

The Price of Anarchy in Cooperative Network Creation Games

ERIK D. DEMAINE

MIT, Cambridge, Massachusetts

and

MOHAMMADTAGHI HAJIAGHAYI

AT&T Labs — Research, Florham Park, New Jersey

and

HAMID MAHINI

Sharif University of Technology, Tehran, Iran

and

MORTEZA ZADIMOGHADDAM

MIT, Cambridge, Massachusetts

1. INTRODUCTION

A fundamental family of problems at the intersection between computer science and operations research is *network design*. This area of research has become increasingly important given the continued growth of computer networks such as the Internet. Traditionally, we want to find a minimum-cost (sub)network that satisfies some specified property such as k -connectivity or connectivity on terminals (as in the classic Steiner tree problem). This goal captures the (possibly incremental) creation cost of the network, but does not incorporate the cost of actually using the network. In contrast, *network routing* has the goal of optimizing the usage cost of the network, but assumes that the network has already been created.

Network creation games attempt to unify the network design and network routing problems by modeling both creation and usage costs. In general, the game is played on a *host graph*, where each node is an independent agent (player), and the goal is to create a network from a subgraph of the host graph. Collectively, the nodes decide which edges of the host graph are worth creating as links in the network. Every link has the same creation cost α . (Equivalently, links have creation costs of

A preliminary version of this paper appeared in *Proceedings of the 26th International Symposium on Theoretical Aspects of Computer Science*, 2008, pages 171–182. This work was done while M. Hajiaghayi was at MIT and CMU and M. Zadimoghaddam was at Sharif University of Technology. H. Mahini was supported in part by Institute for Theoretical Physics and Mathematics (IPM) under grant number CS1384-2-01.

Author's addresses: E. Demaine, M. Zadimoghaddam, Computer Science and Artificial Intelligence Laboratory, MIT, 32 Vassar St., Cambridge, MA 02139, USA, e-mail: {edemaine,morteza}@mit.edu; M. Hajiaghayi, AT&T Labs — Research, 180 Park Ave., Florham Park, NJ 07932, USA, e-mail: hajiagha@research.att.com; H. Mahini, Department of Computer Engineering, Sharif University of Technology, Azadi St., Tehran, Iran, e-mail: mahini@ce.sharif.edu.

α and ∞ , depending on whether they are edges of the host graph.) In addition to these creation costs, each node incurs a usage cost equal to the sum of distances to all other nodes in the network. Equivalently, if we divide the cost (and thus α) by the number n of nodes, the usage cost for each node is its average distance to all other nodes. (This natural cost model has been used in, e.g., contribution games and network-formation games.)

There are several versions of the network creation game that vary how links are purchased. In the *unilateral* model—introduced by Fabrikant, Luthra, Maneva, Papadimitriou, and Shenker [2003]—every node (player) can locally decide to purchase any edge incident to the node in the host graph, at a cost of α . In the *bilateral* model—introduced by Corbo and Parkes [2005]—both endpoints of an edge must agree before they can create a link between them, and the two nodes share the α creation cost equally. In the *cooperative* model—introduced by Albers, Eilts, Even-Dar, Mansour, and Roditty [2006]—any node can purchase any amount of any edge in the host graph, and a link gets created when the total purchased amount is at least α .

To model the dominant behavior of large-scale networking scenarios such as the Internet, we consider the case where every node (player) selfishly tries to minimize its own creation and usage cost [Jackson 2003; Fabrikant et al. 2003; Albers et al. 2006; Corbo and Parkes 2005]. This game-theoretic setting naturally leads to the various kinds of *equilibria* and the study of their structure. Two frequently considered notions are *Nash equilibrium* [Nash 1950; 1951], where no player can change its strategy (which edges to buy) to locally improve its cost, and *strong Nash equilibrium* [Aumann 1959; Andelman et al. 2007; Albers 2008], where no coalition of players can change their collective strategy to locally improve the cost of each player in the coalition. Nash equilibria capture the combined effect of both selfishness and lack of coordination, while strong Nash equilibria separates these issues, enabling coordination and capturing the specific effect of selfishness. However, the notion of strong Nash equilibrium is extremely restrictive in our context, because all players can simultaneously change their entire strategies, abusing the local optimality intended by original Nash equilibria, and effectively forcing globally near-optimal solutions [Andelman et al. 2007].

We consider weaker notions of equilibria, which broadens the scope of equilibria and therefore strengthens our upper bounds, where players can change their strategy on only a single edge at a time. In a *collaborative equilibrium*, even coalitions of players do not wish to change their collective strategy on any single edge; this concept is particularly important for the cooperative network creation game, where multiple players must negotiate their relative valuations of an edge. (This notion is the natural generalization of pairwise stability from [Corbo and Parkes 2005] to arbitrary cost sharing.) Collaborative equilibria are essentially a compromise between Nash and strong Nash equilibria: they still enable coordination among players and thus capture the specific effect of selfishness, like strong Nash, yet they consider more local moves, in the spirit of Nash (but with a different form of locality). In particular, any results about all collaborative equilibria also apply to all strong Nash equilibria. Collaborative equilibria also make more sense computationally: players can efficiently detect equilibrium using a simple bidding procedure

(whereas this problem is NP-hard for strong Nash), and the resulting dynamics converge to such equilibria (see Section 2.2).

The structure of equilibria in network creation games is not very well understood. For example, Fabrikant et al. [2003] conjectured that equilibrium graphs in the unilateral model were all trees, but this conjecture was disproved by Albers et al. [2006]. One particularly interesting structural feature is whether all equilibrium graphs have small *diameter* (say, polylogarithmic), analogous to the small-world phenomenon [Kleinberg 2000; Even-Dar and Kearns 2006]. In the original unilateral version of the problem, the best general lower bound is just a constant and the best general upper bound is polynomial. A closely related issue is the *price of anarchy* [Koutsoupias and Papadimitriou 1999; Papadimitriou 2001; Roughgarden 2002b], that is, the worst possible ratio of the total cost of an equilibrium (found by independent selfish behavior) and the optimal total cost possible by a centralized solution (maximizing social welfare). The price of anarchy is a well-studied concept in algorithmic game theory for problems such as load balancing, routing, and network design; see, e.g., [Papadimitriou 2001; Czumaj and Vöcking 2002; Roughgarden 2002a; Fabrikant et al. 2003; Anshelevich et al. 2003; Anshelevich et al. 2004; Chun et al. 2004; Corbo and Parkes 2005; Albers et al. 2006; Demaine et al. 2007]. Upper bounds on diameter of equilibrium graphs translate to approximately equal upper bounds on the price of anarchy, but not necessarily vice versa. In the unilateral version, for example, there is a general $2^{O(\sqrt{\lg n})}$ upper bound on the price of anarchy.

1.0.0.1 *Previous work.* Network creation games have been studied extensively in the literature since their introduction in 2003.

For the unilateral version and a complete host graph, Fabrikant et al. [2003] prove an upper bound of $O(\sqrt{\alpha})$ on the price of anarchy for all α . Lin [2003] proves that the price of anarchy is constant for two ranges of α : $\alpha = O(\sqrt{n})$ and $\alpha \geq cn^{3/2}$ for some $c > 0$. Independently, Albers et al. [2006] prove that the price of anarchy is constant for $\alpha = O(\sqrt{n})$, as well as for the larger range $\alpha \geq 12n \lceil \lg n \rceil$. In addition, Albers et al. prove a general upper bound of $15 \left(1 + (\min\{\frac{\alpha^2}{n}, \frac{n^2}{\alpha}\})^{1/3}\right)$. The latter bound shows the first sublinear worst-case bound, $O(n^{1/3})$, for all α . Demaine et al. [2007] prove the first $o(n^\epsilon)$ upper bound for general α , namely, $2^{O(\sqrt{\lg n})}$. They also prove a constant upper bound for $\alpha = O(n^{1-\epsilon})$ for any fixed $\epsilon > 0$, and improve the constant upper bound by Albers et al. (with the lead constant of 15) to 6 for $\alpha < (n/2)^{1/2}$ and to 4 for $\alpha < (n/2)^{1/3}$. Andelmen et al. [2007] show that, among strong Nash equilibria, the price of anarchy is at most 2.

For the bilateral version and a complete host graph, Corbo and Parkes [2005] prove that the price of anarchy is between $\Omega(\lg \alpha)$ and $O(\min\{\sqrt{\alpha}, n/\sqrt{\alpha}\})$. Demaine et al. [2007] prove that the upper bound is tight, establishing the price of anarchy to be $\Theta(\min\{\sqrt{\alpha}, n/\sqrt{\alpha}\})$ in this case.

For the cooperative version and a complete host graph, the only known result is an upper bound of $15 \left(1 + (\min\{\frac{\alpha^2}{n}, \frac{n^2}{\alpha}\})^{1/3}\right)$, proved by Albers et al. [2006].

Other variations of network creation games allow nonuniform interests in connectivity between nodes [Halevi and Mansour 2007] and nodes with limited budgets for buying edges [Laoutaris et al. 2008].

1.0.0.2 *Our results.* Our research pursues two important facets of the network creation game.

