

TREE NASH EQUILIBRIA IN THE NETWORK CREATION GAME

Akaki Mamageishvili, Matúš Mihalák, and Dominik Müller

Department of Computer Science, ETH Zurich, Zurich, Switzerland

Abstract *In the network creation game with n vertices, every vertex (player) creates an (adjacent) edge and decides to which other vertices the created edge should go. Each created edge costs a fixed amount $\alpha > 0$. Each player aims to have a good connection with the rest of the vertices and, at the same time, to pay as little as possible. Formally, the cost of each player in the resulting (created) graph is defined as α times the number of edges created by the player plus the sum of the distances to all other vertices. It has been conjectured that for $\alpha \geq n$, every Nash equilibrium of this game is a tree and has been confirmed for every $\alpha \geq 273 \cdot n$. We improve on this bound and show that this is true for every $\alpha \geq 65 \cdot n$. We also show that our approach cannot be used to show the desired bound, but we conjecture that a slightly worse bound $\alpha \geq 1.3 \cdot n$ can be achieved. Toward this conjecture, we show that if a Nash equilibrium has a cycle of length at most 10, then indeed $\alpha < 1.3 \cdot n$. We investigate our approach for a coalitional variant of a Nash equilibrium, which coalitions of two players cannot collectively improve, and show that if $\alpha \geq 41 \cdot n$, then every such Nash equilibrium is a tree.*

1. INTRODUCTION

The network creation game was introduced by [9] as a formal model to study the effects of strategic decisions of economically motivated agents in decentralized networks such as the Internet. In such networks, local decisions, including those about infrastructure, are decided by autonomous systems. Autonomous systems follow their own interest, and as a result, their decisions may be suboptimal for the whole society. Network creation games allow us to formally study the structure of networks created in such a manner, and to compare them with potentially optimal networks (optimal with respect to the whole society).

In the *network creation game*, there are n players $V = \{1, \dots, n\}$, each representing a vertex of an undirected graph. A *strategy* s_i of a player i is a subset of all possible adjacent edges of player i , i.e., $s_i \subseteq \{\{i, j\} \mid j \in V \setminus \{i\}\}$. Chosen strategy s_i expresses that player i creates all edges in s_i . Each created edge costs a fixed amount $\alpha > 0$, which is a parameter of the game. The played (chosen) strategies $s = (s_1, \dots, s_n)$ collectively define an edge set $E_s = \bigcup_i s_i$, and, thus, a graph $G_s = (V, E_s)$. The *goal* of every player is to minimize its

Address correspondence to Akaki Mamageishvili, Department of Computer Science, ETH Universitatstrasse 6, Zurich, Zurich 8092, Switzerland. E-mail: akaki@inf.ethz.ch

Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/uinm.

cost $c_i(s)$, which is the amount paid for the edges (*creation cost*), plus the total distances of the player to every other node of the resulting network G (*usage cost*), i.e.,

$$c_i(s) := \alpha \cdot |s_i| + \sum_{j=1}^n d_s(i, j),$$

where $d_s(i, j)$ denotes the distance between i and j in the resulting network G_s . If the graph is disconnected and vertices i and j belong to different connected components, then the distance between i and j is infinity. We sometimes omit the index s from $d_s(i, j)$ if the graph G_s and the respective strategy profile s are clear from the context.

A *strategy vector* $s = (s_1, \dots, s_n)$ is a *Nash equilibrium* if no player i can change the set s_i of created edges to another set s'_i and lower its cost c_i . Abusing the definition, the resulting graph G_s itself is called a Nash equilibrium, too. The *social cost* $c(s)$ of strategy vector s is the sum of the individual costs, i.e., $c(s) := \sum_{i=1}^n c_i(s)$. It is a trivial observation to see that in any Nash equilibrium G_s , no edge is bought more than once. Also note that any Nash equilibrium graph is connected because, otherwise, the individual cost of each player would be infinity, and any player would be willing to buy new edges toward disconnected vertices and make his/her cost finite. From now on, we consider only such strategy vectors, and observe, then, that

$$c(s) := \sum_{i=1}^n c_i(s) = \alpha \cdot |E_s| + \sum_{i=1}^n \sum_{j=1}^n d_s(i, j).$$

Considering only such strategy vectors s , we define the *social cost of a graph* G_s , denoted by $c(G_s)$, as the cost $c(s)$, i.e., the cost of the corresponding strategy vector s .

A graph $G = (V, E)$ can be created by many strategy vectors s (precisely in $2^{|E|}$ many ways, because every edge in E can be bought by exactly one of its two endpoints), but each of such realizations has the same social cost.

In view of this, $c(G)$ (which is equal to $c(s)$ for any s such that $G = G_s$) can be expressed as

$$c(G) = \alpha \cdot |E| + \sum_{i=1}^n \sum_{j=1}^n d_G(i, j),$$

where $d_G(i, j)$ is the distance between i and j in G .

We call graph $G^* = (V, E)$ an *optimum graph*, if it minimizes the social cost $c(G)$.

Let \mathcal{N} denote the set of all Nash equilibria of a network creation game on n vertices and with an edge price α . The *price of anarchy* (PoA) of the network creation game is the ratio

$$\text{PoA} = \max_{s \in \mathcal{N}} \frac{c(G_s)}{c(G^*)}.$$

PoA expresses the (worst-case) loss of the quality of an optimum network (that the society could hypothetically build), if the society decides to create the network locally by self-interested individuals.

In a series of articles [9, 1, 6, 11] it has been shown that the price of anarchy of the network creation game is $O(1)$, i.e., a constant independent of both n and α , for every value $\alpha > 0$ with the exception of the range $n^{1-\varepsilon} < \alpha < 273 \cdot n$, where $\varepsilon = \Omega(\frac{1}{\log n})$. For the value of α with $n^{1-\varepsilon} < \alpha < 273 \cdot n$, an upper bound of $2^{\frac{2}{\sqrt{\log n}}}$ on the price of anarchy is

known (whereas no Nash equilibrium with a considerably large social cost is known). It is conjectured, however, that the price of anarchy is constant also in this range of α . It remains a major open problem to confirm or disprove this conjecture. It is certainly of interest to note that there are several variants of the network creation game (see, e.g., [2, 4, 8, 3]), but in none of these has the price of anarchy been shown to be constant; a study by [5], however, shows an exception to the other works.

