

Constant Price of Anarchy in Network-Creation Games via Public-Service Advertising

Erik D. Demaine and Morteza Zadimoghaddam

Abstract. Network-creation games have been studied recently in many different settings. These games are motivated by social networks in which selfish agents want to construct a connection graph among themselves. Each node wants to minimize its average or maximum distance to the others, without paying much to construct the network. Many generalizations have been considered, including nonuniform interests between nodes, general graphs of allowable edges, and bounded-budget agents. In all of these settings, there is no known constant bound on the price of anarchy. In fact, in many cases, the price of anarchy can be very large, namely, a constant power of the number of agents. This means that we have no control over the behavior of a network when agents act selfishly. On the other hand, the price of stability in all these models is constant, which means that there is chance that agents act selfishly and we end up with a reasonable social cost.

In this paper, we show how to use an advertising campaign (as introduced in [Balcan et al. 09]) to find efficient equilibria in an (n, k) -uniform bounded-budget connection game [Laoutaris et al. 08]; our result holds for $k = \omega(\log(n))$. More formally, we present advertising strategies such that if an α fraction of the agents agree to cooperate in the campaign, the social cost will be at most $O(1/\alpha)$ times the optimum cost. This is the first constant bound on the price of anarchy that interestingly can be adapted to different settings. We also generalize our method to work in cases in which α is not known in advance. Also, we do not need to assume that the cooperating agents spend

all their budget on the campaign; even a small fraction (β fraction) of their budget is sufficient to obtain a constant price of anarchy.

1. Introduction

In network-creation games, nodes construct an underlying graph in order to have short routing paths among themselves. So each node incurs two types of costs: a network design cost, which is the amount of the contribution of the node in constructing the network; and a network usage cost, which is the sum of the distances to all other nodes. Nodes act selfishly, and every node wants to minimize its own cost, i.e., the network design cost plus the usage cost. The social cost in these games is equal to the sum of the costs of all the nodes.

To study the behavior of social networks, we try to understand how large the social cost can be in the presence of selfish agents. Formally, we have a set of selfish agents $N = \{1, 2, \dots, n\}$. Agent i chooses some strategy $s_i \in S_i$ from its set of possible strategies (actions) S_i . Combining these strategies of players, we get the strategy profile $s = (s_1, s_2, \dots, s_n)$ among the set of all strategy profiles $\times_{i=1}^n S_i$. For each $1 \leq i \leq n$, player i has value function v_i that maps each strategy profile s to some value $v_i(s)$ for player i . Since agents are acting selfishly, they try to maximize their own values. In particular, we are interested in Nash equilibria, which are the stable strategy profiles of this game. A strategy profile $s = (s_1, s_2, \dots, s_n)$ is a Nash equilibrium if and only if for each player i , strategy s_i is $\operatorname{argmax}_{s'_i \in S_i} v_i(s_1, s_2, \dots, s_{i-1}, s'_i, s_{i+1}, \dots, s_n)$.

Nash equilibria are the stable networks in which every agent is acting selfishly. Intuitively, in a Nash equilibrium, every agent has no incentive to change its strategy on the assumption that all other agents keep the same strategies. In this setting, the price of anarchy is the worst ratio of the social cost of Nash equilibria and the optimal social cost of a network that can be designed by a central authority. In our setting, the social cost of a strategy profile s is the sum of players' values for s , i.e., $\sum_{i=1}^n v_i(s)$.

The price of anarchy was introduced in [Koutsoupias and Papadimitriou 99, Papadimitriou 01], and it is used to measure the behavior of games and networks with selfish agents. The small values of price anarchy suggest that allowing agents to be selfish does not increase the social cost by much. On the other hand, large values of the price of anarchy indicate that the selfish behavior of agents can lead the whole game (network) to stable situations with large social cost in comparison with the optimal cases.

1.1. The Model

In a *network-creation game*, there is a set of selfish nodes. Every node can construct an undirected edge to any other node at a fixed given cost.¹ So the strategy set of each node is a subset of other nodes (its neighbors). Each node also incurs a usage cost related to its distance to the other nodes. So the usage cost of a node is the sum of its distances to all other nodes. Clearly, every node is trying to minimize its own total cost, i.e., usage cost plus the construction cost. So every player's value is the negative of its cost.

In another variant of network-creation games, called an (n, k) -*uniform bounded-budget connection game*, we have n nodes in the graph, and each node can construct up to k edges to other nodes. So every node has only the usage cost, but its budget to build edges is limited. The strategy set of each node is a subset of size at most k of other nodes. In both directed and undirected settings, an edge is built if both of its endpoints have the edge in their strategy sets. Edge (i, j) is constructed if and only if $j \in s_i$ and $i \in s_j$. For undirected graphs, every node can have k edges in its strategy, and for directed graphs, every node can have k incoming and k outgoing edges in its strategy.

The advertising-campaign scenario can be applied to different game-theoretic situations. In these scenarios, we can use a public-service advertising campaign to encourage players to follow a specific strategy. We can design the strategy to improve the social cost.