First, we make an extensive study of a natural version of the game—the cooperative model—where the only previous results were simple extensions from unilateral analysis. We substantially improve the bounds in this case, showing that the price of anarchy is polylogarithmic in n for *all values of α* in complete graphs. This is the first result of this type for any version of the network creation game. As mentioned above, this result applies to both collaborative equilibria and strong Nash equilibria. Interestingly, we also show that equilibrium graphs have polylogarithmic diameter for the most natural range of α (at most $n \text{ poly} \lg n$). Note that, because of the locally greedy nature of Nash equilibria, we cannot use the classic probabilistic spanning (sub)tree embedding machinery of [Bartal 1998; Fakcharoenphol et al. 2004; Elkin et al. 2005] to obtain polylogarithmic bounds (although this machinery can be applied to approximate the global social optimum).

Second, we study the impact of the natural assumption that the host graph is a general graph, not necessarily complete, inspired by practical limitations in constructing network links. This model is a simple example of nonuniform creation costs among the edges (effectively allowing weights of α and ∞). Surprisingly, no bounds on the diameter or the price of anarchy have been proved before in this context. We prove several upper and lower bounds, establishing nontrivial tight bounds for many ranges of α , for both the unilateral and cooperative versions. In particular, we establish polynomial lower bounds for both versions and many ranges of α , even for this simple nonuniform cost model. These results are particularly interesting because, by contrast, no superconstant lower bound has been shown for either game in complete (uniform) graphs. Thus, while we believe that the price of anarchy is polylogarithmic (or even constant) for complete graphs, we show a significant departure from this behavior in general graphs.

Our proof techniques are most closely related in spirit to “region growing” from approximation algorithms; see, e.g., [Leighton and Rao 1999]. Our general goal is to prove an upper bound on diameter by way of an upper bound on the expansion of the graph. However, we have not been able to get such an argument to work directly in general. The main difficulty is that, if we imagine building a breadth-first-search tree from a node, then connecting that root node to another node does not necessarily benefit the node much: it may only get closer to a small fraction of nodes in the BFS subtree. Thus, no node is motivated selfishly to improve the network, so several nodes must coordinate their changes to make improvements. The cooperative version of the game gives us some leverage to address this difficulty. We hope that this approach, particularly the structure we prove of equilibria, will shed some light on the still-open unilateral version of the game, where the best bounds on the price of anarchy are $\Omega(1)$ and $2^{O(\sqrt{\lg n})}$.

Table I summarizes our results. Section 4 proves our polylogarithmic upper bounds on the price of anarchy for all ranges of α in the cooperative network creation game in complete graphs. Section 5 considers how the cooperative network creation game differs in general graphs, and proves our upper bounds for this model. Section 6 extends these results to apply to the unilateral network creation game in general graphs. Section 7 proves lower bounds for both the unilateral and coopera-

$\alpha =$	0	n	$n \lg^{0.52} n$	$n \lg^{7.16} n$	$n^{3/2}$	$n^{5/3}$	n^2	$n^2 \lg n$	∞
Cooperative, complete graph	$\Theta(1)$	$\lg^{3.32} n$	$O\left(\lg n + \sqrt{\frac{n}{\alpha}} \lg^{3.58} n\right)$	$\Theta(1)$					
Cooperative, general graph	$O(\alpha^{1/3})$	$O(n^{1/3}), \Omega(\sqrt{\frac{\alpha}{n}})$			$\Theta(\frac{n^2}{\alpha})$	$O(\frac{n^2}{\alpha} \lg n)$	$\Theta(1)$		
Unilateral, general graph	$O(\alpha^{1/2})$	$O(n^{1/2}), \Omega(\frac{\alpha}{n})$			$\Theta(\frac{n^2}{\alpha})$	$\Theta(1)$			

Table I. Summary of our bounds on equilibrium diameter and price of anarchy for cooperative network creation in complete graphs, and unilateral and cooperative network creation in general graphs. For all three of these models, our bounds are strict improvements over the best previous bounds.

tive network creation games in general graphs, which match our upper bounds for some ranges of α .

2. MODELS

In this section, we formally define the different models of the network creation game.

2.1 Unilateral Model

We start with the unilateral model, introduced in [Fabrikant et al. 2003]. The game is played on a *host graph* $G = (V, E)$. Assume $V = \{1, 2, \dots, n\}$. We have n players, one per vertex. The strategy of player i is specified by a subset s_i of $\{j : \{i, j\} \in E\}$, defining the set of neighbors to which player i creates a link. Thus each player can only create links corresponding to edges incident to node i in the host graph G . Together, let $s = \langle s_1, s_2, \dots, s_n \rangle$ denote the joint strategy of all players.

To define the cost of strategies, we introduce a spanning subgraph G_s of the host graph G . Namely, G_s has an edge $\{i, j\} \in E(G)$ if either $i \in s_j$ or $j \in s_i$. Define $d_{G_s}(i, j)$ to be the distance between vertices i and j in graph G_s . (If G_s has no path between i and j , then $d_{G_s}(i, j) = \infty$.) The *cost* (creation plus usage cost) incurred by player i is

$$c_i(s) = \alpha |s_i| + \sum_{j=1}^n d_{G_s}(i, j).$$

The total cost incurred by joint strategy s is $c(s) = \sum_{i=1}^n c_i(s)$.

A (pure) *Nash equilibrium* is a joint strategy s such that $c_i(s) \leq c_i(s')$ for all joint strategies s' that differ from s in only one player i . The *price of anarchy* is then the maximum cost of a Nash equilibrium divided by the minimum cost of any joint strategy (called the *social optimum*).

2.2 Cooperative Model

Next we turn to the cooperative model, introduced in [Fabrikant et al. 2003; Albers et al. 2006]. Again, the game is played on a host graph $G = (V, E)$, with one player per vertex. Assume $V = \{1, 2, \dots, n\}$ and $E = \{e_1, e_2, \dots, e_{|E|}\}$. Now the strategy of player i is specified by a vector $s_i = \langle s(i, e_1), s(i, e_2), \dots, s(i, e_{|E|}) \rangle$, where $s(i, e_j)$ corresponds to the value that player i is willing to pay for link e_j . Together, $s = \langle s_1, s_2, \dots, s_n \rangle$ denotes the strategies of all players.

We define a spanning subgraph $G_s = (V, E_s)$ of the host graph G : e_j is an edge of G_s if $\sum_{i \in V(G)} s(i, e_j) \geq \alpha$. To make the total cost for an edge e_j exactly 0 or α in all cases, if $\sum_{i \in V(G)} s(i, e_j) > \alpha$, we uniformly scale the costs to sum to α : $s'(i, e_j) = \alpha s(i, e_j) / \sum_{k \in V(G)} s(k, e_j)$. (Equilibria will always have $s = s'$.) Then the cost incurred by player i is

$$c_i(s) = \sum_{e_j \in E_s} s'(i, e_j) + \sum_{j=1}^n d_{G_s}(i, j).$$

The total cost incurred by joint strategy s is

$$c(s) = \alpha |E_s| + \sum_{i=1}^n \sum_{j=1}^n d_{G_s}(i, j).$$

In this cooperative model, the notion of Nash equilibrium is less natural because it allows only one player to change strategy, whereas a cooperative purchase in general requires many players to change their strategy. Therefore we use a stronger notion of equilibrium that allows coalition among players, inspired by the strong Nash equilibrium of Aumann [1959], and modeled after the pairwise stability property introduced for the bilateral game [Corbo and Parkes 2005]. Namely, a (pure) joint strategy s is a *collaborative equilibrium* if, for any edge e of the host graph G , for any coalition $C \subseteq V$, for any joint strategy s' differing from s in only $s'(i, e)$ for $i \in C$, some player $i \in C$ has $c_i(s') > c_i(s)$. Note that any such joint strategy must have every sum $\sum_{i \in V(G)} s(i, e_j)$ equal to either 0 or α , so we can measure the cost $c_i(s)$ in terms of $s(i, e_j)$ instead of $s'(i, e_j)$. The *price of anarchy* is the maximum cost of a collaborative equilibrium divided by the minimum cost of any joint strategy (the *social optimum*).

We can define a simple dynamics for the cooperative network creation game in which we repeatedly pick a pair of vertices, have all players determine their valuation of an edge between those vertices (change in $c_i(s)$ from addition or removal), and players thereby bid on the edge and change their strategies. These dynamics always converge to a collaborative equilibrium because each change decreases the total cost $c(s)$, which is a discrete quantity in the lattice $\mathbb{Z} + \alpha\mathbb{Z}$. Indeed, the system therefore converges after a number of steps polynomial in n and the smallest integer multiple of α (if one exists). More generally, we can show an exponential upper bound in terms of just n by observing that the graph uniquely determines $c(s)$, so we can never repeat a graph by decreasing $c(s)$.

3. PRELIMINARIES

In this section, we define some helpful notation and prove some basic results. Call a graph G_s corresponding to an equilibrium joint strategy s an *equilibrium graph*. In such a graph, let $d_{G_s}(u, v)$ be the length of the shortest path from u to v and $\text{Dist}_{G_s}(u)$ be $\sum_{v \in V(G_s)} d_{G_s}(u, v)$. Let $N_k(u)$ denote the set of vertices with distance at most k from vertex u , and let $N_k = \min_{v \in G} |N_k(v)|$. In both the unilateral and cooperative network creation games, the total cost of a strategy consists of two parts. We refer to the cost of buying edges as the *creation cost* and the cost $\sum_{v \in V(G_s)} d_{G_s}(u, v)$ as the *usage cost*.

