

# Anarchy Is Free in Network Creation

RONALD GRAHAM, University of California, San Diego

LINUS HAMILTON, Carnegie Mellon University

ARIEL LEVAVI, University of California, San Diego

PO-SHEN LOH, Carnegie Mellon University

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 article considers one of the most well-studied models, originally proposed by Fabrikant et al. In the model, each of  $n$  agents corresponds to a vertex, which can create edges to other vertices at a cost of  $\alpha$  each, for some parameter  $\alpha$ . 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  $(\alpha, n)$  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 article, we significantly sharpen some of those results, proving that for all constant nonintegral  $\alpha > 2$ , the price of anarchy is in fact  $1 + o(1)$ ; that is, not only is it bounded by a constant, but also 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.

Categories and Subject Descriptors: G.2.2 [Discrete Mathematics]: Graph Theory

General Terms: Performance, Theory

Additional Key Words and Phrases: Network formation, algorithmic game theory, extremal graph theory

## ACM Reference Format:

Ronald Graham, Linus Hamilton, Ariel Levavi, and Po-Shen Loh. 2016. Anarchy is free in network creation. *ACM Trans. Algorithms* 12, 2, Article 15 (February 2016), 10 pages.

DOI: <http://dx.doi.org/10.1145/2729978>

## 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 important question

---

This work is supported by National Science Foundation grants DMS-1041500 and DMS-1201380, an NSA Young Investigators Grant, and a USA-Israel BSF Grant.

Authors' addresses: R. Graham and A. Levavi, Department of Computer Science and Engineering, University of California, San Diego, La Jolla, CA 92093; emails: [graham@ucsd.edu](mailto:graham@ucsd.edu), [alevavi@cs.ucsd.edu](mailto:alevavi@cs.ucsd.edu); L. Hamilton and P.-S. Loh, Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA 15213; emails: [luh@andrew.cmu.edu](mailto:luh@andrew.cmu.edu), [ploh@cmu.edu](mailto:ploh@cmu.edu).

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).

© 2016 ACM 1549-6325/2016/02-ART15 \$15.00

DOI: <http://dx.doi.org/10.1145/2729978>

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 article, we focus on a well-studied game-theoretic model of network creation, which was formulated by Fabrikant et al. [2003]. There are  $n$  agents, each corresponding to a vertex. They form a network (graph) 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  $\alpha$  be an arbitrary real 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 a 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 of Anshelevich et al. [2004], Anshelevich et al. [2003], Czumaj and Vöcking [2002], Papadimitriou [2001], Roughgarden [2002], and Roughgarden [2005] all consider network design issues such as load balancing, routing, and so forth. Numerous papers, including Albers [2008], Albers et al. [2006], Alon et al. [2010], Andelman et al. [2007], Corbo and Parkes [2005], Demaine et al. [2012], Mamageishvili et al. [2013], and Mihalák and Schlegel [2013], and the surveys of Jackson [2005] and Tardos and Wexler [2007] 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 [1999]. 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.

Table I. Previous Work

Regime	Upper Bound on Price of Anarchy
General $\alpha$	$2^{O(\sqrt{\log n})}$
$2 \leq \alpha < \sqrt[3]{n/2}$	4
$\sqrt[3]{n/2} \leq \alpha < \sqrt{n/2}$	6
$\alpha = O(n^{1-\epsilon})$	$O(1)$
$\alpha \geq 65n$	$O(1)$

Previous upper bounds on the price of anarchy. The last bound is due to Mamageishvili et al. [2013], and the other bounds are due to Demaine et al. [2012].

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 that 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. [2003], 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  $\alpha$  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  $\alpha$ , 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 I.

## 1.2. Our Contribution

Much work had been done to achieve constant upper bounds on the price of anarchy in various regimes of  $\alpha$ , 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  $\alpha$ . In other words, the lack of coordination has negligible effect on the social cost as  $n$  grows.

**THEOREM 1.1.** *For a fixed nonintegral  $\alpha > 2$ , the price of anarchy converges to 1 as  $n \rightarrow \infty$ . Quantitatively: when  $\alpha > 2$ ,  $\alpha$  is not an integer, and  $n > \alpha^3$ , the price of anarchy is bounded above by*

$$1 + \frac{200\alpha^6}{(\alpha - \lfloor \alpha \rfloor)^2} \sqrt{\frac{\log n}{n}}.$$

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

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

*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  $v$ .*

*Remark.* We focus on the range  $\alpha \geq 2$  because the price of anarchy is already understood in the range  $\alpha < 2$ , as mentioned earlier in the introduction.




Fig. 1. Partitioning the graph into sets.




Fig. 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  $b$ .

**2. PROOF FOR NONINTEGRAL  $\alpha$**

Fix a nonintegral  $\alpha > 2$ . 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  $v$ , and so forth, as diagrammed in Figure 1. Since the graph in every Nash equilibrium is obviously connected, every vertex falls into one of these 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.