In our model, we find an advertising strategy to reduce the price of anarchy and control the behavior of selfish nodes. We do not need everyone to help us in order to achieve a small price of anarchy. We assume that an α fraction of players are willing to follow our strategy, and each of them agrees to spend β fraction of its budget on the campaign. Formally, we assume that every node agrees to contribute to the campaign with probability α . We call these users *receptive* users, as used in the literature [Balcan et al. 09]. Every receptive user is willing to use βk of its edges for the campaign. At first we assume that α and β are some known parameters, and we present a strategy that leads the network to an equilibrium with small price of anarchy. Then we adapt our strategies to work in cases in which α and β are not known in advance, but some lower bounds on these two parameters are given. To get constant bounds on the price of anarchy, we assume that k is greater than $\frac{c \log n}{\alpha \beta}$ for some sufficiently large constant c .

¹One can get the same results using the same techniques and maintaining two ingoing and outgoing trees from the root for directed graphs as well.

1.2. Previous Work

Network-creation games were introduced in [Fabrikant et al. 03]. The authors studied the price of anarchy in these games, and achieved the first nontrivial bounds on it. They studied the structure of Nash equilibria, and conjectured that only trees can be stable graphs in this model. Later, an interesting class of stable graphs was presented in [Albers et al. 06] that disproved the tree conjecture. The authors of this paper also presented better upper bounds on the price of anarchy. They proved that the price of anarchy cannot be more than $O(n^{1/3})$ in general, and in some cases they gained a constant upper bound on the price of anarchy. In [Corbo and Parkes 05], the authors considered a slightly different model called *bilateral network-formation games*, and they studied the price of anarchy in this model. They were able to prove a $O(\sqrt{c})$ upper bound on the price of anarchy, where c is the cost of constructing one edge in their model. Since c can be as large as n , this bound is also as large as n to the power of a constant.

Demaine et al. studied the sizes of neighborhood sets in stable graphs, and with a recursive technique, they presented the first subpolynomial bounds on the price of anarchy [Demaine et al. 07]. They also studied a variant of these games called *cooperative network-creation games*, and they were able to achieve the first polylogarithmic upper bounds on the price of anarchy [Demaine et al. 09]. This result actually shows that the diameter of stable graphs is polylogarithmic, which implies the small-world phenomenon in these games. For more details about the small-world phenomenon, we refer to [Kleinberg 01, Kleinberg 00].

Laoutaris et al. studied network-creation games in the bounded-budget model [Laoutaris et al. 08]. They claimed that in many practical settings, a selfish agent cannot build an arbitrary number of edges to other nodes even if there is an incentive for the node in building the edge. In this model, every node has a limited budget, and according to the limit, each node can build up to a given number of edges. They call these games *uniform bounded-budget connection games*, and they achieve sublinear upper and lower bounds on the price of anarchy in these games. They prove that the price of anarchy in these games is between

$$\Omega\left(\sqrt{\frac{n/k}{\log_k(n)}}\right) \quad \text{and} \quad O\left(\sqrt{\frac{n}{\log_k(n)}}\right),$$

where n and k are respectively the number of nodes and the maximum number of edges that each node can have. Although this is an interesting model in the sense that each node has a limited number of edges, in these games the price of anarchy can be very large.

In many games including network-creation, selfish routing, and fair cost-sharing, the cost of a stable graph can vary in a large range. In other words,

we have both low-cost and high-cost Nash equilibria. Balcan et al. claim that in such games one can hope to lead the game to low-cost equilibria using a public-service advertising campaign [Balcan et al. 09]. They study the price of anarchy using some advertising strategies. In some cases, such as their fair cost-sharing, they present advertising strategies that reduce the price of anarchy to a constant number, and in some other games, such as scheduling games, they show that no useful advertising strategy exists.

1.3. Our Results

Since the uniform games have a very large price of anarchy [Laoutaris et al. 08], we try to find advertising strategies to reduce the price of anarchy to a constant number in uniform bounded-budget games. In this way, we can be sure that the degree of each node is bounded, so no node is overwhelmed in the network. On the other hand, we also know that the price of anarchy is small, so the behavior of these games is under control.

Formally, we present an advertising strategy that leads the game to equilibria with price of anarchy at most $O(1/\alpha)$, where α is the fraction of nodes that follow our strategy. We do not assume that everyone is willing to contribute to our strategy, and even if α is very small, we still get a small price of anarchy. We also do not assume that every node that contributes to our strategy is willing to spend all its k edges as we suggest. We use only βk edges of a player that contributes to the advertising strategy, where $0 < \beta < 1$ can be a small constant.

In Section 2, we present an advertising strategy that knows the values of α and β in advance. Then in Section 3, we adapt our strategy to work in cases in which these two parameters are not given in the input. However, we assume in Section 3 that two lower bounds on the values of α and β are given.

