

Anarchy is free in network creation

Ronald Graham* Linus Hamilton† Ariel Levavi‡ Po-Shen Loh§

Abstract

The Internet has emerged as perhaps the most important network in modern computing, but rather miraculously, it was created through the individual actions of a multitude of agents rather than by a central planning authority. This motivates the game theoretic study of network formation, and our paper considers one of the most-well studied models, originally proposed by Fabrikant et al. In it, each of n agents corresponds to a vertex, which can create edges to other vertices at a cost of α each, for some parameter α . Every edge can be freely used by every vertex, regardless of who paid the creation cost. To reflect the desire to be close to other vertices, each agent’s cost function is further augmented by the sum total of all (graph theoretic) distances to all other vertices.

Previous research proved that for many regimes of the (α, n) parameter space, the total social cost (sum of all agents’ costs) of every Nash equilibrium is bounded by at most a constant multiple of the optimal social cost. In algorithmic game theoretic nomenclature, this approximation ratio is called the price of anarchy. In our paper, we significantly sharpen some of those results, proving that for all constant non-integral $\alpha > 2$, the price of anarchy is in fact $1 + o(1)$, i.e., not only is it bounded by a constant, but it tends to 1 as $n \rightarrow \infty$. For constant integral $\alpha \geq 2$, we show that the price of anarchy is bounded away from 1. We provide quantitative estimates on the rates of convergence for both results.

1 Introduction

Networks are of fundamental importance in modern computing, and substantial research has been invested in network design and optimization. However, one of the most significant networks, the Internet, was not created “top-down” by a central planning authority. Instead, it was constructed through the cumulative actions of countless agents, many of whom built connections to optimize their individual objectives. To understand the dynamics of the resulting system, and to answer the

*Department of Computer Science and Engineering, University of California, San Diego, La Jolla, CA 92093. E-mail: graham@ucsd.edu.

†Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA 15213. E-mail: luh@andrew.cmu.edu. Research supported by NSF grant DMS-1041500.

‡Department of Computer Science and Engineering, University of California, San Diego, La Jolla, CA 92093. E-mail: alevavi@cs.ucsd.edu.

§Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA 15213. E-mail: ploh@cmu.edu. Research supported by NSF grants DMS-1201380 and DMS-1041500, an NSA Young Investigators Grant and a USA-Israel BSF Grant.

important question of how much inefficiency is introduced through the selfish actions of the agents, it is therefore natural to study it through the lens of game theory.

In this paper, we focus on a well-studied game-theoretic model of network creation, which was formulated by Fabrikant et al. in [10]. There are n agents, each corresponding to a vertex. They form a network by laying down connections (edges) between pairs of vertices. For this, each agent v has an individual strategy, which consists of a subset S_v of the rest of the vertices that it will connect to. The resulting network is the disjoint union of all (undirected) edges between vertices v and vertices in their S_v . Note that in this formulation, an edge may appear twice, if v lays a connection to w and w lays a connection to v . Let α be an arbitrary parameter, which represents the cost of making a connection. In order to incorporate each agent's desire to be near other vertices, the total cost to each agent is defined to be:

$$\text{cost}(v) = \alpha|S_v| + \sum_w \text{dist}(v, w),$$

where the sum is over all vertices in the graph, and $\text{dist}(v, w)$ is the number of edges in the shortest path between v and w in the graph, or infinity if v and w are disconnected. The *social cost* is defined as the total of the individual costs incurred by each agent. This cost function summarizes the fact that v must pay the construction cost for the connections that it initiates, but v also prefers to be graph-theoretically close to the other nodes in the network. This model also encapsulates the fact that, just as in the Internet, once a connection is made, it can be shared by all agents regardless of who paid the construction cost.

The application of approaches from algorithmic game theory to the study of networks is not new. The works [5, 6, 8, 13, 14, 15] all consider network design issues such as load balancing, routing, etc. Numerous papers, including [1, 2, 3, 4, 7, 9] and the surveys [11, 16], have considered network formation itself, by formulating and studying network creation games. From a game-theoretic perspective, a (pure) *Nash equilibrium* is a tuple of deterministic strategies S_v (one per agent) under which no individual agent can strictly reduce its cost by unilaterally changing its strategy assuming all other agents maintain their strategies. If every unilateral deviation strictly increases the deviating agent's cost, then the Nash equilibrium is *strict*.

