

A Bounded Budget Network Creation Game

SHAYAN EHSANI, SABER SHOKAT FADAEI, and MOHAMMADAMIN FAZLI,

Sharif University of Technology

ABBAS MEHRABIAN, University of Waterloo

SINA SADEGHIAN SADEGHABAD, MOHAMMADALI SAFARI,

and MORTEZA SAGHAFIAN, Sharif University of Technology

We introduce a network creation game in which each player (vertex) has a fixed budget to establish links to other players. In this model, each link has a unit price, and each agent tries to minimize its cost, which is either its eccentricity or its total distance to other players in the underlying (undirected) graph of the created network. Two versions of the game are studied: In the MAX version, the cost incurred to a vertex is the maximum distance between the vertex and other vertices, and, in the SUM version, the cost incurred to a vertex is the sum of distances between the vertex and other vertices. We prove that in both versions pure Nash equilibria exist, but the problem of finding the best response of a vertex is NP-hard. We take the social cost of the created network to be its diameter, and next we study the maximum possible diameter of an equilibrium graph with n vertices in various cases. When the sum of players' budgets is $n - 1$, the equilibrium graphs are always trees, and we prove that their maximum diameter is $\Theta(n)$ and $\Theta(\log n)$ in MAX and SUM versions, respectively. When each vertex has a unit budget (i.e., can establish a link to just one vertex), the diameter of any equilibrium graph in either version is $\Theta(1)$. We give examples of equilibrium graphs in the MAX version, such that all vertices have positive budgets and yet the diameter is $\Omega(\sqrt{\log n})$. This interesting (and perhaps counterintuitive) result shows that increasing the budgets may increase the diameter of equilibrium graphs and hence deteriorate the network structure. Then we prove that every equilibrium graph in the SUM version has diameter $2^{O(\sqrt{\log n})}$. Finally, we show that if the budget of each player is at least k , then every equilibrium graph in the SUM version is k -connected or has a diameter smaller than 4.

Categories and Subject Descriptors: G.2.2 [Discrete Mathematics]: Graph Theory—*Network problems*; F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems

General Terms: Algorithms, Design, Theory

Additional Key Words and Phrases: Braess's paradox, eccentricity, game theory, Nash equilibria, network creation games, network design, price of anarchy

A preliminary version of this paper appeared in *Proceedings of the 23rd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA'11)*, pp. 207–214.

Authors' addresses: S. Ehsani, Department of Management Science and Engineering, Huang Engineering Center, 475 Via Ortega, Stanford University, Stanford, CA 94305, USA; email: shayane@stanford.edu; S. S. Fadaei, College of Computer and Information Sciences, Northeastern University, 440 Huntington Ave, #266 WWH, Boston, MA 02115, USA; email: saber@ccs.neu.edu; M. Fazli and M. Safari, Department of Computer Engineering, Sharif University of Technology, Azadi Street, Tehran, Iran; emails: fazli@ce.sharif.edu, msafari@sharif.edu; A. Mehrabian and S. S. Sadeghabad, Department of Combinatorics and Optimization, University of Waterloo, 200 University Avenue West, Waterloo, ON, N2L3G1, Canada; emails: {amehrabi, s3sadegh}@uwaterloo.ca; M. Saghafeian, Department of Computer Science, Sharif University of Technology, Azadi Street, Tehran, Iran; email: saghafeian@ce.sharif.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.

© 2015 ACM 1549-6325/2015/04-ART34 \$15.00

DOI: <http://dx.doi.org/10.1145/2701615>

ACM Reference Format:

Shayan Ehsani, Saber Shokat Fadaee, Mohammadamin Fazli, Abbas Mehrabian, Sina Sadeghian Sadeghabad, Mohammadali Safari, and Morteza Saghaian. 2015. A bounded budget network creation game. *ACM Trans. Algor.* 11, 4, Article 34 (April 2015), 25 pages.
DOI: <http://dx.doi.org/10.1145/2701615>

1. INTRODUCTION

In recent years, a lot of research has been conducted on network design problems because of their importance in computer science and operations research [Christofides 1976; Johnson et al. 2000; Salman et al. 2000]. The aim in this line of research is usually to build a minimum-cost network that satisfies certain properties, and the network structure is usually determined by a central authority. However, this is in contrast to many real-world situations, such as social networks, client-server systems, and peer-to-peer networks, where network structures are determined in a distributed manner by selfish agents [Feldman et al. 2004; Golle et al. 2001; Skyrms and Pemantle 2000]. The formation of these networks can be formulated as a game, which is usually called a *network creation game*. In network creation games, as in any other game, there are selfish players that interact with each other. Each player has its own objective and attempts to minimize the cost it incurs in the network, regardless of how its actions affect other agents. The players are placed at the nodes of the network graph and can create links to other nodes with certain restrictions; for example, there could be an upper bound for the number of links a player constructs. The utility functions of the players should be defined properly to be consistent with their natural interests (e.g., minimizing the cost to communicate with other players).

In network creation games, players interact with each other by adding and removing links between themselves. Many variants of these games arise by defining different utility functions and possible transitions between the strategies of each player. Some authors have studied undirected graphs, whereas others have considered directed graphs. For undirected graphs, there is an issue of “ownership”: When there exists a link between two nodes, but just one of the nodes wants to keep it, is it removed from the network? For directed graphs, the question is whether both endpoints of a link can use it to communicate. In some models, creating a link incurs a cost to the player (i.e., the number of created links appears in utility functions), whereas in other models restrictions for creating links appear in the set of available strategies for players.

Although the creation of a network in such a game is a dynamic process, the structure of the resulting network (if it does converge to a stable structure) provides valuable information about the effectiveness of the game rules. Thus, the existence and structure of *stable networks* has been widely studied. Jackson and Wolinsky [1996] defined a network to be *pairwise stable* if, roughly speaking, none of the nodes is willing to delete an incident link, and no pair of nonadjacent nodes is willing to build a link between themselves. Fabrikant et al. [2003] considered *Nash equilibria* of the game as its stable states. A Nash equilibrium, which is a well-known concept in game theory, is a state of a game in which no player can increase its utility by changing its strategy, assuming the strategies of other players are kept unchanged. The main difference between these two concepts is that, when considering pairwise stability, we are thinking of the players cooperating with each other, whereas we think of noncooperating players when we consider Nash equilibria.

The efficiency of the network formed by a game is measured by different factors rather than the player utilities. We are mostly interested in measuring a global parameter, for instance, the diameter (the largest distance between any pair of nodes) and the vertex connectivity (the minimum number of nodes whose removal disconnects the network) of the network are two of the possible candidates; these are important parameters in

every network. The global parameter, called the *social cost* of the network, quantizes how effective the network is. To measure how the efficiency of a system degrades due to selfish behavior of its agents, we find the social cost achieved by selfish players in a stable state of the network creation game and calculate its ratio to the minimum social cost among all networks. This parameter, called the *price of anarchy* of the game, was introduced by Koutsoupias and Papadimitriou [1999] and has been the focus of study of many works on network creation (and many other) games.

In this article, we introduce and study a new class of network creation games, which is motivated by the work of Laoutaris et al. [2008]. In our model, there is an upper bound on the number of links each player can create, hence the name *bounded budget network creation game*.

1.1. Previous Work

Jackson and Wolinsky [1996] were one of the first to study the stable states of networks created by selfish players. They introduced the notion of pairwise stability and studied the efficiency of pairwise stable networks. Fabrikant et al. [2003] suggested studying Nash equilibria instead of pairwise stable networks. In their model, the network graph is undirected, and there is a link between two nodes if at least one of the nodes wants to create it. There is a cost α for creating a link, and the goal of each node is to minimize the sum of its distances to other nodes plus the amount it pays for creating links. Note that every node wants to build more links to get closer to other vertices, but, on the other hand, the more links it creates, the more money it has to pay. They showed that for large α , the equilibrium graphs are sparse, whereas for small values of α , the equilibrium graphs are dense. This implies that the structure of the equilibria is highly affected by the parameter α . They called a Nash equilibrium *transient* if there exists a sequence of single-player strategy changes that do not alter the changer's cost and leads to a nonequilibrium. They conjectured that there is a universal constant A such that for $\alpha > A$, any nontransient equilibrium graph is a tree. Albers et al. [2006] disproved this conjecture using geometric constructions. The results of Fabrikant et al. were improved in Corbo and Parkes [2005] and Demaine et al. [2007]. In particular, Demaine et al. [2007] found a $2^{O(\sqrt{\log n})}$ upper bound on the diameter of equilibrium graphs, where n is the number of players (nodes), and conjectured that, indeed, the diameter of equilibrium graphs is polylogarithmic. They also considered another version of the game in which the cost of each player is its maximum distance to other nodes plus the amount it pays for creating links. Mihalák and Schlegel [2010] further studied the game and proved that, for the original (sum of distances) version, the price of anarchy is $O(1)$ for $\alpha > 273n$, and, for the latter (maximum distance) version, the price of anarchy is $2^{O(\sqrt{\log n})}$ for all α and is $O(1)$ for $\alpha > 129$. Brandes et al. [2008] studied a variant in which the cost for a pair of disconnected vertices is a finite value, as opposed to previous variants, where this value was infinity.

Laoutaris et al. [2008] introduced another variant of network creation games, in which a budget is dedicated to each player for buying links. That is, each player can build a certain number of links, and there is no cost term in the utility function and thus no parameter α . This is a natural and interesting perspective for formulating peer-to-peer and overlay networks. To eliminate the intricacies with ownership of the links, they assumed that the links are directed, and each link will be used by one of its endpoints (i.e., its tail). They defined the utility function of a player as its average distance to other nodes. If all players have the same budget k , they proved that Nash equilibria always exist and that the price of anarchy is between $c_1 \sqrt{\frac{n \log k}{k \log n}}$ and $c_2 \sqrt{\frac{n \log k}{\log n}}$ for suitable positive constants c_1, c_2 . Our model is mainly motivated by their work,

but we assume that links are bidirectional and can be used by both of their endpoints. However, one of the endpoints of a link is its “owner” and is responsible for creating it.

Alon et al. [2013] simplified Fabrikant et al.’s model by eliminating the parameter α and introduced basic network creation games. In their model, each node locally tries to minimize its maximum distance or average distance to other nodes by swapping one incident edge at a time. Considering swap equilibria seems to be more realistic, since each player can determine its best response in polynomial time and thus the computational needs of each player is decreased. Also, the removal of the parameter α has made the proofs cleaner and more general, and thus we have also used this idea in our model. The main difference between our games and basic network creation games is the ownership of the links. In their model, any of the two endpoints of a link can remove it, whereas in our model, any link is owned by one of its endpoints, and only that node is able to remove it. Two versions of basic creation games are considered, the SUM version and the MAX version, depending on whether the goal of each player is to minimize its average distance or maximum distance to other nodes. Alon et al. found an upper bound of $2^{O(\sqrt{\log n})}$ on the diameter of SUM equilibria, which is stronger than previous bounds on models with a parameter α . We prove that the same upper bound holds for our model. However, in the MAX version, we observe essential differences even in tree equilibria: In basic network creation games, the diameter of equilibria is at most 3, whereas in our model, we have tree equilibria with diameter $\Theta(n)$.