2. How Public-Service Advertising Affects the Price of Anarchy

In this section we present strategies that lead the network to stable graphs with low social costs in settings of undirected graphs and in those of directed graphs. We first present our strategy and its analysis for the undirected setting, which is simpler. We assume that every node follows our strategy with probability α . We call these follower nodes *receptive nodes* because of their interest in the advertised strategy. We also do not ask a node to spend all its budget on our strategy. A receptive node has to spend just βk edges on our strategy ($0 < \beta < 1$), and it can use the rest of its edges arbitrarily. At first we assume that α and β are

some parameters given in advance. In Section 3, we change our strategies to be adaptive and to work when these parameters are not revealed in advance.

2.1. Advertising Strategy for the Undirected Setting

In this part, we assume that the edges are undirected and that every node wants to minimize its total distance to all other nodes. The main idea is to make a low-diameter subgraph using the receptive nodes and to use this low-diameter subgraph as a global hub to route other nodes' traffic as well. Every receptive node is willing to spend βk of its budget on the advertising campaign. We ask every receptive node to use $\beta k/2$ of its budget to form a low-diameter subgraph between receptive nodes, and the other half of its budget on nonreceptive nodes. We explain in our strategy how the receptive nodes should manage the β fraction of their budget in our campaign. The other $1 - \beta$ fraction of budgets of receptive nodes, and the whole budget of nonreceptive nodes, is managed selfishly as in other game-theoretic settings.

The advertising strategy is as follows. Define k' to be $\frac{\alpha\beta}{c \log(n)} k$ for a sufficiently large constant c ; for example, $c \geq 5$ would work. We assume that $k' > 1$, i.e., k is greater than $\frac{c \log(n)}{\alpha\beta}$. We partition the nodes into $l \leq \log_{k'}(n)$ sets S_1, S_2, \dots, S_l such that $|S_1| = \beta k/6$, and $|S_{i+1}|/|S_i| = k'$ for each $1 \leq i < l$. Note that the only important properties of these sets are their sizes. For example, we can set S_1 to be the nodes $1, 2, \dots, |S_1|$, set S_2 to be the nodes $|S_1| + 1, \dots, |S_1| + |S_2|$, and so on.

We ask nodes in the first set S_1 to construct edges to all other nodes in set S_1 . So every receptive node in set S_1 uses $\beta k/6 - 1$ edges to get directly connected to all other nodes in S_1 . For $i > 1$, we ask each node in set S_i to pick $c \log(n)/6\alpha$ nodes randomly from set S_{i-1} and construct edges to them. Note that $c \log(n)/6\alpha$ is at most $\beta k/6$, because k is greater than $\frac{c \log(n)}{\alpha\beta}$.

On the other hand, nodes in set S_{i-1} receive some edges from set S_i . We ask all receptive nodes to accept up to $\beta k/3$ edges coming from their lower sets. So if a receptive node $u \in S_i$ wants to make an edge to a receptive node $v \in S_{i-1}$ that also has not exceeded its $\beta k/3$ budget in our strategy, this edge (u, v) will be accepted.

So we cannot assume that every proposed edge in our strategy is accepted. For example, if a nonreceptive node receives an edge, the node might delete the edge, e.g., the node is not interested in our strategy.

Even if the node in set S_{i-1} is receptive, the edge is not necessarily accepted. Assume that the node $v \in S_{i-1}$ receives more than $\beta k/3$ edges from set S_i , and that node v deletes some of these proposed edges. So a receptive node might become overwhelmed by the nodes in the lower set. So we have to take into

account these overwhelmed receptive nodes in our analysis. So we can assume that if a receptive node receives at most $\beta k/3$ edges from the nodes of the lower set, it will accept these proposed edges.

Lemma 2.1. *The edges built in the above strategy form a hierarchical tree-shaped subgraph with $\log_{k'}(n)$ levels. The diameter of this subgraph is at most $2 \log_{k'} n + 1$, and every receptive node is contained in this subgraph with high probability.² Every receptive node spends at most $\beta k/2$ of its budget on this part of the advertising campaign.*

Proof. We just need to prove that every receptive node v in set S_i gets connected to a receptive node v' in set S_{i-1} , and node v' does not delete the edge (v, v') , i.e., node v' does not become overwhelmed. Node v picks $c \log(n)/6\alpha$ random nodes in set S_{i-1} . There are $c \log(n)/6$ receptive nodes among these nodes in expectation, because every node is receptive with probability α . Using a Chernoff bound, we can say that there are at least $2 \log(n)$ receptive nodes among them with high probability (note that c is sufficiently large).

So every receptive vertex v in level i is connected to at least $\log(n)$ receptive nodes in set $i - 1$ unless they delete their incoming edges because they have been overwhelmed. Now we prove that every node is overwhelmed in this structure with probability at most $1/2$.

Each node in set S_i is receptive with probability α . Each receptive node makes $c \log(n)/6\alpha$ edges to the nodes in set S_{i-1} randomly. So the expected number of incoming edges from set S_i to a node in set S_{i-1} is equal to

$$\frac{\alpha |S_i| (c \log(n)/6\alpha)}{|S_{i-1}|}.$$

We also know that $\frac{|S_i|}{|S_{i-1}|}$ is equal to $k' = \frac{\alpha\beta}{c \log(n)} k$. We conclude that every node u in set S_{i-1} receives $\alpha\beta k/6$ edges in expectation. Using a Markov inequality, we can say that a node can be overwhelmed with probability at most $\alpha/2 < 1/2$ because a receptive node is overwhelmed when it receives more than $\beta k/3$ edges.

