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

Erik D. Demaine and Morteza Zadimoghaddam

MIT Computer Science and Artificial Intelligence Laboratory, Cambridge, MA 02139, USA,
{edemaine,morteza}@mit.edu

Abstract. Network creation games have been studied in many different settings recently. 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 non-uniform interests between nodes, general graphs of allowable edges, bounded budget agents, etc. 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 on the behavior of 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 SODA 2009 [2]) to find such efficient equilibria in (n, k)-*uniform bounded budget connection* game [10]; 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 would 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 that α is not known in advance. Also, we do not need to assume that the cooperating agents spend all their budget in the campaign; even a small fraction (β fraction) of their budget is sufficient to obtain a constant price of anarchy.

Keywords: algorithmic game theory, price of anarchy, selfish agents

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, network design cost which is the amount of the contribution of the node in constructing the network, and network usage cost which is the sum of the distances to all other nodes. Nodes act selfishly, and everyone 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 sum of the costs of all nodes.

To study the behavior of social networks, we try to understand how large the social cost can be in presence of selfish agents. Formally, we have a set of selfish agents $N = \{1, 2, \cdots, n\}$. 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 $s = (s_1, s_2, \cdots, s_n)$ 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 Equilibriums which are the stable strategy profiles of this game. A strategy profile $s = (s_1, s_2, \cdots, s_n)$ is a Nash equilibrium if and only if for each player i, strategy s_i is $argmax_{s'_i \in S_i} v_i(s_1, s_2, \cdots, s_{i-1}, s'_i, s_{i+1}, \cdots, s_n)$.

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 her strategy assuming 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 the network which 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 is introduced by Koutsoupias and Papadimitriou in [9, 11], and is used to measure the behavior of the games and networks with selfish agents. The small values of price anarchy shows that allowing agents to be selfish does not increase the social cost a lot. On the other hand, large values of price of anarchy means 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.

Model In a *network creation game*, there is a set of selfish nodes. Every node can construct an undirected[1] edge to any other node at a fixed given cost. So the strategy set of each node is a subset of other nodes (as 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 negative its cost.

In another variant of network creation games, called (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 only have 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

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

(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 encourage people using a public service advertising 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 to achieve a small price of anarchy. We assume that α fraction of people are willing to follow our strategy, and each of them agrees to spend β fraction of its budget in the campaign. Formally, we assume that every node accepts to contribute in the campaign with probability α. We call these users *receptive* users as used in the literature [2]. Every receptive person 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 an strategy that leads the network to an equilibrium with small price of anarchy. Then we adapt our strategies to work in cases that 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.

Previous Work Fabrikant et al. introduced the network creation games [6]. They studied the price of anarchy in these games, and achieved the first non-trivial bounds on it. They studied the structure of Nash Equilibria, and conjectured that only trees can be stable graphs in this model. Later, Abers et al. came up with an interesting class of stable graphs, and disproved the tree conjecture [1]. They also presented better upper bounds on the price of anarchy. They proved that the price of anarchy can not be more than $O(n^{1/3})$ in general, and in some cases they gained a constant upper bound on the price of anarchy. Corbo and Parkes in [3] considered a slightly different model called bilateral network formation games, and 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 the stable graphs, and with a recursive technique, they presented the first sub-polynomial bounds on the price of anarchy [5]. They also studied a variant of these games called cooperative network creation games, and they were able to achieve the first poly-logarithmic upper bounds on the price of anarchy [4]. This result actually shows that the diameter of stable graphs is poly-logarithmic which implies the small-world phenomenon in these games. For more details about the small world phenomenon, we refer to Kleinberg's works [7, 8].

Laoutaris et al. studied the network creation games in the bounded budget model [10]. They claimed that in many practical settings a selfish agent can not 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 amount of 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, they achieve both sub-linear upper and lower bounds on the price of anarchy in these games. They prove that the price of anarchy is between $\Omega(\sqrt{\frac{n/k}{\log_k(n)}})$, and $O(\sqrt{\frac{n}{\log_k(n)}})$ where n and k are respectively the number of nodes, and the maximum number of edges that each node can have in these games. Although this is an interesting model in the sense that each node has a limited number of edges, the price of anarchy can be very large in these games.

In many games including network creation, selfish routing, fair cost sharing, etc, 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 service [2]. They study the price of anarchy using some advertising strategies. In some cases like fair cost sharing, they present advertising strategies that reduce the price of anarchy to a constant number, and in some other games like scheduling games, they show that there exists no useful advertising strategy.

Our Results Since the uniform games have a very large price of anarchy [10], we try to find advertising strategies to reduce the price of anarchy to a constant number in uniform bounded budget games. This way we can be sure that the degree of each node is bounded, so no one 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 in our strategy, and even if α is very small, we still get small price of anarchy. We also do not assume that every node that contributes in our strategy is willing to spend all its k edges as we say. We just use βk edges of a player that contributes in 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 that these two parameters are not given in the input. However we assume that two lower bounds on values of α and β are given in Section 3.