There are utility functions studied in the literature other than the average or maximum distance of a node to other nodes. We will not pursue them here but mention a few of them for completeness. Bei et al. [2011] investigated a model in which each node aims to maximize its *betweenness*. This notion is introduced originally in social network analysis and, roughly speaking, measures the amount of information passing through a node among all pairwise information exchanges. The *clustering coefficient* of a node is defined as the probability that two of its randomly selected neighbors are directly connected to each other. Brautbar and Kearns [2011] considered the clustering coefficient as an incentive in creation of networks such as social networks, and considered network creation games in which the utility of each node equals its clustering coefficient.

In the model we introduce here, every player has a budget that determines the number of links it is able to build. This simplifies the utility functions, but complicates the strategy set of the players. Nevertheless, we are able to prove many results in both SUM and MAX versions of the game, and we observe that, just like the model of Laoutaris et al. [2008], changing the node budgets significantly changes the structure of equilibrium networks. In contrast to Laoutaris et al. [2008], in our model, once a link is established, both its endpoints can use it equally. This is a natural model in applications where the direction of links does not matter (e.g., computer networks).

Bilò et al. [2012] have proposed a network creation game that is the dual of our model, in the sense that the constraints and the cost functions are swapped. More precisely, they assume that the cost function of each player only consists of the number of bought edges (without any budget on them), but with the constraint that a player needs to connect to the network by staying within a given (either maximum or average) distance to the set of players. Their motivation is that, in some real-world scenarios, a player might have a strong objective about its centrality in the created network. They take the total cost of the players as the social cost and prove several bounds for the price of anarchy for various ranges of parameters.

1.2. Our Model and Notation

Let G be a directed graph on n vertices, and let $V(G)$ denote the vertex set of G . The *underlying graph* of G , which is an undirected graph obtained by ignoring the arc

directions in G , is denoted by $U(G)$. If both arcs \vec{uv} and \vec{vu} are in G , then uv is a multiple edge with multiplicity 2 in $U(G)$, which is viewed as a cycle with two vertices. In the following, whenever we refer to the distance between two vertices of G , we mean their distance in $U(G)$. The distance between two vertices u and v is denoted by $\text{dist}(u, v)$. If u and v are in different connected components of $U(G)$, then it is natural to define their distance as infinity. However, we define their distance to be a large constant C_{inf} , so that the vertices have the incentive to decrease the number of connected components. We choose $C_{\text{inf}} = n^2$ for a reason that will be discussed later. The *diameter* of G , written $\text{diam}(G)$, is the maximum distance between any two vertices of G . For a vertex u and subset $A \subseteq V(G)$, the distance between u and A , written $\text{dist}(u, A)$, is defined as

$$\text{dist}(u, A) = \min\{\text{dist}(u, a) : a \in A\}.$$

The *eccentricity* of a vertex u is the maximum of its distances to other vertices. Note that if the graph is disconnected, then the eccentricity of all vertices is n^2 .

Let n be a positive integer and b_1, b_2, \dots, b_n be nonnegative integers less than n . A *bounded budget network creation game* with parameters b_1, b_2, \dots, b_n , denoted by $(b_1, b_2, \dots, b_n)\text{-BG}$, is the following game. There are n players numbered from 1 to n , and the strategy of player i is a subset $S_i \subseteq \{1, 2, \dots, n\} \setminus \{i\}$ with $|S_i| = b_i$. We may build a directed graph G for every strategy profile (S_1, \dots, S_n) of the game, with vertex set $\{u_1, \dots, u_n\}$ and such that $\vec{u_i u_j}$ is an arc in G if and only if $j \in S_i$. Any such graph G is called a *realization* of $(b_1, b_2, \dots, b_n)\text{-BG}$. We will identify each vertex with its corresponding player. If $\vec{u_i u_j}$ is in G , then we say $\vec{u_i u_j}$ is *owned* by vertex u_i . Note that u_i owns exactly b_i arcs. We think of b_i as the *budget* available to vertex u_i , which it can use to build links to other vertices. If both $\vec{u_i u_j}$ and $\vec{u_j u_i}$ are in G , then the pair $\{u_i, u_j\}$ is called a *brace*.

We consider two versions of bounded budget network creation games, which differ in the definition of the cost function. In the *SUM* version, the cost incurred to each vertex is the sum of its distances to other vertices; that is, for each $u \in V(G)$,

$$c_{\text{SUM}}(u) = \sum_{v \in V(G)} \text{dist}(u, v).$$

By choosing C_{inf} to be at least n^2 , we ensure that for every vertex u , $c_{\text{SUM}}(u)$ decreases whenever u changes its strategy so that the number of vertices in its connected component is increased. In the *MAX* version, if $U(G)$ has κ connected components, then the cost incurred to each vertex is defined as

$$c_{\text{MAX}}(u) = \max\{\text{dist}(u, v) : v \in V(G)\} + (\kappa - 1)n^2.$$

The term $\max\{\text{dist}(u, v) : v \in V(G)\}$ is simply the *eccentricity* of u , and the (artificial) term $(\kappa - 1)n^2$ has been added to the cost function so that when the network is disconnected, the vertices would have the incentive to decrease the number of connected components.

We say a vertex is playing its *best response* if it cannot decrease its cost by changing its strategy while the other vertices' strategies are fixed. Notice that a vertex does not need to have a unique best response. A strategy profile is called a (*pure*) *Nash equilibrium* if, in that profile, all players are playing their best responses. In this case, the graph G is said to be a *Nash equilibrium graph* or simply an *equilibrium graph* for $(b_1, b_2, \dots, b_n)\text{-BG}$. The *price of stability* of $(b_1, \dots, b_n)\text{-BG}$ is defined as

$$\frac{\min\{\text{diam}(G^{NE}) : G^{NE} \text{ is an equilibrium for } (b_1, \dots, b_n)\text{-BG}\}}{\min\{\text{diam}(G) : G \text{ is a realization of } (b_1, \dots, b_n)\text{-BG}\}}.$$

And the *price of anarchy* of (b_1, \dots, b_n) -BG is defined as

$$\frac{\max\{\text{diam}(G^{NE}) : G^{NE} \text{ is an equilibrium for } (b_1, \dots, b_n)\text{-BG}\}}{\min\{\text{diam}(G) : G \text{ is a realization of } (b_1, \dots, b_n)\text{-BG}\}}.$$

Note that both these fractions are at least 1. The price of anarchy measures how the efficiency of the network degrades due to selfish behavior of its agents. In this article, networks with smaller diameter are considered more efficient, and the social cost of a strategy profile is the diameter of the constructed graph. It is worth noting that, if

$$b_1 + b_2 + \dots + b_n \geq n - 1,$$

then the denominator of both of these fractions is $O(1)$ (see Theorem 2.3), and the main challenge is to evaluate the nominators (i.e., the diameters of equilibrium graphs). In this case, the underlying undirected graphs of equilibria are connected (see Lemma 3.1). Instances with

$$b_1 + b_2 + \dots + b_n < n - 1$$

are not very interesting since the constructed networks are always disconnected, and both of the fractions are equal to 1. In this article, all logarithms are in base 2, and generally we do not try to optimize the constant factors.

1.3. Our Results and Organization of the Article

We study various properties of equilibrium graphs for bounded budget network creation games. In particular, we analyze the diameter of equilibrium graphs in various special cases, which results in bounds for the price of anarchy in these cases. First, in Section 2, we prove that for every nonnegative sequence b_1, \dots, b_n , the game (b_1, b_2, \dots, b_n) -BG has a Nash equilibrium in both versions and that the price of stability of this game is $\Theta(1)$. In Section 3, we study the price of anarchy in extreme instances in which the sum of budgets is $n - 1$. Note that this is the smallest sum needed to have a connected network. For these instances, we prove that the price of anarchy is $\Theta(n)$ and $\Theta(\log n)$ in MAX and SUM versions, respectively.

In Section 4, we prove that the price of anarchy in instances in which the budget of each players is equal to 1 is $\Theta(1)$ in either version. One may expect that further increasing the players' budgets will result in equilibrium graphs with even smaller diameters. In Section 5, we show that, interestingly, this is not true and there exist instances in which all players have positive budgets, and the price of anarchy is $\Omega(\sqrt{\log n})$ in the MAX version. Such a counterintuitive behavior had been observed previously in algorithmic game theory, and perhaps the most famous example, known as the *Braess's paradox*, is given by Braess et al. [2005] in network routing games. They observed that adding extra capacity to the links of a road network, which is used by several selfish commuters, in some cases might reduce the overall performance.

In Section 6, we give a general upper bound of $2^{O(\sqrt{\log n})}$ for the price of anarchy in the SUM version. Our bounds on the price of anarchy in various classes of instances in both versions are summarized in Table I. In Section 7, we consider the connectivity of equilibrium graphs and prove that if the budget of each player is at least k , then every equilibrium graph in the SUM version with diameter larger than 3 is k -connected. We conclude with a discussion of our results, and we propose some open problems in Section 8.

2. EXISTENCE OF NASH EQUILIBRIA

In this section, we prove that for every nonnegative b_1, b_2, \dots, b_n , Nash equilibria exist for both MAX and SUM versions of (b_1, b_2, \dots, b_n) -BG. Moreover, we prove that the

Table I. Our Bounds on the Price of Anarchy in Various Classes of Instances (n is the Number of Players)

Budget Constraints	MAX	SUM
$\sum_{i=1}^n b_i = n - 1$	$\Theta(n)$	$\Theta(\log n)$
$b_i = 1, \quad i = 1, 2, \dots, n$	$\Theta(1)$	$\Theta(1)$
$b_i \geq 1, \quad i = 1, 2, \dots, n$	$\Omega(\sqrt{\log n})$	$2^{O(\sqrt{\log n})}$
no constraints	$\Theta(n)$	$2^{O(\sqrt{\log n})}$

price of stability of this game is $\Theta(1)$. Before proving the main result of this section, we show that computing a player's best response in bounded budget network creation games is an intractable problem.

THEOREM 2.1. *The problem of finding a player's best response in both MAX and SUM versions of bounded budget network creation games is NP-Hard.*

PROOF. We reduce the k -CENTER problem to the problem of finding a player's best response in the MAX version of the game. In the k -CENTER problem, an undirected graph H and a positive integer k are given, and the aim is to find a subset S of k vertices so as to minimize the maximum distance from a vertex to S ; that is, we want to find

$$\min_{S \subseteq V(H): |S|=k} \max_{v \in V(H)} \text{dist}(v, S).$$

Assume that we are given an n -vertex graph H and a positive integer k , and we are supposed to find an optimal solution to the k -CENTER problem. Consider a directed graph G such that $U(G) = H$, and a game $(b_1, b_2, \dots, b_n, b_{n+1})$ -BG, where b_i is the outdegree of the i -th vertex in G , and define $b_{n+1} = k$. Now compute a best response of the $(n+1)$ -th player in the MAX version of this game, where the strategies of other players are realized by G . A best response is clearly an optimal solution for the k -CENTER problem in H . The proof is complete by noting that k -CENTER is NP-hard (see Hsu and Nemhauser [1979] for instance).