So every node $v \in S_i$ is connected to at least $\log(n)$ receptive nodes in set S_{i-1} . Each of them is overwhelmed with probability at most $1/2$. Since the overwhelming events for different nodes are negatively correlated, we can say that with high probability, node v is connected to at least one receptive node in set S_{i-1} that is not overwhelmed. This is sufficient to see that with high probability, each receptive node has a path of length at most l to some receptive

²Probability $1 - 1/n^{c'}$ for some large constant c' .

node in set S_1 , where l is the number of levels. Since receptive nodes in set S_1 makes direct edges to all other nodes in set S_1 (and to themselves as well), they form a complete graph. We conclude that the diameter of all receptive nodes is at most $2l + 1 = 2 \log_{k'}(n) + 1$ with high probability.

We also know that each receptive node spends at most $\beta k/6 + \beta k/3 = \beta k/2$ of its budget at this stage to make a low-diameter subgraph between receptive nodes. \square

Up to now, we have shown how receptive nodes should use $\beta k/2$ edges to make a low-diameter subgraph between themselves. Now we show how they should use the other half in the campaign to help other nodes get close to them. If a nonreceptive node v wants to make an edge to a receptive node u , this edge will be accepted by u as one of its $\beta k/2$ extra edges if node u has distance greater than $2 \log_{k'} n + 2$ from v in the current graph and node u has not exceeded its quota of $\beta k/2$ edges. Note that every receptive node has $(1 - \beta)k$ other edges to use selfishly. One useful implication of this strategy is that a receptive node u does not use one of its $\beta k/2$ edges to accept an edge from a node v that has already had an edge to one of the receptive nodes. Because if v has an edge to a receptive node u' , the distance between v and u will be at most $2 \log_{k'} n + 1 + 1$, which is a contradiction. The expected number of receptive nodes is αn , and with high probability this number is not less than $\alpha n/2$. So there is a budget of at least $(\alpha n/2) \cdot (\beta k/2) > n$ for the edges from nonreceptive nodes to receptive ones. We have also proved that a nonreceptive node v cannot consume two of these edges. So there is always a free spot. We formalize this discussion in the following lemma:

Lemma 2.2. *If a nonreceptive node v has distance greater than $2 \log_{k'} n + 2$ from all receptive nodes, then there always exists a receptive node u that has not used its $\beta k/2$ budget completely for accepting incoming edges from far receptive nodes, so node v can make a link to node u and be sure that this edge will be accepted.*

Now we can bound the diameter of the whole graph (not only the subgraph of receptive nodes).

Lemma 2.3. *The diameter of a stable graph after running the advertisement strategy is at most $O(\log_{k'}(n)/\alpha)$.*

Proof. Using Lemma 2.1, we know that with high probability, the diameter of receptive nodes is at most $2l + 1 = O(\log_{k'}(n))$. There are αn receptive nodes in expectation, and with high probability the number of them is not less than $\alpha n/2$.

Consider a nonreceptive vertex v . We prove that v has distance at most $O(l + \log_k(n)/\alpha)$ to some receptive node. Delete all edges in G that are contained in some cycle of length at most $l' = l + 2\log_k(n) + 1$. We prove that if one of the k edges of v is in a cycle of length at most l' , the distance from v to some receptive node is at most $5l'/\alpha$. Let e be an edge owned by v that is in a cycle of length at most l' . Let u be the receptive node in Lemma 2.2 that has a free spot. If the distance between v and u is at most $5l'/\alpha$, the claim is proved. Otherwise, v can delete edge e and make an edge to node u . This edge will be accepted by u , because $5l'/\alpha$ is greater than $2\log_{k'}n + 2$ and node u has some free spot as well. We also show that v has incentive to switch these two edges.

If node v deletes edge e , its distance to other nodes increases by at most $l' \times n$. On the other hand, if v makes an edge to node u , its distance to all receptive nodes decreases by at least

$$\frac{5l'}{\alpha} - (2\log_{k'}n + 2) \geq \frac{3l'}{\alpha}$$

(before adding the edge, its distances to receptive nodes were at least $5l'/\alpha$, and after that the distances are at most $2\log_{k'}n + 2$). So the total decrease in the cost of u would be at least

$$\frac{\alpha n}{2} \cdot \frac{3l'}{\alpha} = \frac{3l'n}{2}$$

because there are at least $\alpha n/2$ receptive nodes with high probability. Therefore, node v has incentive to switch these two edges, which contradicts the fact that we are in the stable graph. So node v has distance at most $5l'/\alpha$ to some receptive node.