Recently, a more detailed study [10] of the exact constant hidden in the $O(1)$ bound on the price of anarchy showed that the price of anarchy converges to 1 for any constant nonintegral α when n tends to infinity, and also provided a lower bound example of $\frac{3}{2}$ for integral α 's, which sounds intriguing and suggests that more needs to be learned about the structure of Nash equilibria.

Understanding the structure of Nash equilibria has proven to be important in bounding the price of anarchy. It was shown [9] that the social cost of any tree G in Nash equilibrium is upper bounded by $O(1) \cdot c(G^*)$. Therefore, the price of anarchy is $O(1)$ for all values of α for which every Nash equilibrium is a tree. It has been shown that every Nash equilibrium is a tree for all values of α greater than n^2 , $12n \log n$, and $273n$, respectively, in [9], [1], and [11]. It has been conjectured that every Nash equilibrium is a tree for every $\alpha \geq n$. Since for $\alpha = n/2$, nontree Nash equilibria are known, this *tree conjecture* is asymptotically tight.

In this study, we take steps to resolve the tree conjecture. We first tighten the tree conjecture and provide a construction of a nontree Nash equilibrium for every $\alpha = n - 3$ (thus, showing that, asymptotically, one cannot hope to show that every Nash equilibrium is a tree for some value $\alpha \geq q \cdot n$, $q < 1$ a fixed constant). We then apply a “linear-programming-like” approach to show that, for $\alpha \geq 65n$, every Nash equilibrium is a tree. To show this, we obtain new structural results on Nash equilibria and combine them with the previous approach of [11]. In the end, we take further steps toward the conjecture. We show that if $\alpha \geq n$, then there is no nontree Nash equilibrium containing exactly one cycle. We then apply the linear-programming-like approach again to show that the girth of every nontree Nash equilibrium (for any $\alpha \geq n$) is at least 8. Using the same ideas, we show that if a nontree Nash equilibrium has girth at most 10, then $\alpha \leq 1.3n$. By further experimental results, we conjecture that this holds for any girth, i.e., that nontree Nash equilibria can appear only for $\alpha \leq 1.3n$. To the end, we investigate the developed techniques for a coalitional variant of a Nash equilibrium, where coalitions of two players cannot collectively improve, and show that if $\alpha \geq 41 \cdot n$, then every such coalitional Nash equilibrium is a tree.

Throughout this article, we will use standard graph-theoretic terms/concepts such as a *path*, a *cycle*, a *subgraph*, a *bridge*, an *articulation*, etc. without defining them explicitly. We refer the reader to standard graph-theory textbooks for details, e.g., [7]. We will often denote the considered Nash equilibrium graph $G_s = (V, E_s)$ simply as $G = (V, E)$, if there is no danger of confusion. Even though the graph G_s is undirected, we will often direct every edge in G_s to express the identity of the player that bought the edge in the strategy vector s : An edge (u, v) directed from u to v denotes the fact that u bought/created the edge in s . It is certainly instructive to note that the orientation of edges in the topology is crucial and can have decisive influence on whether a given graph G_s is a Nash equilibrium. To illustrate this, consider the example in Figure 1. It presents two strategy vectors s and s' of the game with four players, resulting in the same graph $G_s = G_{s'}$ being a path. Interestingly, s is not a Nash equilibrium for any value of α ; s' is a Nash equilibrium for $\alpha = 2$.



Figure 1 A path on four vertices is realized by two strategy profiles: (a) s , which is not a Nash equilibrium, because vertex 1 can improve its cost by “swapping” the edge (1, 2) for edge (1, 3); and (b) s' , which, to the contrary, for $\alpha = 2$, (b) is a Nash equilibrium.

2. PRELIMINARIES

Every nontree graph G contains a cycle. Let C be a shortest cycle of G , let c be its length, and let a_0, a_1, \dots, a_{c-1} be the players that form C where $\{a_i, a_{i+1}\} \in E$ for every $i = 0, 1, \dots, c - 1$ (where indices on vertices of the cycle are, throughout the whole article, to be understood as modulo c). Observe the crucial property of the shortest cycle C : the distance between a_i and a_j in the graph G is equal to the distance between a_i and a_j on the cycle C . Moreover, for any vertex a_k , even if we delete one or both of the adjacent edges of a_k in the cycle C (and in the whole graph G), the distance between any two vertices a_i and a_j , other than a_k in the modified graph G , is the distance between the two vertices on the modified cycle C .

2.1. Main approach

If a Nash equilibrium $s = (s_1, \dots, s_n)$ is not a tree, then it contains a cycle. Let C be a shortest cycle in G_s . We will consider all the players on the cycle C and their strategy changes that involve only the c edges of the cycle. For each considered strategy change s'_{a_i} of player a_i , we obtain an inequality $c_{a_i}(s) \leq c_{a_i}(s_1, \dots, s'_{a_i}, \dots, s_n) =: c_{a_i}(s')$ stating simply the fact (definition) that in a Nash equilibrium s , player a_i cannot improve by changing its strategy. We will often express such an inequality in the form of “SAVINGS” \leq “INCREASE,” where “SAVINGS” denotes the parts (and amount) of $c_{a_i}(s)$ that decreased their value in $c_{a_i}(s')$, and “INCREASE” denotes the parts (and amount) of $c_{a_i}(s)$ that increased their value in $c_{a_i}(s')$. For example, assume that a_i buys the edge $e = (a_i, a_{i+1})$ (i.e., $e \in s_{a_i}$), and let us consider the strategy change where a_i deletes the edge e (i.e., $s'_{a_i} = s_{a_i} \setminus \{e\}$). Recall that $c_{a_i}(s) = \alpha \cdot |s_{a_i}| + \sum_j d(a_i, j)$. Then, in such a strategy change, the “SAVINGS” are clearly on the edge-creation side, i.e., the player a_i saves α for not paying for the edge e . At the same time, some distances of player a_i may have increased; the distance to a vertex v increases if in G_s every shortest path from a_i to v uses the deleted edge e . But the distance from a_i to v could have increased by at most $c - 2$ (because before, a_i needed to go to vertex a_{i+1} but now vertex a_{i+1} can be reached “around” the cycle, making a detour of at most $c - 1$). Because of the Nash equilibrium property of s , we have “SAVINGS” \leq “INCREASE,” which implies $\alpha \leq (c - 2)(n - 1)$ (as the distance to at most $n - 1$ vertices could have increased).