By using exactly the same idea, one can reduce the k -MEDIAN problem (see Lin and Vitter [1992] for the definition) to the problem of finding a best response in the SUM version of the game. Since the former problem is NP-hard, the latter one is NP-hard, too. \square

For proving the main theorem of this section, we need a lemma that gives a sufficient condition for guaranteeing that a vertex is playing its best response.

LEMMA 2.2. *Let u be a vertex in a realization of a bounded budget network creation game. If $c_{\text{MAX}}(u) \leq 2$ and u is not contained in any brace, or $c_{\text{MAX}}(u) = 1$, then u is playing its best response in both MAX and SUM versions of the game.*

PROOF. If $c_{\text{MAX}}(u) = 1$, then it is clear that u cannot decrease its cost. Otherwise, let V^- be the set of vertices that have an arc to u and V^+ be the set of vertices that have an arc from u . Since u is not an endpoint of any brace, $V^+ \cap V^- = \emptyset$. Thus, no matter how u plays, it always has distance one to at most $|V^+| + |V^-|$ vertices and distance at least two to the rest of the vertices. Therefore, regardless of how u plays, its cost in the MAX version will be at least 2, and its cost in the SUM version will be at least $2(n - 1 - |V^-| - |V^+|) + |V^+| + |V^-|$. Hence, u is already playing its best response. \square

We are now ready to prove the main theorem of this section.

THEOREM 2.3. *For every nonnegative b_1, b_2, \dots, b_n , Nash equilibria exist for both MAX and SUM versions of (b_1, \dots, b_n) -BG. Moreover, the price of stability of this game is $\Theta(1)$.*

PROOF. Let z be the number of players with zero budget, and let $\sigma = b_1 + b_2 + \dots + b_n$. Without loss of generality, assume that the b_1, \dots, b_n are in nondecreasing order, that is,

$$0 = b_1 = b_2 = \dots = b_z < b_{z+1} \leq b_{z+2} \leq \dots \leq b_{n-1} \leq b_n. \quad (1)$$

We consider three cases.

Case 1. $\sigma \geq n - 1$ and $b_n \geq z$

We provide an algorithm to build a graph G that realizes (b_1, \dots, b_n) -BG and all of whose vertices satisfy the conditions of Lemma 2.2. Thus, G is an equilibrium graph for both versions. Moreover, G has diameter $O(1)$, which shows that the price of stability is $O(1)$ for instances satisfying $\sigma \geq n - 1$ and $b_n \geq z$. In this case, we can use a single vertex to link to zero-budget vertices and keep G connected.

The graph G has vertex set $\{v_1, \dots, v_n\}$ and is initially empty. Add the arcs $\overrightarrow{v_n v_1}, \overrightarrow{v_n v_2}, \dots, \overrightarrow{v_n v_b_n}$ and then the arcs $\overrightarrow{v_{b_n+1} v_n}, \overrightarrow{v_{b_n+2} v_n}, \dots, \overrightarrow{v_{n-1} v_n}$ to G . Note that G has diameter 2 at this point, but there might be vertices whose outdegrees are less than their budgets. If u is such a vertex, then add arcs from u to arbitrary vertices until its outdegree equals its budget. This operation clearly does not increase the diameter, but may create braces. For every brace $\{u, v\}$ such that u has eccentricity 2 and 3 exists a vertex w not adjacent to u , replace the arc \overrightarrow{uw} with \overrightarrow{uw} . This can be done only a finite number of times since after every replacement the number of braces decreases. The vertices of the obtained graph have the properties of Lemma 2.2, and thus this graph is an equilibrium graph.

Case 2. $\sigma \geq n - 1$ and $b_n < z$

As in Case 1 (but using a more complicated construction), we build a graph that is an equilibrium graph in both versions and has diameter $O(1)$. In this case, we cannot use a single vertex to link to zero-budget vertices and keep the graph connected; hence, we should use several vertices to do this. We would like to use as few vertices as possible, so we will focus on vertices with large degrees. Let t be the largest index with

$$b_n + b_{n-1} + \dots + b_t \geq z + n - t. \quad (2)$$

First, note that such a t exists and is larger than z , since

$$b_n + b_{n-1} + \dots + b_{z+1} = \sigma \geq n - 1 = z + n - (z + 1).$$

Second, note that $t < n$ since $b_n < z = z + n - n$. Define

$$A = \{v_1, v_2, \dots, v_z\}, \quad B = \{v_{z+1}, v_{z+2}, \dots, v_t\}, \quad \text{and} \quad C = \{v_{t+1}, v_{t+2}, \dots, v_{n-1}\}.$$

Note that A is the set of zero-budget vertices, and $\{v_t\} \cup C \cup \{v_n\}$ is the set of vertices that will connect the set A to the rest of the graph. Notice that every vertex in this connecting set has budget greater than 1, since by the definition of t , we have $b_n + b_{n-1} + \dots + b_{t+1} < z + n - t - 1$, which together with Equations (2) and (1) implies $1 < b_t \leq b_{t+1} \leq \dots \leq b_n$.

We start with an empty graph with vertex set $A \cup B \cup C \cup \{v_n\}$ and add arcs to it as described in the four following phases, until the outdegree of each vertex becomes equal to its budget. A concrete example is illustrated in Figure 1, in which $n = 22$, $z = 16$, and $t = 19$.

- (1) Add an arc from every vertex in $B \cup C$ to v_n (the arcs $\overrightarrow{v_{17} v_{22}}, \overrightarrow{v_{18} v_{22}}, \dots, \overrightarrow{v_{21} v_{22}}$ in Figure 1).
- (2) Add arcs from $\{v_n\} \cup C \cup \{v_t\}$ to A :
 - First, add b_n arcs from v_n to the first b_n vertices of A (the arcs $\overrightarrow{v_{22} v_1}, \overrightarrow{v_{22} v_2}, \dots, \overrightarrow{v_{22} v_5}$ in Figure 1);
 - Second, add $b_{n-1} - 1$ arcs from v_{n-1} to the next $b_{n-1} - 1$ vertices of A (the arcs $\overrightarrow{v_{21} v_6}, \dots, \overrightarrow{v_{21} v_9}$ in Figure 1);




Fig. 1. Case 2 in the proof of Theorem 2.3.

- Third, add $b_{n-2} - 1$ arcs from v_{n-2} to the next $b_{n-2} - 1$ vertices of A ; (the arcs $\overrightarrow{v_{20}v_{10}}, \dots, \overrightarrow{v_{20}v_{13}}$ in Figure 1);
- Continue similarly until you add $b_{t+1} - 1$ arcs from v_{t+1} to the next $b_{t+1} - 1$ vertices of A ;
- At last, add s arcs from v_t to the last s vertices of A , where

$$s = z + n - (t + 1) - (b_n + \dots + b_{t+1})$$

is positive by the definition of t (the arcs $\overrightarrow{v_{19}v_{14}}, \overrightarrow{v_{19}v_{15}}, \overrightarrow{v_{19}v_{16}}$ in Figure 1).

When this phase is completed, every vertex in A has exactly one incoming arc. Moreover, the graph is a tree at this stage, and, since v_n has eccentricity 2, its diameter is at most 4.

- (3) Add arcs from B to $C \cup \{v_t\}$: For every vertex u in B whose outdegree is less than its budget, add arcs from u to vertices in $C \cup \{v_t\}$ in reverse order; that is, add the arcs $\overrightarrow{uv_{n-1}}, \overrightarrow{uv_{n-2}}$, and so on, until either arcs to all vertices of $C \cup \{v_t\}$ have been added, or the outdegree of u equals its budget (the arcs $\overrightarrow{v_{17}v_{21}}, \overrightarrow{v_{18}v_{21}}, \overrightarrow{v_{18}v_{20}}, \overrightarrow{v_{18}v_{19}}, \overrightarrow{v_{19}v_{21}}$ in Figure 1).
- (4) Add arcs from B to A : For every vertex u in B whose outdegree is still less than its budget, add arcs from u to vertices in A in order; that is, add the arcs $\overrightarrow{uv_1}, \overrightarrow{uv_2}$, and so on, until the outdegree of u equals its budget (the arc $\overrightarrow{v_{18}v_1}$ in Figure 1).

When phase 4 is completed, for every $u \in B$, since the budget of u is not more than the budget of v_n , the set of neighbors of u in A is a subset of the set of neighbors of v_n in A . Therefore, every vertex in A that is not adjacent to v_n has only one neighbor, which is in $C \cup \{v_t\}$. Let \overrightarrow{wx} be an arc from C to A . This arc could have been added in phase 2 only, so x is not a neighbor of v_n . Thus, we have the following property:

PROPERTY 2.4. *For every arc \overrightarrow{wx} from C to A , w is the only neighbor of x .*

Now, we prove that every vertex is playing its best response, concluding that the obtained graph G is an equilibrium graph. This also implies that the price of stability of this game is $O(1)$ in this case since G has diameter at most 4 when phase 2 finishes. Observe that we create no brace in our construction. Vertices in A are obviously playing their best strategies because their budgets are zero. Since v_n has eccentricity 2, it is playing its best response by Lemma 2.2.

Next, we show that every vertex in C is playing its best response. Let $u \in C$. Every outgoing arc from u is either going to v_n or to some vertex in A . The latter cannot be removed by Property 2.4. Since v_n has the maximum degree and eccentricity 2, u cannot decrease its cost by removing the arc $\overrightarrow{uv_n}$ and adding an arc to another vertex.

Finally, we show that every vertex in B is playing its best response. Let $u \in B$. If in phase 4 some outgoing arcs from u have been added, then, in phase 3, u has already been joined to all vertices in $C \cup \{v_t\}$ and so has eccentricity 2. Thus, in this case, vertex u satisfies the conditions of Lemma 2.2 and is playing its best response. Otherwise, since u is adjacent to v_n , it has eccentricity 3. Assume that in the beginning of phase 3, the budget of u minus its outdegree was p . Note that $p < |C| + 1$. First, since v_n has the maximum degree and eccentricity 2, u has no incentive to replace its arc $\overrightarrow{uv_n}$ with any other arc. For any $w \in C$, at least one arc was added from w to A in phase 2; so, by Property 2.4, there is a vertex $x \in A$ such that w is its only neighbor. Therefore, u cannot make its eccentricity less than 3, and so it is playing its best response in the MAX version. Also, in the SUM version, it is easy to verify that the best strategy for u is to be adjacent to the vertices with largest degrees (i.e., v_{n-1}, \dots, v_{n-p}).