To quantify the cumulative losses incurred by the lack of coordination, the key ratio is called the *price of anarchy*, a term coined by Koutsoupias and Papadimitriou [12]. It is defined as the maximum social cost incurred by any Nash equilibrium, divided by the minimum possible social cost incurred by any tuple of strategies. Note that the minimizer, also known as the *social optimum*, is not necessarily a Nash equilibrium itself. The central questions in this area are thus to understand the price of anarchy, and to characterize the Nash equilibria.

1.1 Previous work

To streamline our discussion, we will represent a tuple of strategies with a directed graph, whose underlying undirected graph is the resulting network, and where each edge vw is oriented from v to w if it was constructed by v 's strategy ($w \in S_v$). This is well-defined because it is clear that the social optimum and all Nash equilibria will avoid multiple edges, and so each edge is either not present at all, or present with a single orientation.

The problem is trivial for $\alpha < 1$, because all Nash equilibria produce complete graphs, as does the social optimum, and therefore the price of anarchy is 1 in this range. For $\alpha \geq 1$, a new Nash equilibrium arises: the star with all edges oriented away from the central vertex. Indeed, the central vertex has no incentive to disconnect any of the edges which it constructed, as its individual cost function would rise to infinity, and no other vertex has incentive to add more connections, because a new connection would cost an additional $\alpha \geq 1$, and reduce at most one of the pairwise distances by 1. Yet, as observed in the original paper of Fabrikant et al. [10], when $\alpha < 2$, the social optimum is a clique, and they calculate the price of anarchy to be $\frac{4}{2+\alpha} + o(1)$, where the error term tends to 0 as $n \rightarrow \infty$. This ranges from $\frac{4}{3}$ to 1 as α varies in that interval.

For $\alpha \geq 2$, the social optimum is the star. Various bounds on the price of anarchy were achieved, with particular interest in constant bounds, which were derived in many ranges of the parameter space. From the point of view of approximation algorithms, these show that in those ranges of α , the Nash equilibria that arise from the framework of selfish agents still are able to approximate the optimal social cost to within a constant factor. The current best bounds are summarized in Table 1.

Regime	Upper bound on price of anarchy
General α	$2^{O(\sqrt{\log n})}$
$2 \leq \alpha < \sqrt[3]{n/2}$	4
$\sqrt[3]{n/2} \leq \alpha < \sqrt{n/2}$	6
$\alpha = O(n^{1-\epsilon})$	$O(1)$
$\alpha \geq 12n \lceil \log_2 n \rceil$	$O(1)$

Table 1: Previous upper bounds on the price of anarchy. The last bound above is due to Albers et al. [2], and the other bounds are due to Demaine et al. [9].

1.2 Our contribution

Much work had been done to achieve constant upper bounds on the price of anarchy in various regimes of α , because those imply the satisfying conclusion that selfish agents fare at most a constant factor worse than optimally coordinated agents. Perhaps surprisingly (or perhaps reassuringly), it turns out that the price of anarchy is actually $1 + o(1)$ for most constant values of α . In other words, the lack of coordination has negligible effect on the social cost as n grows.

Theorem 1.1. *For non-integral $\alpha > 2$, and $n > \alpha^3$, the price of anarchy is at most*

$$1 + \frac{150\alpha^6}{(\alpha - \lfloor \alpha \rfloor)^2} \sqrt{\frac{\log n}{n}} = 1 + o(1).$$

On the other hand, for each integer $\alpha \geq 2$, the price of anarchy is at least

$$\frac{3}{2} - \frac{3}{4\alpha} + \frac{1}{\alpha^2} + o(1),$$

and it is achieved by the following construction. Start with an arbitrary orientation of the complete graph on k vertices. For each vertex v of the complete graph, add $\alpha - 1$ new vertices, each with a single edge oriented from v .

