L. In other words, removing the edge to $L_j^{x_i}$ results in an increase in the shortest path distance cost of at most $s+5$.

Case (b): Point x_j^i is connected to one line from $L_1^{x_1}, \ldots, L_q^{x_1}$ and to one line from $L_1^{x_2}, \ldots, L_q^{x_2}$. From there x_1 and x_2 can be reached. By laying out an edge to x_j^i, line L saves a shortest path distance cost of 5 relative to S_0 and hence a value of at most 5 relative to any other strategy. Again, removing this link can increase the shortest path distance cost by at most 5.

Case (c): Assume w.l.o.g. that an edge to $L_j^{x_1}$ is built. The analysis of a link to $L_j^{x_2}$ is similar. Line $L_j^{x_1}$ is linked to x_1 and to one line from $L_1^{x_2}, \ldots, L_q^{x_2}$. Furthermore $L_j^{x_1}$ is linked to one point from $x_1^i, \ldots, x_q^i$, for any $r+1 \leq i \leq q-1$. Relative to S_0 the shortest path distances decrease by $s+5$. Removing the edge results in an increase of at most $s+5$.

Case (d): Point x_j^i is connected to one line from $L_1^{x_1}, \ldots, L_q^{x_1}$ and to one line from $L_1^{x_2}, \ldots, L_q^{x_2}$. From there x_1 and x_2 can be reached. Building an edge to x_j^i saves a shortest path distance cost of 6 relative to S_0. Not building this edge results in an increase of at most 6.

The last two cases are studied under the condition that the other edges built by L are also of type (e) or (f).

Case (e): If L builds only edges of type (e) and (f), then points $x_1^i, \ldots, x_q^i$ are still at distance 3 and by setting a link to L_i the shortest path distance cost reduces by $q+1$.

Case (f): Again, assume that L builds only edges of type (e) and (f). Without an edge to x_1, lines $L_1^{x_1}, \ldots, L_q^{x_1}$ are a distance of 3 away from L and x_1 is a distance of 4 away. Building an edge to x_1 reduces the shortest path distance cost by $q+3$.

With the above case distinction (a–f) we are able to finish the proof. Recall that L builds at most $s+2$ edges. If S contains edges of types (a–d), then we simultaneously replace all of these edges by edges of type (e) or (f). Any such edge replacement increases the shortest path distance cost by at most 6 or $s+5$ while the decrease is at least $q+1$. Since, for $q > 10$, we have $q+1 > q/2+6 \geq s+5 \geq 6$, strategy S is worse than L's strategy defined by graph G. So suppose that S only builds edges of types (e) or (f). If S builds less than $s+2$ edges, then we introduce additional edges of types (e) or (f) until a total of $s+2$ edges are laid out. For any additional edge, there is an edge building cost of α while the shortest path distance cost decreases by at least $q+1$. If $\alpha < q+1$, there is a net cost saving and S is worse than L's original strategy given by G. If $\alpha = q+1$, then L's original strategy is at least as good. □ □

Lemma 2 *For α in the range $1 < \alpha \leq q+1$, no player associated with a point x has a different strategy that achieves a cost equal to or smaller than that of x's original strategy. For $\alpha = 1$, no player associated with a point has a strategy that achieves a smaller cost.*

The proof is given in Appendix A. The above two lemmata yield the main result of this section.

Theorem 1 *Let $q > 10$. The graph G is a strong Nash equilibrium, for $1 < \alpha < q+1$, and a Nash equilibrium, for $1 \leq \alpha \leq q+1$.*

3 Improved bounds for the price of anarchy

We first consider the case that $\alpha \geq 12n \log n$, proving a constant price of anarchy. Then we address the remaining range of α. In both cases, for a given equilibrium graph $G(\vec{S})$, we need the concept of a shortest path tree rooted at a certain vertex u. The root of $T(u)$ is vertex u and this vertex represents *layer* 0 of the tree. Given vertex layers 0 to $i-1$, layer i is constructed as follows. A node w belongs to layer i if it is not

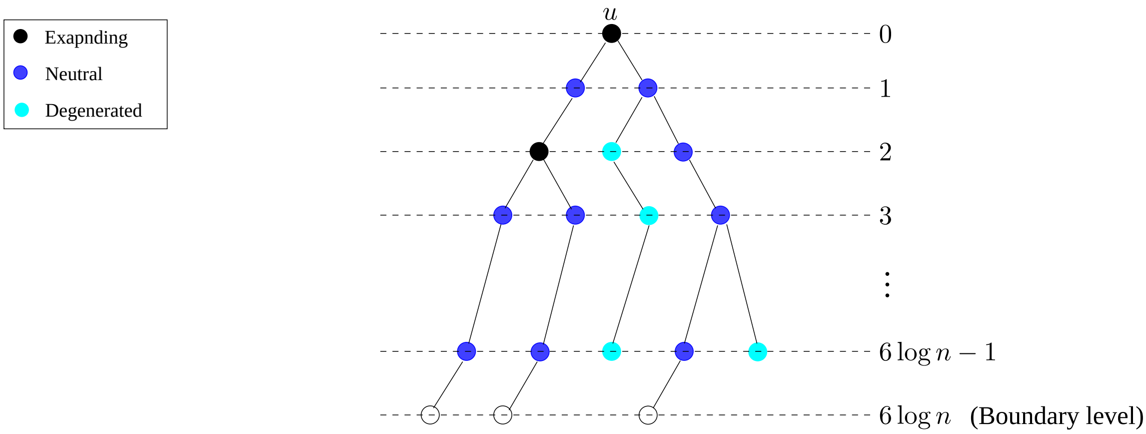

Figure 5: A classification of the vertices of T(u).

yet contained in layers 0 to $i-1$ and there is a vertex v in layer $i-1$ such that there is an edge connecting v and w, i.e. $\{v, w\} \in E$. We add this edge to the shortest path tree. We emphasize that if w is linked to several vertices of layer $i-1$, only one such edge is added to the tree at this point. Suppose that all vertices of V have been added to $T(u)$ in this fashion. The edges inserted to far are referred to as *tree edges*. We now add all remaining edges of E to $T(u)$ and refer to these edges as *non-tree edges*. Essentially, $T(u)$ is just a layered version of G with distinguished tree edges.

3.1 Constant price of anarchy for $\alpha \geq 12n \log n$

In order to establish a constant price of anarchy, we prove that if $\alpha \geq 12n \log n$, then every Nash equilibrium graph is a tree. This implies an upper bound of 5 on the price of anarchy [10]. However, we here give an improved upper bound of 1.5 for the considered range of α.

Our proof has the following structure. Given an equilibrium graph whose girth (i.e., the length of the minimal cycle in the graph) is at least $12 \log n$, we prove that the graph diameter is bounded by $6 \log n$. The proof is by contradiction. We assume that there exists a vertex u with eccentricity at least $6 \log n$ and examine its shortest path tree $T(u)$. We show that the maximal depth of $T(u)$ is less than $6 \log n$. This immediately implies that the equilibrium graph is a tree, given the bound on the girth. Also, since we have chosen an arbitrary vertex this implies that the diameter is at most $6 \log n$. We complete the proof by showing that for high edge costs the graph has a high girth.

We classify the vertices of the equilibrium graph according to their location in the tree $T(u)$. We refer to the vertices at depth exactly $6 \log n$ as vertices in the *Boundary* level. We classify the vertices in the levels before the Boundary level according to the number of descendent their children have in the Boundary level. We have three types of vertices. The first are *Expanding vertices* which lead to an exponential growth, the second, and the most problematic, are *Neutral vertices* that do not lead to a growth but have descendants in the Boundary level, and the third are *Degenerate vertices* that have no descendants in the Boundary level. The vertices of the Boundary level, and at levels of larger depth, are unclassified. We now give the formal definition.

Definition 2 *Let $G(\vec{S})$ be an equilibrium graph and let $u \in V$. Let $T(u)$ be a shortest path tree rooted at u. We say that a vertex $v \in V$, at a depth smaller than $6 \log n$ in $T(u)$, is:*

- **Expanding** *- If v has at least two children with at least one descendent in the Boundary level.*

- **Neutral** - *If v has exactly one child with at least one descendent in the Boundary level.*
- **Degenerate** - *If v does not have any descendent in the Boundary level.*

An example to this classification is given in Figure 5. Note that vertices at level $6 \log n$ (the Boundary level) and higher levels are not classified. Our target is to show that there are n vertices in the Boundary level. This implies that there are no vertices in levels higher than $6 \log n$. It is important to note that since the graph has girth at least $12 \log n$, there is a unique tree $T(u)$ up to level $6 \log n$ (the Boundary level).

In the next Lemma we show that Degenerate children of a Neutral vertex v and their descendants are connected only through v to vertices out of the subtree of v in $T(u)$.