Case 3. $\sigma < n - 1$

Let m be the smallest positive integer that satisfies

$$b_m + b_{m+1} + \dots + b_n \geq n - m.$$

Clearly, $1 < m \leq n$ and $b_1 = b_2 = \dots = b_{m-1} = 0$. Let G be a graph with vertex set $\{v_1, v_2, \dots, v_n\}$ such that the subgraph induced by $\{v_m, v_{m+1}, \dots, v_n\}$ is an equilibrium graph for $(b_m, b_{m+1}, \dots, b_n)$ -BG in the SUM version, and there is no other edge in G . Then, it is easy to verify that G is an equilibrium graph for (b_1, b_2, \dots, b_n) -BG in both versions. Moreover, in this case, any realization of (b_1, \dots, b_n) -BG is disconnected and has diameter n^2 , which shows that the price of stability is 1. \square

3. THE DIAMETER OF EQUILIBRIUM TREES

If the sum of players' budgets is less than $n - 1$, then any realization of the bounded budget network creation game is disconnected and has diameter n^2 . So, the smallest interesting instances of the game are those in which the players' budgets add up to $n - 1$.

LEMMA 3.1. *For any nonnegative b_1, \dots, b_n for which $\sum_{i=1}^n b_i \geq n - 1$, the underlying graphs of Nash equilibria of (b_1, \dots, b_n) -BG are connected.*

PROOF. Let G be an equilibrium graph for (b_1, \dots, b_n) -BG, where $\sum_{i=1}^n b_i \geq n - 1$. If G is not connected, then it has a cycle C , where a brace is also considered a cycle. Pick a vertex v from C that owns at least one arc \overrightarrow{vw} of C , and let u be a vertex that is in a different component. If v replaces \overrightarrow{vw} with \overrightarrow{vu} , then the number of vertices in its connected component increases, and the total number of connected components decreases. Thus, by definition of cost functions (and since $C_{\inf} = n^2$), the cost of v decreases in either version. So v is not playing its best response in G ; that is, G is not an equilibrium graph, which is a contradiction. \square

When $\sum_{i=1}^n b_i = n - 1$, it can be easily seen that every equilibrium graph is a tree. We write *Tree-BG* for the set of instances of bounded budget network creation games in which the sum of budgets equals $n - 1$.

In this section, we study the price of anarchy of games in *Tree-BG*. We prove that, in the MAX version, there exist equilibrium graphs with diameter $\Theta(n)$, so the price of anarchy is $\Theta(n)$. In the SUM version, we prove that equilibrium graphs have diameter $O(\log n)$, and this bound is asymptotically tight, so the price of anarchy is $\Theta(\log n)$.




Fig. 2. Illustration of the proof of Theorem 3.2.

THEOREM 3.2. *In the MAX version, there are Tree-BG instances that have equilibrium graphs with diameter $\Omega(n)$.*

PROOF. Let k be a positive integer, and let $n = 3k + 1$. Define

$$X = \{x_1, x_2, \dots, x_k\}, Y = \{y_1, y_2, \dots, y_k\}, \text{ and } Z = \{z_1, z_2, \dots, z_k\}.$$

Let G be the tree with vertex set $X \cup Y \cup Z \cup \{w\}$ and arc set

$$\{\overrightarrow{x_1 x_2}, \dots, \overrightarrow{x_{k-1} x_k}, \overrightarrow{y_1 y_2}, \dots, \overrightarrow{y_{k-1} y_k}, \overrightarrow{z_1 z_2}, \dots, \overrightarrow{z_{k-1} z_k}, \overrightarrow{x_1 w}, \overrightarrow{y_1 w}, \overrightarrow{z_1 w}\}.$$

See Figure 2. Then, G is a realization of a Tree-BG instance (in which the budget of each vertex equals its outdegree in G) and has diameter $2k = \Theta(n)$. To complete the proof, we need to show that G is an equilibrium graph.

Since the vertices w, x_k, y_k, z_k have zero budgets, they are playing their best responses. By symmetry, it is sufficient to prove that for all $1 \leq i < k$, x_i is playing its best response. For $2 \leq i < k$, note that x_i has a unit budget and currently has an arc to x_{i+1} . If it replaces its outgoing arc with $\overrightarrow{x_i x_j}$ for some $j > i + 1$, then its cost does not change at all. If it replaces its outgoing arc with any other arc, then the graph gets disconnected, and the cost of x_i becomes $2n^2$. Thus, x_i is playing its best response.

Now, let $i = 1$. Note that x_1 has budget 2. Clearly, in order to keep the graph connected, x_1 should have arcs to one vertex from each of the two disjoint paths $x_2 x_3 \dots x_k$ and $z_k z_{k-1} \dots z_1 w y_1 y_2 \dots y_k$. Therefore, to minimize its eccentricity, its best response is to choose the middle of the second path (which is w) and an arbitrary vertex in the first path. Thus, x_1 is playing its best response, and the proof is complete. \square

Next, we show that the diameters of equilibrium graphs in the SUM version are much smaller.

THEOREM 3.3. *Any equilibrium graph for a Tree-BG instance in the SUM version has diameter $O(\log n)$.*




Fig. 3. Illustration of the proof of Theorem 3.3.

PROOF. Let G be an equilibrium graph for a Tree-BG instance in the SUM version. Let d be the diameter of G and $P = v_0 v_1 \dots v_d$ be a longest path in G . At least half of the arcs in P are in the same direction along P . By symmetry, we may assume that these are the arcs $\overrightarrow{v_{i_1} v_{i_1+1}}, \overrightarrow{v_{i_2} v_{i_2+1}}, \dots, \overrightarrow{v_{i_t} v_{i_t+1}}$, where $t \geq d/2$. Every vertex not in P is connected to P via a unique path. Let A_i be the set of vertices that are connected to P through v_i (including v_i itself), and let $a(i) = |A_i|$. See Figure 3 for an example. Notice that all $a(i)$'s are positive since $v_i \in A_i$, and that each vertex appears in exactly one of the sets A_i .

For $1 \leq j < t$, if v_{i_j} replaces its arc $\overrightarrow{v_{i_j} v_{i_j+1}}$ with the arc $\overrightarrow{v_{i_j} v_{i_j+2}}$, then its distances to vertices in $A_{i_{j+1}}$ increase by one, and its distances to vertices in A_k , $k > i_j + 1$ decrease by one, and its distances to other vertices do not change. Since v_{i_j} is playing its best response, we have

$$a(i_j + 1) \geq \sum_{k=i_j+2}^d a(k) \geq \sum_{l=j+1}^t a(i_l + 1) \quad \forall 1 \leq j < t. \quad (3)$$

Setting $j = t - 1$ in Equation (3) gives

$$a(i_{t-1} + 1) \geq a(i_t + 1) \geq 1.$$

Setting $j = t - 2$ in Equation (3) gives

$$a(i_{t-2} + 1) \geq a(i_{t-1} + 1) + a(i_t + 1) \geq 1 + 1 \geq 2.$$

Setting $j = t - 3$ in Equation (3) gives

$$a(i_{t-3} + 1) \geq a(i_{t-2} + 1) + a(i_{t-1} + 1) + a(i_t + 1) \geq 2 + 1 + 1 \geq 4.$$

Continuing similarly, we find that $a(i_j + 1) \geq 2^{t-j-1}$ for $1 \leq j < t$. Therefore,

$$\sum_{j=1}^{t-1} a(i_j + 1) \geq \sum_{j=1}^{t-1} 2^{t-j-1} = 2^{t-1} - 1.$$

On the other hand, since all vertices appear in exactly one of the sets A_i , we have

$$n \geq \sum_{i=1}^d a(i) \geq \sum_{j=1}^{t-1} a(i_j + 1) \geq 2^{t-1} - 1.$$

Therefore, $d \leq 2t = O(\log n)$. \square

The bound $O(\log n)$ proved in the above theorem is tight up to constant factors, as we next prove that there exist Tree-BG instances having equilibrium graphs with diameter $\Theta(\log n)$.

THEOREM 3.4. *In the SUM version, there exist instances of Tree-BG that have an equilibrium graph with diameter $\Theta(\log n)$.*

PROOF. Let k be a positive integer, and let $n = 2^{k+1} - 1$. Let G be a perfect binary tree on n vertices; that is, G has vertex set $\{u_1, \dots, u_n\}$, and for all $1 \leq i < n/2$, vertex u_i has arcs to vertices u_{2i} and u_{2i+1} . Then, G is a realization of a Tree-BG instance and has diameter $\Theta(\log n)$. To complete the proof, we just need to show that G is an equilibrium graph in the SUM version.

For each i , let T_i be the tree rooted at vertex i . For each $1 \leq i < n/2$, vertex u_i has budget 2. In order for the graph to be connected, u_i must have an arc to a vertex in T_{2i} and an arc to a vertex in T_{2i+1} . Observe that for every $1 \leq j \leq n$, vertex u_j has less total distance to vertices in T_j than any other vertex in T_j . So, the best response for u_i is to have arcs to vertices u_{2i} and u_{2i+1} , and thus it is already playing its best response. For $i > n/2$, vertex u_i has zero budget, so it is obviously playing its best response. Therefore, all vertices are playing their best responses, and G is an equilibrium graph. \square

4. THE STRUCTURE OF EQUILIBRIUM GRAPHS FOR $(1, 1, \dots, 1)$ -BG

In the previous section, we considered bounded budget network creation games in which the sum of players' budgets is $n - 1$. In these instances, there is at least one player with zero budget. One might expect that if all players have positive budgets, then the diameter of equilibrium graphs drop significantly. In this section, we consider an extreme case for which this expectation is realistic, and, in fact, the diameter is $O(1)$. More precisely, we study games in which all vertices have unit budgets (i.e., $b_i = 1$ for all i), and prove that the equilibrium graphs for these instances have a special structure. In particular, we prove that all equilibrium graphs for $(1, \dots, 1)$ -BG in SUM and MAX versions have diameters of less than 5 and 8, respectively, and therefore these instances have bounded price of anarchy in both versions.

THEOREM 4.1. *Any equilibrium graph for $(1, \dots, 1)$ -BG in the SUM version is connected and has a unique cycle with at most 5 vertices, and any vertex is either on the cycle or has a neighbor in the cycle.*

PROOF. Let G be an equilibrium graph for $(1, \dots, 1)$ -BG in the SUM version. If the number of players is two, then the only realization of the game consists of a 2-cycle, and the proof is complete. So we may assume that n is larger than 2. First, we show that G does not have a brace. Assume that there is a brace $\{u, v\}$. Because $n > 2$, there is a third vertex w such that at least one of u or v , say u , is not adjacent to w in $U(G)$. So, if u replaces its arc \overrightarrow{uw} with the arc \overrightarrow{uw} , its cost decreases, which is a contradiction. Hence, G does not have a brace.

Second, $U(G)$ is connected by Lemma 3.1, and has n edges, thus it has exactly one cycle. Indeed, since every vertex in G has outdegree 1, G has a unique directed cycle. Let $C = (v_1, v_2, \dots, v_k)$ be the unique directed cycle in G . For the ease of notation, define $v_0 = v_k$. Every vertex not in C is connected to C via a unique path. Let A_i be the set of vertices that are connected to C through v_i (including v_i itself). We may assume, by relabeling the vertices of C if necessary, that $|A_k| \geq |A_{k-1}|$.