2 Proof for non-integral α

Assume that we are given a Nash equilibrium. In this section, we prove that its total social cost is bounded by $1 + o(1)$ times the social optimum, as stated in Theorem 1.1. Throughout this proof, we impose a structure on the graph as follows: select a vertex v , and partition the remainder of the graph into sets based on their distance from v . Let N_1 denote the set of vertices at distance 1 from v , let N_2 denote the set of vertices at distance 2 from v , etc., as diagrammed in Figure 1. Since the graph in every Nash equilibrium is obviously connected, every vertex falls into one of these sets.

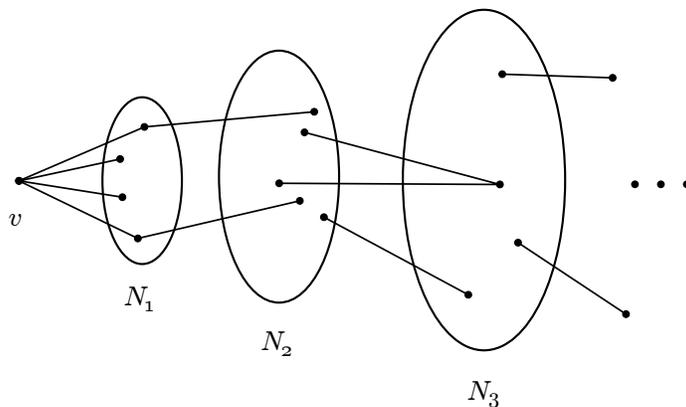


Figure 1: Partitioning the graph into sets.

Consider any vertex $v_i \in N_i$ where $i \geq 3$. Since the graph is connected, we can always find a path $v_i v_{i-1} v_{i-2} \dots v_2 v_1 v$, where $v_j \in N_j$ for all $1 \leq j \leq i$. In this case, we will call v_i a *child* of v_2 . (Note that v_i may be a child of more than one vertex, but is always a child of at least one vertex.) This is diagrammed in Figure 2.

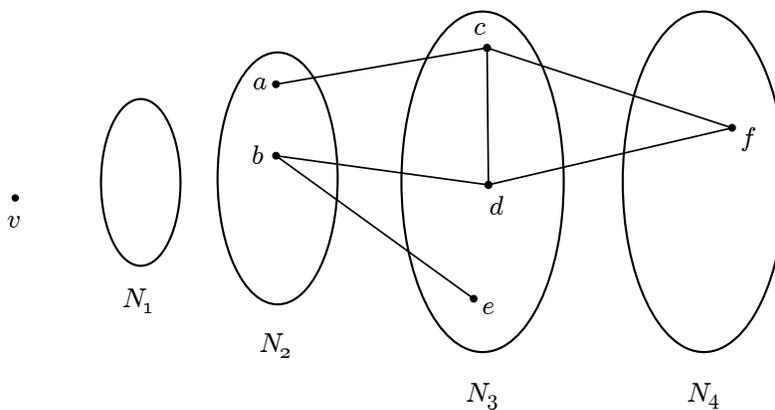


Figure 2: Here, d and e are children of b ; c is a child of a , but not a child of b ; and f is a child of both a and b .

Lemma 2.1. *No matter which vertex is used as v to construct the vertex partition, every vertex in N_2 has at most $\lfloor \alpha - 1 \rfloor$ children.*

Proof. Suppose $w \in N_2$ has more than $\alpha - 1$ children. Consider what happens if v buys an edge to w . Although v pays α for the edge, it gets one step closer to w and all of its children, and so the distance component of v 's cost function reduces by more than $1 + (\alpha - 1) = \alpha$. Therefore, buying the edge is a net positive gain for v . But we assumed the graph was a Nash equilibrium—contradiction. Therefore, w has at most $\alpha - 1$ children, and since its number of children is an integer, we may round the bound down as in the statement of the lemma. \square

Lemma 2.2. *Regardless of the choice of v , the resulting parts N_i satisfy:*

$$|N_1| + |N_2| + 1 \geq \frac{n}{\alpha}.$$

Proof. Since every vertex in $N_3 \cup N_4 \cup \dots$ is a child of at least one vertex of N_2 , but Lemma 2.1 bounds the number of children per N_2 -vertex by $\alpha - 1$, we must have

$$\begin{aligned} (\alpha - 1)|N_2| &\geq |N_3 \cup N_4 \cup \dots| = (n - 1 - |N_1| - |N_2|) \\ \alpha|N_2| + |N_1| + 1 &\geq n, \end{aligned}$$

which implies the desired result. \square

Lemma 2.3. *If x has degree at least α , then every vertex is at most distance 3 from it.*