We call a vertex *incomplete* if at least one of its edges is deleted (in the process of removing edges in short cycles). As proved above, each incomplete vertex is at distance at most $O(l'/\alpha)$ from some receptive node. We also note that the remaining graph does not have a cycle of length at most l' . We claim that each vertex is either incomplete or has distance at most l' from one of the incomplete vertices. So the distances of all vertices from v is at most $l' + O(l'/\alpha) = O(l'/\alpha)$.

Consider a vertex v' and all walks of length $l'/2$ starting from v' in the remaining graph. If one of these walks passes over an incomplete vertex, the claim is proved. Otherwise, we have $k^{l'/2}$ walks starting from the same vertex u . The endpoints of these walks are also different. Otherwise, we find a cycle of length at most l' in the remaining graph. So there are at least $k^{l'/2} > n$ different vertices in the graph, which is a contradiction because l' is greater than $2\log_k(n)$.

So every vertex in the graph has distance at most $O(l'/\alpha)$ from some receptive node. Note that $O(l'/\alpha) = O((l + \log_k(n))/\alpha)$, that l is equal to $\log_{k'}(n)$, and

that k' is at most k . So the diameter of the whole graph is simply at most $O(\log_{k'}(n)/\alpha)$. \square

Theorem 2.4. *The price of anarchy is at most*

$$O\left(\frac{\log_{k'}(n)}{\alpha \log_k(n)}\right) = O\left(\frac{\log_{k'}(k)}{\alpha}\right)$$

using our advertising strategy, where k' is $\frac{\alpha\beta}{c \log(n)}k$ for a constant c .

Proof. Using Lemma 2.3, the diameter of a stable graph is at most $O(\log_{k'}(n)/\alpha)$. Since each vertex has degree at most k , the average distance in the optimal graph is at least $\Omega(\log_k(n))$. Combining these two facts completes the proof of this lemma. \square

Corollary 2.5. *For $k = \Omega(\log^{1+\epsilon}(n))$, the price of anarchy is $O(1/\alpha\epsilon)$.*

Proof. Note that α and β are some constant parameters. So k/k' is $O(\log(n))$. Since k is at least $\Omega(\log^{1+\epsilon}(n))$, we can say that k is at most $O(k^{1/\epsilon})$. This shows that $\log_{k'}(k)$ is $O(1/\epsilon)$, which completes the proof. \square

Corollary 2.6. *For $k = \Omega(\log(n))$, the price of anarchy is at most $O(\log \log(k)/\alpha)$.*

Proof. One has simply to set k' to an appropriate constant. The rest is similar to the above line of argument. \square

2.2. Advertising Strategy for Directed Setting

We now consider the situation that all edges are directed and the cost of a node v is the sum of distances from v to all other nodes plus the sum of distances from all other nodes to v . Every node can have k incoming edges and k outgoing edges, i.e., a budget of k for incoming edges and another k for outgoing ones. An edge is constructed from node u to node v if and only if they are both willing to have this edge. Note that this edge costs both u and v . It is clear that the average distance from any node and to any node is $\Omega(\log_k n)$, because of the degree (budget) limitations. Here we show how to achieve average distance $O(\log_k n)$ and therefore a constant price of anarchy using an advertising campaign.

The advertising strategy is very similar to the undirected strategy. In the undirected setting, we had two main goals:

1. Construct a low-diameter subgraph on receptive nodes.
2. Make sure that edges coming from nonreceptive nodes far from the receptive nodes will be accepted.

Our directed strategy is as follows. In the first part, receptive nodes make edges in both directions as in the undirected part. A receptive node v in set S_i chooses $c \log n / 6\alpha$ nodes in set S_{i-1} and makes edges from itself (v) to these nodes, and edges from these selected nodes to itself (v). Since we make edges in both directions, these edges work as undirected edges. Therefore, we can say that the receptive nodes form a low-diameter subgraph, and they don't spend more than $\beta k / 2$ of their budget of incoming edges and the same amount of their outgoing budget. In other words, the proof of Lemma 2.1 works for the following lemma as well.

Lemma 2.7. *The edges built in the above directed strategy form a hierarchical tree-shaped subgraph with $\log_{k'}(n)$ levels. The (directed) diameter of this subgraph is at most $2 \log_{k'} n + 1$, and every receptive node is contained in this subgraph with high probability.³ Every receptive node spends at most $\beta k / 2$ of its incoming budget and at most $\beta k / 2$ of its outgoing budget in this part of the advertising campaign.*

So we can build a low-diameter subgraph on receptive nodes successfully. The second part of the strategy is to make sure that if a nonreceptive node v wants to make an edge to some receptive node, and the distance of v to receptive nodes is more than a threshold, there exists a receptive node u that has a free spot to accept an edge from v to u (the same is true for the other direction— incoming edges to nonreceptive nodes). We ask every receptive node u to save $\beta k / 2$ incoming edges and $\beta k / 2$ outgoing edges for distant nonreceptive nodes. The strategy is that if a nonreceptive node v wants to make an edge to a receptive node u , this edge will be accepted by node u if u has not exceeded its $\beta k / 2$ incoming budget and the distance from v to u in the current graph is more than $1 + 2 \log_{k'} n + 1$. The same is true for the other direction. Note that with high probability, there are $\alpha n / 2$ receptive nodes, and each of them has $\beta k / 2$ budget for incoming and $\beta k / 2$ budget for outgoing edges from/to far nonreceptive nodes. In total, there is at least $\alpha \beta k n / 4 > n$ budget for each direction. We make sure that no vertex can use more than one edge from each direction, so there is

³Probability $1 - 1/n^{c'}$ for some large constant c' .

always some free spot if a far receptive node wants to make an edge from/to some receptive node. In particular, the following lemma can be proved with an argument similar to the proof of Lemma 2.2.