Third, we show that k is at most 5. Suppose that $k > 5$. If v_{k-2} replaces its arc $\overrightarrow{v_{k-2}v_{k-1}}$ with the arc $\overrightarrow{v_{k-2}v_k}$, then its distances to vertices in A_k and to v_1 decrease by one (as $k > 5$), and the only increment in the cost of v_{k-2} would be because its distances to vertices in A_{k-1} are increased by one. Recall that $|A_{k-1}| < |A_k| + 1$, so this swap would

decrease the cost of v_{k-2} , which contradicts the assumption that G is an equilibrium graph. Hence, $k \leq 5$.

Fourth, to complete the proof, we show that for all $1 \leq i \leq k$, every vertex in A_i is either equal or adjacent to v_i . Assume that this is not the case for some i , and let v be a vertex in A_i with maximum distance from v_i . Let $l = \text{dist}(v, v_i)$. So, by the assumption, $l > 1$. Note that the subgraph of G induced by A_i is a tree T in which all arcs are directed toward v_i . Make T rooted by setting v_i as the root. Let p be the parent of v , and p' be the parent of p in T . Since v is not adjacent to v_i , we have $p \neq v_i$, and p' is well defined. Let W be the set of children of p in T , and let $F = W \cup \{p\} \setminus \{v\}$. If v replaces its arc \overrightarrow{vp} with the arc $\overrightarrow{vp'}$, then its distances to vertices in F increase by one, and its distances to all other vertices decrease by one since vertices of W have no children. As v is playing its best response, we have

$$|F| \geq n - |F| - 1. \quad (4)$$

If v_{i-1} replaces its arc $\overrightarrow{v_{i-1}v_i}$ with the arc $\overrightarrow{v_{i-1}p}$, its distances to vertices in $F \cup \{v\}$ decrease by $l - 1$, and its distance to any other vertex increases by at most $l - 1$. Since v_{i-1} is playing its best response, we have

$$(n - |F| - 2)(l - 1) \geq (|F| + 1)(l - 1),$$

which contradicts Equation (4) as $l > 1$. \square

Next, we prove a similar structure theorem for equilibria of $(1, \dots, 1)$ -BG in the MAX version.

THEOREM 4.2. *Any equilibrium graph for $(1, \dots, 1)$ -BG in the MAX version is connected, has a unique cycle with at most 7 vertices, and all vertices are within distance 2 of the cycle.*

PROOF. Let G be an equilibrium graph for $(1, \dots, 1)$ -BG in the MAX version. By Lemma 3.1, $U(G)$ is connected, so it has exactly one cycle. In fact, G has a unique directed cycle. (G may have a brace, which is simply a cycle with two vertices.)

Let $C = (v_1, v_2, \dots, v_k)$ be the unique directed cycle in G . We show that k is at most 7. Suppose that $k > 7$. Every vertex not in C is connected to C via a unique path. Let A_i be the set of vertices that are connected to C through v_i (including v_i itself). Define

$$m_i = \max\{\text{dist}(v, v_i) : v \in A_i\}.$$

We may assume, by relabeling the vertices of C if necessary, that $m_{1+\lfloor k/2 \rfloor}$ is the largest m_i . Then the eccentricity of v_1 is exactly $\lfloor k/2 \rfloor + m_{1+\lfloor k/2 \rfloor}$. It can be verified that if v_1 replaces its arc $\overrightarrow{v_1v_2}$ with the arc $\overrightarrow{v_1v_4}$, then because $k > 7$, the distance between v_1 and any vertex $u \in A_j$ becomes at most $\lfloor k/2 \rfloor - 1 + \text{dist}(v_j, u)$. We have

$$\left\lfloor \frac{k}{2} \right\rfloor - 1 + \text{dist}(v_j, u) \leq \left\lfloor \frac{k}{2} \right\rfloor - 1 + m_j < \left\lfloor \frac{k}{2} \right\rfloor + m_{1+\lfloor k/2 \rfloor},$$

which contradicts the assumption that v_1 is playing its best response in G . Hence, $k \leq 7$.

Finally, to complete the proof, we show that for all $1 \leq i \leq k$, every vertex in A_i is within distance at most 2 from v_i . Assume that this is not the case for some i , and let v be a vertex in A_i with maximum distance from v_i . Note that the subgraph of G induced by A_i is a tree T in which all arcs are directed toward v_i . Make T rooted by setting v_i as the root. Let p be the parent of v and p' be the parent of p in T . Since v is not adjacent to v_i , we have $p \neq v_i$, and p' is well defined. Since $\text{dist}(v, v_i) \geq 3$, the eccentricity of v is larger than 3. If v replaces its arc \overrightarrow{vp} with the arc $\overrightarrow{vp'}$, then its distances to all

vertices except the children of p decreases, and its distance to each child of p (other than v itself) increases by 1 and becomes at most 3. Consequently, the eccentricity of v_1 decreases, contradicting the fact that v is playing its best response in G . \square

5. A LOWER BOUND FOR THE PRICE OF ANARCHY IN THE MAX VERSION

We saw in the previous section that if all players have budget 1, then the equilibrium graphs have diameter $O(1)$. It appears intuitive that increasing the budgets would decrease the diameter of the equilibrium graphs. However, this is not true, and, in this short section, we prove that for some positive budget values, there exist equilibrium graphs in the MAX version with diameter $\Omega(\sqrt{\log n})$. This surprising phenomenon resembles Braess's paradox in network routing games [Braess et al. 2005]. This result implies that the price of anarchy of bounded budget network creation games when all players have positive budgets is $\Omega(\sqrt{\log n})$ in the MAX version.

LEMMA 5.1. *Let U be an undirected graph with n vertices, diameter d , and maximum degree Δ satisfying $\Delta^d - 1 < n(\Delta - 1)$. Then, for any vertex v and any subset A of vertices having size at most Δ , there exists a vertex u , different from v , with $\text{dist}(u, A) > d - 2$.*

PROOF. There are at most $|A|\Delta$ vertices whose distance from A is exactly 1. Similarly, there are at most $|A|\Delta^2$ vertices with distance exactly 2 from A . Continuing in the same way, we find that there are at most $|A|\Delta^{d-2}$ vertices with distance exactly $d - 2$ from A . If there is no $u \neq v$ with $\text{dist}(u, A) > d - 2$, then we must have

$$n \leq 1 + |A| + |A|\Delta + \cdots + |A|\Delta^{d-2} \leq 1 + \Delta + \Delta^2 + \cdots + \Delta^{d-1} = \frac{\Delta^d - 1}{\Delta - 1},$$

which contradicts the assumption $\Delta^d - 1 < n(\Delta - 1)$. Thus, there exists a vertex u , different from v , with $\text{dist}(u, A) > d - 2$. \square

LEMMA 5.2. *For every integer $t, k > 3$ satisfying $(2t)^k - 1 < t^k(2t - 1)$, there exists an undirected graph U with t^k vertices, minimum degree at least 2, and diameter k , such that every G with $U = U(G)$ is an equilibrium graph in the MAX version.*

PROOF. Let U be the graph with vertex set $\{1, 2, \dots, t\}^k$ and with vertices (x_1, x_2, \dots, x_k) and (y_1, y_2, \dots, y_k) being adjacent if at least one of the following happens:

- (1) $x_i = y_{i+1}$ for all $1 \leq i \leq k - 1$,
- (2) $y_i = x_{i+1}$ for all $1 \leq i \leq k - 1$.

Note that we want U to be a simple graph, so we only add edges between distinct vertices, and we add at most one edge between any pair. Then, U has minimum degree at least $t - 1$, maximum degree $2t$, and t^k vertices. The eccentricity of every vertex is k : For an arbitrary $(x_1, \dots, x_k) \in V(U)$ choose $y_1, \dots, y_k \notin \{x_1, \dots, x_k\}$. Note that this is always possible, since $(2t)^k - 1 < t^k(2t - 1)$ and $t, k > 3$ imply $t > k$. Then the distance between (x_1, \dots, x_k) and (y_1, \dots, y_k) is k .

Let G be a directed graph such that $U = U(G)$. Assume for the sake of contradiction that v is a vertex of G that is not playing its best response. Let A be the set of neighbors of v (vertices with an incoming arc from v or an outgoing arc to v) if it had changed its strategy and played its best response. Because v has degree at most $2t$, we have $|A| \leq 2t$. Since $(2t)^k - 1 < t^k(2t - 1)$, by Lemma 5.1, there exists a vertex u , different from v , with $\text{dist}(u, A) \geq k - 1$.

Now, suppose that v changes its strategy so that its neighborhood becomes A . Then it is not hard to see that for any vertex $w \neq v$, the new distance between w and A is not less than their old distance. In particular, the new distance between u and A is at

least $k - 1$. Hence, the new distance between u and v is at least k (i.e., the eccentricity of v has not decreased, a contradiction). Therefore, all vertices were playing their best responses in G , and G is an equilibrium graph. \square

THEOREM 5.3. *For infinitely many n , there exist positive integers b_1, b_2, \dots, b_n , such that there exists an equilibrium graph for (b_1, b_2, \dots, b_n) -BG in the MAX version with diameter $\sqrt{\log n}$.*

PROOF. Let $k > 3$ and $t = 2^k$. It is easy to check that we have $(2t)^k - 1 < t^k(2t - 1)$. Let U be the graph given by Lemma 5.2, which has $n = (2^k)^k = 2^{k^2}$ vertices, minimum degree at least 2, and diameter $k = \sqrt{\log n}$. Now, let G be a directed graph with $U(G) = U$ and such that the outdegree of all vertices of G is at least 1. Such a G exists because the minimum degree of U is larger than 1. Then G is an equilibrium graph by Lemma 5.2, and the proof is complete. \square

6. AN UPPER BOUND FOR THE PRICE OF ANARCHY IN THE SUM VERSION

In this section, we prove a general bound of $2^{O(\sqrt{\log n})}$ for the price of anarchy of bounded budget network creation games in the SUM version. The proof, which is long and consists of several steps, follows the line of the proof of Alon et al. [2013, Theorem 9], but the first step is more involved.

In the following, we focus on the SUM version. For a vertex u and a nonnegative integer r , define

$$B_r(u) = \{v : \text{dist}(u, v) \leq r\}.$$

The first step is to prove the following theorem:

THEOREM 6.1. *Let u be a vertex of an equilibrium graph G , and let r be a positive integer. Assume that the subgraph induced by $B_r(u)$ is a tree T . Then, $r = O(\log n)$.*

Consider T rooted at u . If every vertex in T has at least two children, then $r = O(\log |V(T)|)$, and since $|V(T)| \leq n$, this implies $r = O(\log n)$, and the theorem is proved. Hence, the problematic vertices are those with zero or one child. Roughly speaking, in the next three lemmas, we will prove that those vertices cannot increase the height of the tree significantly.