2.2. Partitioning of vertices

In the following, we will use slightly more involved forms of the previously described inequalities.

Definition 2.1. Strategy s'_{a_i} of player a_i of a shortest cycle C is called a *local strategy change* with respect to C and s , if

1. Every deleted edge $e \in s_{a_i} \setminus s'_{a_i}$ belongs to C .
2. There is at most one newly created edge e , i.e., $|s'_{a_i} \setminus s_{a_i}| \leq 1$, and this edge connects a_i with a vertex a_j of the cycle C .
3. a_i remains connected by at least one edge to a vertex a_j , $j \neq i$.

We are interested in the changes of the distances from a_i , $i = 0, \dots, c - 1$, i.e., in the values $\Delta_v := d_{s'}(a_i, v) - d_s(a_i, v)$, $v \in V$. For this reason, it does not matter what the exact distance $d_s(a_i, v)$ is, but rather what the relative difference is. This relative difference lies, as we will see, between 0 and $c - 1$ (because of the detour we can always make around the cycle C).

Let v be an arbitrary vertex $v \in V$. Let $G \setminus C$ be the graph G without the c edges of the cycle C . Let us denote the distances of v to the vertices a_0, a_1, \dots, a_{c-1} in $G \setminus C$ by the vector $d(v) = (d_0, d_1, \dots, d_{c-1})$, respectively, where $d_i = \infty$ if a_i and v are disconnected in $G \setminus C$. Observe that for at least one i , $d_i \neq \infty$. We call d_i the *outer distance* of v to a_i in the Nash equilibrium G , and d the *vector of outer distances* of v in G .

Proposition 2.2. For any nontree Nash equilibrium s with a shortest cycle C , for any $a_i \in C$ and any $v \in V$:

$$d_s(a_i, v) = \min_j (d_s(a_i, a_j) + d_j).$$

Furthermore, the distance $d_s(a_i, a_j)$ is achievable via the edges of C only.

Proof. The distance between a_i and v in G_s is equal to $\min_j (d_s(a_i, a_j) + d_j)$, because there always is a shortest path from a_i to v that first uses a part of the cycle C (until vertex a_j , where a_j can be a_i), leaves C and never comes back to C (because, recall, C is a shortest cycle of G). □

For a strategy profile $s' = (s_1, \dots, s'_i, \dots, s_n)$, where s'_i is a local strategy change with respect to s and C , we have a similar property about the distances $d_{s'}(a_i, v)$.

Proposition 2.3. For any nontree Nash equilibrium s with a shortest cycle C , for any $a_i \in C$ and any $v \in V$, for any local strategy change s'_i with respect to s and C :

$$d_{s'}(a_i, v) = \min_j (d_{s'}(a_i, a_j) + d_j).$$

Proof. Obviously, $d_{s'}(a_i, v) = \min_j (d_{s'}(a_i, a_j) + d_{s'}(a_j, v))$. Observe also, because s'_i is a local strategy change, that the distance d_j (i.e., the distance from a_j to v in $G_s \setminus C$) does not change in $G_{s'} \setminus C$, i.e., $d_j := d_{G_s \setminus C}(a_j, v) = d_{G_{s'} \setminus C}(a_j, v)$. Among all $j \in \arg \min_j (d_{s'}(a_i, a_j) + d_{s'}(a_j, v))$, there needs to exist one j^* for which a shortest path from a_{j^*} to v does not go via any of the a_0, \dots, a_{c-1} . For this j^* we have $d_{j^*} = d_{s'}(a_{j^*}, v)$ and the claim follows. □

As a consequence of the two propositions, we obtain:

$$\min_j d_j \leq d_s(a_i, v) \leq (c - 1) + \min_j d_j,$$

and

$$\min_j d_j \leq d_{s'}(a_i, v) \leq (c - 1) + \min_j d_j.$$

Because we are interested in the changes of the distances from a_i , i.e., in the values of $\Delta_v = d_{s'}(a_i, v) - d_s(a_i, v)$, $v \in V$, we can normalize the vector $d(v)$ by subtracting $\min_j d_j$ from each of the elements d_0, d_1, \dots, d_{c-1} (which, according to Proposition 2.2 and Proposition 2.3 does not change the value of Δ_v). Observe that after the normalization, there is an entry d_i equal to zero. We will “normalize” the entries more further on. Because we are interested in the value Δ_v , we can handle all entries $d_j \geq c - 1$ in the same way: they do not have any influence on Δ_v at all. Indeed, observe that there is no shortest path, neither in G_s nor in $G_{s'}$, from a_i to v that goes via a_j with $d_j > c - 1$, because the vertex a_{j^*} with $d_{j^*} = 0$ can be reached in at most $c - 1$ steps along the cycle C . Similarly, there always is a shortest path from a_i to v that does not go via a_j with $d_j = c - 1$. We will, therefore, further modify the vector d by substituting every entry $d_j \geq c - 1$ with the value $c - 1$ (with the meaning that a value $d_j = c - 1$ is actually useless for considering shortest paths from a_i to v).

This gives partition of all vertices into groups V_d , where each group has associated vectors of “normalized” outer distances $d = (d_0, \dots, d_{c-1})$, one of the distances is necessarily equal to 0 and all the distances are upper bounded by $c - 1$. In this way, there are $t = c^c - (c - 1)^c$ groups. We denote the set of all “normalized” distance vectors by D . Trivially, as $V_d, d \in D$, form a partition of V , $\sum_{d \in D} |V_d| = n$.