First we prove the existence of collaborative equilibria for complete host graphs. Similar results are known in the unilateral case [Fabrikant et al. 2003; Andelman et al. 2007].

Lemma 1. *In the cooperative network creation game, any complete graph is a collaborative equilibrium for $\alpha \leq 2$, and any star graph is a collaborative equilibrium for $\alpha \geq 2$.*

Next we show that, in the unilateral version, a bound on the usage cost suffices to bound the total cost of an equilibrium graph G_s , similarly to [Demaine et al. 2007, Lemma 1].

Lemma 2. *The total cost of any equilibrium graph in the unilateral game is at most $\alpha n + 2 \sum_{u,v \in V(G_s)} d_{G_s}(u,v)$.*

Proof: Let $v = \operatorname{argmin}_{u \in V(G_s)} \operatorname{Dist}_{G_s}(u)$. Therefore $\operatorname{Dist}_{G_s}(v) \leq \frac{1}{n} \sum_{u,v \in V(G_s)} d_{G_s}(u,v)$. Let T be the BFS tree rooted at v . For every vertex x , if x changes its strategy in order to keep only edges in T that x bought, the sum of its distance to the other vertices would be at most $\sum_{y \in V(G_s)} [d_{G_s}(x,v) + d_{G_s}(v,y)] \leq n d_{G_s}(x,v) + \operatorname{Dist}_{G_s}(v)$. On the other hand, x 's creation cost would be at most αt_x , where t_x is the number of edges in T bought by x . Thus the total cost of vertex x would be at most $\alpha t_x + n d_{G_s}(x,v) + \operatorname{Dist}_{G_s}(v)$. But x did not choose this strategy, so $c_x(s) \leq \alpha t_x + n d_{G_s}(x,v) + \operatorname{Dist}_{G_s}(v)$. By summing all these costs, the total cost of the equilibrium graph is at most

$$\begin{aligned} c(s) &\leq \sum_{x \in V(G_s)} (n d_{G_s}(x,v) + \operatorname{Dist}_{G_s}(v) + \alpha t_x) \\ &\leq 2n \operatorname{Dist}_{G_s}(v) + \alpha(n-1) \\ &\leq 2 \sum_{u,v \in V(G_s)} d_{G_s}(u,v) + \alpha n. \end{aligned}$$

□

Next we prove a more specific bound for the cooperative version, using the following bound on the number of edges in a graph of large girth:

Lemma 3. [Dutton and Brigham 1991] *The number of edges in an n -vertex graph of odd girth g is $O(n^{1+2/(g-1)})$.*

Lemma 4. *For any integer g , the total cost of any equilibrium graph G_s is at most $\alpha O(n^{1+2/g}) + g \sum_{u,v \in V(G_s)} d_{G_s}(u,v)$.*

Proof: Let $g' = g + 1$. Consider an edge x in at least one cycle of length at most g' . We know $\sum_{v \in V(G_s)} c(v,x) \geq \alpha$. For every vertex v , consider the shortest path from v to other vertices, and let $f(v,x)$ denote the number of such paths that contain x . If we delete x from graph G_s , then the length of these $f(v,x)$ shortest paths increase by at most $g' - 2$. Because the edge x is in the equilibrium graph, we conclude that $(g' - 2) \sum_{v \in V(G_s)} f(v,x) \geq \sum_{v \in V(G_s)} c(v,x) \geq \alpha$, which implies that $\sum_{v \in V(G_s)} f(v,x) \geq \alpha/(g' - 2)$. Thus edge x is in at least $\alpha/(g' - 2)$ shortest path in the equilibrium graph. On the other hand, $\sum_{u,v \in V(G_s)} d_{G_s}(u,v)$ equals the number

of edges in all shortest path. Therefore the number of edges like x that are in at least one cycle of length at most g' is at most $\sum_{u,v \in V(G_s)} d_{G_s}(u,v)/(\alpha/(g'-2))$.

If we delete all such edges in at least one cycle of length at most g' , then the girth of the remaining graph is at least $g'+1$. By Lemma 3, the number of edges in the remaining graph is $O(n^{1+2/(g'-1)})$ (because $g'+1$ maybe an even number). Thus the number of edges in the equilibrium graph is at most $\sum_{u,v \in V(G_s)} d_{G_s}(u,v)/(\alpha/(g'-2)) + O(n^{1+2/(g'-1)})$ and the cost of buying these edges is at most $(g'-2) \sum_{u,v \in V(G_s)} d_{G_s}(u,v) + \alpha O(n^{1+2/(g'-1)})$. Therefore the total cost is at most $\alpha O(n^{1+2/g}) + g \sum_{u,v \in V(G_s)} d_{G_s}(u,v)$. \square

4. COOPERATIVE VERSION IN COMPLETE GRAPHS

In this section, we study the price of anarchy when any number of players can cooperate to create any link, and the host graph is the complete graph.

We start with two lemmata that hold for both the unilateral and cooperative versions of the problem. The first lemma bounds a kind of “doubling radius” of large neighborhoods around any vertex, which the second lemma uses to bound the usage cost.

Lemma 5. [Demaine et al. 2007, Lemma 4] *For any vertex u in an equilibrium graph G_s , if $|N_k(u)| > n/2$, then $|N_{2k+2\alpha/n}(u)| \geq n$.*

Proof: We prove the contrapositive. Suppose $|N_{2k+2\alpha/n}(u)| < n$. Then there is a vertex v with $d_{G_s}(u,v) \geq 2k+1+2\alpha/n$. For every vertex $x \in N_k(u)$, $d_{G_s}(u,x) \leq k$. By the triangle inequality, $d_{G_s}(u,x) + d_{G_s}(x,v) \geq d_{G_s}(u,v)$, so $d_{G_s}(x,v) \geq k+1+2\alpha/n$. If vertex v bought the edge $\{v,u\}$, then the distance between v and x would decrease by at least $2\alpha/n$, so $\text{Dist}_{G_s}(v)$ would decrease by at least $|N_k(u)| \cdot 2\alpha/n$. Because v has not bought the edge $\{v,u\}$, we have $\alpha \geq |N_k(u)| \cdot 2\alpha/n$, i.e., $|N_k(u)| \leq n/2$. \square

Lemma 6. *If we have $N_k(u) > n/2$ for some vertex u in an equilibrium graph G_s , the usage cost is at most $O(n^2k + \alpha n)$.*

Proof: By Lemma 5, the sum of the distances from u to all other vertices is at most $2k+2\alpha/n$. Thus the distance between any pair of vertices is at most $4k+4\alpha/n$, so $\sum_{u,v \in V(G_s)} d_{G_s}(u,v) = O(n^2(k+\alpha/n)) = O(n^2k + \alpha n)$. \square

Next we show how to improve the bound on “doubling radius” for large neighborhoods in the cooperative game:

Lemma 7. *For any vertex u in an equilibrium graph G_s , if $|N_k(u)| > n/2$, then $|N_{2k+4\sqrt{\alpha/n}}(u)| \geq n$.*

Proof: We prove the contrapositive. Suppose $|N_{2k+4\sqrt{\alpha/n}}(u)| < n$. Then there is a vertex v with $d_{G_s}(u,v) \geq 2k+1+4\sqrt{\alpha/n}$. For every vertex $x \in N_k(u)$ and $y \in N_{\sqrt{\alpha/n}}(v)$, we have $d_{G_s}(u,x) \leq k$ and $d_{G_s}(v,y) \leq \sqrt{\alpha/n}$; see Figure 1. By the triangle inequality $d_{G_s}(u,x) + d_{G_s}(x,y) + d_{G_s}(y,v) \geq d_{G_s}(u,v)$, we have $d_{G_s}(x,y) \geq k+1+3\sqrt{\alpha/n}$.




Fig. 1. The vertices x and y have incentive to buy edge (u, v) .

Adding edge $\{v, u\}$ decreases the distance between x and y by at least $2\sqrt{\alpha/n}$, so $\text{Dist}_{G_s}(y)$ would decrease by at least $N_k(u) \cdot 2\sqrt{\alpha/n}$. Every node $y \in N_{\sqrt{\alpha/n}}(v)$ can pay at least $N_k(u) \cdot 2\sqrt{\alpha/n}$ for edge $\{v, u\}$. Because the edge $\{v, u\}$ is not bought, we have $\alpha \geq |N_k(u)| \cdot 2\sqrt{\alpha/n} \cdot |N_{\sqrt{\alpha/n}}(v)|$. Note that $|N_{\sqrt{\alpha/n}}(v)| \geq \sqrt{\alpha/n}$. Therefore $|N_k(u)| \leq n/2$, which is a contradiction. \square

Next we consider what happens with arbitrary neighborhoods, using techniques similar to [Demaine et al. 2007, Lemma 5].