To prove Theorem 6.1, we need to consider weighted graphs. We denote a weighted directed graph by $G = (V, A, w)$, where V and A are the vertex set and the arc set of G , respectively, and $w : V \rightarrow \mathbb{Z}^+$ assigns a weight to each vertex. For every vertex u , the cost of u is defined as

$$c(u) = \sum_{v \in V} w(v) \text{dist}(u, v).$$

Note that if all vertices have unit weights, then this reduces to the original (unweighted) model. For a subgraph H of G , define

$$w(H) = \sum_{u \in V(H)} w(u).$$

We say that G is a *weak equilibrium graph* if no vertex can decrease its cost by swapping exactly one of its edges; more precisely, for every arc $\vec{uv} \in A$ and $x \in V$ with $\vec{ux} \notin A$, the cost of u does not decrease if the arc \vec{uv} is replaced with the arc \vec{ux} . Clearly, every equilibrium graph is also a weak equilibrium graph.

Remark 6.2. Mihalák and Schlegel [2012] studied weak equilibrium graphs under the name of “asymmetric swap-equilibria.” Among other results, they showed that the

underlying undirected graph of every weak equilibrium graph has at most one 2-edge-connected component.

A vertex of G with degree 1 is called a *leaf*. It turns out that one should distinguish between two types of leaves: a *poor* leaf is a leaf with outdegree 0, and a *rich* leaf is a leaf with outdegree 1. The poor leaves cause the most trouble, and they are the reason for introducing the weights. Let l be a poor leaf in G , and let $\vec{ul} \in A$. Define $G_0 = (V_0, A_0, w_0)$ to be a weighted directed graph with

$$V_0 = V \setminus \{l\}, \quad A_0 = A \setminus \{\vec{ul}\}, \quad w_0(v) = \begin{cases} w(v) & \text{if } v \neq u \\ w(u) + w(l) & \text{if } v = u. \end{cases}$$

Then it can be verified that if G is a weak equilibrium graph, then so is G_0 . We say that G_0 is obtained by *folding* the poor leaf l into u . The following lemma is used for handling the poor leaves:

LEMMA 6.3. *Let G be a weak equilibrium graph and T be an induced rooted subtree of G with root z . Assume that*

- every arc of T is oriented away from z , and
- no nonroot vertex of T is adjacent to a vertex outside T .

Then, the height of T is at most $1 + \log w(T)$.

PROOF. For every vertex v of T , let T_v be the subtree of T rooted at v . We will prove that for every $v \in V(T)$ that is not the root, if T_v has height k , then $w(T_v) \geq 2^k$. This shows that if the height of T is h , then $w(T) \geq 2^{h-1}$, or equivalently, $h \leq \log w(T) + 1$.

The proof is by induction on k . Correctness of the case $k = 0$ follows from the fact that all weights are positive integers. Assume that the induction hypothesis is true for k , and let $v \in V(T)$ be such that T_v has height $k + 1$. Let p be the parent and x_1, x_2, \dots, x_m be the children of v . At least one of $T_{x_1}, T_{x_2}, \dots, T_{x_m}$ has height k . We may assume without loss of generality that the height of T_{x_1} is k . By the induction hypothesis, $w(T_{x_1}) \geq 2^k$. We have

$$\sum_{i=2}^m w(T_{x_i}) + w(v) \geq w(T_{x_1}),$$

otherwise, the vertex p could decrease its cost by replacing the arc \vec{pv} with the arc $\vec{px_1}$. Thus, we find

$$w(T_v) = w(T_{x_1}) + w(T_{x_2}) + \dots + w(T_{x_m}) + w(v) \geq 2w(T_{x_1}) \geq 2^{k+1},$$

which completes the proof. \square

Note that if the conditions of Lemma 6.3 hold, then one can fold the whole subtree T into the vertex z . An induced rooted subtree T of G is called *foldable* if it satisfies the conditions of Lemma 6.3. Observe that if T_1 and T_2 are two foldable subtrees that have a nonempty intersection, then their union is also a foldable subtree. Hence, the maximal foldable subtrees of G are disjoint. Moreover, if z is the root of a maximal foldable subtree of G , then z is not a poor leaf in the graph obtained from G by folding T . So, if we fold all maximal foldable subtrees of G , we obtain a graph G' with no poor leaves. Using Lemma 6.3, we obtain the following corollary:

COROLLARY 6.4. *Let G be a weak equilibrium graph, and let G' be the graph obtained from folding all maximal foldable subtrees of G . Then G' is a weak equilibrium graph with no poor leaves, and any vertex in $V(G) \setminus V(G')$ has distance $O(\log w(G))$ to some vertex in $V(G')$.*

Handling rich leaves is easy, as shown by the following lemma:

LEMMA 6.5. *Let G be a weak equilibrium graph. Then the distance between any two rich leaves of G is at most 2.*

PROOF. Let u, v be two rich leaves. By symmetry, we may assume that

$$c(u) - \text{dist}(u, v)w(v) \leq c(v) - \text{dist}(u, v)w(u).$$

Since u is a rich leaf, it owns an arc $\overrightarrow{u\hat{p}}$. If v changes its strategy by replacing its outgoing arc with the arc $\overrightarrow{v\hat{p}}$, then its cost becomes

$$c(u) - \text{dist}(u, v)w(v) + 2w(u) \leq c(v) - \text{dist}(u, v)w(u) + 2w(u) = c(v) + (2 - \text{dist}(u, v))w(u).$$

Since v is already playing its best response, $\text{dist}(u, v) \leq 2$, and the proof is complete. \square

To handle the vertices of degree 2, which have one child, the following lemma will be used.

LEMMA 6.6. *Let G be a weak equilibrium graph and P be a path in $U(G)$ such that for every two vertices u and v in P , the (u, v) -path along P is the unique shortest (u, v) -path (which implies, in particular, that P is an induced subgraph of $U(G)$). Then the number of edges $uv \in E(P)$ such that both u and v have degree 2 is $O(\log w(P))$.*

PROOF. The proof is similar to the proof of Theorem 3.3. Let $P = v_0v_1 \dots v_d$. Suppose that the set of edges of P whose endpoints have degree 2 is

$$S = \{v_{i_1}v_{i_1+1}, v_{i_2}v_{i_2+1}, \dots, v_{i_m}v_{i_m+1}\}.$$

At least half of the arcs in S are in the same direction along P . By symmetry, we may assume that these are the arcs $\overrightarrow{v_{i_1}v_{i_1+1}}, \overrightarrow{v_{i_2}v_{i_2+1}}, \dots, \overrightarrow{v_{i_t}v_{i_t+1}}$, where $t \geq m/2$. For $1 \leq j < t$, if v_{i_j} replaces its arc $\overrightarrow{v_{i_j}v_{i_j+1}}$ with the arc $\overrightarrow{v_{i_j}v_{i_j+2}}$, then its distances to vertices in $\{v_{i_j+2}, \dots, v_d\}$ decrease by one, and its distance to v_{i_j+1} increases by one. Since v_{i_j} is playing its best response, we have

$$w(v_{i_j+1}) \geq \sum_{k=i_j+2}^d w(v_k) \geq \sum_{l=j+1}^t w(v_{i_l+1})$$

for all $1 \leq j < t$. Now, by the same reason as the one in the last part of the proof of Theorem 3.3, we have $m \leq 2t = O(\log w(P))$. \square

Now we prove Theorem 6.1.

PROOF OF THEOREM 6.1. First, make G weighted by setting $w(x) = 1$ for all vertices x . Thus, $w(G) = n$. Second, fold the maximal foldable subtrees of G . By Corollary 6.4, the height of T changes by $O(\log w(G))$, and the obtained graph is a weak equilibrium graph with no poor leaves.

Third, for each edge $xy \in E(T)$ such that both x and y have degree 2 in T , contract the edge, and repeat until no such edge exists. We claim that doing all these contractions changes the height of T by $O(\log w(G))$. Indeed, let v be any vertex in T and P be the unique (u, v) -path in T . Then, P satisfies the conditions of Lemma 6.6; so, by this lemma, the change in the distance between u and v after doing all these contractions is $O(\log w(P))$, which is $O(\log w(G))$.

Note that since the graph was a weak equilibrium right before doing the contractions, by Lemma 6.5 the distance between any two rich leaves was at most 2. The contractions do not increase the distances and do not create new leaves, so, in the final graph, there is at most one vertex in G that is adjacent to leaves. Hence, the height of the obtained

tree T' is $O(\log |V(T')|)$, which is $O(\log w(G))$. Therefore, the height of the original tree T is $O(\log w(G))$, which is $O(\log n)$ as well. \square

For the rest of the section, all graphs are unweighted. The rest of the proof is similar to the proof of Alon et al. [2013, Theorem 9].

LEMMA 6.7. *Let u, v, x be vertices of a graph G such that the arc \vec{uv} is not in G . Assume that adding the arc \vec{uv} to G decreases the cost of u by s , where $s > n \text{dist}(x, u)$. Then, adding the arc \vec{xv} to G decreases the cost of x by at least $s - n \text{dist}(x, u)$.*

PROOF. For every vertex w , let $\text{improve}_u(w)$ be the amount u gets closer to w by adding the arc \vec{uv} . Similarly, let $\text{improve}_x(w)$ be the amount x gets closer to w by adding the arc \vec{xv} . Let $\text{dist}^{\text{new}}(x, w)$ be the distance between x and w in $G \cup \vec{xv}$. Let W be the set of vertices w with $\text{improve}_u(w) > 0$. For all $w \in W$, we have

$$\begin{aligned} \text{improve}_x(w) &= \text{dist}(x, w) - \text{dist}^{\text{new}}(x, w) \geq (\text{dist}(u, w) - \text{dist}(u, x)) - (1 + \text{dist}(v, w)) \\ &= \text{dist}(u, w) - (1 + \text{dist}(v, w)) - \text{dist}(u, x) \\ &= \text{improve}_u(w) - \text{dist}(u, x). \end{aligned}$$

Thus,

$$\begin{aligned} \sum_{w \in V} \text{improve}_x(w) &\geq \sum_{w \in W} \text{improve}_x(w) \geq \sum_{w \in W} [\text{improve}_u(w) - \text{dist}(u, x)] \\ &= s - |W| \text{dist}(u, x) \geq s - n \text{dist}(u, x), \end{aligned}$$

and the proof is complete. \square

LEMMA 6.8. *Let G be a connected equilibrium graph in the SUM version that is not a tree. Given any vertex u , there is an arc \vec{xy} with $\text{dist}(x, u) = O(\log n)$ and whose removal increases the cost of x by at most $O(n \log n)$.*

PROOF. Let r be the smallest positive integer such that the subgraph induced by $B_{r+1}(u)$ has a cycle. Note that r is well defined since G is connected and is not a tree. Theorem 6.1 gives $r = O(\log n)$. Consider a breadth-first search from u in $U(G)$, and let T denote the top $r + 2$ levels of the BFS tree, from level 0 (just u) to level $r + 1$. Since $B_{r+1}(u)$ has a cycle, there is an edge $xy \notin E(T)$ with $x, y \in V(T)$. Assume by symmetry that the arc direction is from x to y . Clearly, $\text{dist}(x, u) \leq r + 1 = O(\log n)$. If the arc \vec{xy} is deleted, then the distance between x and any vertex increases by at most $1 + 2r$, since the shortest path can use the alternate path in T from x to the lowest common ancestor of x and y in T , and then to y , instead of using \vec{xy} . Therefore, removing \vec{xy} increases the cost of x by at most $n(1 + 2r) = O(n \log n)$. \square