3. BOUNDS ON α FOR EXISTENCE OF CYCLES

We first give in Figure 2 a construction of a nontree Nash equilibrium graph for $n = 2q + 3$ vertices and $\alpha = 2q = n - 3$, for any integer q . There are three vertices z_1, z_2 , and z_3 in the middle. Each of these three vertices buys an edge toward L and an edge toward R . Vertices z_1, z_2 , and z_3 have the same creation cost, and they are symmetric, thus having the same usage cost. Furthermore, there are $q - 1$ vertices $\ell_2, \ell_3, \dots, \ell_q$, to each of which vertex L buys an edge. Similarly, there are $q - 1$ vertices r_2, r_3, \dots, r_q , to each of which vertex R buys an edge. The vertices l_i and $r_i, i = 2, \dots, q$, do not buy any edge. Thus, L and R are symmetric, having the same creation and usage costs. Similarly, all l_i are symmetric, and all r_i are symmetric, all having the same creation and usage cost. By an extensive search for improving strategies for each player type, we can check that no player wants to deviate from its strategy and therefore the graph is a Nash equilibrium. This thus shows that the conjecture “for $\alpha \geq n$, all Nash equilibria are trees” cannot be improved to “for $\alpha \geq (1 - \varepsilon)n$, all Nash equilibria are trees.”

Theorem 3.1. *For any integer q , there is a nontree Nash equilibrium on $n = 2q + 3$ vertices for $\alpha = n - 3$. In particular, for every constant ε , there is n and a nontree Nash equilibrium on n vertices for $\alpha = (1 - \varepsilon)n$.*




Figure 2 Nontree Nash equilibrium for $n = 2q + 3$ players and $\alpha = n - 3$. An edge directed from a node u to a node v denotes that u buys the edge.

We now proceed to show that for any α large enough, only tree Nash equilibria exist. We first derive an auxiliary claim providing a lower bound on the length of a shortest cycle in any Nash equilibrium.

Theorem 3.2. *The length c of a shortest cycle C in any Nash equilibrium is at least $\frac{2\alpha}{n} + 2$.*

Proof. We distinguish two cases. First, assume that there is a player, which buys both its adjacent edges on the cycle C . Without loss of generality assume that this player is a_0 . Consider the strategy change by which a_0 deletes both these edges (a_0, a_1) and (a_0, a_{c-1}) and buys an edge toward player a_i on the cycle, $i = 2, \dots, c - 2$. The player cannot improve by such a change, and therefore “SAVINGS” \leq “INCREASE.” Here, the player saves at least α (by buying one edge less). Let us denote the increase of distances of player a_0 to the players of the group V_d by $c_{i,d}$. Then, we get that $\alpha \leq \sum_{d \in D} \delta_{i,d} |V_d|$. Summing up all the $c - 3$ inequalities, one for every i , we get $(c - 3)\alpha \leq \sum_{i=2}^{c-2} \sum_d \delta_{i,d} |V_d|$.

We now show that for every d , the coefficient $\sum_i \delta_{i,d}$ at $|V_d|$ is at most $(c - 2)(c - 3)/2$. Consider arbitrary $d = (d_0, d_1, \dots, d_{c-1})$ of the outer distances of the vertices in V_d . Clearly, the strategy change of a_0 increases its distances to V_d iff every shortest path from a_0 to V_d goes through the deleted edges. Thus, we can assume (for the worst case with respect to the coefficient $\sum_d \delta_{i,d}$) that $d_0 = c - 1$. Assume that one shortest path (in G_s) leaves the cycle at a_e , $e \in \{1, \dots, c - 2\}$. In the new graph $G_{s'}$, player a_0 can always use the new edge (a_0, a_i) and then go to a_e for the remainder of the cycle C . Thus, the increase of distances $\delta_{i,d}$ is at most $(1 + |i - e|) - 1 = |i - e|$. In total, we obtain $\sum_{i=2}^{c-2} \delta_{i,d} \leq \sum_i |i - e| \leq \sum_i (i - 1) = (c - 3)(c - 2)/2$, as claimed. Now, since $\sum_{d \in D} |V_d| = n$, we finally get that $\alpha \leq \frac{(c-2)}{2}n$, which gives the claimed $c \geq \frac{2\alpha}{n} + 2$.

Consider now the second case in which no player buys two of its adjacent edges in C , i.e., every player buys exactly one edge. Without loss of generality assume that every player a_i buys the edge (a_i, a_{i+1}) . For each player i , we consider the strategy change of deleting the edge (a_i, a_{i+1}) . Similarly to the previous case, we obtain $\alpha \leq \sum_{d \in D} \delta_{i,d} |V_d|$. Summing for every i , we get $c\alpha \leq \sum_{i=0}^{c-1} \sum_d \delta_{i,d} |V_d|$. We show this time that $\sum_{i=0}^{c-1} \delta_{i,d}$, the coefficient at $|V_d|$, is upper bounded by $1 + 2 + \dots + (c - 2) = (c - 2)(c - 1)/2$. Consider an arbitrary $d = (d_0, \dots, d_{c-1}) \in D$, and assume without loss of generality that $d_0 = 0$. For every player a_i , $\delta_{i,d}$ is at most $i - 1$, because the worst-case increase in a distance of player a_i to vertices V_d happens when all shortest paths from a_i used the deleted edge (a_i, a_{i+1}) . But because, after the deletion, there is an alternative path from a_i to V_d using




Figure 3 The 5-Neighborhood $N_5(u)$ of vertex u .

a_0 , the increase is at most $i - 1$. Thus, summing over all i , the total increase in distances to V_d is at most $0 + 1 + 2 + \dots + (c - 2) = (c - 2)(c - 1)/2$ as claimed. Plugging this into our inequality, $c\alpha \leq \sum_i \sum_d \delta_{i,d} |V_d|$ and using the fact that $\sum_d |V_d| = n$, we obtain that $c > \frac{2\alpha}{n} + 2$. \square

Let H be a nontrivial biconnected component of a nontree Nash equilibrium, i.e., an induced subgraph of H of at least three vertices containing no articulation point (vertex removal of which leaves the graph disconnected). For any vertex $v \in H$, let $S(v)$ be the set of vertices that do not belong to H , and that have v as the closest vertex among all vertices in H . For any vertex $u \in H$, we define $deg_H(u)$ to be the degree of vertex u in the graph induced by H . Furthermore, we define $N_k(u)$ to be the k -th neighborhood of u in H , i.e., $N_k(u) := \{w \in H \mid d(u, w) \leq k\}$. The following lemma has been shown in [11]. We will use it to prove the subsequent lemma.