Lemma 3 *Let $G(\vec{S})$ be an equilibrium graph whose girth at least $12 \log n$. Let v be a Neutral vertex in $T(u)$ and let $D_u(v)$ be the set of its Degenerate children and their descendants at $T(u)$. Every path from $x \in D_u(v)$ to $y \in V \setminus D_u(v)$ in $G(\vec{S})$ must go through v.*

Proof. We show that any path from x to y must go through v. Suppose that there is a path that does not go through v then either it goes through a vertex z from the Boundary level or the entire path does not cross the Boundary level. However, x is Degenerate and wlog z is its descendant and can not be in the Boundary level since it violates the definition Degenerate vertex. Thus, it must be that $\delta(u, z) < 6 \log n$. Now if every vertex z on the path from x to y satisfies that $\delta(u, z) < 6 \log n$ then there is a cycle of length less than $12 \log n$. We conclude that any path from x to y must go through v. □

The above Lemma shows that Neutral vertices have a crucial role in connecting Degenerate vertices. The next Lemma will use this property to show that although many Neutral vertices can be found in the tree, the number of times that two Neutral vertices can appear consecutively on a path from u is limited.

Lemma 4 *Let $G(\vec{S})$ be an equilibrium graph whose girth is at least $12 \log n$. Let $u = w_0, w_1, \ldots, w_l = v$ be a shortest path from u to v. An edge on the path is said to be a Neutral edge if both of its endpoints are Neutral vertices. The total number of Neutral edges is $2 \log n$.*

Proof. Let (w_{i-1}, w_i) be a Neutral edge on the path from u to v. There are two possible types of Neutral edges. Edges which are bought by their tail (i.e. w_{i-1}) or edges which are bought by their head (i.e. w_i). We assume w.l.o.g that the number of edges which are bought by their tail is larger than the number of edges which are bought by their head. We bound the total number of such Neutral edges with $\log n$. This gives the desired bound of $2 \log n$.

Let $(w_{i_1-1}, w_{i_1}), (w_{i_2-1}, w_{i_2}), \ldots, (w_{i_m-1}, w_{i_m})$ be the Neutral edges on the path which are bought by their tail. We show that $m \le \log n$. Let $D_u(w_{i_j})$ be the set of all the Degenerate children of w_{i_j} and their descendants. By Lemma 3 every path from a vertex in $V \setminus D_u(w_{i_j})$ to a vertex in $D_u(w_{i_j})$ goes through w_{i_j}. Let n_j denote the size of $D_u(w_{i_j})$. Now since we are in equilibrium the benefit of w_{i_j-1} from buying the edge (w_{i_j-1}, w_{i_j}) is larger than the benefit from buying the edge (w_{i_j-1}, w_{i_j+1}). Thus, $n_j \ge \sum_{k=j+1}^{m} n_k$. As a result $n_j \ge 2^{m-j-1}$ and m is bounded by $\log n$. □

Based on the above Lemma we prove the main result of this section. We show that every equilibrium graph whose girth is at least $12 \log n$ must be a tree whose maximal depth is $6 \log n$.

Proposition 1 *If $G(\vec{S})$ is an equilibrium graph whose girth is at least $12 \log n$ then the diameter of $G(\vec{S})$ is at most $6 \log n$ and $G(\vec{S})$ is a tree.*

Proof. For the sake of contradiction, we start by assuming that the diameter is at least $6\log n$. Let $u \in V$ be a vertex on one of the endpoints of the diameter. We look on a shortest path tree rooted at u. Since u is one of the diameter endpoints our assumption implies that u is either Neutral or Expanding vertex. We show that the number of descendants at the Boundary level (i.e. vertices at a depth of **exactly** $6\log n$) is at least n. As it is not possible to have n vertices in the Boundary level we reach to a contradiction. This obviously implies that the maximal depth is at most $6\log n$ and that there are no cycles. Let $v \in V$, we denote with d the depth of v in $T(u)$ and with b the number of Neutral edges on the path from u to v. We label a vertex by (d,b). For example, the label for the root u is $(0,0)$ because $d=0$ and $b=0$. Let v be a non-Degenerate vertex whose label is (d,b), and let $N(d,b)$ be a lower bound on the number of its descendants at the Boundary level. (Note that two vertices might have the same label, but have different number of descendants at the boundary level.) We claim that $N(d,b) \geq 2^{\frac{6\log n-d}{2}-(2\log n-b)}$. This implies for the root that $N(0,0) \geq 2^{\frac{6\log n-0}{2}-(2\log n-0)} = n$, thus proving the claim will lead to the desired contradiction.

The proof will be by a backwards induction on d and b. As for the induction basis we show that $N(6\log n, b) \geq 2^{-(2\log n-b)}$ and $N(d, 2\log n) \geq 2^{\frac{6\log n-d}{2}}$. We first show that $N(6\log n, b) \geq 2^{-(2\log n-b)}$. The only descendent at the Boundary level is the vertex itself and $N(6\log n, b) = 1$. Thus, we need to show that $2^{-(2\log n-b)} \leq 1$. This follows directly from Lemma 4 since $b \leq 2\log n$. Next, we prove that $N(d, 2\log n) \geq 2^{\frac{6\log n-d}{2}}$. The proof here is a bit more subtle and a secondary induction on d is needed. The basis for the secondary induction, $N(6\log n, 2\log n) \geq 1$, trivially holds. We assume that $N(d', 2\log n) \geq 2^{\frac{6\log n-d'}{2}}$ for every $d' > d$ and prove it for d. Let v be a vertex at depth d with $b = 2\log n$ which may be either Expanding or Neutral. We show that in either case v has at least two descendants at depth $d+2$ which are either Expanding or Neutral. For the case that v is Expanding it follows from the definition of Expanding vertex that v has at least two descendants at depth $d+2$ which are either Expanding or Neutral. For the case that v is Neutral it follows that v cannot have a Neutral child since $b = 2\log n$ and there are at most $2\log n$ Neutral edges by Lemma 4. Thus, v must have an Expanding child which again has by definition at least two children which are either Expanding or Neutral. We conclude that in both cases, i.e. v is Expanding or Neutral, it has at least two descendants at depth $d+2$ which are either Expanding or Neutral. The induction hypothesis holds for these descendants of v and we get that:

$$N(d, 2\log n) \geq N(d+2, 2\log n) + N(d+2, 2\log n) \geq 2^{\frac{6\log n-d-2}{2}} + 2^{\frac{6\log n-d-2}{2}} = 2^{\frac{6\log n-d}{2}}$$

This completes the proof of the basis of the primary induction. We assume the induction hypothesis holds for every $d' \geq d$ and $b' \geq b$ (note that one inequality must be sharp). Let v be a vertex at depth d with b Neutral edges on the path from v. Let w be a child of v. There are four possibilities: both v and w are Expanding, v is Expanding and w is Neutral, v is Neutral and w is Expanding and both v and w are Neutral. In the first three possibilities, as we already discussed above, v has at least two descendants at depth $d+2$ which are either Expanding or Neutral and thus the induction hypothesis holds for them and we have:

$$N(d,b) \geq N(d+2,b)+N(d+2,b) \geq 2^{\frac{6\log n-d-2}{2}-(2\log n-b)}+2^{\frac{6\log n-d-2}{2}-(2\log n-b)} = 2^{\frac{6\log n-d}{2}-(2\log n-b)}$$

In the fourth case in which both v and w are Neutral there is one more Neutral edge and we have

$$N(d,b) = N(d+2, b+1) = 2^{\frac{6\log n-d-2}{2}-(2\log n-b-1)} = 2^{\frac{6\log n-d}{2}-(2\log n-b)}$$

□

So far the only assumption that we used in our proofs on the equilibrium graph is that its girth is of length at least $12\log n$. The next lemma connects between the girth of an equilibrium graph and the edge cost α.

Lemma 5 *Let $G(\vec{S})$ be an equilibrium graph and c be any positive constant. If $\alpha > cn \log n$ then the length of the girth of $G(\vec{S})$ is at least $c \log n$.*

Proof. Suppose for the sake of contradiction that the size of the minimal cycle is $c \log n$, and look on a vertex u on the cycle that buys a cycle edge. The benefit of u from this edge is at most $(c \log n - 1)n$, which is strictly less than $cn \log n = \alpha$ the cost of an edge. Therefore, this is not an equilibrium graph and we reach to a contradiction. □

We are ready to state our main results, which is a characterization of every Nash equilibrium and a constant price of anarchy whenever $\alpha \geq 12n \log n$.

Theorem 2 *For $\alpha \geq 12n \log n$ the price of anarchy is bounded by $1 + \frac{6n \log n}{\alpha} \leq 1.5$ and any equilibrium graph is a tree.*

Proof. The fact that the graph is a tree follows form Lemma 5 and Proposition 1. The *social cost* of the optimum, a star graph, is $\alpha(n-1) + 2(n-1)^2$. By Proposition 1 we know that every Nash equilibrium graph is a tree whose maximal depth is $6 \log n$. Therefore, the cost of every equilibrium graph is bounded by $\alpha(n-1) + 6n^2 \log n$ and the price of anarchy is bounded by

$$\frac{\alpha(n-1) + 6n^2 \log n}{\alpha(n-1) + 2(n-1)^2} \leq 1 + \frac{6n^2 \log n}{\alpha n + 2(n-1)^2 - \alpha} \leq 1 + \frac{6n \log n}{\alpha}$$

□

3.2 Improved upper bound for $\alpha < 12n \log n$

We give a new upper bound for $\alpha < 12n \log n$. In fact, the following theorem holds for any α and is stated in this general form so that it can be generalized to a weighted game in Section 5. Furthermore, it implies a constant upper bound for $\alpha \leq O(\sqrt{n})$. The proof is given in Appendix B.