Let p and q be positive constants such that if G is a connected nontree equilibrium graph, then for any $u \in V$, there is an arc \vec{xy} with $\text{dist}(x, u) \leq p \log n$ and whose removal increases the cost of x by at most $qn \log n$. These constants exist by Lemma 6.8. Using Lemma 6.7, we obtain the following corollary:

COROLLARY 6.9. *In a connected nontree equilibrium graph G in the SUM version, the addition of any arc \vec{uv} decreases the cost of u by at most $(p + q + 1)n \log n$.*

PROOF. Suppose for the sake of contradiction that the addition of \vec{uv} decreases the cost of u by more than $(p + q + 1)n \log n$. By the definition of p and q , there is an arc \vec{xy} with $\text{dist}(x, u) \leq p \log n$, and whose removal increases the cost of x by at most $qn \log n$. We show that if x replaces the arc \vec{xy} with the arc \vec{xv} , then its cost decreases, which contradicts the fact that G is an equilibrium graph. By Lemma 6.7, inserting the arc \vec{xv} decreases the cost of x by at least $(p + q + 1)n \log n - pn \log n$. Now, deleting the arc

$\vec{x}\vec{y}$ from the graph $G \cup \vec{x}\vec{v}$ increases the cost of x by at most $qn \log n$. This completes the proof since

$$-(p+q+1)n \log n + pn \log n + qn \log n < 0. \quad \square$$

Now we are ready to prove the main theorem of this section, which implies that the price of anarchy in the SUM version is $2^{O(\sqrt{\log n})}$.

THEOREM 6.10. *Let G be a connected equilibrium graph in the SUM version. Then the diameter of G is $2^{O(\sqrt{\log n})}$.*

PROOF. If G is a tree, then, by Theorem 3.3, its diameter is $O(\log n)$, which is $2^{O(\sqrt{\log n})}$, so we may assume that G is not a tree.

To prove the theorem, we first prove a certain “high-expansion property” for G and use it to deduce that the diameter cannot be too large. Define $f(k) = \min_u |B_k(u)|$. First, we show that the following high-expansion property holds:

$$f(4k) \geq \min \left\{ \frac{n+1}{2}, \frac{kf(k)}{4(p+q+1)\log n} \right\}. \quad (5)$$

This is true for $k=1$ because $f(4) \geq f(1)$, so we may assume that $k \geq 2$. Fix a vertex u and assume that $f(4k) \leq n/2$. Then, certainly, $f(3k) \leq n/2$. Let T be a maximal set of vertices at distance exactly $3k$ from u subject to the distance between any pair of vertices in T being at least $2k+1$. We claim that, for every vertex v of distance more than $3k$ from u , the distance of v from the set T is at most $\text{dist}(u, v) - k$. Indeed, v has distance $\text{dist}(u, v) - 3k$ to some vertex at distance exactly $3k$ from u , and any such vertex is within distance $2k$ of some vertex of T , by the maximality of T .

Because we assumed that at least $n/2$ vertices have distance of more than $3k$ from u , by the pigeonhole principle, there are at least $n/(2|T|)$ such vertices v whose distance from the same $t \in T$ is at most $\text{dist}(u, v) - k$. Adding the arc \vec{ut} decreases the distances between u and such vertices v by $k-1$, and so improves the cost of u by at least

$$(k-1) \frac{n}{2|T|} \geq \frac{kn}{4|T|}.$$

By Corollary 6.9, this improvement is at most $(p+q+1)n \log n$, so we find that

$$|T| \geq \frac{k}{4(p+q+1)\log n}.$$

Now, the sets $\{B_k(t) : t \in T\}$ are all pairwise disjoint, all lie within distance $4k$ of u , and each of them has at least $f(k)$ vertices (by the definition of f). Thus,

$$f(4k) \geq f(k) \frac{k}{4(p+q+1)\log n}$$

and Equation (5) holds.

Now we use the high-expansion property Equation (5) to prove the theorem. Note that

$$f(2^{\sqrt{\log n}}) \geq 2^{\sqrt{\log n}}$$

simply because G is connected. Let k be the smallest nonnegative integer for which $f(2^{\sqrt{\log n}} 4^k) > n/2$. By Equation (5), for every $1 \leq i < k$, we have

$$\frac{f(2^{\sqrt{\log n}} 4^i)}{f(2^{\sqrt{\log n}} 4^{i-1})} \geq \frac{2^{\sqrt{\log n}} 4^{i-1}}{4(p+q+1)\log n} = 2^{\Omega(\sqrt{\log n})}.$$

One can prove by induction on i that for all $1 \leq i < k$,

$$f(2^{\sqrt{\log n}} 4^i) = 2^{\Omega(i \sqrt{\log n})}.$$

But, since

$$f(2^{\sqrt{\log n}} 4^{k-1}) \leq n/2 = 2^{\log n - 1},$$

we have $k = O(\sqrt{\log n})$. Recall that we have $f(2^{\sqrt{\log n}} 4^k) > n/2$. Thus, for any two vertices u and v ,

$$B_{2^{\sqrt{\log n}} 4^k}(u) \cup B_{2^{\sqrt{\log n}} 4^k}(v) \neq \emptyset,$$

which means that there is a vertex x within distance $2^{\sqrt{\log n}} 4^k$ of both u and v . So the distance between u and v is at most

$$2 \times 2^{\sqrt{\log n}} 4^k = 2^{O(\sqrt{\log n})},$$

and the proof is complete because u and v were chosen arbitrarily. \square

COROLLARY 6.11. *The price of anarchy of any bounded budget network creation game in the SUM version is $2^{O(\sqrt{\log n})}$.*

PROOF. Consider a bounded budget network creation game. If the sum of players' budgets is less than $n - 1$, then clearly any realization of the game has diameter n^2 , so the price of anarchy is 1. Otherwise, by Lemma 3.1, every equilibrium graph is connected and so by Theorem 6.10 has diameter $2^{O(\sqrt{\log n})}$, and this completes the proof. \square

7. VERTEX CONNECTIVITY OF EQUILIBRIUM GRAPHS IN THE SUM VERSION

One of the most important issues in designing stable networks is the connectivity of the resulting network. In this section, we find a direct connection between the budget limits and the connectivity of an equilibrium graph in the SUM version, which shows that we can guarantee strong connectivity or a small diameter for our network when all players have large enough budgets. Assume that all players have budgets at least k for some positive integer k , and let G be any equilibrium graph in the SUM version. We show that if G has diameter larger than 3, then it is k -connected; that is, removing any $k - 1$ vertices does not disconnect the graph. By Menger's theorem (see, e.g., Diestel [2005, Theorem 3.3.6]), this also means that if the diameter of G is larger than 3, then for any two vertices u and v , there exist k internally disjoint paths connecting u and v . For a subset A of vertices, let $G - A$ denote the subgraph induced by $V(G) \setminus A$.

Before proving this, we remark that it is not generally true that whenever all players have budgets at least k , any equilibrium graph is k -connected. Consider the graph illustrated in Figure 4, which is a realization of $(2, 2, 2, 2, 2, 2)$ -BG that is not 2-connected. The central vertex has eccentricity 1, and all other vertices have eccentricity 2. The graph has no brace, so it is an equilibrium graph by Lemma 2.2.

First, we prove a lemma that will be used in the proof of the main result of this section.

LEMMA 7.1. *Let G be an equilibrium graph in the SUM version. Let C be a subset of vertices of G and A be the vertex set of a connected component of $U(G - C)$. Assume that, for all $a \in A$, $\text{dist}(a, C) = 1$ and the budget of vertex a is larger than $|C|$. Then, every vertex in A has eccentricity at most 2.*

PROOF. Fix a vertex $a \in A$ and let C' be the set of vertices in C that are not adjacent to a . Note that C' might be empty. Assume for the sake of contradiction that there




Fig. 4. An equilibrium graph for $(2,2,2,2,2,2,2,2)$ -BG that is not 2-connected.

exists a vertex x with $\text{dist}(a, x) > 2$. Since the budget of a is larger than $|C|$, and all neighbors of a are in $A \cup C$, vertex a owns at least $|C'| + 1$ arcs to vertices in A . Let A' be a set of $|C'| + 1$ vertices in A to which a has an arc. If vertex a deletes all these arcs and adds arcs to the vertices $C' \cup \{x\}$ instead, then the following happens to the distances between a and other vertices:

- The distance between a and every vertex in A' changes from 1 to at most 2 (recall that every vertex in A' has a neighbor in C);
- If a vertex in $A \setminus A'$ had distance 1 to a , its distance to a does not change.
- If a vertex in $A \setminus A'$ had distance more than 1 to a , its distance to a becomes at most 2 and hence does not increase.
- The distances between a and vertices outside A do not increase; moreover, its distances to vertices in C' decrease by at least 1, and its distance to x decreases by at least 2.

Since $|A'| = |C'| + 1$, the cost of vertex a decreases by this strategy change, which contradicts the assumption that a is playing its best response. Therefore, every $a \in A$ has eccentricity at most 2 in G . \square

Now we prove the main result of this section.

THEOREM 7.2. *Suppose that G is an equilibrium graph in the SUM version and all vertices have budgets at least k . If G has diameter greater than 3, then it is k -connected.*

PROOF. Let C be a minimum cardinality subset of vertices such that $U(G - C)$ is disconnected and let k' be the size of C . If $k' \geq k$ then G is k -connected and the theorem is proved, so we may assume that $k' < k$, and we need to show that the diameter of G is at most 3. Let A be the vertex set of the connected component of $U(G - C)$ with minimum size, and let $B = V(G) \setminus (A \cup C)$.

We claim that every vertex in A is adjacent to some vertex in C . Let $u \in A$ be a vertex with maximum distance from C , and let $l = \text{dist}(u, C)$. If $l = 1$, then the claim is proved, so assume that $l > 1$. Since the budget of u is larger than k' and u has no neighbor in C , there is a set $Q \subseteq A$ of size k' to each vertex of which u has an arc. Let us study what happens if u changes its strategy by removing the arcs to Q and adding arcs to all vertices in C . Note that after this change, the distance between u and any vertex in A becomes at most $l + 1$.

- The distance between u and every vertex in $C \cup B$ decreases by at least $l - 1$, giving a total decrease of at least $(|B| + k')(l - 1)$.
- The distance between u and every vertex in Q changes from 1 to at most $l + 1$, giving a total increase of at most $k'l$.

—The distance between u and every vertex $v \in A \setminus (Q \cup \{u\})$ increases by at most $l - 1$: if v was adjacent to u , it remains adjacent, and if v had distance at least 2, its distance becomes at most $l + 1$. Therefore, the total increase is at most $(|A| - k' - 1)(l - 1)$.