Lemma 8. *If $|N_k(u)| \geq Y$ for every vertex u in an equilibrium graph G_s , then either $|N_{4k+2}(u)| > n/2$ for some vertex u or $|N_{5k+3}(u)| \geq Y^2 n/\alpha$ for every vertex u .*

Proof: If there is a vertex u with $|N_{4k+2}(u)| > n/2$, then the claim is obvious. Otherwise, for every vertex u , $|N_{4k+2}(u)| \leq n/2$. Let u be an arbitrary vertex. Let S be the set of vertices whose distance from u is $4k + 3$. We select a subset of S , called *center points*, by the following greedy algorithm. We repeatedly select an unmarked vertex $z \in S$ as a center point, mark all unmarked vertices in S whose distance from z is at most $2k$, and assign these vertices to z .

Suppose that we select l vertices x_1, x_2, \dots, x_l as center points. We prove that $l \geq |N_k(u)|n/\alpha$. Let C_i be the vertices in S assigned to x_i ; see Figure 2. By construction, $S = \bigcup_{i=1}^l C_i$. We also assign each vertex v at distance at least $4k + 4$ from u to one of these center points, as follows. Pick any one shortest path from v to u that contains some vertex $w \in S$, and assign v to the same center point as w . This vertex w is unique in this path because this path is a shortest path from v to u . Let T_i be the set of vertices assigned to x_i and whose distance from u is more than $4k + 2$. By construction, $\bigcup_{i=1}^l T_i$ is the set of vertices at distance more than $4k + 2$ from u . The shortest path from $v \in T_i$ to u uses some vertex $w \in C_i$. For any vertex x whose distance is at most k from u and for any $y \in T_i$, adding the edge $\{u, x_i\}$ decreases the distance between x and y at least 2, because the shortest path from $y \in T_i$ to u uses some vertex $w \in C_i$, as shown in Figure 2. By adding edge $\{u, x_i\}$, the distance between u and w would become at most $2k + 1$ and the distance between x and w would become at most $3k + 1$, where x is any vertex whose distance from u is at most k . Because the current distance between x and




Fig. 2. Center points.

w is at least $4k + 3 - k = 3k + 3$, adding the edge $\{u, x_i\}$ decreases this distance by at least 2. Consequently the distance between x and any $y \in T_i$ decreases by at least 2. Note that the distance between x and y is at least $d_{G_s}(u, y) - k$, and after adding edge (u, x_i) , this distance becomes at most

$$\begin{aligned} 3k + 1 + d_{G_s}(w, y) &= 3k + 1 + d_{G_s}(u, y) - d_{G_s}(u, w) \\ &= 3k + 1 + d_{G_s}(u, y) - (4k + 3) \\ &= d_{G_s}(u, y) - k - 2. \end{aligned}$$

Thus any vertex $y \in T_i$ has incentive to pay at least $2|N_k(u)|$ for edge $\{u, x_i\}$. Because the edge $\{u, x_i\}$ is not in equilibrium, we conclude that $\alpha \geq 2|T_i||N_k(u)|$. On the other hand, $|N_{4k+2}(u)| \leq n/2$, so $\sum_{i=1}^l |T_i| \geq n/2$. Therefore, $l\alpha \geq 2|N_k(u)| \sum_{i=1}^l |T_i| \geq n|N_k(u)|$ and hence $l \geq n|N_k(u)|/\alpha$.

According to the greedy algorithm, the distance between any pair of center points is more than $2k$; hence, $N_k(x_i) \cap N_k(x_j) = \emptyset$ for $i \neq j$. By the hypothesis of the lemma, $|N_k(x_i)| \geq Y$ for every vertex x_i ; hence $|\bigcup_{i=1}^l N_k(x_i)| = \sum_{i=1}^l |N_k(x_i)| \geq lY$. For every $i \leq l$, we have $d_{G_s}(u, x_i) = 4k + 3$, so vertex u has a path of length at most $5k + 3$ to every vertex whose distance to x_i is at most k . Therefore, $|N_{5k+3}(u)| \geq |\bigcup_{i=1}^l N_k(x_i)| \geq lY \geq Yn|N_k(u)|/\alpha \geq Y^2n/\alpha$. \square

Now we are ready to prove bounds on the price of anarchy. We start with the case when α is a bit smaller than n :

Theorem 9. *For $1 \leq \alpha < n^{1-\varepsilon}$, the price of anarchy is at most $O(1/\varepsilon^{1+\lg 5})$.*

Proof: Consider an equilibrium graph G_s . Let $X = n/\alpha > n^\varepsilon$. Define $a_1 = 2$ and $a_i = 4a_{i-1} + 3$, or equivalently $a_i = 3 \cdot 4^{i-1} - 1 < 4^i$, for all $i > 1$. By Lemma 8, for each $i \geq 1$, either $N_{3a_i+2}(v) > n/2$ for some vertex v or $N_{a_{i+1}} \geq (n/\alpha) N_{a_i}^2 = X N_{a_i}^2$. Let j be the least number for which $|N_{3a_j+2}(v)| > n/2$ for some vertex v . By this definition, for each $i < j$, $N_{a_{i+1}} > (n/\alpha) N_{a_i}^2 = X N_{a_i}^2$. Because $N_{a_1} > 1$, we obtain that $N_{a_i} > X^{2^{i-1}-1}$ for every $i \leq j$. On the other hand, $X^{2^{j-1}-1} < N_{a_j} \leq n$, which implies $2^{j-1} - 1 < 1/\varepsilon$. Thus $j < \lg(1/\varepsilon + 1) + 1 \leq \lg(1/\varepsilon) + 2$ and $a_j < 4^{\lg(1/\varepsilon)+2} = 16/\varepsilon^2$. Therefore $N_{3 \cdot 16/\varepsilon^2+2} > n/2$ and using Lemma 6, we conclude that the usage cost is at most $O(\frac{n^2}{\varepsilon^2} + n^2 + \alpha n)$. Now set $g = 2/\varepsilon$ in Lemma 4. Then the total cost is $O(\alpha n^{1+\varepsilon}) + \frac{2}{\varepsilon} O(\frac{n^2}{\varepsilon^2} + n^2 + \alpha n)$. The cost of the social optimum is at least $\Omega(\alpha n + n^2)$, and the value of α is less than $n^{1-\varepsilon}$. Therefore the price of anarchy is $O(1/\varepsilon^3)$. \square

Next we prove a polylogarithmic bound on the price of anarchy when α is close to n .

Theorem 10. *For $\alpha = O(n)$, the price of anarchy is $O(\lg^{1+\lg 5} n)$ and the diameter of any equilibrium graph is $O(\lg^{1+\lg 5} n)$.*

Proof: Consider an equilibrium graph G_s . The proof is similar to the proof of Theorem 9. Define $a_1 = \max\{2, 2\alpha/n\} + 1$ and $a_i = 5a_{i-1} + 3$, or equivalently $a_i = \frac{4a_1+3}{20} \cdot 5^i - \frac{3}{4} < a_1 5^i$, for all $i > 1$. By Lemma 8, for each $i \geq 1$, either $N_{4a_i+2}(v) > n/2$ for some vertex v or $N_{a_{i+1}} \geq (n/\alpha) N_{a_i}^2$. Let j be the least number for which $|N_{4a_j+2}(v)| > n/2$ for some vertex v . By this definition, for each $i < j$, $N_{a_{i+1}} \geq (n/\alpha) N_{a_i}^2$. Because $N_{a_1} > 2 \max\{1, \alpha/n\}$, we obtain that $N_{a_i} > 2^{2^{i-1}} \max\{1, \alpha/n\}$ for every $i \leq j$. On the other hand, $2^{2^{j-1}} \leq 2^{2^{j-1}} \max\{1, \alpha/n\} < N_{a_j} \leq n$, so $j < \lg \lg n + 1$ and $a_j < a_1 5^{\lg \lg n+1} < (2 + 2\alpha/n + 1 + 1) 5 \lg^{1+\lg 5} n = 10(2 + \alpha/n) \lg^{1+\lg 5} n$. Therefore $N_{4 \cdot [10(2+\alpha/n) \lg^{1+\lg 5} n]+2}(v) > n/2$ for some vertex v and using Lemma 5, we conclude that the distance of v to all other vertices is at most $2[40(2 + \alpha/n) \lg^{1+\lg 5} n + 2] + 2\alpha/n$. Thus the diameter of G_s is at most $O((1 + \alpha/n) \lg^{1+\lg 5} n)$. Setting $g = \lg n$ in Lemma 4, the cost of G_s is at most $\alpha O(n) + (\lg n) O(n^2(1 + \alpha/n) \lg^{1+\lg 5} n) = O((\alpha n + n^2) \lg^{1+\lg 5} n)$. Therefore the price of anarchy is at most $O(\lg^{1+\lg 5} n)$. \square

When α is a bit larger than n , we can obtain a constant bound on the price of anarchy. First we need a somewhat stronger result on the behavior of neighborhoods:

Lemma 11. *If $|N_k(u)| \geq Y$ for every vertex u in an equilibrium graph G_s , then either $|N_{5k}(u)| > n/2$ for some vertex u or $|N_{6k+1}(u)| \geq Y^2 kn/2\alpha$ for every vertex u .*

Proof: The proof is similar to the proof of Lemma 8. If there is a vertex u with $|N_{5k}(u)| > n/2$, then the claim is obvious. Otherwise, for every vertex u , $|N_{5k}(u)| \leq n/2$. Let u be an arbitrary vertex. Let S be the set of vertices whose distance from u is $5k+1$. We select a subset of S as we did in the proof of Lemma 8.