Theorem 3 *Let $\alpha > 0$. For any Nash equilibrium N, the price of anarchy is bounded by $15(1 + (\min\{\frac{\alpha^2}{n}, \frac{n^2}{\alpha}\})^{1/3})$.*

The next theorem implies that the only critical part in bounding the price of anarchy is the sum of the shortest path distances between players. The proof is given in Appendix B.

Theorem 4 *In any Nash equilibrium N, the total cost incurred by the players in building edges is bounded by twice the cost of the social optimum. There exists a shortest path tree such that, for any player v, the number of non-tree edges built by v is bounded by $1 + \lfloor (n-1)/\alpha \rfloor$.*

4 Characterizations of Nash equilibria

We give further characterization of Nash equilibria. Our first contribution is to show that, for any n and any $\alpha < n/2$, there exist transient Nash equilibria which are not trees. We then show that every Nash equilibrium which is chordal graph is a transient Nash equilibrium. An undirected graph is chordal if every cycle of length at least four has a chord, i.e. has an edge connecting two non-adjacent vertices on the cycle. Chordal graphs play a very important role in graph theory, see e.g. [9]. Finally, we show that for $\alpha < n/2$ every Nash equilibrium which is a tree must be star. The proofs of the results are given in Appendix C.

Theorem 5 *For any integer n and for any integer cost $\alpha \leq n/2$, there exists a Nash equilibrium forming a non-tree chordal graph on n vertices.*

Theorem 6 *Let $\alpha > 1$ and N be a Nash equilibrium that has a cycle in the associated graph $G = (V, E)$. If G is chordal, then N is transient.*

Theorem 7 *For $\alpha < n/2$, the star is the only Tree which is an equilibrium graph.*

We note that for $\alpha = n/2$, the construction of Theorem 5 is an equilibrium graph which is also a tree with diameter 3, and as a result Theorem 7 is tight.

5 A weighted network creation game

So far, we have considered an unweighted network creation game in which all players incur the same traffic. We now study a weighted game in which player u sends a traffic amount of $w_{uv} > 0$ to player v, with $u \neq v$. In the cost of player u, the shortest path distance between u and v is multiplied by w_{uv}. Let $\mathcal{W} = (w_{uv})_{u,v}$ be the resulting $n \times n$ traffic matrix. We use $w_{\min} = \min_{u \neq v} w_{uv}$ to denote smallest traffic entry and $w_{\max} = \max_{u \neq v} w_{uv}$ to denote the largest one. Let $W = \sum_{u=1}^{n} \sum_{v=1}^{n} w_{uv}$ be the sum of the traffic values. We extend the upper bounds of Section 3 to the weighted case. Again we assume that there are at least $n \geq 2$ players. The following theorem is a generalization of Theorem 3. In the unweighted case we have $w_{\min} = 1$ and the bounds given in the next theorem are identical to that of Theorem 3, up to constant factors. The proof is given in Appendix D.

Theorem 8 *a) Let $0 < \alpha \leq w_{\min} n^2$. For any Nash equilibrium N, the price of anarchy is bounded by $60(1 + \min\{(\alpha^2/(w_{\min}^2 n))^{1/3}, W/(w_{\min} n^4 \alpha)^{1/3}, n\})$.*

b) Let $w_{\min} n^2 < \alpha < w_{\max} n^2$. Then the price of anarchy is bounded by $12 + 3 \min\{\sqrt{\alpha / w_{\min}}, W/(\sqrt{\alpha w_{\min}}(n-1)), n\}$.

c) Let $w_{\max} n^2 \leq \alpha$. Then the price of anarchy is bounded by 4.

6 Cost sharing

We study the effect of cost sharing where players can pay for a fraction of an edge. An edge exists if the total contribution is at least α. We first show that the bounds on the price of anarchy developed in Section 3 and 5 essentially carry over. We then prove that there exist strong Nash equilibria containing cycles in which the cost is split evenly among players. We present the proofs in Appendix E.

Theorem 9 *a) In the unweighted scenario the bounds of Theorem 3 hold. b) In the weighted scenario the bound of Theorem 8 hold.*

Theorem 10 *For $n > 6$ and α in the range $\frac{1}{6}n^2 + n < \alpha < \frac{1}{2}n^2 - n$, there exist strong Nash equilibria with n players that contain cycle an in which the cost is split evenly among players.*

References

[1] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *45th Annual IEEE Symposium on Foundations of Computer Science*, pages 295–304, 2004.

[2] Elliot Anshelevich, Anirban Dasgupta, Eva Tardos, and Tom Wexler. Near-optimal network design with selfish agents. In *STOC '03: Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 511–520, New York, NY, USA, 2003. ACM Press.

[3] Venkatesh Bala and Sanjeev Goyal. A non-cooperative theory of network formation. *Econometrica*, 68(5):1181–1229, 2000.

[4] A. Blokhuis and A.E. Brouwer. Geodetic graphs of diameter two. *Geometricae Dedicata*, 25:527–533, 1988.

[5] Jacomo Corbo and David C. Parkes. The price of selfish behavior in bilateral network formation. In *Proc. 24rd ACM Symp. on Principles of Distributed Computing (PODC'05)*, 2005. To appear.

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

[7] A. Czumaj, P. Krysta, and B. Vocking. Selfish traffic allocation for server farms. In *Proceedings of the 34th Symposium on Theory of Computing*, pages 287–296, 2002.

[8] A. Czumaj and B. Vocking. Tight bounds on worse case equilibria. In *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 413–420, 2002.

[9] R. Distel. *Graph Theory. Graduate Texts in Mathematics*. Springer, 2000.

[10] Alex Fabrikant, Ankur Luthra, Elitza Maneva, Christos H. Papadimitriou, and Scott Shenker. On a network creation game. In *PODC '03: Proceedings of the twenty-second annual symposium on Principles of distributed computing*, pages 347–351. ACM Press, 2003.

[11] D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. Spirakis. The structure and complexity of nash equilibria for a selfish routing game. In *Proceedings of the 29th ICALP*, pages 123–134, 2002.

[12] A. Gupta, A. Srinivasan, and E. Tardos. Cost-sharing mechanisms for network design. In *Proc. 7th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, pages 139–150, 2004.

[13] H. Haller and S. Sarangi. Nash networks with heterogeneous agents, 2003.

[14] M. Jackson. A survey of models of network formation:stability and efficiency. In G. Demange and M. Wooders, editors, *Group Formation in Economics: Networks, Clubs and Coalitions*. 2003.

[15] K. Jain and V. Vazirani. Applications of approximation algorithms to cooperative games. In *Proc. 33rd Annual ACM Symposium on Theory of Computing*, pages 364–372, 2001.

[16] Ramesh Johari, Shie Mannor, and John N. Tsitsiklis. A contract-based model for directed network formation, 2003.

[17] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Proceedings of 16th STACS*, pages 404–413, 1999.

[18] J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, 1951.

[19] M. Pal and E. Tardos. Strategy proof mechanisms via primal-dual algorithms. In *Proc. 44th Annual IEEE Symposium on Foundations of Computer Science*, pages 584–593, 2003.

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

[21] F.J. Mac Williams and N.J.A. Sloane. *The Theory of Error-Correcting Codes*. North Holland Publishing Company, 2nd edition edition, 1978.

Appendix A

Proof of Lemma 2. Consider an arbitrary point x. We study all possible strategy changes. If x builds $l > q-1$ edges then at best there are $l+2$ vertices at distance 1 and the remaining vertices at distance 2. In x's original strategy, there are $q+1$ vertices at distance 1 while the other vertices are at distance 2. The cost difference between the new and old strategy is $(l-(q-1))\alpha + l - (q-1)$, and this value is stricly positive if $\alpha > 1$ and zero if $\alpha = 1$.

In the following we assume that x builds at most $q-1$ edges and first investigate the strategy S_0 in which x does not build any edges. The new graph relative to x is shown in Figure 6. Any line L^i_j, with $3 \le i \le q+1$ and $1 \le j \le q-1$, is at distance 3 from x because these lines are not connected to L^x_1 or L^x_2 but are each connected to one point from $x^1_1, \ldots, x^1_{q-1}$ and to one point from $x^2_1, \ldots, x^2_{q-1}$. Similarly, any point x^i_j, with $3 \le i \le q+1$ and $1 \le j \le q-1$ is at distance 3 because the point is not contained in L^x_1 or L^x_2 but is contained in one line from $L^1_1, \ldots, L^1_{q-1}$ and in one line from $L^2_1, \ldots, L^2_{q-1}$. Any line L^x_i, $3 \le i \le q+1$, is at distance 4 from x. This is because this line does not contain points x^1_j or x^2_j, for $j = 1, \ldots, q-1$, and is not parallel to lines $[L^x_1]$ and $[L^x_2]$. In Figure 6, $\mathbf{L}$ denotes the lines $L \neq L^x_i$, $i = 1, \ldots, q+1$, $L \notin [L^x_1] \cup [L^x_2]$. Symbol $\mathbf{x}$ denotes the points not equal to x, x^1_i and x^2_i, for $i = 1, \ldots, q-1$. Symbol $\mathbf{L^x}$ denotes the lines $L^x_i, 3 \le i \le q+1$. The cost difference between S_0 and the original strategy of x in G is $-(q-1)\alpha + 2(q-1)^2 + 3(q-1) = (q-1)(2q+1-\alpha) > 0$ and hence S_0 is worse.