2 How the 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 both undirected and directed graphs settings. At first, we present our strategy and its analysis for 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 person to spend all its budget in our strategy.

A receptive node just has to spend βk edges in our strategy ($0 < \beta < 1$), and can use the rest of its edges arbitrarily. At first we assume that α and β are some given parameters in advance. In Section 3, we change our strategies to be adaptive and work when these parameters are not revealed in advance.

2.1 Advertising Strategy for Undirected Setting

In this part, assume that the edges are undirected and 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 use this low diameter sungraph as a global hub to route other nodes' traffic as well. Every receptive node is willing to spend βk of its budget for 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 for non-receptive 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 non-receptive nodes is managed selfishly like all 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, i.e. $c \geq 5$ 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 $S_1, S_2, \cdots, S_l$ such that $|S_1| = \beta k/6$, and $\frac{|S_{i+1}|}{|S_i|} = k'$ 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 $1, 2, \cdots, |S_1|$, set S_2 to be the nodes $|S_1|+1, \cdots, |S_1|+|S_2|$, 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 $v \in S_{i-1}$ which also has not exceeded its $\beta k/3$ budget in our strategy, this edge (u, v) will be accepted.

So we can not assume that every proposed edge in our strategy is accepted. For example if a non-receptive 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, node v deletes some of this proposed edges. So a receptive node might get 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 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 in 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 get 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 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 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 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.

Up to now, we showed 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 non-receptive 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

[2] probability $1 - 1/n^{c'}$ for some large constant c'

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 which has already had an edge to one of the receptive node. 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 are at least $(\alpha n/2)\cdot(\beta k/2) > n$ budget for the edges from non-receptive nodes to receptive ones. We also proved that a non-receptive node v can not consume two of these edges. So there is always a free spot. We can formalize this discussion with the following lemma:

Lemma 2. *If a non-receptive node v has distance greater than $2\log_{k'} n+2$ from all receptive nodes, there exists always 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 3. *The diameter of a stable graph after running the advertisement strategy is at most $O(\log_{k'}(n)/\alpha)$.*

Proof. Using Lemma 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 non receptive 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 which is in a cycle of length at most l'. Let u be the receptive node in lemma 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 $5l'/\alpha-(2\log_{k'} n+2) \geq 3l'/\alpha$ (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 $(\alpha n/2)\cdot(3l'/\alpha) = 3l'n/2$ because there are at least $\frac{\alpha n}{2}$ 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 in 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 $O(l'/\alpha) = O((l+\log_k(n))/\alpha)$, 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 $O(\log_{k'}(n)/\alpha)$.

Theorem 1. *The price of anarchy is at most $O(\frac{\log_{k'}(n)}{\alpha\log_k(n)}) = O(\frac{\log_{k'}(k)}{\alpha})$ using our advertising strategy where k' is $\frac{\alpha\beta}{c\log(n)}k$ for a constant c.*

Proof. Using Lemma 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.

Corollary 1. *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.

Corollary 2. *For $k = \Omega(\log(n))$, the price of anarchy is at most $O(\log\log(k)/\alpha)$.*

Proof. One just need to set k' to an appropriate constant. The rest is similar to above.

2.2 Advertising Strategy for Directed Setting

In this setting, 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 non-receptive nodes far from the receptive nodes will be accepted.

Our directed strategy is as follows. At the first part, receptive nodes make edges in both directions like the undirected part. A receptive node v in set S_i chooses $c \log n / 6\alpha$ nodes in set S_{i-1}, and make edges from itself (v) to these nodes, and edge 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 incoming edges budget, and their outgoing budget. In other words, the proof of Lemma 1 works for the following lemma as well.

Lemma 4. *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 non-receptive 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 free spot to accept an edge from v to u, (the same is true for the other direction (incoming edges to non-receptive nodes)). We ask every receptive node u to save $\beta k/2$ incoming edges, and $\beta k/2$ outgoing edges for far non-receptive nodes. The strategy is that if a non-receptive 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 than 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 non-receptive 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 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 similarly to the proof of Lemma 2.

Lemma 5. *If a non-receptive node v has distance greater than $2\log_{k'} n + 2$ from/to all receptive nodes, there exists always 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 the all non-receptive nodes are close from/to some receptive node, otherwise something contradicts the stability of

[3] probability $1 - 1/n^{c'}$ for some large constant c'

the Equilibria graph. Here we want to prove a lemma similar to Lemma 3, but proof of Lemma 3 can not be turned into 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 different. Because unlike the undirected setting, two different walks of length $l'/2$ from a vertex to another one in a directed graph does 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 Lemma 1 in [10]. Although the directed setting of [10] and in particular cost function of players is slightly different from our setting, the same proof for following lemma works in our setting as well.

Lemma 6. *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 the main lemma which can help us to obtain constant bounds on the price of anarchy in the directed setting.

Lemma 7. *The diameter of a stable graph after running the directed advertisement strategy is at most* $O(\log_{k'}(n)/\alpha)$.