Suppose that we select l vertices x_1, x_2, \dots, x_l as center points. We prove that $l \geq |N_k(u)| kn/2\alpha$. Let C_i be the vertices in S assigned to x_i . By construction,

$S = \bigcup_{i=1}^l C_i$. We also assign each vertex v at distance at least $5k + 2$ from u to one of these center points, as follows. Pick any one shortest path from v to u that contains a vertex $w \in S$, and assign v to the same center point as w . This vertex w is unique in this path because this path is a shortest path from v to u . Let T_i be the set of vertices assigned to x_i and whose distance from u is more than $5k$. By construction, $\bigcup_{i=1}^l T_i$ is the set of vertices at distance more than $5k$ from u . The shortest path from $v \in T_i$ to u uses some vertex $w \in C_i$. For any vertex x whose distance is at most k from u and any $y \in T_i$, adding the edge $\{u, x_i\}$ decreases the distance between x and y at least k , because the shortest path from $y \in T_i$ to x uses some vertex $w \in C_i$. By inserting edge $\{u, x_i\}$, the distance between u and w would become at most $2k + 1$ and the distance between x and w would become at most $3k + 1$, where x is any vertex whose distance from u is at most k . Because the current distance between x and w is at least $5k + 1 - k = 4k + 1$, adding edge $\{u, x_i\}$ decreases this distance by at least k . Consequently the distance between x and any $y \in T_i$ decreases by at least k .

Thus any vertex $y \in T_i$ has incentive to pay at least $k|N_k(u)|$ for edge $\{u, x_i\}$. Because the edge $\{u, x_i\}$ is not in equilibrium, we conclude that $\alpha \geq k|T_i||N_k(u)|$. On the other hand, $|N_{4k}(u)| \leq n/2$, so $\sum_{i=1}^l |T_i| \geq n/2$. Therefore, $l\alpha \geq k|N_k(u)| \sum_{i=1}^l |T_i| \geq kn|N_k(u)|/2$ and hence $l \geq kn|N_k(u)|/2\alpha$.

According to the greedy algorithm, the distance between any pair of center points is more than $2k$; hence, $N_k(x_i) \cap N_k(x_j) = \emptyset$ for $i \neq j$. By the hypothesis of the lemma, $|N_k(x_i)| > Y$ for every vertex x_i ; hence $|\bigcup_{i=1}^l N_k(x_i)| = \sum_{i=1}^l |N_k(x_i)| > lY$. For every $i \leq l$, we have $d_{G_s}(u, x_i) = 5k + 1$, so vertex u has a path of length at most $6k + 1$ to every vertex whose distance to x_i is at most k . Therefore, $|N_{6k+1}(u)| \geq |\bigcup_{i=1}^l N_k(x_i)| > lY \geq Ykn|N_k(u)|/2\alpha > Y^2kn/2\alpha$. \square

Theorem 12. *For any $\alpha > n$, the price of anarchy is $O(\sqrt{n/\alpha} \lg^{1+\lg 6} n)$ and the diameter of any equilibrium graph is $O(\lg^{\lg 6} n \cdot \sqrt{\alpha/n})$.*

Proof: Consider an equilibrium graph G_s . The proof is similar to the proof of Theorem 10. Define $a_1 = 2\sqrt{\alpha/n} + 1$ and $a_i = 6a_{i-1} + 1$, or equivalently $a_i = \frac{5a_1+1}{30} \cdot 6^i - \frac{1}{5} < a_1 6^i$, for all $i > 1$. By Lemma 11, for each $i \geq 1$, either $|N_{5a_i}(v)| > n/2$ for some vertex v or $N_{a_{i+1}} \geq (n/2\alpha) a_i N_{a_i}^2$. Let j be the least number for which $|N_{5a_j}(v)| > n/2$ for some vertex v . By this definition, for each $i < j$, $N_{a_{i+1}} > (n/2\alpha) a_i N_{a_i}^2$. Because $N_{a_1} > 2\sqrt{\alpha/n}$ and $a_i \geq a_1$, we obtain that $N_{a_i} > 2^{2^{i-1}} \sqrt{\alpha/n}$ for every $i \leq j$. On the other hand, $2^{2^{j-1}} \leq 2^{2^{j-1}} \sqrt{\alpha/n} < N_{a_j} \leq n$, so $j < \lg \lg n + 1$ and $a_j < a_1 6^{\lg \lg n + 1} = 6a_1 \lg^{\lg 6} n < 18\sqrt{\alpha/n} \lg^{\lg 6} n$. Therefore $|N_{5 \cdot 18\sqrt{\alpha/n} \lg^{\lg 6} n}(v)| > n/2$ for some vertex v and using Lemma 7, we conclude that the distance of v to other vertices is at most $10[18\sqrt{\alpha/n} \lg^{\lg 6} n] + 4\sqrt{\alpha/n}$. Thus the diameter of G_s is at most $O(\lg^{\lg 6} n \cdot \sqrt{\alpha/n})$. Setting $g = \lg n$ in Lemma 4, the cost of G_s is at most $\alpha O(n) + (\lg n)O(n^2 \lg^{\lg 6} n \cdot \sqrt{\alpha/n})$. Therefore the price of anarchy is at most $\frac{\alpha O(n) + O(n^2 \lg^{1+\lg 6} n \sqrt{\alpha/n})}{\alpha n} = O(\sqrt{n/\alpha} \lg^{1+\lg 6} n)$. \square

By Theorem 12, we conclude the following:

Corollary 13. *For $\alpha = \Omega(n \lg^{2+2\lg 6} n) \approx \Omega(n \lg^{7.16} n)$, the price of anarchy*

is $O(1)$.

5. COOPERATIVE VERSION IN GENERAL GRAPHS

In this section, we study the price of anarchy when only some links can be created, e.g., because of physical limitations. In this case, the social optimum is no longer simply a clique or a star.

We start by bounding the growth of distances from the host graph G to an arbitrary equilibrium graph G_s :

Lemma 14. *For any two vertices u and v in any equilibrium graph G_s , $d_{G_s}(u, v) = O(d_G(u, v) + \alpha^{1/3}d_G(u, v)^{2/3})$.*

Proof: Let $u = v_0, v_1, \dots, v_k = v$ be a shortest path in G between u and v , so $k = d_G(u, v)$. Suppose that the distance between v_0 and v_i in G_s is d_i , for $0 \leq i \leq k$. We first prove that $d_{i+1} \leq d_i + 1 + \sqrt{9\alpha/d_i}$ for $0 \leq i < k$. If edge $\{v_i, v_{i+1}\}$ already exists in G_s , the inequality clearly holds. Otherwise, adding this edge decreases the distance between x and y by at least $\frac{d_{i+1}-d_i}{3}$, where x is a vertex whose distance is at most $\frac{d_{i+1}-d_i}{3} - 1$ from v_{i+1} and y is a vertex in a shortest path from v_i to v_0 . Therefore any vertex x whose distance is at most $\frac{d_{i+1}-d_i}{3} - 1$ from v_{i+1} can pay $\frac{d_{i+1}-d_i}{3}d_i$ for this edge. Because this edge does not exist in G_s and because there are at least $\frac{d_{i+1}-d_i}{3}$ vertices of distance at most $\frac{d_{i+1}-d_i}{3} - 1$ from v_{i+1} , we conclude that $\left(\frac{d_{i+1}-d_i}{3}\right)^2 d_i \leq \alpha$. Thus we have $d_{i+1} \leq d_i + 1 + \sqrt{9\alpha/d_i}$ for $0 \leq i < k$. Next we prove that $d_{i+1} \leq d_i + 1 + 5\alpha^{1/3}$. If edge $\{v_i, v_{i+1}\}$ already exists in G_s , the inequality clearly holds. Otherwise, adding this edge decreases the distance between z and w by at least $\frac{d_{i+1}-d_i}{5}$, where z and w are two vertices whose distances from v_{i+1} and v_i , respectively, are less than $\frac{d_{i+1}-d_i}{5}$. There are at least $\left(\frac{d_{i+1}-d_i}{5}\right)^2$ pair of vertices like (z, w) . Because the edge $\{v_i, v_{i+1}\}$ does not exist in G_s , we conclude that $\left(\frac{d_{i+1}-d_i}{5}\right)^3 \leq \alpha$. Therefore $d_{i+1} \leq d_i + 1 + 5\alpha^{1/3}$.

Combining these two inequalities, we obtain $d_{i+1} \leq d_i + 1 + \min\{\sqrt{9\alpha/d_i}, 5\alpha^{1/3}\}$.