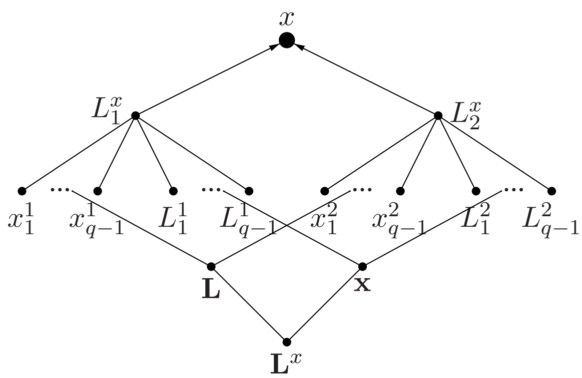

Figure 6: Strategy change S_0.

Next consider a strategy S that builds edges to vertices not equal to L^x_i, $3 \le i \le q+1$. These edges can be of four different types: x builds an edge to (a) a point x^1_j or x^2_j, for some $1 \le j \le q-1$; (b) a line L^1_j or L^2_j, for some $1 \le j \le q-1$; (c) a line L' with $L' \neq L^x_i$, for $3 \le i \le q+1$, and $L' \notin [L^x_1] \cup [L^x_2]$; (d) a point x' with $x' \neq x^1_i$ and $x' \neq x^2_i$, for $1 \le i \le q-1$. The different cases are depicted in Figure 7. We investigate how many additional vertices at distance 2 point x can reach compared to S_0. We remark that in x's original strategy each link to a line L^x_i, $3 \le i \le q+1$, gives $2(q-1)$ such vertices.

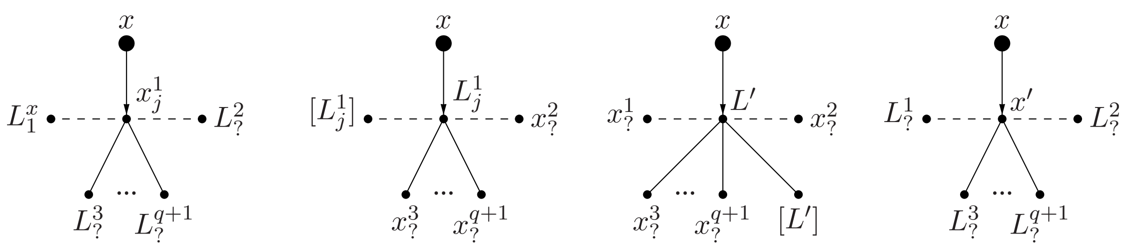

Figure 7: The effect of edges of types (a – d).

Case (a): We analyze an edge to x_j^1. This point is connected to exactly one line from $L_1^i, \ldots, L_{q-1}^i$, for any $3 \le i \le q+1$. Thus at best $q-1$ additional vertices at distance 2 are reached by x.

Case (b): We consider an edge to L_j^1. This line is connected to exactly one point from $x_1^i, \ldots, x_{q-1}^i$, for any $3 \le i \le q+1$. Thus at best $q-1$ additional vertices at distance 2 can be reached.

Case (c): Suppose that $L' \in [L_j^x]$, with $3 \le j \le q+1$. Then L' is connected to lines in $[L_j^x]$ and to exactly one point from $x_1^i, \ldots, x_{q-1}^i$, for any $3 \le i \le q+1$ with $i \ne j$. This gives a total of at most $2q-3$ extra vertices at distance 2.

Case (d): Suppose that x' belongs to L_j^x, $3 \le j \le q+1$. Thus x' is connected L_j^x and to exactly one line from $L_1^i, \ldots, L_{q-1}^i$, for any $3 \le i \le q+1$ with $i \ne j$. The number of new vertices is $q-1$.

We conclude that if S builds k edges of types (a–d), then, compared to S_0, less than $2(q-1)k$ additional edges at distance 2 can be reached by x. Now, if S builds a total of l, $l < q-1$ edges, then there must be at least $(q-1-l)2(q-1)$ edges at distance 3 from x. The cost difference relative to the original strategy of x in G is $-(q-1-l)\alpha + (q-1-l)(2q-2) > 0$ and hence S is worse. If S builds $l = q-1$ edges, then S has a cost as low as x's original strategy and only if all edges are built to L_i^x, $3 \le i \le q+1$. □

Appendix B

Proof of Theorem 3. Consider an arbitrary Nash equilibrium $N = \vec{S}$ and let $G(\vec{S}) = (V, E)$ be the corresponding equilibrium graph. We assume that $|V| = n > 1$ since otherwise, if $n = 1$, the edge set is empty and the price of anarchy is 1. Given a shortest path tree $T(u)$ and a vertex v, let $\ell(v)$ be the index of the layer v belongs to in $T(u)$. We need the following lemma.

Lemma 6 *For any $T(u)$ and any $v, w \in V$, the shortest path between v and w in G consists of at least $|\ell(v) - \ell(w)|$ edges.*

Proof. We first observe that any non-tree edge connects vertices of the same layer or of adjacent layers: If there was an edge linking a vertex x of layer i to a vertex x' of layer j, with $j \ge i+2$, then x' would rather belong to layer $i+1$. Clearly, tree edges link vertices of adjacent layers. Now, consider a shortest path $v = v_0, v_1, \ldots, v_k = w$ in G. For any i with $0 \le i \le k-1$, we have $|\ell(v_i) - \ell(v_{i+1})| \le 1$. Thus, in traversing the shortest path, each edge can reduce the layer difference between v and w by at most 1. □

Let $Cost(N)$ be the cost of N and $Cost(\text{OPT})$ be the cost of a social optimum. For the analysis of $Cost(N)$, let $Cost(v)$ be the cost paid by player $v \in V$ in N. We have $Cost(N) = \sum_{v \in V} Cost(v)$. The cost incurred by v consists of the cost for building edges and $Dist(v)$, the sum of the shortest path distances from v to all the other vertices in the equilibrium graph. Fix an arbitrary $v_0 \in V$. We prove

$$Cost(N) \le 2\alpha(n-1) + nDist(v_0) + (n-1)^2. \tag{1}$$

Consider the shortest path tree $T(v_0)$. For any vertex $v \in V$, let E_v be the number of tree edges built by v in $T(v_0)$. Vertex v_0 built only tree edges while the other vertices may have built tree as well as non-tree edges. To prove (1), we show for $v \in V$, $v \ne v_0$,

$$Cost(v) \le \alpha(E_v + 1) + Dist(v_0) + n - 1. \tag{2}$$

To verify this inequality, we modify v's strategy as follows. Vertex v discards the non-tree edges it built formerly; it only builds the tree edges it laid out before and, additionally, builds an edge to v_0. The new cost for building edges is $\alpha(E_v + 1)$. Since only non-tree edges were deleted, $Dist(v_0)$ is not affected by v's new strategy. The new edge between v and v_0 ensures that the shortest path distance between v and any other vertex w is by at most 1 larger than the shortest path distance between v_0 and w. This gives

$Dist(v) \le Dist(v_0) + n - 1$ and (2) is established. Summing (2) over all $v \ne v_0$ and adding $Cost(v_0)$ we obtain (1). This is because v_0 built only tree edges and the total number of tree edges in $T(v_0)$ is $n-1$.

It remains to analyze $Dist(v_0)$. If $\alpha < 1$, then there is a direct link between any pair of vertices and hence $Dist(v_0) \le n-1$. We obtain $Cost(N) \le 2\alpha(n-1) + 2n(n-1)$ and the price of anarchy is bounded by 2 because $Cost(\mathrm{OPT}) \ge \alpha(n-1) + n(n-1)$. If $\alpha > n^2$, then we use the trivial bound $Dist(v_0) \le (n-1)^2$ and $Cost(N) \le 2\alpha(n-1) + 2n(n-1)^2$ and the price of anarchy is bounded by 4 because $Cost(\mathrm{OPT}) > \alpha(n-1) > n^2(n-1)$.