Proof. If some vertex w is distance at least 4 from x , then w can buy an edge to x . Vertex w will pay α for the edge, and get 3 steps closer to x , as well as at least 1 step closer to all of x 's immediate neighbors, for a net gain. Hence this cannot appear in a Nash equilibrium. \square

Corollary 2.1. *If n is sufficiently large ($n > \alpha^3$), then the graph has diameter at most 4.*

Proof. Consider an arbitrary pair of vertices v, w . Lemma 2.2 implies that for n sufficiently large ($n > \alpha^3$ suffices), either v has degree at least α , or one of v 's neighbors has degree at least α . In either case, we can travel from v to a vertex with degree at least α in at most one step, and then by Lemma 2.3, travel to w in at most 3 more steps. Therefore, v and w are at distance at most 4. \square

Remark. From now on, we will assume $n > \alpha^3$, and so for any initial choice of v , the resulting partition will only have N_1, N_2, N_3 , and N_4 .

Lemma 2.4. *Consider the partition constructed from an arbitrary initial vertex v . Select any $w \in N_2$, and let d be the number of edges w pays for which connect to other vertices in N_2 . Then $d \leq |N_1| \cdot \frac{\alpha}{\alpha - \lfloor \alpha \rfloor}$.*

Proof. Consider the following strategy for w : disconnect those d edges, and instead connect to every vertex in N_1 . We will carefully tally up the potential gain for this amendment.

- *Paying for edges:* w **saves at least $(d - |N_1|)\alpha$** in terms of paying for edges. (The “at least” is because w might already be connected to some vertices in N_1 .)
- *Connectedness to v and N_1 :* w obviously can’t get farther away from v or any vertices in N_1 .
- *Connectedness within N_2 :* w gets farther away from all d vertices it disconnected from, but remains at distance 2 from all of N_2 , since every vertex in N_2 is connected to some vertex in N_1 . This results in a maximum **increased cost of d** in terms of distances to other vertices within N_2 .
- *Connectedness to N_3 and N_4 :* When disconnecting from a vertex $x \in N_2$, w might get farther away from all of x ’s children in N_3 and N_4 . However, remember that w is still distance 2 from all of N_2 . Hence, w is still distance 3 from all of N_3 and distance 4 from all of N_4 . Therefore, w can only get 1 step farther from x ’s children, and doesn’t get any farther from vertices in N_3 and N_4 that aren’t x ’s children. By Lemma 2.1, every N_2 -vertex has at most $\lfloor \alpha - 1 \rfloor$ children. Therefore, in disconnecting from d vertices, w gets 1 step farther from at most $d\lfloor \alpha - 1 \rfloor$ vertices in N_3 and N_4 , for a **cost increase of at most $d\lfloor \alpha - 1 \rfloor$** .

Adding, w ’s net cost savings total to at least $(d - |N_1|)\alpha - d - d\lfloor \alpha - 1 \rfloor$, which must be ≤ 0 since we are at a Nash equilibrium. Rearranging, $d \leq |N_1| \cdot \frac{\alpha}{\alpha - \lfloor \alpha \rfloor}$, as desired. \square

Lemma 2.5. *If $|N_1|$ is $o(n)$, then so is $|N_3 \cup N_4|$. Quantitatively, $|N_3 \cup N_4| < |N_1| \cdot \frac{5\alpha^3}{\alpha - \lfloor \alpha \rfloor}$.*

Proof. Let P be the number of pairs of vertices (x, y) , such that $x \in N_3 \cup N_4$ and y is at most distance 2 from x . We will bound this number in two ways. First, Lemma 2.2 tells us that for any vertex in the graph, the number of vertices at most distance 2 from it is at least $\frac{n}{\alpha}$. Therefore, $P \geq |N_3 \cup N_4| \cdot \frac{n}{\alpha}$.