Lemma 3.3. (From [11]) *If $u, v \in V(H)$ are two vertices of a biconnected component H of a Nash equilibrium graph G with $d(u, v) \geq 3$ such that u buys the edge to its adjacent vertex x in a shortest $u - v$ -path and v buys the edge to its adjacent vertex y in that path, then $deg_H(x) \geq 3$ or $deg_H(y) \geq 3$.*

Lemma 3.4. *If H is a biconnected component of G , then for any vertex u , its neighborhood $N_5(u)$ in H contains a vertex v with $deg_H(v) \geq 3$.*

Proof. Assume that this is not true. Then the 5-neighborhood $N_5(u)$ of vertex u is formed by two disjoint paths. (The case that the 5-neighborhood forms a cycle is excluded by Theorem 3.2 for big enough α). We consider two cases. First, we will assume that at least one of the two paths starting at u is directed away from u (see Figure 3(a)). In the second case, in each of the two paths, there has to be a vertex that buys an edge toward u . It follows from Lemma 3 that these two vertices are the two neighbors of u in $N_5(u)$ (see Figure 3(b)).

In the first case, there is a sequence of five edges directed away from u , with the naming as in Figure 3(a)). Let $s_u := |S(u)|$, $s_i = |S(i)|$ for $0 \leq i \leq 4$. Then,

$$s_0 \geq s_1 + s_2 + s_3 + s_4, s_1 \geq s_2 + s_3 + s_4, s_2 \geq s_3 + s_4, s_3 \geq s_4, s_4 \geq k, \tag{3.1}$$

where k is the number of vertices that which are descendants of vertex 5 in the breadth-first-search (BFS) tree rooted at vertex 3. We can obtain these inequalities by considering the following strategy changes of the players u and i , $0 \leq i \leq 3$: delete the edge directed away from u and buy a new edge to the next vertex in the sequence; now simply apply the “SAVINGS” \leq “INCREASE” principle.

We first assume that vertex 5, the neighbor of vertex 4 in H , has degree at least 3 in H (i.e., it has at least two children in the BFS tree rooted at vertex 3). The case when the degree-3 vertex appears later along the path is easier and will be discussed later. We now distinguish two cases. First, we assume that one of the children of vertex 5 in the considered BFS tree buys an edge to vertex 5. Let us call it vertex 6. The other case is when vertex 5 buys all the edges to its children.

Consider the following strategy change: vertex 6 deletes an edge toward vertex 5 and buys a new edge toward vertex u . This decreases its distance cost, at least to vertices in $S(0)$ by 4, and to vertices in $S(1)$ by 2, while it increases distances to vertices in the set of descendants of 5 in the BFS tree rooted at 3 by at most 6, to the vertices in $S(4)$ by 4, and to the vertices in $S(3)$ by 2. By this strategy change, distance from vertex 6 to any other vertex is not increased, because vertex u is located deeper than vertex 6 in the BFS tree rooted at vertex 3. But then, according to the chain of inequalities (3.1) we get $4s_0 + 2s_1 > 6k + 4s_4 + 2s_3$, and thus, the player 6 can improve, which is a contradiction.

In the case where vertex 5 buys all edges toward its children, consider the following strategy change of vertex 5: delete all the edges to its children (in the considered BFS tree) and buy one edge to vertex u . By this, the “SAVINGS” are at least α . Furthermore, since H is biconnected, the graph remains connected. Distances from vertex 5 are increased only to vertices in the set K – the set of the vertices that are descendants of vertex 5 in the BFS tree rooted at vertex 3. This “INCREASE” is at most $2 \cdot \text{diam}(H)$, where $\text{diam}(H)$ is the diameter of H . By the “SAVINGS” \leq “INCREASE” principle, we get that $\alpha \leq 2 \cdot \text{diam}(H)k$. At the same time, $\alpha \geq (\text{rad}(H) - 1)s_0$, where $\text{rad}(H)$ is the radius of H , because otherwise, a vertex at distance $\text{rad}(H)$ from vertex 0 could buy an edge toward vertex 0 and decrease its cost. Combining these two inequalities with the inequality $s_0 \geq 8k$, which is obtained from (3.1), we get that $8(\text{rad}(H) - 1)k \leq 2 \cdot \text{diam}(H)k \leq 4 \cdot \text{rad}(H)k$, which is a contradiction.

The second case depicted in Figure 3(b) is analyzed in the very same way, the only change is that, according to a new chain of inequalities, now the heaviest component in the sense of the number of vertices is $S(u)$ (in the previous case it was $S(0)$). The chain of inequalities is similar to (3.1):

$$s_u \geq s_0 + s_1 + s_2 + s_3, s_1 \geq s_2 + s_3 + s_4, s_2 \geq s_3 + s_4, s_3 \geq s_4, s_4 \geq k, \quad (3.2)$$

where the notation is the same as in the first case. We obtain that $s_u \geq 7k$, and subsequently, arguing about the vertex at distance $\text{rad}(H)$ from u , we get a contradiction $7(\text{rad}(H) - 1)k \leq 2 \cdot \text{diam}(H)k \leq 4\text{rad}(H)k$.

Finally, if there is a longer sequence of vertices with degree 2 than the considered sequence of length 5 of edges directed away from u , then we can only consider the last 5 edges (all directed away from u) and apply the very same reasoning. \square

We can strengthen the result if we consider a stronger version of a Nash equilibrium in which no coalition of two players can change their strategies and improve their overall cost. We call such an equilibrium a *2-coalitional Nash equilibrium*.




Figure 4 The 3-neighborhood $N_3(u)$ of vertex u .

Lemma 3.3. *The 3-neighborhood $N_3(u)$ of any vertex u of a biconnected component H of a 2-coalitional Nash equilibrium has a vertex of degree at least 3.*