In the remainder of this proof we assume $1 \le \alpha \le n^2$. In this case a social optimum is given by the star graph, which incurs a cost of $Cost(\mathrm{OPT}) = \alpha(n-1) + 2(n-1)^2 > \alpha(n-1) + n^2$, for $n \ge 2$ players. Let d be the depth of $T(v_0)$, i.e. d is the maximum layer number $\max_{v\in V} \ell(v)$. If $d \le 9$, we are easily done. We have $Dist(v_0) \le 9n$ and $Cost(N) \le 2\alpha(n-1) + 10n^2$ and the desired price of anarchy holds because $Cost(\mathrm{OPT}) > \alpha(n-1) + n^2$. Thus, in the following we restrict ourselves to the case $d \ge 10$.

Determine c, $1/3 \le c \le 1$, such that $\alpha = n^{3c-1}$. Let $V' = \{v \in V \mid \ell(v) \le \lfloor \frac{2}{5}d \rfloor \text{ in } T(v_0)\}$ be the set of vertices of depth at most $\lfloor \frac{2}{5}d \rfloor$ in $T(v_0)$. We distinguish two cases depending on whether $|V'| \ge \frac{2}{3}n^c$ or $|V'| < \frac{2}{3}n^c$.

If $|V'| \ge \frac{2}{3}n^c$, then consider a vertex w_0 at depth d in $T(v_0)$, i.e. in $\ell(w_0) = d$ in $T(v_0)$. By Lemma 6, the shortest path distance between w_0 and any vertex $v \in V'$ is at least $\lceil \frac{3}{5}d \rceil$. If there was an edge between w_0 and v_0, then the distance between w_0 and v would be at most $\lfloor \frac{2}{5}d \rfloor + 1$. Since w_0 did not build an edge to v_0 we have

$$\alpha > |V'| \left(\left\lceil \frac{3}{5}d \right\rceil - \left\lfloor \frac{2}{5}d \right\rfloor - 1 \right) \ge \frac{2}{3}n^c \left(\frac{1}{5}d - 1 \right) \ge \frac{2}{3}n^c \frac{1}{10}d$$

and hence

$$d \le \frac{15\alpha}{n^c}. \tag{3}$$

Next assume $|V'| < \frac{2}{3}n^c$. For any i with $\lfloor \frac{1}{5}d \rfloor + 1 \le i \le \lfloor \frac{2}{5}d \rfloor$ let $V'_i = \{v \in V' \mid \ell(v) = i \text{ in } T(v_0)\}$ be the vertices at depth i in $T(v_0)$. There must exist an i_0 with $|V'_{i_0}| < \frac{2}{3}n^c / \lfloor \frac{1}{5}d \rfloor$ since otherwise

$$|V'| \ge \sum_{i=\lfloor \frac{1}{5}d \rfloor + 1}^{\lfloor \frac{2}{5}d \rfloor} |V'_i| \ge \lfloor \frac{1}{5}d \rfloor \cdot \frac{2}{3}n^c / \lfloor \frac{1}{5}d \rfloor = \frac{2}{3}n^c,$$

contradicting the assumption that $|V'| < \frac{2}{3}n^c$. There are at least $n - \frac{2}{3}n^c \ge \frac{1}{3}n$ vertices in $V \setminus V'$. Each such vertex is decendent of one vertex in V_{i_0}. Thus, there is one vertex $v_{i_0} \in V_{i_0}$ having at least

$$\frac{n/3}{\frac{2}{3}n^c / \lfloor \frac{1}{5}d \rfloor} = \frac{1}{2}n^{1-c} \left\lfloor \frac{1}{5}d \right\rfloor \ge \frac{1}{2}n^{1-c} \left(\frac{1}{5}d - 1 \right) \ge \frac{d}{20}n^{1-c}$$

decendents. If there was an edge from v_0 to v_{i_0}, then the shortest path distance from v_0 to these decendents would be reduced by at least $\left\lfloor \frac{1}{5}d \right\rfloor \frac{d}{20}n^{1-c} \ge \frac{d^2}{100}n^{1-c}$. Since v_0 did not build such an edge, $\alpha \ge \frac{d^2}{200}n^{1-c}$, which gives

$$d \le 15\sqrt{\frac{\alpha}{n^{1-c}}}. \tag{4}$$

The bounds on d shown in (3) and (4) are identical because $\frac{15\alpha}{n^c} = 15\sqrt{\alpha/n^{1-c}}$ is equivalent to $\alpha = n^{3c-1}$ and this holds by the choice of c.

We finally determine the price of anarchy. We have $Dist(v_0) \le (n-1)15\alpha/n^c \le 15\alpha n^{1-c}$. Using (1) we obtain $Cost(N) \le 2\alpha(n-1) + 15\alpha n^{2-c} + n^2$. The price of anarchy is bounded by

$$\frac{2\alpha(n-1) + 15\alpha n^{2-c} + n^2}{\alpha(n-1) + n^2} \le 3 + \frac{15\alpha n^{2-c}}{\alpha(n-1) + n^2}.$$

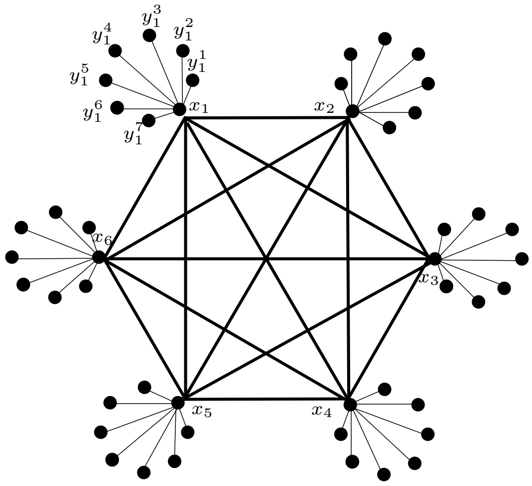

Figure 8: A $(6, 8)$ clique of stars graph, an equilibrium graph which is not a tree.

If $\alpha \leq n$, then the price of anarchy is bounded by $3 + 15\alpha/n^c < 15(1 + \alpha/n^c) = 15(1 + n^{2c-1}) = 15(1 + (\alpha^2/n)^{1/3})$ because the definition of c implies that $n^c = (\alpha n)^{1/3}$. If $\alpha > n$, then we use the fact that $\alpha(n-1) + n^2 > \alpha n$. This hold because $\alpha \leq n^2$. The price of anarchy is bounded by $3 + 15n^{1-c} < 15(1 + n^{1-c}) = 15(1 + (n^2/\alpha)^{1/3})$, using again the fact that $n^c = (\alpha n)^{1/3}$. $\square$

Proof of Theorem 4. Consider the graph $G = (V, E)$ associated with N. Again, for $v \in V$, let $Cost(v)$ be the cost incurred by v and let $Dist(v)$ be the sum of the shortest path distances from v to all the other vertices in V. Choose a vertex v_0 with minimum $Dist$-value among all vertices, i.e. $Dist(v_0) = \min_{v \in V} Dist(v)$ and consider the shortest path tree $T(v_0)$. For any $v \in V$, let E_v be the number of tree edges and let E'_v be the number of non-tree edges built by v in $T(v_0)$. The total cost incurred by the players in building edges is $\sum_{v \in V}(E_v + E'_v)$.

Suppose that player v's strategy, $v \neq v_0$, is modified as follows. Agent v deletes its E'_v non-tree edges. It only builds the E_v tree edges it laid out before and, additionally, build an edge to v_0. With this additional edge, the shortest path distance from v to any vertex w is by at most one larger then the shortest path distance from v_0 to w. Since v does not follow this strategy, $Cost(v) = \alpha(E_v + E'_v) + Dist(v) \leq \alpha E_v + \alpha + Dist(v_0) + n - 1$, which by the minimality of $Dist(v_0)$ implies

$$E'_v \leq 1 + \lfloor (n-1)/\alpha \rfloor. \tag{5}$$

There is a total of $n - 1$ tree edges in $T(v_0)$ and $E'_{v_0} = 0$. Thus the total cost paid by the players in building edges is bounded by $\alpha(n-1) + \alpha(n-1) + (n-1)^2$ and this is at most twice the cost $Cost(\text{OPT})$ of a social optimum because $Cost(\text{OPT}) \geq \alpha(n-1) + n(n-1)$. $\square$

Appendix C

Proof of Theorem 5. We start by describing our non-tree chordal equilibrium graph. A (k, ℓ) *clique of stars* is a clique with k vertices, where each vertex of the clique is a root of a star with ℓ vertices. A $(6, 8)$ clique of stars is depicted in Figure 8.

We next prove that a (k, ℓ) clique of stars is a Nash equilibrium when $\alpha = \ell$ and the edges of each star are bought only by its root, and clique edges are bought arbitrarily by one of their vertices.

Lemma 7 *Let $G(\vec{S})$ be a (k, ℓ) clique of stars. If the cost of an edge equals to ℓ and all the edges are bought by the clique vertices (and no edge is bought twice), then $G(\vec{S})$ is an equilibrium graph.*

Proof. We prove that a (k, ℓ) clique of stars is an equilibrium in this setting by showing that no player has an incentive to deviate form her strategy. We denote with $x_1, \ldots, x_k$ the vertices of the clique and with $y_i^1, \ldots, y_i^{\ell-1}$ the vertices of the star rooted at x_i.