For the second way, we will find an upper bound for the number of ways to start at a vertex $x \in N_3 \cup N_4$, and then travel along at most two edges in some way. This is an overcount for P , so it will give an upper bound. To count the number of these paths, we do casework on the various ways to start at a vertex in $N_3 \cup N_4$ and then travel along at most two edges.

Case 1: The path stays inside $N_3 \cup N_4$. Any vertex in $N_3 \cup N_4$ can be connected to at most $\alpha - 1$ other vertices in $N_3 \cup N_4$ (otherwise v would gain from connecting to it directly), so the number of paths for us to count for each starting vertex is at most $1 + (\alpha - 1) + (\alpha - 1)^2 \leq \alpha^2$. Therefore, the total number of paths of this type is at most $|N_3 \cup N_4|\alpha^2$.

Case 2: The path travels from $N_3 \cup N_4$ to N_3 to N_2 , or is a length 1 path traveling from N_3 to N_2 . We count these backwards, starting from N_2 . The number of edges from N_2 to N_3 is at most $\alpha|N_2|$ by Lemma 2.1, and again, every vertex in N_3 is connected to at most α vertices in $|N_3 \cup N_4|$, if including itself. Therefore, the number of paths here is at most $\alpha^2|N_2|$.

Case 3: The path travels from N_3 to N_2 to N_3 . We can count these by looking at the vertex in N_2 first, and then picking 2 of its children in N_3 . Thus, the number of such paths is at most $|N_2|\alpha^2$.

Case 4: The path travels from N_3 to N_2 to N_1 . Similarly to Case 2, the number of such paths is at most $\alpha|N_2||N_1|$.

Case 5: The path travels from N_3 to N_2 to N_2 . By Lemma 2.4, the number of edges inside N_2 is at most $|N_2||N_1|\frac{\alpha}{\alpha-\lfloor\alpha\rfloor}$. Each such path consists of one of these edges, together with an edge to N_3 from one of its two endpoints. Therefore, the number of paths for us to count is at most $|N_2||N_1|\frac{2\alpha^2}{\alpha-\lfloor\alpha\rfloor}$.

Total: summing over all cases, we have:

$$\begin{aligned} P &\leq |N_3 \cup N_4|\alpha^2 + 2\alpha^2|N_2| + |N_1||N_2|\alpha + |N_1||N_2|\frac{2\alpha^2}{\alpha-\lfloor\alpha\rfloor} \\ &< 2\alpha^2n + |N_1|n\left(\alpha + \frac{2\alpha^2}{\alpha-\lfloor\alpha\rfloor}\right) \\ &< |N_1|n\left(\frac{5\alpha^2}{\alpha-\lfloor\alpha\rfloor}\right). \end{aligned}$$

But $P \geq |N_3 \cup N_4|\frac{n}{\alpha}$ from above, so:

$$\begin{aligned} |N_3 \cup N_4|\frac{n}{\alpha} &< |N_1|n\left(\frac{5\alpha^2}{\alpha-\lfloor\alpha\rfloor}\right) \\ |N_3 \cup N_4| &< |N_1|\cdot\frac{5\alpha^3}{\alpha-\lfloor\alpha\rfloor}. \end{aligned}$$

□

Lemma 2.6. *If every vertex has degree $> \sqrt{n \log n}$, then the graph is asymptotically socially optimal: the total social cost is at most $2n^2 + \alpha n^{3/2}\sqrt{\log n}$.*

Proof. Suppose we have a Nash equilibrium where all vertices have degree greater than $\sqrt{n \log n}$. We give a strategy for an arbitrary vertex to achieve an individual cost of at most $\alpha\sqrt{n \log n} + 2n$, by changing only its own behavior. Since this is a Nash equilibrium, we will then be able to conclude that every vertex must have had individual cost at most $\alpha\sqrt{n \log n} + 2n$, proving this claim.