Lemma 2.8. *If a nonreceptive node v has distance greater than $2 \log_{k'} n + 2$ from/to all receptive nodes, there always exists a receptive node u that has not used its $\beta k/2$ budget completely for accepting outgoing/incoming edges to/from far receptive nodes, so node v can make a link from/to node u and be sure that this edge will be accepted.*

The only left thing is to prove that all nonreceptive nodes are close from/to some receptive node, since otherwise, something will contradict the stability of the equilibrium graph. Here we want to prove a lemma similar to Lemma 2.3, but the proof of Lemma 2.3 cannot be applied to the directed setting, because in that proof, we are somehow using the fact that undirected graphs without short cycles (large girth) are sparse. This is not necessarily true for directed graphs. In particular, the endpoints of the walks of length $l'/2$ starting from v' in that proof are not necessarily distinct. Because in contrast to the undirected setting, two different walks of length $l'/2$ from a vertex to another one in a directed graph do not provide a cycle of length at most $2 \cdot l'/2 = l'$. So we have to cope with this problem in another way. We use the following lemma, which has been proved as [Laoutaris et al. 08, Lemma 1]. Although the directed setting of [Laoutaris et al. 08], and in particular the cost function of players, is slightly different from our setting, the same proof for the following lemma works in our setting as well.

Lemma 2.9. *In any stable graph, the outgoing (incoming) cost of any node is at most $n + n \lfloor \log_k n \rfloor$ more than the cost of any other node. The outgoing/incoming cost of a node is the sum of distances to/from all other nodes.*

Now we can prove our main lemma, which can help us to obtain constant bounds on the price of anarchy in the directed setting.

Lemma 2.10. *The diameter of a stable graph after running the directed advertisement strategy is at most $O(\log_{k'}(n)/\alpha)$.*

Proof. Using Lemma 2.7, we know that with high probability the diameter of receptive nodes is at most $2l + 1 = O(\log_{k'}(n))$, where we recall that l is the number of levels in the tree construction of receptive nodes. There are αn receptive nodes in expectation, and with high probability the number of them is not less than $\alpha n/2$.

Let d be the diameter of the stable graph G . We want to prove that d is $O(\log_{k'}(n)/\alpha)$. Fix a receptive node v_r as a center. We know that either the distance of some nonreceptive node to v_r or the distance of v_r to some nonreceptive node is at least $d/2$. Otherwise, the distances between all nonreceptive nodes will also be less than d , and the diameter d will be at most the diameter between receptive nodes, which is $O(\log_{k'}(n))$, so the claim is proved in that case. Without loss of generality we can assume that the distance from a nonreceptive node v to the receptive node v_r is at least $d/2$.

Using Lemma 2.8, we know that there exists a receptive node u that has a free spot for incoming edges. So if node v (or any other node like v) is willing to make an edge to u , and the distance from v to u is greater than $2 + 2 \log_{k'} n$, which is the case here, the edge (v, u) will be accepted by u . Consider the arborescence tree T rooted at u that contains shortest paths of u to all vertices; i.e., we can find a directed tree T in the stable graph in which every vertex has incoming degree 1 except u (which has incoming degree zero in T), and the shortest paths of u to all other vertices are contained in T .

We call a vertex x *half-complete* if x has not used more than $k/2$ outgoing edges in this tree T , i.e., the outgoing degree of x in T is at most $k/2$. It is not hard to prove that there exists an incomplete vertex v' such that the distance from v to v' is at most $\log_{k/2} n + 1$, because if we start at v in tree T and go down $1 + \log_{k/2} n$ levels, we will see an incomplete vertex. Otherwise, there will be at least $(k/2)^{1 + \log_{k/2} n} > n$ vertices in T , which is a contradiction. We consider this half-complete vertex v' , which has at least $k/2$ outgoing edges $e_1, e_2, \dots, e_{k/2}$ that are not present in tree T . We make $k/2$ disjoint groups of vertices as follows. For every vertex x , consider one of the shortest paths from v' to x arbitrarily, and if this shortest path is using edge e_i for some $1 \leq i \leq k/2$, put vertex x in set S_i . Since at most one of these $k/2$ edges is used in each shortest path, these sets are disjoint. So one of these sets, like S_i , has size at most