We start by showing that the star vertices have no incentive to deviate form their strategy of not buying any edge. We look on an arbitrary star vertex y_i^j. The edge connecting it to the graph is bought by x_i. The benefit from buying the edge (y_i^j, x_p) for $p \neq i$ is ℓ, since y_i^j is getting closer by one only to the vertices of the star rooted at x_p. The cost of an edge is also ℓ therefore the player y_i^j is indifferent and will not deviate. The benefit from buying the edge $(y_i^j, y_{i'}^{j'})$ is only one and thus y_i^j will have no incentive to buy it. Since buying a set of edges is at most as beneficial as the sum of their benefits in a connected graph, y_i^j will not deviate.

We now turn our attention to the clique vertices. We take an arbitrary vertex x_i. Its star vertices are connected with an edge of the form (x_i, y_i^j). If x_i does not buy one of these edges the graph get disconnected and the cost of x_i becomes infinity. Thus, these edges are necessary. Suppose that the edge (x_j, x_i) is bought by x_j, then x_j is indifference of buying or not buying the edge, since without the edge the distance to the star rooted at x_i is at least 2 while it is 1 with the edge. The benefit from buying the edge is ℓ which is also the cost of an edge. Clearly x_j can not benefit from buying an edge to a leaf of another star, say y_p^k, since $\alpha \geq 1$ and the benefit is exactly 1. Thus, x_j has no incentive to change its strategy and we conclude that $G(\vec{S})$ is an equilibrium graph. □

For every n we have a family of (k, ℓ) clique of stars with $k \cdot \ell = n$ and $\alpha = \ell$. This implies that we can build a non-tree equilibrium for $\alpha = n/3, n/4, \ldots, 1$. By a slightly more complicated construction it is possible to extend the (k, ℓ) clique of stars construction and to derive the desired theorem. Details are given in the full version of the paper. □

Proof of Theorem 6. Consider an arbitrary cycle of length three in G. On this cycle, considering directed edges, either (a) each of the three cycle vertices has exactly one incoming and one outgoing cycle edge or (b) there exists one vertex that has two outgoing edges. In case (a) we name the vertices on the cycle v_0, v_1 and v_2, starting at an arbitrary vertex and then following the cycle orientation. In case (b), let v_0 be the vertex with two outgoing cycle edges and name the remaining two vertices such that there are oriented edges (v_0, v_1) and (v_1, v_2). This leads to the configuration shown in Figure 9. The edge between v_0 and v_2 can be oriented in two ways.

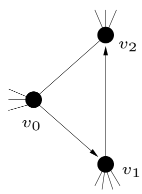

Figure 9: The cycle of vertices v_0, v_1 and v_2.

Let V_{12} be the set of vertices v, $v \neq v_0$, that are directly linked to both v_1 and v_2, i.e. $V_{12} = \{v \in V \mid v \neq v_0 \text{ and } \{v, v_i\} \in E \text{ for } i = 1, 2)\}$. Furthermore, let W be the set of vertices $w \in V$ such that a shortest path from v_1 to w uses edge (v_1, v_2) and any other path from v_1 to w that does not use (v_1, v_2) is strictly longer than a shortest path. Obviously, (v_1, v_2) is the first edge on the shortest paths from v_1 to vertices w. Furthermore, W and V_{12} are disjoint. Set W must contain at least α vertices since otherwise v_1 could delete edge (v_1, v_2) and instead use the edges between v_1 and v_0 and between v_0 and v_2 to reach v_2 on the path to $w \in W$. The would lower v_1's cost for building edges by α while its shortest paths cost would increase by less than α.

Let V_1 be the set of vertices $v \in V$, $v \notin V_{12} \cup \{v_0, v_2\}$, that are directly linked to v_1. Formally, $V_1 = \{v \in V \mid v \notin V_{12} \cup \{v_0, v_2\}$ and edge $\{v, v_1\} \in E\}$. We next prove that, for any $v \in V_1$ and $w \in W$, a shortest path from v to w is by at least 1 longer than a shortest path from v_1 to w. Assume that this were not the case. Let $v \in V_1$ be a vertex such that the desired statement is violated for some vertices in W. Among those candidates, let $w \in W$ be the one having the smallest distance from v_2. Let P_v be a shortest path from v to w and P_{v_1} be a shortest path from v_1 to w. Path P_v does not use (v_1, v_2) since otherwise P_v would be one edge longer than P_{v_1}. Path P_{v_1} does use (v_1, v_2) by the definition of W. Path P_v cannot be shorter then P_{v_1}; otherwise the path consisting of the edge between v_1 and v, followed by P_v would be a shortest path from v_1 to w, contradicting the fact that $w \in W$. Hence P_{v_1} and P_v have the same length. All the vertices of P_{v_1}, except for v_1, belong to W. Therefore P_{v_1} and P_v are edge disjoint. If they was a common suffix S, then the first vertex of S would be a vertex in W closer to v_2 violating the desired statement. Paths P_{v_1} and P_v each have a length of at least two, since otherwise $w = v_2$ and hence $v \in V_{12}$.

Consider the following cycle C that has a length of the least five. Starting at v_1 we follow the edge to v, then traverse the path P_v to w and finally traverse the edges of P_{v_1} to reach v_1. We argue that neither v_1 nor v has a chord to any other vertex on C. A chord between v_1 and another vertex on C would imply a shortest path between v_1 and w that does not use (v_1, v_2), contradicting the definition of W. If there was a chord between v and v_2, then $v \in V_{12}$. If there was a chord between v and any other vertex on C, this would imply the existence of a path form v to w that is shorter then P_v. Using this property of v and v_1, we are able to identify a cycle C' of length at least four that that has no chord. We start at vertex v_1, follow the edge to v and traverse the first edge of P_v. Let w_1 be the vertex reached. ¿From w_1 we traverse the chord that skips the largest number of edges on the arc of C between w_1 and v_2. If there is no chord at w_1, we traverse the next edge of C leaving w_1. Let w_2 be the vertex reached. We proceed in the same way as in vertex w_1. In general, when at vertex w_i we follow the chord that skips the largest number of edges on the cycle arc between w_i and v_2. If there is no such chord, we traverse the next cycle edge. Eventually we reach v_2 and can complete C' by traversing the edge between v_2 and v_1. The existence of C' is a contradiction to the fact that the undirected graph underlying our Nash equilibrium is chordal.

We conclude that, indeed, for any $v \in V_1$ and $w \in W$ a shortest path from v to w is at least one edge longer than a shortest path from v_1 to w. Using this property we can show that N is transient. If vertex $v \in V$ builds an edge to v_2, its cost can only decrease because the shortest path distances between v and $w \in W$ decrease by at least $|W| \geq \alpha$ while the cost for building edges increases by α. The fact that v did not build this edge in N implies that $|W| = \alpha$ and N is transient because v can alter his strategy without changing his cost. An edge (v, v_2) does not change the shortest path distances from other vertices $v' \in V_1$, $v' \neq v_1$, to vertices $w \in W$: If v' uses (v, v_2) on a shortest path, it needs at least two edges to reach v_2 and this was also the number of edges to reach v_2 in N.

The single player changes are now as follows. Agents $v \in V_1$ one after the other introduce an edge (v, v_2). The changer's cost does not change. At this point we have reached a non-equilibrium state N': Agent v_0 can delete edge (v_0, v_1), saving a cost of α. We finally show that only the shortest path distance to v_1 increases by one. In the original equilibrium N, consider a shortest path from v_0 to some vertex $w \neq v_1$ that uses edge (v_0, v_1). After v_1, the shortest path visits a vertex $v' \in V_{12} \cup V_1$. The subpath (v_0, v_1) followed by the edge between v_1 and v' in N can be replaced by the edges between v_0 and v_2 and between v_2 and v' in N'. If $v' \in V_1$, the last edge was newly introduced. □

Proof of Theorem 7. Suppose for the sake of contradiction that there is an equilibrium graph which is a tree but not a star. It is well known that any tree has a centroid vertex whose removal leaves the tree with components of size smaller than $n/2$. Let v be such a centroid vertex and let u be a leaf at depth $d \geq 2$. It is easy to see that since the removal of v leaves the tree with components of size at most $n/2$, there must be at least $n/2$ vertices whose shortest path to u passes through v. Buying the edge (u, v) would save $n(d-1)/2$ to u and thus we get that $\alpha \geq n(d-1)/2 \geq n/2$, a contradiction. □

Appendix D

Proof of Theorem 8. Let N be any Nash equilibrium. We extend the proof of Theorem 3 and first develop a modified bound on $Cost(N)$. Consider the equilibrium graph $G = (V, E)$ given by N and fix an arbitrary player $v_0 \in V$. We use the shortest path tree $T(v_0)$ rooted at v_0, which is defined in the same way as in the unweighted case. We simply ignore traffic weights and just consider the edges in E to identify the structure of $T(v_0)$. Again, let E_v be the number of edges built by player $v \in V$ and let d be the depth of $T(v_0)$. We have

$$Cost(v_0) \leq \alpha E_{v_0} + d \sum_{\substack{u \in V \\ u \neq v_0}} w_{v_0 u}$$

because v_0 builds only tree edges and the number of edges between v_0 and any other $u \in V$ is bounded by d. We next show

$$Cost(v) \leq \alpha(E_v + 1) + (d+1) \sum_{\substack{u \in V \\ u \neq v}} w_{uv}.$$