Inductively we prove that $d_j \leq 3j + 7\alpha^{1/3} + 5\alpha^{1/3}j^{2/3}$. For $j \leq 2$, the inequality is clear. Now suppose by induction that $d_j \leq 3j + 7\alpha^{1/3} + 5\alpha^{1/3}j^{2/3}$. If $d_j \leq 2\alpha^{1/3}$, we reach the desired inequality using the inequality $d_{j+1} \leq d_j + 1 + 5\alpha^{1/3}$. Otherwise, we know that $d_{j+1} \leq d_j + 1 + \sqrt{9\alpha/d_j} = f(d_j)$ and to find the maximum of the function $f(d_j)$ over the domain $d_j \in [2\alpha^{1/3}, j + 7\alpha^{1/3} + 5\alpha^{1/3}j^{2/3}]$, we should check f 's critical points, including the endpoints of the domain interval and where f 's derivative is zero. We reach three values for d_j : $2\alpha^{1/3}$, $j + 7\alpha^{1/3} + 5\alpha^{1/3}j^{2/3}$, and $\left(\frac{9\alpha}{4}\right)^{1/3}$. Because the third value is not in the domain, we just need to check the first two values. The first value is also checked, so just the second value remains. For the second value, we have

$$\begin{aligned} d_{j+1} &\leq d_j + 1 + \sqrt{9\alpha/d_j} \\ &\leq j + 7\alpha^{1/3} + 5\alpha^{1/3}j^{2/3} + 1 + \sqrt{\frac{9\alpha}{j + 7\alpha^{1/3} + 5\alpha^{1/3}j^{2/3}}} \\ &\leq j + 1 + 7\alpha^{1/3} + 5\alpha^{1/3}j^{2/3} + \sqrt{\frac{10\alpha}{5\alpha^{1/3}j^{2/3}}} \end{aligned}$$

$$\leq j + 1 + 7\alpha^{1/3} + 5\alpha^{1/3}j^{2/3} + \frac{\alpha^{1/3}\sqrt{2}}{j^{1/3}}.$$

Because $(j+1)^{2/3} - j^{2/3} = \frac{(j+1)^2 - j^2}{(j+1)^{4/3} + (j+1)^{2/3}j^{2/3} + j^{4/3}} \geq \frac{2j}{3(j+1)^{4/3}}$, we have

$$\begin{aligned} & j + 1 + 7\alpha^{1/3} + 5\alpha^{1/3}j^{2/3} + \frac{\alpha^{1/3}\sqrt{2}}{j^{1/3}} \\ & \leq j + 1 + 7\alpha^{1/3} + 5\alpha^{1/3}(j+1)^{2/3} - 5\alpha^{1/3}\frac{2j}{3(j+1)^{4/3}} + \frac{\alpha^{1/3}\sqrt{2}}{j^{1/3}} \\ & \leq j + 1 + 7\alpha^{1/3} + 5\alpha^{1/3}(j+1)^{2/3} - \frac{10\alpha^{1/3}j}{3j^{4/3}} + \frac{\alpha^{1/3}\sqrt{2}}{j^{1/3}} \\ & \leq j + 1 + 7\alpha^{1/3} + 5\alpha^{1/3}(j+1)^{2/3}. \end{aligned}$$

Note that $j+1 > 2$ and $d_k = d_{G_s}(u, v)$. Therefore $d_{G_s}(u, v)$ is at most $O(d_G(u, v) + \alpha^{1/3}d_G(u, v)^{2/3})$ and the desired inequality is proved. \square

Using this Lemma 14, we prove two different bounds relating the sum of all pairwise distances in the two graphs:

Corollary 15. *For any equilibrium graph G_s , $\sum_{u,v \in V(G)} d_{G_s}(u, v) = O(\alpha^{1/3}) \cdot \sum_{u,v \in V(G)} d_G(u, v)$.*

Theorem 16. *For any equilibrium graph G_s , $\sum_{u,v \in V(G)} d_{G_s}(u, v) \leq \min\{n^3, O(n^{1/3})(\alpha n + \sum_{u,v \in V(G)} d_G(u, v))\}$.*

Proof: We partition pairs of vertices into two parts. The first part contains pairs with distance at most $\frac{\alpha}{n}$ in G . The second part contains pairs with distance more than $\frac{\alpha}{n}$ in G .

$$\begin{aligned} & \sum_{u,v \in V(G_s)} d_{G_s}(u, v) \\ & \leq \sum_{u,v \in V(G)} O(d_G(u, v) + \alpha^{1/3}d_G(u, v)^{2/3}) \\ & \leq \sum_{d_G(u,v) \leq \frac{\alpha}{n}} O(d_G(u, v) + \alpha^{1/3}d_G(u, v)^{2/3}) + \sum_{d_G(u,v) \geq \frac{\alpha}{n}} O(d_G(u, v) + \alpha^{1/3}d_G(u, v)^{2/3}) \\ & \leq O(n^2\alpha^{1/3}(\alpha/n)^{2/3}) + \sum_{u,v \in V(G)} O\left(\frac{d_G(u, v)\alpha^{1/3}}{(\alpha/n)^{1/3}}\right) \\ & \leq O(n^{1/3})\alpha n + O(n^{1/3}) \sum_{u,v \in V(G)} d_G(u, v) \end{aligned}$$

On the other hand, we know that $\sum_{u,v \in V(G_s)} d_{G_s}(u, v)$ is at most n^3 for every connected graph G_s . Therefore we have the desired property for the sum of distances in G_s . \square

Now we can bound the price of anarchy for the various ranges of α , combining Corollary 15, Theorem 16, and Lemma 4, with different choices of g .

Theorem 17. *In the cooperative network creation game in general graphs, the price of anarchy is at most*

(a) $O(\alpha^{1/3})$ for $\alpha < n$,

- (b) $O(n^{1/3})$ for $n \leq \alpha \leq n^{5/3}$,
 (c) $O(\frac{n^2}{\alpha})$ for $n^{5/3} \leq \alpha < n^{2-\varepsilon}$, and
 (d) $O(\frac{n^2}{\alpha} \lg n)$ for $n^2 \leq \alpha$.

Proof: (a) By setting $g = 6$ in Lemma 4, the total cost is at most $\alpha O(n^{4/3}) + 6 \sum_{u,v \in V(G_s)} d_{G_s}(u,v) \leq O(\alpha^{1/3} n^2) + 6 \sum_{u,v \in V(G_s)} d_{G_s}(u,v)$. Using Corollary 15, $\sum_{u,v \in V(G_s)} d_{G_s}(u,v) = O(\alpha^{1/3}) \sum_{u,v \in V(G)} d_G(u,v)$. Thus the total cost is at most $O(\alpha^{1/3} n^2) + O(\alpha^{1/3}) \sum_{u,v \in V(G)} d_G(u,v)$, which is at most $O(\alpha^{1/3})$ times the optimum cost.

(b) By setting $g = 6$ in Lemma 4, the total cost is at most $\alpha O(n^{4/3}) + 6 \sum_{u,v \in V(G_s)} d_{G_s}(u,v)$. Using Theorem 16, $\sum_{u,v \in V(G_s)} d_{G_s}(u,v) = O(n^{1/3})(\alpha n + \sum_{u,v \in V(G)} d_G(u,v))$. Thus the total cost is at most $\alpha O(n^{4/3}) + O(n^{1/3})(\alpha n + \sum_{u,v \in V(G)} d_G(u,v))$. The cost of the social optimum is $\Omega(\alpha n + \sum_{u,v \in V(G)} d_G(u,v))$, so the price of anarchy is at most $O(n^{1/3})$.

(c) By setting $g = 2/\varepsilon$ in Lemma 4, the total cost is at most $\alpha O(n^{1+\varepsilon}) + \frac{2}{\varepsilon} \sum_{u,v \in V(G_s)} d_{G_s}(u,v)$. Using Theorem 16, $\sum_{u,v \in V(G_s)} d_{G_s}(u,v) = O(n^3)$. Thus the total cost is at most $\alpha O(n^{1+\varepsilon}) + \frac{2}{\varepsilon} O(n^3)$. Because the cost of the social optimum is $\Omega(\alpha n + n^2)$, the price of anarchy is at most $O(\max\{n^\varepsilon, \frac{2}{\varepsilon} \frac{n^2}{\alpha}\}) = O(\frac{2}{\varepsilon} \frac{n^2}{\alpha}) = O(\frac{n^2}{\alpha})$.

(d) By setting $g = \lg n$ in Lemma 4, the total cost is at most $\alpha O(n) + \lg n \sum_{u,v \in V(G_s)} d_{G_s}(u,v)$. Using Theorem 16, $\sum_{u,v \in V(G_s)} d_{G_s}(u,v) = O(n^3)$. Thus the total cost is at most $\alpha O(n) + (\lg n) O(n^3)$. Because the cost of the social optimum is $\Omega(\alpha n + n^2)$, the price of anarchy is at most $O(\frac{n^2}{\alpha} \lg n)$. \square

6. UNILATERAL VERSION IN GENERAL GRAPHS

Next we consider how a general host graph affects the unilateral version of the problem.