$$\frac{n}{k/2} = \frac{2n}{k}.$$

We know that vertex u has shortest paths to vertices in S_i in tree T , and therefore u does not need edge e_i to reach the vertices of set S_i . On the other hand, the distance of u to these vertices is not more than the diameter of the graph d . Now if vertex v' chooses to remove edge e_i and make a directed edge to u instead, its incoming cost will be increased, because e_i is an outgoing edge for v' . The distances of v' to vertices in $V(G) \setminus S_i$ are also not increased. Recall that e_i is used only in some shortest paths from v' to set S_i . So the total cost of v' might

be increased by at most

$$|S_i| \cdot (d + 1) \leq \frac{n(d + 1)}{k}.$$

Delete all edges in G that are contained in some cycle of length at most $l' = l + 2 \log_k(n) + 1$. We prove that if one of the k edges of v is in a cycle of length at most l' , the distance from v to some receptive node is at most $5l'/\alpha$. Let e be an edge owned by v that is in a cycle of length at most l' . Let u be the receptive node in Lemma 2.2 that has a free spot. If the distance between v and u is at most $5l'/\alpha$, the claim is proved. Otherwise, v can delete edge e and make an edge to node u . This edge will be accepted by u because $5l'/\alpha$ is greater than $2 \log_{k'} n + 2$, and node u has some free spot as well. We also show that v has an incentive to switch these two edges.

If node v deletes edge e , its distance to other nodes increases by at most $l' \times n$. On the other hand, if v makes an edge to node u , its distance to all receptive nodes decreases by at least

$$\frac{5l'}{\alpha} - (2 \log_{k'} n + 2) \geq \frac{3l'}{\alpha}$$

(before adding the edge, its distances to receptive nodes were at least $5l'/\alpha$, and after that, the distances are at most $2 \log_{k'} n + 2$). So the total decrease in the cost of u would be at least

$$\frac{\alpha n}{2} \cdot \frac{3l'}{\alpha} = \frac{3l'n}{2}$$

because there are at least $\alpha n/2$ receptive nodes with high probability. Therefore node v has incentive to switch these two edges, which contradicts the fact that we are in the stable graph. So node v has distance at most $5l'/\alpha$ to some receptive node.

We call a vertex *incomplete* if at least one of its edges is deleted (in the process of removing edges in short cycles). As proved above, each incomplete vertex is at distance at most $O(l'/\alpha)$ from some receptive node. We also note that the remaining graph does not have a cycle of length at most l' . We claim that each vertex is either incomplete or has distance at most l' from one of the incomplete vertices. So the distances of all vertices from v is at most $l' + O(l'/\alpha) = O(l'/\alpha)$.

Consider a vertex v' and all walks of length $l'/2$ starting from v' in the remaining graph. If one of these walks passes over an incomplete vertex, the claim is proved. Otherwise, we have $k^{l'/2}$ walks starting from the same vertex u . The endpoints of these walks are also different, for otherwise, we find a cycle of length at most l' in the remaining graph. So there are at least $k^{l'/2} > n$ different vertices in the graph, which is a contradiction because l' is greater than $2 \log_k(n)$.

So every vertex in the graph has distance at most $O(l/\alpha)$ from some receptive node. Note that $O(l/\alpha) = O((l + \log_k(n))/\alpha)$, and l is equal to $\log_{k'}(n)$, and k' is at most k . So the diameter of the whole graph is simply at most $O(\log_{k'}(n)/\alpha)$. \square

3. How to Deal with Unknown α and β

In Section 2, we presented an advertising strategy that given two parameters α and β , led the network to some equilibria with a small price of anarchy. Here we try to make our strategy adaptive for the cases in which the parameters are not known in advance, i.e., sometimes a large number of agents contribute in the campaign, and sometimes only a small fraction of them participate. So in these cases, we know that an $\alpha > \epsilon$ fraction of agents are willing to spend $\beta > \epsilon'$ fraction of their budget in the campaign, where ϵ and ϵ' are two given lower bounds on these two parameters. We note that these two lower bounds are two constants that can be very small.

Define m and m' to be the two smallest integers such that $\epsilon > 1/2^m$ and $\epsilon' > 1/2^{m'}$. So there exist two integers i and j such that $1/2^i \leq \alpha \leq 1/2^{i-1}$ and $1/2^j \leq \beta \leq 1/2^{j-1}$, where $1 \leq i \leq m$ and $1 \leq j \leq m'$.

Note that we do not need to know the exact values of parameters α and β in the advertising strategy, just an estimate would work. For example, if we know two integers i and j such that $1/2^i \leq \alpha \leq 1/2^{i-1}$ and $1/2^j \leq \beta \leq 1/2^{j-1}$, we can run the above strategy with parameters $1/2^i$ and $1/2^j$ instead of α and β . The same probabilistic bounds would work in the same way, and we can prove the same claims as proved in Section 2. But we do not even have good estimates of these two parameters. The only thing we know is that they are in the range $[\epsilon, 1]$ and $[\epsilon', 1]$, respectively.

