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. 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) would give us 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 behaviour of social networks, we try to understand how large the social cost can be in presence of selfish agents. Nash Equilibria are the stable networks in which every agent is acting selfishly. More formally, in a Nash Equilibria 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. The price of anarchy is introduced by Koutsoupias and Papadimitriou in [9, 11], and is used to measure the behaviour 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 behaviour 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. 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.

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 k edges to other nodes. So every node only have the usage cost, but its budget to build edges is limited.

The advertising campaign scenario can be applied to different game theoretic situations. In this 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 behaviour 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. 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

[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.

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 sub-linear both 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 reduces 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 At first we prove a tight upper bound in the uniform bounded budget games. We prove that the price of anarchy is $O(\sqrt{\frac{n/k}{\log_k(n)}})$. According to the lower bound in [10], this is a tight result, and shows that the price of anarchy is in fact $\Theta(\sqrt{\frac{n/k}{\log_k(n)}})$.

Since the uniform games have a very large price of anarchy, 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 3, we present an advertising strategy that knows the values of α and β in advance. Then in Section 4, we adapt our strategy to work in cases that these two parameters are not given in the input, and we should find out about their values as well.

2 A Tight Upper Bound for Price of Anarchy in Uniform Games

Here we show that the price of anarchy can not be more than $O(\sqrt{\frac{n/k}{\log_k(n)}})$ in the (n,k)-uniform BBC game[2]. According to the $\Omega(\sqrt{\frac{n/k}{\log_k(n)}})$ lower bound for the price of anarchy presented in [10], this is the best upper bound that can be achieved in this game. This means that setting any limit on the budget of each node implies very large values of price of anarchy.

We prove that the diameter of any stable graph in this model is bounded by $O(\sqrt{n\log_k(n)/k})$.

Lemma 1. *The diameter of any stable graph in (n,k)-uniform game is at most $O(\sqrt{n\log_k(n)/k})$.*

Proof. We just need to show that there is a vertex v whose distance to any other vertex is at most $O(\sqrt{n\log_k(n)/k})$. Let G be a stable graph. Define g to be $c\log_k(n)$ for a sufficiently large constant $c \geq 1$. Delete all edges that are contained in a cycle of length at most g. Let G' be the remaining graph. Clearly G' has no cycle of length at most g. We claim that G' has a vertex with degree at most $k/2$. If all degrees are at least $k/2$, we have at least $(k/2)^{g/2}$ walks with length $g/2$ starting from an arbitrary vertex u. The endpoint of these walks are different. Otherwise we would find two walks with length $g/2$ starting from u and ending at the same vertex, we can also find a cycle of length at most $2(g/2) = g$ in the remaining graph which is a contradiction. So there are at least $(k/2)^{g/2}$ different endpoints for these walks. On the other hand there are

[2] Bounded Budget Connection game

at most n vertices in the graph. For any k, there is a sufficiently large value of c for which $(k/2)^{g/2}$ is greater than n which is a contradiction.

So there is a vertex v with degree at most $k/2$ in G'. This means that v has at least $k/2$ edges like $e_1, e_2, \cdots, e_{k/2}$ each of which is contained in a cycle of length at most g. For each vertex $u \neq v$ in G consider a shortest path from v to u in graph G. So we have $n-1$ shortest paths, and each of these paths might use at most one of these $k/2$ edges. So there is an edge e_i that is used in at most $\frac{n}{k/2} = 2n/k$ paths. If vertex v deletes edge e_i, its distance to at most $2n/k$ other vertices might increase by at most $g-1$ because edge e_i is contained in a cycle of length at most g. So the cost of v is increased by at most $2ng/k$.

Now let d be the maximum distance from v in the stable graph G. Assume that d is the distance between v and v'. If v deletes edge e_i, and adds edge (v, v') its cost increases by at most $2ng/k$, and decreases by at least $(d/3)^2 = d^2/9$. To see this, we just need to consider the shortest path from v to v', there are at least $d/3$ vertices whose distance to v is at least $2d/3$ now, and by adding edge (v, v') their distances to v would be at most $d/3$. So the cost of v decreases by at least $d^2/9$ by adding edge (v, v'). Since we are in a stable graph, $2ng/k$ should be at least $d^2/9$. This means that d is at most $O(\sqrt{ng/k})$. Note that g is $c \log_k(n)$, and this completes the proof.

Theorem 1. *The price of anarchy in a (n,k)-uniform BBC game is at most $O(\sqrt{\frac{n/k}{\log_k(n)}})$.*

Proof. Using Lemma 1, we know that the diameter of any stable graph in this model is bounded by $O(\sqrt{n \log_k(n)/k})$. As mentioned in the proof of Theorem 3 in [10], the average distance in the optimum solution is at least $\Omega(\log_k(n))$. This shows that the price of anarchy in these uniform games is not more than $\frac{O(\sqrt{n \log_k(n)/k})}{\Omega(\log_k(n))} \leq O(\sqrt{\frac{n/k}{\log_k(n)}})$.