Specifically, we show that for any vertex w , the strategy “undo all edges you’re currently paying for, and connect to $\sqrt{n \log n}$ vertices at random” has a positive probability of bringing it within distance ≤ 2 from every other vertex in the graph. Indeed, if w does this, then for any other vertex x ,

$$\begin{aligned} &\mathbb{P}[x \text{ is now distance } > 2 \text{ from } w] \\ &\leq \mathbb{P}[w \text{ didn't choose any of } x\text{'s neighbors}] \\ &\leq \left(1 - \frac{\sqrt{n \log n}}{n}\right)^{\sqrt{n \log n}} \leq e^{-\log n} = \frac{1}{n}. \end{aligned}$$

Since there are only $n-1$ other vertices $x \neq w$ to consider, a union bound shows that the probability of failure is at most $(n-1)\frac{1}{n} < 1$, and therefore there is a way for w to attain an individual cost of at most $\alpha\sqrt{n \log n} + 2n$, as desired. □

Lemma 2.7. *Even if there is a vertex of degree at most $\sqrt{n \log n}$, the graph is still asymptotically socially optimal: the total social cost is at most $2n^2 + n^{3/2} \sqrt{\log n} \cdot \frac{290\alpha^6}{(\alpha - \lfloor \alpha \rfloor)^2}$.*

Proof. Let v be a vertex of degree at most $\sqrt{n \log n}$, and construct the vertex partition N_1, N_2, N_3, N_4 . We already know $|N_1|$ is at most $\sqrt{n \log n} = o(n)$, so by Lemma 2.5, $|N_3 \cup N_4|$ is at most $\sqrt{n \log n} \cdot \frac{5\alpha^3}{\alpha - \lfloor \alpha \rfloor} = o(n)$. By Lemma 2.4, the total number of edges inside N_2 is at most $|N_2| |N_1| \frac{\alpha}{\alpha - \lfloor \alpha \rfloor} \leq n^{3/2} \sqrt{\log n} \cdot \frac{\alpha}{\alpha - \lfloor \alpha \rfloor} = o(n^2)$. Also, the total number of edges not completely inside N_2 is at most $n \cdot (1 + |N_1| + |N_3 \cup N_4|) \leq n^{3/2} \sqrt{\log n} \cdot \frac{6\alpha^3}{\alpha - \lfloor \alpha \rfloor} = o(n^2)$. Therefore, the total number of edges in the whole graph is at most $n^{3/2} \sqrt{\log n} \cdot \frac{7\alpha^3}{\alpha - \lfloor \alpha \rfloor} = o(n^2)$.

Next, we calculate a bound on the total sum of distances in the graph. Using Lemma 2.5 on every vertex in the graph, and the fact that all distances are at most 4 (Corollary 2.1), we get:

$$\begin{aligned} & \text{[total sum of distances in the graph]} \\ & \leq 2n^2 + 4[\# \text{ of distances in the graph that are 3 or 4}] \\ & = 2n^2 + 4 \sum_w [\# \text{ of vertices at distance 3 or 4 from } w] \\ & < 2n^2 + 4 \sum_w \deg(w) \cdot \frac{5\alpha^3}{\alpha - \lfloor \alpha \rfloor}. \end{aligned}$$

The degree sum is precisely twice the total number of edges in the graph, a quantity which we just bounded above. Putting everything together, the total sum of distances is at most:

$$2n^2 + 8n^{3/2} \sqrt{\log n} \cdot \frac{7\alpha^3}{\alpha - \lfloor \alpha \rfloor} \cdot \frac{5\alpha^3}{\alpha - \lfloor \alpha \rfloor} = 2n^2 + n^{3/2} \sqrt{\log n} \cdot \frac{280\alpha^6}{(\alpha - \lfloor \alpha \rfloor)^2}.$$

Adding α times the number of edges to compute the total social cost, we obtain the desired bound. \square

Lemmas 2.6 and 2.7 cover complementary cases, so we now conclude that the total social cost of every Nash equilibrium is at most the bound obtained in Lemma 2.7. As was observed by previous authors [10], the social optimum for $\alpha \geq 2$ is the star, achieving a social cost of at least $2n(n-1)$. Dividing, we find that the price of anarchy is at most

$$1 + \frac{150\alpha^6}{(\alpha - \lfloor \alpha \rfloor)^2} \sqrt{\frac{\log n}{n}} = 1 + o(1),$$

proving the first part of Theorem 1.1.