With this in mind, the main idea of the proof is as follows. We will show that for almost all vertices  $v$ , the third neighborhood  $N_3$  and beyond are small; the first neighborhood  $N_1$  is also small; and the second neighborhood  $N_2$  is almost the entire graph. The proof proceeds as a series of lemmas.

**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  $\alpha$  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  $\alpha$ , 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  $\alpha$  for the edge and get three steps closer to  $x$ , as well as at least one step closer to all of  $x$ 's immediate neighbors, for a net gain. Hence, this cannot appear in a Nash equilibrium.  $\square$

COROLLARY 2.4. *If  $n$  is sufficiently large ( $n > \alpha^3$ ), then the graph has a 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  $\alpha$  or one of  $v$ 's neighbors has degree at least  $\alpha$ . In either case, we can travel from  $v$  to a vertex with degree at least  $\alpha$  in at most one step and then, by Lemma 2.3, travel to  $w$  in at most three 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 2.5. *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 that 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 one 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 one 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  $\leq 0$  since we are at a Nash equilibrium. Rearranging,  $d \leq |N_1| \cdot \frac{\alpha}{\alpha - \lfloor \alpha \rfloor}$ , as desired.  $\square$

**LEMMA 2.6.** *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 backward, 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  $\alpha$  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 two 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 2.5, 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

$$P \leq |N_3 \cup N_4|\alpha^2 + 2\alpha^2|N_2| + |N_1||N_2|\alpha + |N_1||N_2|\frac{2\alpha^2}{\alpha - \lfloor \alpha \rfloor}.$$

Using the trivial bounds  $|N_3 \cup N_4| + |N_2| < n$  and  $|N_2| < n$ , we bound the previous expression:

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



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

$$|N_3 \cup N_4| \frac{n}{\alpha} < |N_1| n \left( \frac{5\alpha^2}{\alpha - \lfloor \alpha \rfloor} \right)$$

$$|N_3 \cup N_4| < |N_1| \cdot \frac{5\alpha^3}{\alpha - \lfloor \alpha \rfloor}. \quad \square$$

**LEMMA 2.7.** *If every vertex has degree more than  $\sqrt{n \log n}$ , 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 are currently paying for, and connect to  $\sqrt{n \log n}$  vertices at random” has a positive probability of bringing it within distance  $\leq 2$  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.  $\square$

**LEMMA 2.8.** *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  $N_1, N_2, N_3, N_4$  with respect to  $v$ . We already know  $|N_1|$  is at most  $\sqrt{n \log n} = o(n)$ , so by Lemma 2.6,  $|N_3 \cup N_4|$  is at most  $\sqrt{n \log n} \cdot \frac{5\alpha^3}{\alpha - \lfloor \alpha \rfloor} = o(n)$ . By Lemma 2.5, 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 2.6 on every vertex in the graph, and the fact that all distances are at most 4

(Corollary 2.4), 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 we just bounded earlier. Putting everything together, the total sum of distances is at most

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

The total cost of edges is  $\alpha$  times the number of edges. By adding this in to compute the total social cost, we obtain the desired bound. Remark: the constant 290 in the statement of the theorem is perhaps overly generous, but it suffices for our purposes.  $\square$

Lemmas 2.7 and 2.8 cover complementary cases, so we now conclude that the total social cost of every Nash equilibrium is at most the bound obtained in Lemma 2.8. As was observed by previous authors [Fabrikant et al. 2003], 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

$$\left(1 + \frac{1}{n-1}\right) \left(1 + \frac{145\alpha^6}{(\alpha - \lfloor \alpha \rfloor)^2} \sqrt{\frac{\log n}{n}}\right) < 1 + \frac{200\alpha^6}{(\alpha - \lfloor \alpha \rfloor)^2} \sqrt{\frac{\log n}{n}} = 1 + o(1),$$

where we merged the  $(1 + \frac{1}{n-1})$  factor into the bound by using our assumptions that  $n > \alpha^3$  and  $\alpha > 2$ . This proves the first part of Theorem 1.1.

### 3. INTEGRAL $\alpha$

There is one catch in our bound earlier. Namely, when  $\alpha$  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  $\alpha$  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  $\alpha$  is an integer. Let  $k \gg \alpha$ , and 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 Albers et al. [2006], as an example of a Nash equilibrium that does not correspond to a tree, but its social cost is not calculated there.

We therefore calculate its social cost now. 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 cliques, 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

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

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

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

giving a price of anarchy at least  $\frac{3}{2} - \frac{3}{4\alpha} + o(1)$ , 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 nonintegral  $\alpha > 2$  but is bounded away from 1 for integer  $\alpha \geq 2$ . Our convergence rate is nonuniform in the sense that it slows down substantially when  $\alpha$  is slightly more than an integer. On the other hand, when  $\alpha$  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 nonintegral  $\alpha$ .