But we know that α is in one of the m ranges

$$[1/2, 1], \quad [1/4, 1/2], \quad \dots, \quad [1/2^m, 1/2^{m-1}],$$

and the same holds for β . We should run the strategy for different estimates of α and β in a parallel manner. So there are $m \times m'$ different pairs of estimates for our parameters. But a receptive agent contributes only βk edges to the campaign. We can ask a receptive node to spend $\frac{\beta k}{m \times m'}$ in each of these runs. Note that in order to run a strategy, we need to set four parameters: α , β , k , and n . Here we want to use the strategy for $m \times m'$ parallel runs. So for each pair (i, j) , we run the strategy with parameters $1/2^i$, $1/2^j$, $\frac{k}{m \times m'}$, and n (instead of α , β , k , and n) for each $1 \leq i \leq m$ and $1 \leq j \leq m'$. Each receptive node spends at most βk edges in all the runs. The only thing that changes our upper bounds on the price

of anarchy is the new value of k in each run. In fact, we are using $\frac{k}{m \times m'}$ edges to reduce the price of anarchy. So we have the following theorem for cases in which the parameters are not known in advance.

Theorem 3.1. *When the parameters $\alpha > \epsilon$ and $\beta > \epsilon'$ are not known in advance, the price of anarchy is at most*

$$O\left(\frac{\log_{k'}(n)}{\alpha \log_k(n)}\right) = O\left(\frac{\log_{k'}(k)}{\alpha}\right)$$

using the above advertising strategy (updated version), where k' is $\frac{\alpha\beta}{c \log(n)} \times \frac{k}{m \times m'}$ for a constant c . Integers m and m' are $\lceil \log(1/\epsilon) \rceil$ and $\lceil \log(1/\epsilon') \rceil$, respectively.

Proof. When we run the original strategy for different pairs (i, j) , one of these pairs is a good estimate for α and β . Using the edges constructed by the receptive nodes in this specific run of the strategy and Theorem 2.4, we can attain this bound. The only difference is that we can use $\frac{k}{m \times m'}$ in each run, and that is why the value of k' is divided by a factor of $m \times m'$. \square

Since ϵ and ϵ' are two (probably very small) constants, we can say that m and m' are also some (probably large) constants. We conclude that Corollaries 2.5 and 2.6 are also true in this case (unknown α and β).

References

- [Albers et al. 06] S. Albers, S. Eilts, E. Even-Dar, Y. Mansour, and L. Roditty. “On Nash Equilibria for a Network Creation Game.” In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms. Miami, FL*, pp. 89–98, 2006.
- [Balcan et al. 09] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. “Improved Equilibria via Public Service Advertising.” In *Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms. New York, NY*, pp. 728–737, 2009.
- [Corbo and Parkes 05] J. Corbo and D. Parkes. “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*, pp. 99–107, 2005.
- [Demaine et al. 07] Erik D. Demaine, MohammadTaghi Hajiaghayi, Hamid Mahini, and Morteza Zadimoghaddam. “The Price of Anarchy in Network Creation Games.” In *Proceedings of the 26th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing*, pp. 292–298, 2007.
- [Demaine et al. 09] Erik D. Demaine, MohammadTaghi Hajiaghayi, Hamid Mahini, and Morteza Zadimoghaddam. “The Price of Anarchy in Cooperative Network Creation Games.” In *SIGecom Exchanges 8.2*, December 2009. (A preliminary version of this

- paper appeared in *Proceedings of the 26th International Symposium on Theoretical Aspects of Computer Science, 2009*, pp. 171–182.)
- [Fabrikant et al. 03] A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. Shenker. “On a Network Creation Game.” In *Proceedings of the 22nd Annual Symposium on Principles of Distributed Computing, Boston, Massachusetts*, pp. 347–351, 2003.
- [Kleinberg 00] Jon Kleinberg. “The Small-World Phenomenon: An Algorithmic Perspective.” In *Proceedings of the 32nd ACM Symposium on Theory of Computing*, 2000.
- [Kleinberg 01] Jon Kleinberg. “Small-World Phenomena and the Dynamics of Information.” *Advances in Neural Information Processing Systems (NIPS)* 14, 2001.
- [Koutsoupias and Papadimitriou 99] E. Koutsoupias and C. Papadimitriou. “Worst-Case Equilibria.” In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science, Trier, Germany*, Lecture Notes in Computer Science 1563, pp. 404–413. New York: Springer, 1999.
- [Laoutaris et al. 08] N. Laoutaris, L. J. Poplawski, R. Rajaraman, R. Sundaram, and S.-H. Teng. “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*, pp. 165–174, 2008.
- [Papadimitriou 01] C. Papadimitriou. “Algorithms, Games, and the Internet.” In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing, Hersonissos, Greece*, pp. 749–753, 2001.

Erik D. Demaine, MIT Computer Science and Artificial Intelligence Laboratory, Cambridge, MA 02139, USA (edemaine@mit.edu)

Morteza Zadimoghaddam, MIT Computer Science and Artificial Intelligence Laboratory, Cambridge, MA 02139, USA (morteza@mit.edu)

Received April 1, 2011; accepted July 19, 2011.