Proof. Assume the converse. Similarly to the proof of Lemma 3.4, there are two different cases of how the neighborhood of vertex u looks (see Figure 4(a) and (b); notation is also the same as that in Lemma 3.4). In both cases, consider the coalition of players 0 and 2. Consider the following strategy changes: player 0 deletes edge $(0, 1)$ and instead buys edge $(0, 3)$; player 2 deletes edge $(2, 3)$ and buys edge $(2, 0)$. This strategy change does not change the player coalition’s creation cost (in terms of α). Among the vertices $S(0)$, $S(1)$, $S(2)$, and $S(u)$ this strategy change decreases the coalition’s usage cost by $s_u + s_0 + s_2$ and it increases by s_1 . Other vertices are partitioned by their shortest distances to vertices 0 and 2; let’s assume that for any vertex v that does not belong to $S(0)$, $S(1)$, $S(2)$, or $S(u)$, shortest distance to vertex 0 is x and the shortest distance to vertex 2 is y . Obviously, $|x - y| \leq 2$. If $|x - y| > 0$, then there is no increase in the usage cost of coalition toward vertex v by this strategy change. The only possibility of increase is when $x = y$, but in that case v is the descendant of vertex 3 in the BFS tree rooted at vertex 1. Similarly to Lemma 3.4, we denote k to be the number of vertices that are descendants of vertex 3 in the BFS tree rooted at vertex 1. Analogously to the proof of Lemma 3.4, the following inequalities hold for the case depicted in Figure 4(a): $s_0 \geq s_1 + s_2$, $s_1 \geq s_2$ and $s_2 \geq k$, whereas, for the case depicted in Figure 4(b), we have $s_u \geq s_0 + s_1 + s_2$, $s_1 \geq s_2$, $s_2 \geq k$. In both cases, $s_u + s_0 + s_2 > s_1 + k$, which results in a contradiction. \square

The following two lemmas are crucial for proving the main result of the paper. The first lemma has been proven [11]. The second lemma strengthens a similar lemma from [11]. Its proof uses the result of Theorem 3.2.

Lemma 3.4. (From [11]) *If the t -neighborhood of every vertex of a biconnected component H of a Nash equilibrium contains a vertex of degree at least 3, then the average degree of H is at least $2 + \frac{1}{3t+1}$.*

Lemma 3.5. *If $\alpha > n$, then the average degree of a biconnected component H of a Nash equilibrium graph is at most $2 + \frac{4n}{\alpha-n}$.*

Proof. Among all vertices of the equilibrium graph G , consider a vertex with the smallest usage cost and let this vertex be v . Consider a BFS tree T rooted in v . Let $T' = T \cap H$. Then the average degree of H is $deg(H) = \frac{2|E(T')|+2|E(H)E(T')|}{|V(T')|} \leq 2 + \frac{2|E(H)E(T')|}{|V(T')|}$ (note that $|V(H)| = |V(T')|$). We now bound $|E(H)E(T')|$. We consider vertices that buy an edge in $E(H)E(T')$ and call them *shopping vertices*. It is easy to see that no shopping vertex buys more than one edge, because if any of them buys two or more edges, it is better for it to delete all of the edges and buy one new edge toward v : this decreases its creation cost by at least α , whereas increases its usage cost increase by at most n . It is, thus, enough to bound the number of shopping vertices. For this, we prove that the distance in the tree T' between any two shopping vertices is lower bounded by $\frac{\alpha-n}{n}$, which then implies that there cannot be too many shopping vertices. In other words, the number of shopping vertices is at most $\frac{2nV(T')}{\alpha-n}$. Assigning every node from H to the closest shopping vertex according to the distance in T' (breaking ties arbitrarily) forms a partition of H , wherein every part contains exactly one shopping vertex. As the distance in T' between shopping vertices is at least $\frac{\alpha-n}{n}$, the size of every part is at least $\frac{\alpha-n}{2n}$.

We assume for contradiction that there is a pair of shopping vertices u_1 and u_2 such that $d_{T'}(u_1, u_2) < \frac{\alpha-n}{n}$. Let $u_1 = x_1, \dots, x_k = u_2$ be the unique path from u_1 to u_2 in T' , and (u_1, v_1) and (u_2, v_2) be the edges bought by u_1 and u_2 in $E(H)E(T')$. Observe first that vertices v_1 and v_2 are not descendants of any vertex x_i , otherwise paths $v_j - x_i$ and $x_i - u_j$ together with an edge (u_j, v_j) form a cycle of length at most $2(d_{T'}(u_1, u_2) + 1) < \frac{2\alpha}{n} + 2$, which contradicts Theorem 3.2. Thus, $x_0 := v_1, x_1, \dots, x_k, x_{k+1} := v_2$ is a path. Since x_1 buys edge (x_0, x_1) , and x_k buys edge (x_k, x_{k+1}) , there is a vertex x_i such that x_i buys both of its adjacent edges (x_{i-1}, x_i) and (x_i, x_{i+1}) . Consider the following strategy change for player x_i : delete the two adjacent edges and buy a new edge to vertex v . In this way, x_i decreases its creation cost by α .

We now show that $U_{new}(x_i)$, the usage cost of x_i in the new graph, is less than $U_G(x_i)$, the usage cost in the original graph, plus α , which gives a contradiction. It is easy to observe that $U_{new}(x_i) \leq n + U_{new}(v)$, because x_i can always go through v in the new strategy to any vertex. We now consider $U_{new}(v)$. Note that only the vertices in the path $u_1 - u_2$ and their descendants can increase their distance to v by the strategy change of x_i . Let y be any such vertex. If the closest ancestor of y on the path is x_i , then $d_{new}(v, y) \leq d_G(v, y)$, so there is no increase. We assume, without loss of generality, that the closest ancestor of y is x_j with $j < i$. Then, the following chain of inequalities and equalities holds: $d_{new}(v, y) \leq d_{new}(v, x_0) + d_{new}(x_0, x_j) + d_{new}(x_j, y) = d_G(v, x_0) + d_G(x_0, x_j) + d_G(x_j, y)$ (the inequality is a triangle inequality, and the equality holds because x_0 is not a descendant of any vertex on the path in the new graph). Since $d_G(v, y) = d_G(v, x_j) + d_G(x_j, y)$, the difference between new and initial distances is $d_{new}(v, y) - d_G(v, y) = d_G(v, x_0) + d_G(x_0, x_j) - d_G(v, x_j) \leq 2d_G(x_0, x_j) \leq d_G(u_1, u_2) \leq 2 \cdot d_{T'}(u_1, u_2) \leq \frac{2(\alpha-n)}{n}$ (where the latter inequality is implied by our assumption).

We need to bound the number of possible y 's. Path $u_1 - u_2$ does not go through vertex v , so the number of possible y 's is bounded by the size of the subtree of T' of a child of v that contains this path. We prove that the size of any subtree of a child of v in T is at most $\frac{n}{2}$.