3 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. 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 4, we change our strategies to be adaptive and work when these parameters are not revealed in advance.

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 $k' > 1$. We partition the nodes into $l \leq \log_{k'}(n)$ sets $S_1, S_2, \cdots, S_l$ such that $|S_1| = \beta k/2$, 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 $\beta k/2 - 1$ 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 $c \log(n)/2\alpha$ nodes randomly from set S_{i-1} and construct edges to them. Note that $c \log(n)/2\alpha$ is at most $\beta k/2$ because k' is greater than one, and it is also equal to $\alpha\beta k/c \log(n)$.

On the other hand nodes in set S_{i-1} receive some incoming edges from set S_i. We do not assume that every such an edge is accepted by nodes in set S_{i-1}. For example if a non-receptive node receives an edge, the node might delete the edge, i.e. the node is not interested in our strategy or it is a malicious player. This assumption just makes our work harder because we have to find a way to take care of deleted edges.

Even if the node in set S_{i-1} is receptive we might have a problem. Assume that the node receives more than $\beta k/2$ edges from set S_i, it might delete some edges. Because a receptive node is not necessarily willing to contribute in the strategy with more than βk edges. So the node might get overwhelmed by the nodes in lower set. In these cases we just do not rely on these edges in our analysis. So we assume that if a receptive node receives at most $\beta k/2$ edges from the nodes of the lower set, it does not delete these edges. This assumption is true because we are basically asking a receptive node in set S_i to handle at most $\beta k/2$ edges from the set S_{i+1}, and build $c \log(n)/2\alpha \leq \beta k/2$ edges to the nodes of set S_{i-1} which is at most βk edges in total.

Lemma 2. *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 $2\log_{k'} n$, and every receptive node is contained in this subgraph with high probability*[3].

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 $c \log(n)/2\alpha$ random nodes in set S_{i-1}. There are $c \log(n)/2$ 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 $\log(n)$ 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 $c \log(n)/2\alpha$ 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 $\frac{\alpha |S_i| (c \log(n)/2\alpha)}{|S_{i-1}|}$. 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

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

u in set S_{i-1} receives $\alpha\beta k/2$ edges in expectation. Using Markov inequality, we can say that a node can be overwhelmed with probability at most $\alpha/2 < 1/2$.

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 $2l = 2\log_{k'}(n)$ with high probability.

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 2, we know that with high probability the diameter of receptive nodes is at most $2l$. There are αn receptive nodes in expectation, and with high probability the number of them is not less than $\alpha n/2$.

Consider a receptive vertex v. Let d be the maximum distance of other nodes from v. We prove that d is $O(l + \log(n)/\alpha)$.

Delete all edges in G that are contained in at least a cycle of length at most $l' = l + 2\log_k(n) + 1$. Consider a non-receptive vertex u. We prove that if one of the k edges of u is in a cycle of length at most l', the distance from u to v is at most l'/α. Let e be an edge owned by u which is in a cycle of length at most l'. Let x be the distance between u and v. If vertex u deletes edge e, its distance to other nodes increases by at most $l' \times n$. On the other hand, if u makes an edge to vertex v, its distance to all receptive nodes decreases by at least $x - 4l - 1$ (before adding the edge its distances to receptive nodes were at least $x - 2l$, and after that the distances are at most $2l + 1$). So the total decrease in the cost of u would be at least $\frac{\alpha n}{2}(x - 4l' - 1)$ because there are at least $\frac{\alpha n}{2}$ receptive nodes with high probability. Since we are in a stable graph, $\frac{\alpha n}{2}(x - 4l - 1)$ should not be greater than $l' \times n$. So x is $O(l'/\alpha + l) = O(l'/\alpha)$ in this case.

We call a vertex incomplete if at least one of its edges is deleted. As proved above, each incomplete vertex is in distance at most $O(l'/\alpha)$ from v. 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 an incomplete vertex u, and all walks of length $l'/2$ starting from u 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 the distances of all vertices from a receptive vertex v are at most $O(l'/\alpha) = O((l + \log_k(n))/\alpha)$. Note that 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 2. *The price of anarchy is at most* $O(\frac{\log_{k'}(n)}{\alpha \log_k(n)}) = O(\frac{\log_{k'}(k)}{\alpha})$ *using the 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 $O(\log_{k'}(n)/\alpha)$. On the other hand as mentioned in proof of Theorem 3 in [10], 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.

4 How to deal with unknown α and β

In Section 3, 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 3. 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 3. *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 2, 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.