To verify this inequality we simply observe that if v_0 decides to build only its tree edges, deleting the non-tree edges, and additionally builds an edge to v_0, its cost is given by the right-hand side of the inequality. Summing the costs over all vertices, we obtain

$$Cost(N) \leq 2\alpha(n-1) + (d+1)W.$$

It remains to analyze d. Obviously, $d \leq n-1$ and hence $Cost(N) \leq 2\alpha(n-1) + nW$. Since $Cost(\mathrm{OPT}) \geq \alpha(n-1) + W$, this establishes the upper bounds of $60(1+n)$ and $12+3n$ in parts a) and b) of the theorem. We can also establish part c) of the theorem because, if $\alpha \geq w_{\max} n^2$, we have $Cost(N) \leq 2\alpha(n-1) + n^3 w_{\max}$ and $Cost(\mathrm{OPT}) \geq n^2(n-1)w_{\max}$. If $\alpha < w_{\min}$, then there is a direct link between any pair of players and the price of anarchy is bounded by 1 because $Cost(\mathrm{OPT}) \geq \alpha n(n-1)/2 + W$.

In the following we assume $w_{\min} \leq \alpha \leq w_{\max} n^2$ and develop a refined bound on d. If $d \leq 9$, then the price of anarchy is bounded by 12. Therefore, we assume $d \geq 10$. To prove part a) of the theorem, we determine c, $1/3 \leq c \leq 1$ such that $\alpha = w_{\min} n^{3c-1}$ and let $V' = \{v \in V \mid \ell(v) \leq \lfloor \frac{2}{5} d \rfloor \text{ in } T(v_0)\}$. If $|V'| \geq \frac{2}{3} n^c$, then a vertex w_0 at depth $T(v_0)$ could save a cost of $w_{\min}|V'|(\lceil \frac{3}{5} d \rceil - \lfloor \frac{2}{5} d \rfloor - 1)$ by building an edge to v_0. Since w_0 does not build such an edge, α is at least as large as the latter expression, implying

$$d \leq \frac{15\alpha}{w_{\min} n^c}. \tag{6}$$

If $|V'| < \frac{2}{3} n^c$, then, as in the proof of Theorem 3, there must exist a vertex v_{i_0} at depth d_0 with $\lfloor \frac{1}{5} d \rfloor + 1 \leq d_0 \leq \lfloor \frac{2}{5} d \rfloor$ having at least $dn^{1-c}/20$ decendents. Building an edge to v_0, vertex v_{i_0} would save a cost of at least $w_{\min} \lfloor \frac{1}{5} d \rfloor \frac{d}{20} n^{1-c} \geq w_{\min} \frac{d^2}{100} n^{1-c}$. This cost saving must be upper bounded by α since v_{i_0} does not build such an edge. We obtain

$$d \leq 15 \sqrt{\frac{\alpha}{w_{\min} n^{1-c}}}. \tag{7}$$

By the choice of c, the bounds on d given in (6) and (7) are identical. Using these bound (6) we derive

$$Cost(N) \leq 2\alpha(n-1) + \left(\frac{15\alpha}{w_{\min} n^c} + 1\right) W \leq 2\alpha(n-1) + 2W \frac{15\alpha}{w_{\min} n^c}$$

We recall that $Cost(\mathrm{OPT}) \geq \alpha(n-1) + W$. Thus, if $\alpha(n-1) \leq W$, the price of anarchy is bounded by

$$2 + \frac{30\alpha}{w_{\min} n^c} = 2 + 30 \left(\frac{\alpha^2}{w_{\min}^2 n} \right)^{1/3}$$

because $n^c = (\alpha n / w_{\min})^{1/3}$. If $\alpha(n-1) > W$, the price of anarchy is bounded by

$$2 + \frac{30W}{w_{\min} n^c (n-1)} \leq 2 + \frac{60W}{(w_{\min} n^4 \alpha)^{1/3}},$$

using again $n^c = (\alpha n / w_{\min})^{1/3}$.

To prove part b) of the theorem, we finally study the case that α is in the range $w_{\min} < \alpha < w_{\max} n^2$. Here we use a different estimate on d. We have that d is upper bounded by $3\sqrt{\alpha/w_{\min}}$, since otherwise v_0 could build an edge to a vertex that is $\lceil \sqrt{\alpha/w_{\min}} \rceil + 1$ edges away on a path of length d. This would reduce the shortest distance cost by at least $w_{\min} \lceil \sqrt{\alpha/w_{\min}} \rceil (3\lceil \sqrt{\alpha/w_{\min}} \rceil - \lceil \sqrt{\alpha/w_{\min}} \rceil) > \alpha$. Thus

$$Cost(N) \leq 2\alpha(n-1) + 3\sqrt{\alpha/w_{\min}} W.$$

If $\alpha(n-1) \leq W$, then the price of anarchy is bounded by $2 + 3\sqrt{\alpha/w_{\min}}$. If $\alpha(n-1) > W$, the price of anarchy is bounded by $2 + 3W/(\sqrt{\alpha w_{\min}}(n-1))$. □

Appendix E

Proof of Theorem 9. We first show part a). Using the terminology of the proof of Theorem 3, we can show that for any $v \in V$, $Cost(v) \leq \alpha(E_v + 1) + Dist(v_0) + n - 1$. To see this inequality, we modify v's strategy such that it removes its cost contributions to non-tree edges. Agent v only maintains its contributions to tree edges and, additionally, builds an edge to v_0, the vertex for which we consider the corresponding shortest path tree. The cost under this modified strategy is bounded by the expression given above. Summing over all v we obtain $Cost(N) \leq 2\alpha(n-1) + nDist(v_0) + (n-1)^2$. We can then bound $Dist(v_0)$ in exactly the same way as in the proof of Theorem 3.

For the proof of part b), using the terminology of the proof of Theorem 8, we can show $Cost(N) \leq 2\alpha(n-1) + (d+1)W$. We can extend the arguments presented for the scenario without cost sharing to bound d in a similar way. □

Proof of Theorem 10. Consider a cycle of n vertices $v_1, \ldots, v_n$. There is an edge between v_i and v_{i+1}, $1 \leq i \leq n-1$, and an edge between v_n and v_1. We associate a player with each of the n vertices. Every player pays a cost of $\alpha/2$ for each of the two edges adjacent to him, incurring a total cost of α for building edges. We show that this cycle represents a strong Nash equilibrium for the given range of α. Since the strategies of players v_i, $1 \leq i \leq n$, are symmetric in i, it suffices to prove that there is no strictly better strategy for v_1. We first analyze the cost of v_1. There are two vertices at each of the distances 1 up to $\lfloor \frac{n}{2} \rfloor - 1$. If n is even, there is one vertex at distance $\lfloor \frac{n}{2} \rfloor$; otherwise there are two such vertices. We have

$$\begin{aligned} Cost(v_0) &= \alpha + 2\left(1 + \ldots + \left\lfloor \frac{n}{2} \right\rfloor\right) - \left\lfloor \frac{n}{2} \right\rfloor ((n+1) \bmod 2) && (8) \\ &= \alpha + \left\lfloor \frac{n}{2} \right\rfloor \left(\left\lfloor \frac{n}{2} \right\rfloor + 1 \right) - \left\lfloor \frac{n}{2} \right\rfloor ((n+1) \bmod 2). && (9) \end{aligned}$$

We investigate the following strategy changes.

(a) Agent v_1 maintains its cost contributions to the two adjacent edges and, additionally, builds new edges to other vertices.

(b) Agent v_1 removes its cost contribution to one of the adjacent edges and does not build any new edges.

(c) Agent v_1 removes its cost contribution to one of the adjacent edges and builds does build new edges to other vertices.

(d) Agent v_1 removes its cost contributions to the two adjacent edges and, instead, builds new edges to other vertices.