Consider any child t of v in T , and consider the subtree of T rooted in t . Let b be the number of vertices in this subtree, and let a be the number of other vertices of T . Let c_1 be the usage cost of t in the subtree, and let c_2 be the usage cost of v in the other part of the tree T ; see Figure 5. Then, the usage cost of t in G is upper bounded by $c_1 + a + c_2$,




Figure 5 Sizes of the subtrees.

whereas the usage cost of v is exactly $b + c_1 + c_2$. Since v is the vertex with the minimal usage cost, we have $c_1 + a + c_2 \geq b + c_1 + c_2$. Since $a + b = n$, we get that $b \leq \frac{n}{2}$.

Note that y was chosen arbitrarily, so the increase of the usage cost for v is less than $\frac{n}{2} \frac{2(\alpha-n)}{n} = \alpha - n$, and, therefore, $U_{\text{new}}(v) < U_G(v) + \alpha - n$, which is a contradiction. \square

Combining Lemmas 3.2 , 3.4 and 3.5 gives the main result.

Theorem 3.6. *For $\alpha \geq 65n$, every Nash equilibrium graph is a tree.*

Note that results obtained in Lemmas 3.4 and 3.5 hold for 2-coalitional Nash equilibrium graphs as well, so combining Lemmas 3.3, 3.4 and 3.5 gives the following result.

Theorem 3.7. *For $\alpha \geq 41n$, every 2-coalitional Nash equilibrium graph is a tree.*

4. SMALL CYCLES AND EXPERIMENTAL RESULTS

In this section, we consider equilibrium graphs that have small girth c , and show that they exist only for small values of α . We start with an observation that limits the girth of equilibrium graphs containing exactly one cycle.

Proposition 4.1. *Let G be a Nash equilibrium graph containing a k -cycle $C = \{v_0, v_1, \dots, v_{k-1}\}$, and F the graph where the edges of C are removed from G . If F consists of k connected components, then $k < 6$.*

Proof. Assume for contradiction that $k \geq 6$. For $0 \leq i < k$, let $s_i > 0$ denote the number of vertices in the connected component of F that contains v_i . If the edge (v_0, v_{k-1}) is bought by the player v_0 , then it could replace (v_0, v_{k-1}) by (v_0, v_{k-2}) . By doing this, the creation cost will remain the same, the distances to $s_{k-3} + s_{k-2}$ vertices decrease by 1, but the distances to s_{k-1} vertices increase by 1. If the edge (v_0, v_{k-1}) is bought by the player v_{k-1} , this player could replace (v_{k-1}, v_0) by (v_{k-1}, v_1) . By this change of strategy, the distances to s_0 vertices would increase, but distances could be decreased to $s_1 + s_2$ vertices.

Because we consider a Nash equilibrium, we deduce that $s_{k-3} + s_{k-2} \leq s_{k-1}$ or $s_0 \geq s_1 + s_2$. Applying this reasoning for every edge of C , we get that for every i ,

$$s_{i-3} + s_{i-2} \geq s_{i-1} \text{ or } s_i \geq s_{i+1} + s_{i+2}, \quad (4.1)$$

where $0 \leq i < k$ (recall that indexes are considered modulo c). The two inequalities $s_i \geq s_{i+1} + s_{i+2}$ and $s_{i-1} + s_i \leq s_{i+1}$ cannot hold simultaneously. Yet, (4.1) forces one of the inequalities $s_{i-1} + s_i \leq s_{i+1}$ or $s_{i+2} \geq s_{i+3} + s_{i+4}$ to be true, so we have that inequality $s_i \geq s_{i+1} + s_{i+2}$ implies $s_{i+2} \geq s_{i+3} + s_{i+4}$ for any $0 \leq i < k$. Without loss of generality, we can assume that the edge (v_{k-1}, v_0) was bought by v_0 . Then, we get the chain of inequalities $s_{2i} \geq s_{2i+1} + s_{2i+2}$ for every i , which is obviously a contradiction, because if we sum up all these inequalities we get $n \geq 2n$. \square

We now describe our computer-aided approach for upper bounding α in case of the existence of small cycles in Nash equilibrium graphs. In our approach, we consider a nontree Nash equilibrium whose smallest cycle has a fixed length c , and we construct a linear program asking for a maximum α , while satisfying inequalities of the type “SAVINGS” \leq “INCREASE,” which we create by considering various strategy changes of the players of the cycle. The partition of vertices of a Nash equilibrium graph into vertices V_d , $d \in D$, gives a variable $|V_d|$ for every d . The number of variables is $t = c^c - (c-1)^c$. We enumerate over all possible (meaningful) directions of the edges on the considered cycle and solve the linear program, which gives us an upper bounds on α for every direction of edges. The largest such value (among all directions) is then obviously an upper bound on α for any direction, and thus, for any Nash equilibrium containing a cycle of the fixed size.

The number of all possible directions is equal to 2^c , but this number can be decreased to at most $2^{c-3} + 2$ by simple observations that all hold without loss of generality. We can assume that the number of *right* edges is at least the number of *left* edges, where an edge (v_i, v_{i+1}) is called a *right* edge, and (v_{i+1}, v_i) is called a *left* edge. Furthermore, we can also assume that the edge (v_0, v_1) is a right edge. If c is even, every considered cycle can be made (by renaming arguments) to fall into one of the following three classes: (i) the edges along the cycle alternate between right and left, or (ii) all edges are right edges, or (iii) the first two edges are right edges and the last edge is a left edge. The same holds when c is odd, with the exception of the alternating edges.

Our linear program contains all inequalities implied by the strategy changes described in Theorem 3.2. We furthermore add inequalities for strategy changes of buying one extra edge, and for swapping an edge of the cycle with a new edge toward a vertex of the cycle. We add the equality $\sum_{d \in D} |V_d| = 1$ (which expresses the fact that the variables should sum to n). Then, the value of a variable $|V_d|$ expresses the fraction of all vertices (instead of the absolute number of vertices).

We used the Gurobi linear-programming solver to maximize α for every generated linear program. The largest such value, thus, gives an upper bound on α for which a cycle of size c can exist. Due to the huge number of variables, we could not solve the linear program

for $c > 7$, because already for $c = 8$, the number of variables was more than 10^7 , and while the number of constraints is $\Theta(c^2)$. We have made further tweaks to the code, which allowed us to speed up the computation. We observed that many variables had the same coefficients in every generated constraint, and thus, at most one such variable is relevant for obtaining the solution of the linear program. We have considered the variables one by one, and added only those having unique coefficients in the considered constraints. To check for uniqueness, we used hashing, because creating the matrix of the linear program was too slow. The obtained compression of the number of variables was huge: for $c = 10$, instead of nearly 10^{10} variables, we obtained only around 10^5 .