Lemma 18. *For any two vertices u and v in any equilibrium graph G_s , $d_{G_s}(u,v) = O(d_G(u,v) + \alpha^{1/2} d_G(u,v)^{1/2})$.*

Proof: Similar to the proof of Lemma 14, we define the sequence d_i , $0 \leq i \leq k$. We prove that $d_{i+1} \leq d_i + 1 + \frac{\alpha}{d_i}$ for $0 \leq i < k$. If edge $\{v_i, v_{i+1}\}$ already exists in G_s , the inequality clearly holds. Otherwise, adding this edge decreases the distance between v_{i+1} and x by at least $d_{i+1} - d_i - 1$, where x is a vertex in a shortest path from v_i to v_0 . Therefore v_{i+1} can pay $(d_{i+1} - d_i - 1)d_i$ for this edge. Because this edge does not exist in G_s , we conclude that $(d_{i+1} - d_i - 1)d_i \leq \alpha$. Thus we have $d_{i+1} \leq d_i + 1 + \frac{\alpha}{d_i}$ for $0 \leq i < k$. On the other hand, we can prove that $d_{i+1} \leq d_i + 1 + 3\alpha^{1/2}$. If edge $\{v_i, v_{i+1}\}$ already exists in G_s , clearly the inequality holds. Otherwise, adding this edge decreases the distance between v_{i+1} and y by at least $\frac{d_{i+1} - d_i}{3}$ where y is a vertex whose distance is less than $\frac{d_{i+1} - d_i}{3}$ from v_i . There are at least $\frac{d_{i+1} - d_i}{3}$ vertices like y . Because the edge $\{v_i, v_{i+1}\}$ does not exist in G_s , $\left(\frac{d_{i+1} - d_i}{3}\right)^2 \leq \alpha$. Therefore $d_{i+1} \leq d_i + 3\alpha^{1/2}$. Combining these two inequalities, we obtain $d_{i+1} \leq d_i + 1 + \min\{1 + \frac{\alpha}{d_i}, 3\alpha^{1/2}\}$.

Inductively we prove that $d_j \leq j + 4\alpha^{1/2} + 2\alpha^{1/2}j^{1/2}$. For $j = 0$, the inequality is clear. Now suppose by induction that $d_j \leq j + 4\alpha^{1/2} + 2\alpha^{1/2}j^{1/2}$. If $d_j \leq \alpha^{1/2}$, we reach the desired inequality using the inequality $d_{j+1} \leq d_j + 1 + 3\alpha^{1/2}$. Otherwise, we know that $d_{j+1} \leq d_j + 1 + \frac{\alpha}{d_j} = f(d_j)$ and to find the maximum of the function $f(d_j)$ over the domain $d_j \in [\alpha^{1/2}, j + 4\alpha^{1/2} + 2\alpha^{1/2}j^{1/2}]$, we should check its critical points including endpoints of the domain interval and where its derivative is zero. We reach three values for d_j : $\alpha^{1/2}$, $j + 4\alpha^{1/2} + 2\alpha^{1/2}j^{1/2}$, and $\alpha^{1/2}$. The first and third values are checked, so just the second value remains. For the second value, we have

$$\begin{aligned} d_{j+1} &\leq d_j + 1 + \frac{\alpha}{d_j} \\ &\leq j + 4\alpha^{1/2} + 2\alpha^{1/2}j^{1/2} + 1 + \frac{\alpha}{j + 4\alpha^{1/2} + 2\alpha^{1/2}j^{1/2}} \\ &\leq j + 1 + 4\alpha^{1/2} + 2\alpha^{1/2}j^{1/2} + \frac{\alpha^{1/2}}{2j^{1/2}}. \end{aligned}$$

Because $(j+1)^{1/2} - j^{1/2} = \frac{1}{(j+1)^{1/2} + j^{1/2}} \geq \frac{1}{2j^{1/2}}$, we have

$$\begin{aligned} &j + 1 + 4\alpha^{1/2} + 2\alpha^{1/2}j^{1/2} + \frac{\alpha^{1/2}}{2j^{1/2}} \\ &\leq j + 1 + 4\alpha^{1/2} + 2\alpha^{1/2}(j+1)^{1/2} - 2\alpha^{1/2} \frac{1}{2j^{1/2}} + \frac{\alpha^{1/2}}{2j^{1/2}} \\ &\leq j + 1 + 4\alpha^{1/2} + 2\alpha^{1/2}(j+1)^{1/2}. \end{aligned}$$

Note that $d_k = d_{G_s}(u, v)$. Therefore $d_{G_s}(u, v)$ is at most $O(d_G(u, v) + \alpha^{1/2}d_G(u, v)^{1/2})$ and the desired inequality is proved. \square

Again we relate the sum of all pairwise distances in the two graphs:

Corollary 19. *For any equilibrium graph G_s , $\sum_{u, v \in V(G)} d_{G_s}(u, v) = O(\alpha^{1/2}) \cdot \sum_{u, v \in V(G)} D_G(u, v)$.*

Theorem 20. *For any equilibrium graph G_s , $\sum_{u, v \in V(G_s)} d_{G_s}(u, v) \leq \min\{O(n^{1/2})(\alpha n + \sum_{u, v \in V(G)} D_G(u, v)), n^3\}$.*

Proof: We partition pairs of vertices into two parts. The first part contains pairs with distance at most $\frac{\alpha}{n}$ in G . The second part contains pairs with distance more than $\frac{\alpha}{n}$ in G .

$$\begin{aligned} \sum_{u, v \in V(G_s)} D_{G_s}(u, v) &\leq \sum_{u, v \in V(G)} O(D_G(u, v) + \alpha^{1/2}D_G(u, v)^{1/2}) \\ &\leq \sum_{D_G(u, v) \leq \frac{\alpha}{n}} O(D_G(u, v) + \alpha^{1/2}D_G(u, v)^{1/2}) \\ &\quad + \sum_{D_G(u, v) \geq \frac{\alpha}{n}} O(D_G(u, v) + \alpha^{1/2}D_G(u, v)^{1/2}) \\ &\leq O(n^2\alpha^{1/2}(\frac{\alpha}{n})^{1/2}) + \sum_{u, v \in V(G)} O\left(\frac{D_G(u, v)\alpha^{1/2}}{(\alpha/n)^{1/2}}\right) \\ &\leq O(n^{1/2})\alpha n + O(n^{1/2}) \sum_{u, v \in V(G)} D_G(u, v) \end{aligned}$$

On the other hand, we know $\sum_{u,v \in V(G_s)} d_{G_s}(u,v)$ is at most n^3 for every connected graph G_s . Therefore we have the desired property for the sum of distances in G_s . \square

To conclude bounds on the price of anarchy, we now use Lemma 2 in place of Lemma 4, combined with Corollary 19 and Theorem 20.

Theorem 21. *For $\alpha \geq n$, the price of anarchy is at most $\min\{O(n^{1/2}), \frac{n^2}{\alpha}\}$.*

Theorem 22. *For $\alpha < n$, the price of anarchy is at most $O(\alpha^{1/2})$.*

7. LOWER BOUNDS IN GENERAL GRAPHS

In this section, we prove polynomial lower bounds on the price of anarchy for general host graphs, first for the cooperative version and second for the unilateral version.

Theorem 23. *The price of anarchy in the cooperative game is $\Omega(\min\{\sqrt{\frac{\alpha}{n}}, \frac{n^2}{\alpha}\})$.*

Proof: For $\alpha = O(n)$ or $\alpha = \Omega(n^2)$, the claim is clear. Otherwise, let $k = \sqrt{\frac{\alpha}{12n}} \geq 2$. Thus $k = O(\sqrt{n})$. We construct graph $G_{k,l}$ as follows; see Figure 3. Start with $2l$ vertices v_1, v_2, \dots, v_{2l} connected in a cycle. For any $1 \leq i \leq 2l$, insert a path P_i of k edges between v_i and v_{i+1} (where we define $v_{2l+1} = v_1$). For any $1 \leq i \leq l$, insert a path Q_i of k edges between v_{2i} and v_{2i+2} (where we define $v_{2l+2} = v_2$). Therefore there are $n = (3k - 1)l$ vertices and $(3k + 2)l$ edges in $G_{k,l}$, so $l = n/(3k - 1)$.

For simplicity, let G denote $G_{k,l}$ in the rest of the proof. Let G_1 be a spanning connected subgraph of G that contains exactly one cycle, namely, $(v_1, v_2, \dots, v_{2l}, v_1)$; in other words, we remove from G exactly one edge from each path P_i and Q_i . Let G_2 be a spanning connected subgraph of G that contains exactly one cycle, formed by the concatenation of Q_1, Q_2, \dots, Q_l , and contains none of the edges $\{v_i, v_{i+1}\}$, for $1 \leq i \leq 2l$; for example, we remove from G exactly one edge from every P_{2i} and every edge $\{v_i, v_{i+1}\}$.

Next we prove that G_2 is an equilibrium. For any $1 \leq i \leq l$, removing any edge of path Q_i increases the distance between its endpoints and at least $n/6$ vertices by at least $\frac{lk}{3} \geq n/6$. Because $\alpha = o(n^2)$, we have $\alpha < \frac{n}{6} \frac{n}{6}$, so if we assign this edge to be bought solely by one of its endpoints, then this owner will not delete the edge. Removing other edges makes G_2 disconnected. For any $1 \leq i \leq l$, adding an edge of path P_{2i} or path P_{2i+1} or edge $\{v_{2i}, v_{2i+1}\}$ or edge $\{v_{2i+1}, v_{2i+2}\}$ to G_2 decreases only the distances from some vertices of paths P_{2i} or P_{2i+1} to the other vertices.