Proof. Using Lemma 4, we know that with high probability the diameter of receptive nodes is at most $2l + 1 = O(\log_{k'}(n))$, remember 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 non-receptive node to v_r or the distance of v_r to some non-receptive node is at least $d/2$. Otherwise the distances between all non-receptive nodes will be also less than d, 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 non-receptive v to the receptive node v_r is at least $d/2$.

Using lemma 5, we know that there exists a receptive node u that has 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 shortests 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 (that has incoming degree zero in T), and the shortest paths of u to all other vertices is contained in T.

Now, we call a vertex x half-complete if and only 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 for $1+\log_{k/2} n$ 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 $e_1, e_2, \cdots, e_{k/2}$ 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 paths, these sets are disjoint. So one of these sets like S_i has size at most $n/(k/2) = 2n/k$. 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 cost 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. Remember 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 $|S_i| \cdot (d+1) \leq n(d+1)/k$.

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 which is in a cycle of length at most l'. Let u be the receptive node in lemma 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 $5l'/\alpha - (2\log_{k'} n + 2) \geq 3l'/\alpha$ (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 $(\alpha n/2) \cdot (3l'/\alpha) = 3l'n/2$ because there are at least $\frac{\alpha n}{2}$ 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 in 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 $O(l'/\alpha) = O((l+\log_k(n))/\alpha)$, 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 $O(\log_{k'}(n)/\alpha)$.

3 How to deal with unknown α and β

In Section 2, we presented an advertising strategy that lead the network to some equilibria with small price of anarchy given two parameters α and β. Here we try to make our strategy adaptive for the cases that the parameters are not known in advance, i.e. some times a lot of agents contribute in the campaign, and sometimes a small fraction of them participate. So in these cases, we know that $\alpha > \epsilon$ 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 exists 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 estimation 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 estimations of these two parameters. The only thing we know is that they are in range $[\epsilon, 1]$ and $[\epsilon', 1]$ respectively.

But we know that α is in one of these m ranges: $[1/2, 1]$, $[1/4, 1/2]$, $\cdots$, $[1/2^m, 1/2^{m-1}]$, and the same for β. We should run the strategy for different estimations of α and β in a parallel manner. So there are $m \times m'$ different pairs of estimations for our parameters. But a receptive agent contributes in the campaign with only βk edges. 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 nodes 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 that parameters are not known in advance.

Theorem 2. *When the parameters $\alpha > \epsilon$ and $\beta > \epsilon'$ are not known in advance, the price of anarchy is at most $O(\frac{\log_{k'}(n)}{\alpha \log_k(n)}) = O(\frac{\log_{k'}(k)}{\alpha})$ 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 of (i, j), one of these pairs is a good estimation for α and β. Using the constructed edges by the receptive nodes in this specific run of the strategy and Theorem 1, we can have this bound. The only different thing 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 $m \times m'$.

Since ϵ and ϵ' are two constant (and probably very small) constants, we can say that m and m' are also some constant (and probably large) numbers. We conclude that the Corollaries 1 and 2 are also true in this case (unknown α and β).

References

1. Albers, S., Eilts, S., Even-Dar, E., Mansour, Y., and Roditty, L. *On Nash Equilibria for a Network Creation Game*. In Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms. Miami, FL, 89-98, 2006.
2. 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, 728-737, 2009.
3. Corbo, J. and Parkes, D. *The price of selish behavior in bilateral network formation*. In Proceedings of the 24th Annual ACM Symposium on Principles of Distributed Computing. Las Vegas, Nevada, 99-107, 2005.
4. Erik D. Demaine, MohammadTaghi Hajiaghayi, Hamid Mahini, and Morteza Zadimoghaddam. *The Price of Anarchy in Cooperative Network Creation Games*. Appeared 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, pages 171-182.
5. 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. 292-298, 2007. To appear in ACM Transactions on Algorithms.
6. Fabrikant, A., Luthra, A., Maneva, E., Papadimitriou, C. H., and Shenker, S. *On a network creation game*. In Proceedings of the 22nd Annual Symposium on Principles of Distributed Computing. Boston, Massachusetts, 347-351.
7. Jon Kleinberg. *Small-World Phenomena and the Dynamics of Information*. Advances in Neural Information Processing Systems (NIPS) 14, 2001.
8. Jon Kleinberg. *The small-world phenomenon: An algorithmic perspective*. In Proceedings of the 32nd ACM Symposium on Theory of Computing, 2000.
9. Koutsoupias, E. and Papadimitriou, C. *Worst-case equilibria*. In Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science. Lecture Notes in Computer Science, vol. 1563. Trier, Germany, 404-413.
10. Laoutaris, N., Poplawski, L. J., Rajaraman, R., Sundaram, R., and Teng, S.-H. *Bounded budget connection (BBC) games or how to make friends and influence people, on a budget*. In Proceedings of the 27th ACM Symposium on Principles of Distributed Computing. 165-174, 2008.
11. Papadimitriou, C. *Algorithms, games, and the internet*. In Proceedings of the 33rd Annual ACM Symposium on Theory of Computing, Hersonissos, Greece, 749-753.