The obtained upper bounds on α are quite close to n . For girth $c \leq 7$, we obtain $\alpha \leq 1$, which corresponds to $\alpha \leq n$ if we required that $\sum_{d \in D} |V_d| = n$ (instead of $\sum_{d \in D} |V_d| = 1$). For girth $c = 8$, α is upper bounded by $\frac{191}{185}$ for only one orientation of the edges in the shortest cycle, for all the other directions, α is upper bounded by 1. Regarding the girth $c = 9$, α is upper bounded by $\frac{13}{12}$, whereas for girth $c = 10$, α is bounded by 1.2. It may seem that the upper bound on α grows with c . This, however, might not be the case, as the following experiments suggest.

We have performed further experiments with larger values of c , but did not consider all orientations of edges (this was out of our computational power). Furthermore, since the number of variables is increasing superexponentially, instead of considering all variables, for larger values of c we have considered only variables $|V_d|$ that have only 0's and $(c - 1)$'s as distances in vector d , that is, we have considered 2^c variables. Additionally, we have taken extra 2^c random variables. We have checked all values of c up to 15. Upper bounds for α obtained using only these variables are very close to the real bounds for $c \leq 10$ (the difference for $k \leq 10$ is between 0 and 0.01). The largest upper bound of $1.3n$ on α appears for $c = 13$, and then only decreases, which is why we conjecture: the upperbound of $\alpha \leq 1.3n$ can be proved by the considered strategy changes.

Constructing lower bound examples is easy for $c = 3$ and $c = 4$, but already complicated for $c = 5$; we managed to provide an example for one direction of edges on the minimal cycle, which matches with the upper bound for α . Possible explanation is that we care only about fractional parts of n in our linear program, but for the real lower bound examples, even constants matter, which are assumed to be 0 by our linear program. Regarding the minimal cycles with length at least 6, there is no example of a Nash equilibrium graph having these, which tempts us to conjecture that it is impossible to construct such an example.

5. CONCLUSIONS

In this article, we have described steps toward resolving the conjecture “Every Nash equilibrium of a Network Creation Game with $\alpha \geq n$ is a tree.” Previously, it has been known that for every $\alpha \geq 273 \cdot n$, every Nash equilibrium is a tree. We have improved this bound and showed that for every $\alpha \geq 65 \cdot n$, every Nash equilibrium is a tree. At the same time, we have provided a construction of a nontree Nash equilibrium for every $\alpha = n - 3$. Previously, nontree Nash equilibria were known only for $\alpha \leq n/2$. We have also considered a specific collaborative notion of a Nash equilibrium, and for these we have shown slightly improved bounds on α for the existence of nontree Nash equilibria by a straightforward adaptation of the techniques.

To gain an insight into the structure of the problem, we have considered Nash equilibria with small cycles. We have shown that for $\alpha > n$, there must be more than

one cycle of size at most 6. Furthermore, using LP-solvers, we showed that for $\alpha > 1.2n$, the girth of a nontree Nash equilibrium is at least 11. To go beyond girth 10, we have considered only several topologies and simplified the constraints by stripping off some of the variables. These experiments suggest that with the increasing girth, the upper bound on α that allows nontree Nash equilibria is not increasing. This, as well as the main considered conjecture, remains to be formally proven.

FUNDING

This work has been partially supported by the Swiss National Science Foundation (SNF) under the grant number 200021_143323/1.

REFERENCES

- [1] S. Albers, S. Eilts, E. Even-Dar, Y. Mansour, and L. Roditty. “On Nash Equilibria for a Network Creation Game.” In *Proc. 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 89–98. New York, NY, USA: ACM, 2006.
- [2] N. Alon, E. D. Demaine, M. T. Hajiaghayi, and T. Leighton. “Basic Network Creation Games.” *SIAM Journal on Discrete Mathematics* 27:2 (2013), 656–668.
- [3] D. Bilò, L. Gualà, and G. Proietti. “Bounded-Distance Network Creation Games.” In *Proc. 8th International Workshop on Internet and Network Economics (WINE)*, pp. 72–85. Springer, 2012.
- [4] M. Brautbar and M. Kearns. “A Clustering Coefficient Network Formation Game.” In *Proc. Fourth International Symposium on Algorithmic Game Theory (SAGT)*, pp. 224–235. Berlin, Heidelberg: Springer, 2011.
- [5] E. Demaine and M. Zadimoghaddam. “Constant Price of Anarchy in Network Creation Games via Public Service Advertising.” In *Proc. Seventh International Workshop on Algorithms and Models for the Web-Graph (WAW)*, pp. 122–131. Berlin, Heidelberg: Springer, 2010.
- [6] E. D. Demaine, M. Hajiaghayi, H. Mahini, and M. Zadimoghaddam. “The Price of Anarchy in Network Creation Games.” *ACM Trans. Algorithms*, 8:2 (2012), 1–13.
- [7] R. Diestel. *Graph Theory*. Graduate Texts in Mathematics, 4th edition, 173. Berlin: Springer, 2010.
- [8] S. Ehsani, M. Fazli, A. Mehrabian, S. S. Sadeghabad, M. Safari, M. Saghafian, and S. S. Fadaee. “On a Bounded Budget Network Creation Game.” In *Proc. 23rd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pp. 207–214. New York, NY: ACM, 2011.
- [9] A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. Shenker. “On a Network Creation Game.” In *Proc. 22nd Annual Symposium on Principles of Distributed Computing (PODC)*, pp. 347–351. New York, NY, USA: ACM, 2003.
- [10] R. L. Graham, L. Hamilton, A. Levavi, and P. S. Loh. “Anarchy is Free in Network Creation.” In *Proc. 10th International Workshop on Algorithms and Models for the Web Graph (WAW)*, volume 8305 of *LNCS*, pp. 220–231. Berlin, Heidelberg: Springer, 2013.
- [11] M. Mihalák and J. C. Schlegel. “The Price of Anarchy in Network Creation Games is (Mostly) Constant.” *Theory Comput. Syst.* 53:1 (2013), 53–72.