Case (a): We first assume that v_1 builds one additional edge and then consider the scenario that more edges are built. If one extra edge is added, then the best strategy is to connect to vertex v_i with $i = \lfloor \frac{n}{2} \rfloor + 1$. With this new link, v_1 has three vertices at distance 1 and four vertices at each of the distances 2 up to $\lfloor \frac{n}{4} \rfloor$. If $n \bmod 4 = 1$, then there is one additional vertex at distance $\lfloor \frac{n}{4} \rfloor + 1$. If $n \bmod 4 = 2$, there are two additional vertices at this distance. Three such vertices exist if $n \bmod 4 = 3$. Thus, v_1's new shortest path distance cost is

$$\begin{aligned} & 3 \cdot 1 + 4(2 + \ldots + \left\lfloor \frac{n}{4} \right\rfloor) + (n \bmod 4) \left(\left\lfloor \frac{n}{4} \right\rfloor + 1 \right) \\ = \; & 2 \left(\left\lfloor \frac{n}{4} \right\rfloor + 1 \right) \left(\left\lfloor \frac{n}{4} \right\rfloor + \frac{1}{2}(n \bmod 4) \right) - 1 \\ \geq \; & \frac{n^2}{8} - 1. \end{aligned}$$

The difference in v_1's shortest path distance cost is

$$\begin{aligned} & \left\lfloor \frac{n}{2} \right\rfloor \left(\left\lfloor \frac{n}{2} \right\rfloor + 1 \right) - \left\lfloor \frac{n}{2} \right\rfloor ((n+1) \bmod 2) - \frac{n^2}{8} + 1 \\ \leq \; & \frac{n^2}{8} + \frac{n}{2} + 1 \\ \leq \; & \alpha - \frac{n^2}{6} - n + \frac{n^2}{8} + \frac{n}{2} + 1 \\ < \; & \alpha \end{aligned}$$

and it does not pay to build an additional edge since the extra cost for that edge is α.

Next assume that there was a strategy in which v_1 builds two or more additional edges, incurring a total cost bounded by (9). Consider the strategy with the smallest number of additional edges and suppose that there are at least two such links. The removal of any extra link to a vertex v_{i_0}, $2 < i_0 < n$, would increase the shortest path distance cost by more than α. In other words, the addition of the link to v_{i_0} leads to a decrease in the shortest path distance cost by more than α. This implies that if v_1 maintained its original strategy and only added one link to v_{i_0}, this would lead to a smaller total cost. This contradicts the calculations of the last paragraph where we showed that an extra link to an optimal vertex v_i, $i = \lfloor \frac{n}{2} \rfloor + 1$, does not pay off.

Case (b): We assume w.l.o.g. that v_1 removes its cost contribution to the edge connecting to v_2, saving a cost of $\alpha/2$. Vertices v_i, for $i = 2, \ldots, \lceil \frac{n}{2} \rceil$, must now be reached by traversing the cycle arc through v_n. The shortest path distance cost of v_1 increases by

$$\left\lfloor \frac{n}{2} \right\rfloor \left(\left\lceil \frac{n}{2} \right\rceil - 1 \right) \geq \frac{n}{2} \left(\frac{n}{2} - 1 \right).$$

Since $\alpha/2$ is smaller than the latter expression, v_1 does not perform the considered strategy change.

Case (c): Again we assume that v_1 removes its cost contribution to the edge connecting to v_2. We first study the scenario that v_1 builds one new edge and then address the case that more new edges are built. If one additional edge ist built, then the best strategy is to connect to vertex v_i with $i = \lfloor \frac{n}{3} \rfloor + 1$. Then v_1 can reach two vertices at distance 1 and three vertices at each of the distances 2 up to $\lfloor \frac{n}{3} \rfloor$. If $n \bmod 3 = 1$, there is one additional vertex at distance $\lfloor \frac{n}{3} \rfloor + 1$. If $n \bmod 3 = 2$, there are two additional vertices at this distance. Thus the new path distance cost of v_1 is

$$\begin{aligned}
&2 \cdot 1 + 3\left(1 + \ldots + \left\lfloor \frac{n}{3} \right\rfloor\right) + \left(\left\lfloor \frac{n}{3} \right\rfloor + 1\right)(n \bmod 3) \\
&= \frac{3}{2}\left(\left\lfloor \frac{n}{3} \right\rfloor + 1\right)\left(\left\lfloor \frac{n}{3} \right\rfloor + \frac{2}{3}(n \bmod 3)\right) - 1 \\
&\geq \frac{3}{2}\left(\left\lfloor \frac{n}{3} \right\rfloor + 1\right)\frac{n}{3} - 1 \\
&\geq \frac{3}{2}\left(\frac{n}{3} + \frac{1}{3}\right)\frac{n}{3} - 1 \\
&\geq \frac{n^2}{6}.
\end{aligned}$$

Hence v_1's saving in the shortest path distance cost is at most

$$\left\lfloor \frac{n}{2} \right\rfloor \left(\left\lfloor \frac{n}{2} \right\rfloor + 1\right) - \left\lfloor \frac{n}{2} \right\rfloor ((n+1) \bmod 2) - \frac{n^2}{6} \leq \frac{n}{2}\left(\frac{n}{2} + 1\right) - \frac{n^2}{6}$$

and this is less than $\alpha/2$, which is the extra cost incurred by v_1 in building edges.

Next assume that there was a strategy in which v_1 builds more than one additional edge, leading to a cost bounded by that given in (9). Consider the strategy with the smallest number of additional edges and suppose that there are at least two such links. Let i_0, $i_0 < n$, be the largest index such that v_1 builds an additional edge to v_{i_0}. As in case (a) it follows that the deletion of the link to v_{i_0} would increase the shortest path distance cost by more than α. Equivalently, the addition of the link to v_{i_0} leads to a decrease of the shortest path distance cost by more than α. This implies that the following strategy leads to a cost smaller than (9): Vertex v_1 maintains its cost contribution to the edges connecting to v_2 and v_n and builds an additional edge to v_{i_0}. This contradicts the fact that, as argued above, strategy changes of type (a) lead to strictly higher cost.

Case (d): We first study the scenario that v_1 builds one new edge and then investigate the case that two or more new edges are built. If one new edge is built, then v_1's total cost for building edges remains the same. The best strategy is to build a link to the vertex v_i with $i = \lfloor \frac{n}{2} \rfloor + 1$. With respect to v_1's shortest path distance cost, there is one vertex at distance 1 and two vertices at each of the distances 2 up to $\lfloor \frac{n}{2} \rfloor$. If n is odd, there is one vertex at distance $\lceil \frac{n}{2} \rceil$. Thus the new shortest path distance cost is

$$1 + 2\left(2 + \ldots + \left\lfloor \frac{n}{2} \right\rfloor\right) + \left\lceil \frac{n}{2} \right\rceil (n \bmod 2).$$

The cost difference with respect to v_1's original strategy is

$$\left\lceil \frac{n}{2} \right\rceil (n \bmod 2) - 1 + \left\lfloor \frac{n}{2} \right\rfloor ((n+1) \bmod 2)$$

and this is strictly positive for $n > 6$.

Next suppose that two new edges are built. The best strategy for v_1 is to connect to v_{i_1}, with $i_1 = \lceil \frac{n}{4} \rceil + 1$, and to v_{i_2}, with $i_1 = \lfloor \frac{3n}{4} \rfloor + 1$. Vertex v_1 has two vertices at distance 1 and four vertices at each of the distances 2 up to $\lfloor \frac{n}{4} \rfloor$. If n is divisible by 4, then there is one additional vertex at distance $\lfloor \frac{n}{4} \rfloor + 1$. If

$n \bmod 4 = 1$, then there are two additional vertices at distance $\lfloor \frac{n}{4} \rfloor + 1$. If $n \bmod 4 = 2$, then there are three additional vertices at that distance. Thus the new shortest path distance cost is

$$\begin{aligned} & 2 \cdot 1 + 4\left(2 + \ldots + \left\lfloor \frac{n}{4} \right\rfloor\right) + \left(\left\lfloor \frac{n}{4} \right\rfloor + 1\right)((n+1) \bmod 4) \\ = \; & 2\left(\left\lfloor \frac{n}{4} \right\rfloor + 1\right)\left(\left\lfloor \frac{n}{4} \right\rfloor + \frac{1}{2}(n+1) \bmod 4\right) - 2 \\ \geq \; & \frac{n}{2}\left(\frac{n}{4} - \frac{3}{4}\right) - 2. \end{aligned}$$

The difference in the shortest path distance cost is upper bounded by

$$\begin{aligned} & \left\lfloor \frac{n}{2} \right\rfloor \left(\left\lfloor \frac{n}{2} \right\rfloor + 1\right) - \left\lfloor \frac{n}{2} \right\rfloor ((n+1) \bmod 2) - \frac{n}{2}\left(\frac{n}{4} - \frac{3}{4}\right) + 2 \\ < \; & \alpha - \frac{1}{6}n^2 - n + \frac{1}{8}n^2 + \frac{7}{8}n + 2 \\ < \; & \alpha. \end{aligned}$$

Hence it does not pay to build two additional edges.

Finally assume that there was a strategy in which v_1 builds three or more additional edges, leading to a cost bounded by that given in (9). As usual, consider the strategy with the smallest number of additional edges and suppose that there are at least three such links. Let i_0, $i_0 < n$, be the second to largest index such that v_1 builds an additional edge to v_{i_0}. We can now argue as in case (c). Removing the link to v_{i_0} increases the shortest path distance cost of v_1 by more than α, i.e. the addition of the link to v_{i_0} leads to a decrease of the shortest path distance cost by more than α. This implies that the following strategy has a cost smaller than (9): Vertex v_1 maintains its cost contribution to the edges connecting to v_2 and v_n and builds an additional edge to v_{i_0}. As before, this contradicts the fact that strategy changes of type (a) lead to strictly higher cost. □