#### ACKNOWLEDGMENTS

We would like to thank the referees for useful suggestions that improved the clarity of this article.

#### REFERENCES

- Susanne Albers. 2008. On the value of coordination in network design. In *Proc. 19th ACM-SIAM Symposium on Discrete Algorithms (SODA'08)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 294–303.
- Susanne Albers, Stefan Eilts, Eyal Even-dar, Yishay Mansour, and Liam Roditty. 2006. On Nash equilibria for a network creation game. In *Proc. 16th ACM-SIAM Symposium on Discrete Algorithms (SODA'06)*. 89–98.
- Noga Alon, Erik D. Demaine, Mohammad Taghi Hajiaghayi, and Tom Leighton. 2010. Basic network creation games. In *Proc. 22nd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA'10)*. ACM, New York, NY, 106–113. DOI: <http://dx.doi.org/10.1145/1810479.1810502>
- Nir Andelman, Michal Feldman, and Yishay Mansour. 2007. Strong price of anarchy. In *Proc. 18th ACM-SIAM Symposium on Discrete Algorithms (SODA'07)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 189–198. <http://dl.acm.org/citation.cfm?id=1283383.1283404>
- Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Eva Tardos, Tom Wexler, and Tim Roughgarden. 2004. The price of stability for network design with fair cost allocation. In *Proc. 45th IEEE Symposium on Foundations of Computer Science (FOCS'04)*. IEEE Computer Society, 295–304. DOI: <http://dx.doi.org/10.1109/FOCS.2004.68>
- Elliot Anshelevich, Anirban Dasgupta, Eva Tardos, and Tom Wexler. 2003. Near-optimal network design with selfish agents. In *Proc. 35th ACM Symposium on Theory of Computing (STOC'03)*. ACM, New York, NY, 511–520. DOI: <http://dx.doi.org/10.1145/780542.780617>
- Jacomo Corbo and David Parkes. 2005. The price of selfish behavior in bilateral network formation. In *Proc. 24th ACM Symposium on Principles of Distributed Computing (PODC'05)*. ACM, New York, NY, 99–107. DOI: <http://dx.doi.org/10.1145/1073814.1073833>

- Artur Czumaj and Berthold Vöcking. 2002. Tight bounds for worst-case equilibria. In *Proc. 13th ACM-SIAM Symposium on Discrete Algorithms (SODA'02)*. 413–420.
- Erik D. Demaine, Mohammad Taghi Hajiaghayi, Hamid Mahini, and Morteza Zadimoghaddam. 2012. The price of anarchy in network creation games. *ACM Transactions on Algorithms* 8, 2 (April 2012), Article 13.
- Alex Fabrikant, Ankur Luthra, Elitza Maneva, Christos H. Papadimitriou, and Scott Shenker. 2003. On a network creation game. In *Proc. 22nd Symposium on Principles of Distributed Computing (PODC'03)*. 347–351.
- Matthew O. Jackson. 2005. A survey of network formation models: Stability and efficiency. In *Group Formation in Economics; Networks, Clubs, and Coalitions*, Gabrielle Demange and Myrna Wooders (Eds.). Cambridge University Press, 11–57.
- Elias Koutsoupias and Christos Papadimitriou. 1999. Worst-case equilibria. In *Proc. 16th Symposium on Theoretical Aspects of Computer Science (STACS'99)*. Springer-Verlag, Berlin, 404–413. <http://dl.acm.org/citation.cfm?id=1764891.1764944>
- Akaki Mamageishvili, Matúš Mihalák, and Dominik Müller. 2013. Tree Nash equilibria in the network creation game. In *Proc. 10th Workshop on Algorithms and Models for the Web Graph (WAW'13)*. 118–129.
- Matúš Mihalák and Jan Christoph Schlegel. 2013. The price of anarchy in network creation games is (mostly) constant. *Theory of Computing Systems* 53, 1 (2013), 53–72.
- Christos Papadimitriou. 2001. Algorithms, games, and the internet. In *Proc. 33rd ACM Symposium on Theory of Computing (STOC'01)*. ACM, New York, NY, 749–753. DOI : <http://dx.doi.org/10.1145/380752.380883>
- Tim Roughgarden. 2002. The price of anarchy is independent of the network topology. In *Proc. 34th ACM Symposium on Theory of Computing (STOC'02)*. 428–437.
- Tim Roughgarden. 2005. *Selfish Routing and the Price of Anarchy*. MIT Press.
- Éva Tardos and Tom Wexler. 2007. Network formation games and the potential function method. In *Algorithmic Game Theory*, Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani (Eds.). Cambridge University Press, 487–516.

Received November 2013; revised January 2015; accepted January 2015