Fig. 3. Lower bound graph.

There are at most $n(2k - 1)$ such pairs. Adding such an edge can decrease each of these distance by at most $3k - 1$. But we know that $\alpha \geq 12nk^2 > 2n(2k - 1)(3k - 1)$, so the price of the edge is more than its total benefit among all nodes, and thus the edge will not be created by any coalition.

The cost of G_1 is equal to $O(\alpha n + n^2(k + l)) = O(\alpha n + n^2(k + \frac{n}{k}))$ and the cost of G_2 is $\Omega(\alpha n + n^2(k + lk)) = \Omega(\alpha n + n^3)$. The cost of the social optimum is at most the cost of G_1 , so the price of anarchy is at least $\Omega(\frac{n^3}{\alpha n + n^3/k + kn^2}) = \Omega(\min\{\frac{n^2}{\alpha}, k, \frac{n}{k}\})$. Because $k = O(\sqrt{n})$, the price of anarchy is at least $\Omega(\min\{\frac{n^2}{\alpha}, k\}) = \Omega(\min\{\frac{n^2}{\alpha}, \sqrt{\frac{\alpha}{n}}\})$. \square

Theorem 24. *The price of anarchy in unilateral games is $\Omega(\min\{\frac{\alpha}{n}, \frac{n^2}{\alpha}\})$.*

Proof: The proof is similar to the proof of Theorem 23. For $\alpha = O(n)$ or $\alpha = \Omega(n^2)$, the claim is clear. Otherwise, let k be the biggest number for which the inequality $3nk < \alpha < \frac{n}{6}(\frac{n}{9k} - 1)$ holds. Therefore $k = O(\sqrt{n})$. Let $l = n/(3k - 1)$. Again consider the host graph $G_{k,l}$ and the subgraphs G_1 and G_2 , as defined in the proof of Theorem 23.

Next we prove that G_2 is an equilibrium. For any $1 \leq i \leq l$, removing any edge of path Q_i increases the distance between its endpoints and at least $n/6$ vertices by at least $\frac{lk}{3} \geq \frac{n}{9} - 1$. Because $\alpha < \frac{n}{6}(\frac{n}{9k} - 1)$, the owner of such an edge will not delete it. Removing any other edge disconnects G_2 . Adding any edge to G_2 decreases the distance of its endpoints to other vertices at most $3k - 2$ because any edge in $G - G_2$ forms a cycle of length at most $3k$ with edges in G_2 . But we know that $\alpha > 3nk$, so neither endpoint will create this edge.

The cost of G_1 is equal to $O(\alpha n + n^2(k + l)) = O(\alpha n + n^2(k + \frac{n}{k}))$ and the cost of G_2 is $\Omega(\alpha n + n^3)$. The cost of the social optimum is at most the cost of G_1 , so the price of anarchy is at least $\Omega(\frac{n^3}{\alpha n + n^3/k + n^2k}) = \Omega(\min\{\frac{n^2}{\alpha}, k, \frac{n}{k}\})$. Because $k = O(\sqrt{n})$ and $k = \Theta(\frac{\alpha}{3n})$, the price of anarchy is at least $\Omega(\min\{\frac{n^2}{\alpha}, k\}) = \Omega(\min\{\frac{n^2}{\alpha}, \frac{\alpha}{n}\})$. \square

8. OPEN PROBLEMS

For cooperative network creation games, the main direction for future research is to determine whether the price of anarchy for the complete graph is constant or polylogarithmic (or somewhere in between). For network creation games in general, the main open question remains to determine the price of anarchy for the unilateral network creation game in complete graphs, in particular, whether it is polylogarithmic.

REFERENCES

- ALBERS, S. 2008. On the value of coordination in network design. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms*. 294–303.
- ALBERS, S., EILTS, S., EVEN-DAR, E., MANSOUR, Y., AND RODITTY, L. 2006. On Nash equilibria for a network creation game. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*. Miami, FL, 89–98.
- ANDELMAN, N., FELDMAN, M., AND MANSOUR, Y. 2007. Strong price of anarchy. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms*. 189–198.
- ACM SIGecom Exchanges, Vol. 8, No. 2, December 2009.

- ANSHELEVICH, E., DASGUPTA, A., KLEINBERG, J., TARDOS, E., WEXLER, T., AND ROUGHGARDEN, T. 2004. The price of stability for network design with fair cost allocation. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*. 295–304.
- ANSHELEVICH, E., DASGUPTA, A., TARDOS, E., AND WEXLER, T. 2003. Near-optimal network design with selfish agents. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*. San Diego, CA, 511–520.
- AUMANN, R. J. 1959. Acceptable points in general cooperative n -person games. In *Contributions to the Theory of Games*. Vol. 4. Princeton University Press.
- BARTAL, Y. 1998. On approximating arbitrary metrics by tree metrics. In *Proceedings of the 13th Annual ACM Symposium on Theory of Computing*. Dallas, Texas, 161–168.
- CHUN, B.-G., FONSECA, R., STOICA, I., AND KUBIATOWICZ, J. 2004. Characterizing selfishly constructed overlay routing networks. In *Proceedings of the 23rd Annual Joint Conference of the IEEE Computer and Communications Societies*. Hong Kong, China.
- CORBO, J. AND PARKES, D. 2005. The price of selfish behavior in bilateral network formation. In *Proceedings of the 24th Annual ACM Symposium on Principles of Distributed Computing*. Las Vegas, Nevada, 99–107.
- CZUMAJ, A. AND VÖCKING, B. 2002. Tight bounds for worst-case equilibria. In *Proceedings of the 13th Annual ACM-SIAM symposium on Discrete Algorithms*. San Francisco, California, 413–420.
- DEMAINE, E. D., HAJIAGHAYI, M., MAHINI, H., AND ZADIMOGHADDAM, M. 2007. The price of anarchy in network creation games. In *Proceedings of the 26th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing*. 292–298. To appear in *ACM Transactions on Algorithms*.
- DUTTON, R. D. AND BRIGHAM, R. C. 1991. Edges in graphs with large girth. *Graphs and Combinatorics* 7, 4 (December), 315–321.
- ELKIN, M., EMEK, Y., SPIELMAN, D. A., AND TENG, S.-H. 2005. Lower-stretch spanning trees. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*. Baltimore, MD, 494–503.
- EVEN-DAR, E. AND KEARNS, M. 2006. A small world threshold for economic network formation. In *Proceedings of the 20th Annual Conference on Neural Information Processing Systems* (December 4-7, 2006). 385–392.
- FABRIKANT, A., LUTHRA, A., MANEVA, E., PAPADIMITRIOU, C. H., AND SHENKER, S. 2003. On a network creation game. In *Proceedings of the 22nd Annual Symposium on Principles of Distributed Computing*. Boston, Massachusetts, 347–351.
- FAKCHAROENPHOL, J., RAO, S., AND TALWAR, K. 2004. A tight bound on approximating arbitrary metrics by tree metrics. *Journal of Computer and System Sciences* 69, 3, 485–497.
- HALEVI, Y. AND MANSOUR, Y. 2007. A network creation game with nonuniform interests. In *Proceedings of the 3rd International Workshop on Internet and Network Economics*. Lecture Notes in Computer Science, vol. 4858. San Diego, CA, 287–292.
- JACKSON, M. O. 2003. A survey of models of network formation: Stability and efficiency. In *Group Formation in Economics; Networks, Clubs and Coalitions*, G. Demange and M. Wooders, Eds. Cambridge University Press.
- KLEINBERG, J. M. 2000. Navigation in a small world. *Nature* 406, 845.
- KOUTSOPIAS, E. AND PAPADIMITRIOU, C. H. 1999. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*. Lecture Notes in Computer Science, vol. 1563. Trier, Germany, 404–413.
- LAOUTARIS, N., POPLAWSKI, L. J., RAJARAMAN, R., SUNDARAM, R., AND TENG, S.-H. 2008. Bounded budget connection (BBC) games or how to make friends and influence people, on a budget. In *Proceedings of the 27th ACM Symposium on Principles of Distributed Computing*. 165–174.
- LEIGHTON, T. AND RAO, S. 1999. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM* 46, 6, 787–832.
- LIN, H. 2003. On the price of anarchy of a network creation game. Class final project.

- NASH, J. 1950. Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences* 36, 1, 48–49.
- NASH, J. 1951. Non-cooperative games. *Annals of Mathematics (2)* 54, 286–295.
- PAPADIMITRIOU, C. 2001. Algorithms, games, and the internet. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*. Hersonissos, Greece, 749–753.
- ROUGHGARDEN, T. 2002a. The price of anarchy is independent of the network topology. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*. Montréal, Canada, 428–437.
- ROUGHGARDEN, T. 2002b. Selfish routing. Ph.D. thesis, Cornell University. Published as *Selfish Routing and the Price of Anarchy*, MIT Press, 2005.

RECEIVED SEPTEMBER 2009; REVIEWED DECEMBER 2009; ACCEPTED DECEMBER 2009