3 Integral α

There is one catch in our bound above. Namely, when α is only slightly greater than an integer (e.g. 4.0001), the terms of the form $\frac{1}{\alpha - \lfloor \alpha \rfloor}$ all blow up, giving the final $o(n^2)$ terms for our bound a large constant factor. Even worse, when α is an exact integer, the proof fails completely. Perhaps surprisingly, this is not an artifact of the proof. In this section, we construct a counterexample

when α is an integer. Let v_1, v_2, \dots, v_k be a large clique with edges oriented arbitrarily. In addition, each vertex v_i in the clique also pays for edges to $\alpha - 1$ separate leaves $l_{i:1}, l_{i:2}, \dots, l_{i:\alpha-1}$. This graph also appears in [2], but we provide a full analysis here for completeness.

Lemma 3.1. *In this graph, no single vertex has a better strategy than the one it's currently using.*

Proof. First, consider any leaf, say $l_{1:1}$. This leaf is not currently paying for any edges, so its only option is to pay for some set of edges. Notice that purely choosing some set of edges to pay for, without being able to delete any edges, is an instance of convex optimization. Therefore, by convexity, if any vertex in any graph can improve its station purely by adding some set of edges S , then it can also do this by adding some single edge $s \in S$. By observation, the leaf can only break even by adding one edge, so it can only break even overall.

Next, consider a members of the clique, say v_1 . This vertex cannot delete its connections to its leaves, because that would disconnect the graph, making the distance component of its cost infinite. If v_1 remains neighbors with v_i and also buys an edge to some leaf $v_{i:j}$, then this is suboptimal: the edge to $v_{i:j}$ costs α but only gets v closer to one vertex. If v_1 deletes its edge to v_i but buys an edge to some leaf $v_{i:j}$, this is unnecessary: v_1 can move the edge from $v_{i:j}$ to v_i , switching its distances to those two vertices and not increasing the distance to any other vertex. Therefore, it is unnecessary for v_1 to consider strategies involving connecting to other vertices' leaves.

Thus, similarly to the previous case, v_1 only needs to consider strategies involving purely deleting edges. Again, by convexity, this reduces to considering strategies involving deleting a single edge. But again, v_1 can only break even by deleting an edge, so it can only break even overall. \square

Therefore, the graph is indeed a weak Nash equilibrium. Let n be the number of vertices in the graph. The size of the clique is $k = \frac{n}{\alpha}$, and so the cost of all of the edges is

$$\alpha \left[\binom{k}{2} + (\alpha - 1)k \right] = \alpha \left[(1 + o(1)) \frac{n^2}{2\alpha^2} + \frac{\alpha - 1}{\alpha} n \right] = (1 + o(1)) \frac{n^2}{2\alpha}.$$

Every clique vertex is distance 1 from the rest of the clique, as well as its leaves, and distance 2 from every other vertex; therefore, each clique vertex sees a distance sum of

$$(1 + o(1))n \left(2 - \frac{1}{\alpha} \right).$$

Since there are $\frac{n}{\alpha}$ clique vertices, these contribute a total of

$$(1 + o(1))n^2 \left(\frac{2}{\alpha} - \frac{1}{\alpha^2} \right).$$

Every leaf vertex is distance 2 from almost all of the clique, and distance 3 from almost all of the leaves, and so it sees a distance sum of

$$(1 + o(1))n \left(3 - \frac{1}{\alpha} \right).$$

Since there are $n(1 - \frac{1}{\alpha})$ leaves, these contribute a total distance sum of

$$(1 + o(1))n^2 \left(3 - \frac{4}{\alpha} + \frac{3}{\alpha^2} \right).$$

Putting everything together, we find that this graph has a total social cost of

$$(1 + o(1))n^2 \left(3 - \frac{3}{2\alpha} + \frac{2}{\alpha^2} \right),$$

giving a price of anarchy at least $\frac{3}{2} - \frac{3}{4\alpha} + \frac{1}{\alpha^2} + o(1)$, as claimed in the second part of Theorem 1.1.

4 Concluding remarks

It is interesting that the price of anarchy converges to 1 for non-integral $\alpha > 2$, but is bounded away from 1 for integer $\alpha \geq 2$. Our convergence rate is non-uniform in the sense that it slows down substantially when α is slightly more than an integer. On the other hand, when α is slightly less than an integer, the convergence rate is still relatively rapid. It would be nice to prove a uniform convergence rate for all non-integral α .