Therefore, the total change in the cost of u is less than or equal to

$$-(|B| + k')(l - 1) + k'l + (|A| - k' - 1)(l - 1) = (|A| - |B|)(l - 1) + k'(2 - l) + 1 - l < 0,$$

which contradicts the fact that u is playing its best response. Thus, $l = 1$, and every vertex in A is adjacent to some vertex in C .

Now, by Lemma 7.1, every vertex in A has eccentricity 2. Let D be any connected component of $U(G - C)$ other than A . Every vertex in D is at distance 2 from A and so is adjacent to some vertex in C . Therefore, by Lemma 7.1, every vertex in D has eccentricity 2. Since D was arbitrary, we conclude that every vertex in $U(G - C)$ has eccentricity 2. To complete the proof, we just need to show that every vertex in C has eccentricity at most 3. Let $x \in C$ be arbitrary. Since the budget of x is larger than $|C|$, it has a neighbor u outside C . Since the eccentricity of u is 2, the eccentricity of x is at most 3 and the proof is complete. \square

8. CONCLUDING REMARKS

We considered network creation games in which every vertex has a specific budget for the number of vertices it can establish links to and the distances are computed based on the underlying graph of the resulting network. Two versions of this game were defined, depending on whether each vertex wants to minimize its maximum distance or sum of distances to other vertices. We proved that Nash equilibria exist in both versions by explicitly constructing them. A natural question is, if the game starts from an arbitrary position and the players keep on improving their strategies, does the game converge to an equilibrium? If yes, then how quickly does it converge? Note that Laoutaris et al. [2008] demonstrated that converging to a pure Nash equilibrium is not guaranteed in their game by constructing an explicit loop. (Since the time of our original writeup, Kawald and Lenzner [2013] have studied the dynamics of many variants of network creation games. In particular, they demonstrated that in both the SUM and MAX versions of our model, even if all vertices have unit budgets, convergence to equilibrium is not guaranteed.)

We studied the diameter of equilibrium graphs and the price of anarchy of this game and found asymptotically tight bounds in two extreme cases: namely, when the sum of players' budgets is $n - 1$ (the threshold for connectivity) and when all players have unit budgets. There are other special cases that might be interesting; for example, the cases in which all players have the same budget $B > 1$.

We observed a nonmonotone property of this game (in the MAX version): When all budgets are 1, the diameter of equilibrium graphs (and the price of anarchy) is $O(1)$; but when all budgets are at least 1, there are instances of the game in which the diameter of equilibrium graphs is $\Omega(\sqrt{\log n})$, and so the price of anarchy is also $\Omega(\sqrt{\log n})$. We proved that this happens in the MAX version, and it is interesting to know if it can also happen in the SUM version. More precisely, when all players have positive budgets, is the diameter of equilibrium graphs in the SUM version bounded?

In the SUM version, we proved a general upper bound of $2^{O(\sqrt{\log n})}$ for the diameter of equilibrium graphs (and the price of anarchy). The proof is complicated and based on showing a certain expansion property for equilibrium graphs. The bound $2^{O(\sqrt{\log n})}$ seems strange and probably is not tight. Finally, we showed an interesting connection between the minimum budget available to the players and the connectivity of

the resulting network. More precisely, if each vertex has budget at least k , then any equilibrium graph in the SUM version has diameter less than 4, or is k -connected.

ACKNOWLEDGMENTS

The authors are thankful to Nastaran Nikparto for her useful comments throughout the preparation of this article and to Laura Poplawski Ma for suggesting the idea of defining the distance of vertices in different components to be $C_{\text{inf}} = n^2$, which resulted in slightly cleaner arguments. We also thank the anonymous referees for carefully reading the manuscript and their many useful comments.

REFERENCES

- Susanne Albers, Stefan Eilts, Eyal Even-Dar, Yishay Mansour, and Liam Roditty. 2006. On Nash equilibria for a network creation game. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'06)*. ACM, New York, 89–98. DOI: <http://dx.doi.org/10.1145/1109557.1109568>
- Noga Alon, Erik D. Demaine, MohammadTaghi Hajiaghayi, and Tom Leighton. 2013. Basic network creation games. *SIAM Journal on Discrete Mathematics* 27, 2 (2013), 656–668. DOI: <http://dx.doi.org/10.1137/090771478>
- Xiaohui Bei, Wei Chen, Shang-Hua Teng, Jialin Zhang, and Jiajie Zhu. 2011. Bounded budget betweenness centrality game for strategic network formations. *Theoretical Computer Science* 412, 52 (Dec. 2011), 7147–7168. DOI: <http://dx.doi.org/10.1016/j.tcs.2011.09.028>
- Davide Bilò, Luciano Gualà, and Guido Proietti. 2012. Bounded-distance network creation games. In *Proceedings of the 8th International Conference on Internet and Network Economics (WINE'12)*. Springer-Verlag, Berlin, Heidelberg, 72–85. DOI: http://dx.doi.org/10.1007/978-3-642-35311-6_6
- Dietrich Braess, Anna Nagurney, and Tina Wakolbinger. 2005. On a paradox of traffic planning. *Transportation Science* 39, 4 (2005), 446–450.
- Ulrik Brandes, Martin Hoefer, and Bobo Nick. 2008. Network creation games with disconnected equilibria. In *Proceedings of the 4th International Workshop on Internet and Network Economics (WINE'08)*. Springer-Verlag, Berlin, 394–401. DOI: http://dx.doi.org/10.1007/978-3-540-92185-1_45
- Michael Brautbar and Michael Kearns. 2011. A clustering coefficient network formation game. In *Proceedings of the 4th International Symposium on Algorithmic Game Theory (SAGT'11)*. Springer-Verlag, Berlin, 224–235.
- Nicos Christofides. 1976. *Worst-case Analysis of a New Heuristic for the Travelling Salesman Problem*. Technical Report CS-93-13. Carnegie Mellon University, Graduate School of Industrial Administration, Pittsburgh, PA.
- Jacomo Corbo and David Parkes. 2005. The price of selfish behavior in bilateral network formation. In *Proceedings of the 24th Annual ACM Symposium on Principles of Distributed Computing (PODC'05)*. ACM, New York, NY, 99–107. DOI: <http://dx.doi.org/10.1145/1073814.1073833>
- Erik D. Demaine, Mohammad Taghi Hajiaghayi, Hamid Mahini, and Morteza Zadimoghaddam. 2007. The price of anarchy in network creation games. In *Proceedings of the 26th Annual ACM Symposium on Principles of Distributed Computing (PODC'07)*. ACM, New York, NY, 292–298. DOI: <http://dx.doi.org/10.1145/1281100.1281142>
- Reinhard Diestel. 2005. *Graph Theory* (3rd ed.). Graduate Texts in Mathematics, Vol. 173. Springer-Verlag, Berlin.
- Alex Fabrikant, Ankur Luthra, Elitza Maneva, Christos H. Papadimitriou, and Scott Shenker. 2003. On a network creation game. In *Proceedings of the 22nd Annual Symposium on Principles of Distributed Computing (PODC'03)*. ACM, New York, NY, 347–351. DOI: <http://dx.doi.org/10.1145/872035.872088>
- Michal Feldman, Kevin Lai, Ion Stoica, and John Chuang. 2004. Robust incentive techniques for peer-to-peer networks. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC'04)*. ACM, New York, NY, USA, 102–111. DOI: <http://dx.doi.org/10.1145/988772.988788>
- Philippe Golle, Kevin Leyton-Brown, Ilya Mironov, and Mark Lillibridge. 2001. Incentives for sharing in peer-to-peer networks. In *Proceedings of the 2nd International Workshop on Electronic Commerce (WELCOM'01)*. Springer-Verlag, London, UK, 75–87.
- Wen Lian Hsu and George L. Nemhauser. 1979. Easy and hard bottleneck location problems. *Discrete Applied Mathematics* 1, 3 (1979), 209–215. DOI: [http://dx.doi.org/10.1016/0166-218X\(79\)90044-1](http://dx.doi.org/10.1016/0166-218X(79)90044-1)
- Matthew O. Jackson and Asher Wolinsky. 1996. A strategic model of social and economic networks. *Journal of Economic Theory* 71, 1 (1996), 44–74. DOI: <http://dx.doi.org/10.1006/jeth.1996.0108>

- David S. Johnson, Maria Minkoff, and Steven Phillips. 2000. The prize collecting Steiner tree problem: Theory and practice. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'00)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 760–769.
- Bernd Kawald and Pascal Lenzner. 2013. On dynamics in selfish network creation. In *Proceedings of the 25th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA'13)*. ACM, New York, NY, 83–92. DOI: <http://dx.doi.org/10.1145/2486159.2486185>
- Elias Koutsoupias and Christos Papadimitriou. 1999. Worst-case equilibria. In *Proceedings of the 16th Annual Conference on Theoretical Aspects of Computer Science (STACS'99)*. Springer-Verlag, Berlin, Heidelberg, 404–413. <http://dl.acm.org/citation.cfm?id=1764891.1764944>
- Nikolaos Laoutaris, Laura J. Poplawski, Rajmohan Rajaraman, Ravi Sundaram, and Shang-Hua Teng. 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 (PODC'08)*. ACM, New York, NY, 165–174. DOI: <http://dx.doi.org/10.1145/1400751.1400774>
- Jyh-Han Lin and Jeffrey Scott Vitter. 1992. ϵ -approximations with minimum packing constraint violation. In *Proceedings of the 24th Annual ACM Symposium on Theory of Computing (STOC'92)*. ACM, New York, NY, 771–782. DOI: <http://dx.doi.org/10.1145/129712.129787>
- Matúš Mihalák and Jan Christoph Schlegel. 2010. The price of anarchy in network creation games is (mostly) constant. In *Proceedings of the 3rd International Symposium on Algorithmic Game Theory (SAGT'10)*. Springer-Verlag, Berlin, 276–287.
- Matúš Mihalák and Jan Christoph Schlegel. 2012. Asymmetric swap-equilibrium: A unifying equilibrium concept for network creation games. In *Proceedings of the 37th International Conference on Mathematical Foundations of Computer Science (MFCS'12)*. Springer-Verlag, Berlin, 693–704. DOI: http://dx.doi.org/10.1007/978-3-642-32589-2_60
- F. Sibel Salman, Joseph Cheriyan, Ramamoorthi Ravi, and Sairam Subramanian. 2000. Approximating the single-sink link-installation problem in network design. *SIAM Journal on Optimization* 11 (March 2000), 595–610. Issue 3. DOI: <http://dx.doi.org/10.1137/S1052623497321432>
- Brian Skyrms and Robin Pemantle. 2000. A dynamic model of social network formation. In *Proceedings of the National Academy of Sciences* 97, 16 (2000), 9340–9346. DOI: <http://dx.doi.org/10.1073/pnas.97.16.9340>

Received June 2012; revised November 2014; accepted December 2014