References

- [1] ALBERS, S. On the value of coordination in network design. In *Proc. 19th ACM-SIAM Symposium on Discrete Algorithms* (Philadelphia, PA, USA, 2008), SODA '08, Society for Industrial and Applied Mathematics, pp. 294–303.
- [2] ALBERS, S., EILTS, S., EVEN-DAR, E., MANSOUR, Y., AND RODITTY, L. On Nash equilibria for a network creation game. In *Proc. 16th ACM-SIAM Symposium on Discrete Algorithms* (2006), SODA '06, pp. 89–98.
- [3] ALON, N., DEMAINE, E. D., HAJIAGHAYI, M., AND LEIGHTON, T. Basic network creation games. In *Proc. 22nd ACM Symposium on Parallelism in Algorithms and Architectures* (New York, NY, USA, 2010), SPAA '10, ACM, pp. 106–113.
- [4] ANDELMAN, N., FELDMAN, M., AND MANSOUR, Y. Strong price of anarchy. In *Proc. 18th ACM-SIAM Symposium on Discrete Algorithms* (Philadelphia, PA, USA, 2007), SODA '07, Society for Industrial and Applied Mathematics, pp. 189–198.
- [5] ANSHELEVICH, E., DASGUPTA, A., KLEINBERG, J., TARDOS, E., WEXLER, T., AND ROUGHGARDEN, T. The price of stability for network design with fair cost allocation. In *Proc. 45th IEEE Symposium on Foundations of Computer Science* (Washington, DC, USA, 2004), FOCS '04, IEEE Computer Society, pp. 295–304.
- [6] ANSHELEVICH, E., DASGUPTA, A., TARDOS, E., AND WEXLER, T. Near-optimal network design with selfish agents. In *Proc. 35th ACM Symposium on Theory of Computing* (New York, NY, USA, 2003), STOC '03, ACM, pp. 511–520.

- [7] CORBO, J., AND PARKES, D. The price of selfish behavior in bilateral network formation. In *Proc. 24th ACM Symposium on Principles of Distributed Computing* (New York, NY, USA, 2005), PODC '05, ACM, pp. 99–107.
- [8] CZUMAJ, A., AND VÖCKING, B. Tight bounds for worst-case equilibria. In *Proc. 13th ACM-SIAM Symposium on Discrete Algorithms* (2002), SODA '02.
- [9] DEMAINE, E. D., HAJIAGHAYI, M., MAHINI, H., AND ZADIMOGHADDAM, M. The price of anarchy in network creation games. *ACM Transactions on Algorithms* 8, 2 (April 2012), Paper 13.
- [10] FABRIKANT, A., LUTHRA, A., MANEVA, E., PAPADIMITRIOU, C. H., AND SHENKER, S. On a network creation game. In *Proc. 22nd Symposium on Principles of Distributed Computing* (2003), PODC '03, pp. 347–351.
- [11] JACKSON, M. O. A survey of network formation models: stability and efficiency. In *Group Formation in Economics; Networks, Clubs, and Coalitions*, G. Demange and M. Wooders, Eds. Cambridge University Press, 2005.
- [12] KOUTSOUPIAS, E., AND PAPADIMITRIOU, C. Worst-case equilibria. In *Proc. 16th Symposium on Theoretical Aspects of Computer Science* (Berlin, Heidelberg, 1999), STACS'99, Springer-Verlag, pp. 404–413.
- [13] PAPADIMITRIOU, C. Algorithms, games, and the internet. In *Proc. 33rd ACM Symposium on Theory of Computing* (New York, NY, USA, 2001), STOC '01, ACM, pp. 749–753.
- [14] ROUGHGARDEN, T. The price of anarchy is independent of the network topology. In *Proc. 34th ACM Symposium on Theory of Computing* (2002), STOC '02, pp. 428–437.
- [15] ROUGHGARDEN, T. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.
- [16] TARDOS, E., AND WEXLER, T. Network formation games and the potential function method. In *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, Eds. Cambridge University Press, 2007, ch. 19, pp. 487–516.