The ring design game with fair cost allocation [Extended Abstract]*

Angelo Fanelli[1] Dariusz Leniowski[2] Gianpiero Monaco[3] Piotr Sankowski[4]

[1] Division of Mathematical Sciences, School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore. Email: angelo.fanelli@ntu.edu.sg
[2] Institute of Informatics, University of Warsaw. Email: dariusz.leniowski@mimuw.edu.pl
[3] Dipartimento di Ingegneria e Scienze dell'Informazione e Matematica, University of L'Aquila and Institute of Informatics, University of Warsaw. Email: gianpiero.monaco@univaq.it
[4] Institute of Informatics, University of Warsaw and Department of Computer and System Science, Sapienza University of Rome. Email: sank@mimuw.edu.pl

Abstract. In this paper we study the network design game when the underlying network is a ring. In a network design game we have a set of players, each of them aims at connecting nodes in a network by installing links and sharing the cost of the installation equally with other users. The ring design game is the special case in which the potential links of the network form a ring. It is well known that in a ring design game the price of anarchy may be as large as the number of players. Our aim is to show that, despite the worst case, the ring design game always possesses good equilibria. In particular, we prove that the price of stability of the ring design game is at most 3/2, and such bound is tight. We believe that our results might be useful for the analysis of more involved topologies of graphs, e.g., planar graphs.

1 Introduction

In a network design game, we are given an undirected graph $G = (V, E)$ and edge costs given by a function $c : E \to \mathbb{R}^+$. The edge cost function naturally extends to any subset of edges, that is $c(B) = \sum_{e \in B} c(e)$ for any $B \subseteq E$. We define $c(\emptyset) = 0$. There is a set of n players $[n] = \{1, \ldots, n\}$; each player $i \in [n]$ wishes to establish a connection between two nodes $s_i, t_i \in V$ called the source and destination node of player i, respectively. The set of strategies available to player i consists of all paths connecting nodes s_i and t_i in G. We call a *state* of the game a set of strategies $\sigma \in \Sigma$ (where Σ is the set of all the states of the game), with one strategy per player, i.e., $\sigma = (\sigma_1, \ldots, \sigma_n)$ where σ_i denotes the strategy of player i in σ . Given a state σ, let $n_\sigma(e)$ be the

* This work was partially supported by the ERC StG project PAAl (Practical Approximation Algorithms) no. 259515, by the grant NRF-RF2009-08 "Algorithmic aspects of coalitional games" and by PRIN 2008 research project COGENT (COmputational and GamE-theoretic aspects of uncoordinated NeTworks).

number of players using edge e in σ. Then, the cost of player i in σ is defined as $c_\sigma(i) = \sum_{e\in\sigma_i} \frac{c(e)}{n_\sigma(e)}$. Let $E(\sigma)$ be the set of edges that are used by at least one player in state σ. The social cost $C(\sigma)$ is simply the total cost of the edges used in state σ which coincides with the sum of the costs of the players, i.e., $C(\sigma) = \sum_{e\in E(\sigma)} c(e) = \sum_{i\in[n]} c_\sigma(i) = c(E(\sigma))$.

Let (σ_{-i}, σ'_i) denote the state obtained from σ by changing the strategy of player i from σ_i to σ'_i. Given a state $\sigma = (\sigma_1, \dots, \sigma_n)$, an *improving move* of player i in σ is a strategy σ'_i such that $c_{(\sigma_{-i},\sigma'_i)}(i) < c_\sigma(i)$. A state of the game is a *Nash equilibrium* if and only if no player can perform any improving move. An *improvement dynamics* (shortly *dynamics*) is a sequence of improving moves. A game is said to be *convergent* if, given any initial state σ, any dynamics leads to a Nash equilibrium. It is well known, as it has been proved by Rosenthal [6] for the more general class of congestion games, that any network design game is convergent. We denote by NE the set of states that are Nash equilibria. A Nash equilibrium can be different from the socially optimal solution. Let $\textsc{Opt}$ be a state of the game minimizing the social cost. The *price of anarchy* (PoA) of a network design game is defined as the ratio of the maximum social cost among all Nash equilibria over the optimal cost, i.e., $\text{PoA} = \frac{\max_{\sigma\in\text{NE}} C(\sigma)}{C(\textsc{Opt})}$. It is trivial to observe that the PoA in a network design game may be as large as the number of players n, and such bound is tight. The *price of stability* (PoS) is defined as the ratio of the minimum social cost among all Nash equilibria over the optimal cost, i.e., $\text{PoS} = \frac{\min_{\sigma\in\text{NE}} C(\sigma)}{C(\textsc{Opt})}$. Anshelevich *et al.* [1] proved that the price of stability is at most $H_n = 1 + 1/2 + \ldots + 1/n$. Although the upper bound proof has been shown to be tight for directed networks, the problem is still open for undirected networks. There have been several attempts to give a significant lower bound for the undirected case, e.g., [5, 4, 2, 3]. The best known lower bound so far of $348/155 \approx 2.245$, has been recently shown in [2].

The aim of the current paper is to analyze the network design game when the underlying graph is a ring. We refer to this special case as *ring design game.* For the sake of clarity, by a ring we mean an undirected graph $G = (V, E)$ where $V = \{v_1, v_2, \dots, v_k\}$, $E = \{e_1, e_2, \dots, e_k\}$, and $e_i = v_i v_{i+1}$, $i = 1, \dots, k$ (where $v_{k+1} = v_1$). Note that this simple case captures the whole spectra of interesting behavior, i.e., PoA remains equal the number of players. Moreover, the ring is crucial in the sense that it is the first non-trivial topology to analyze in the context of network design and it is the first step in order to cope with more involved topologies, like planar graphs. Hence, we believe that giving a tight bounds here could give some insight for studying more general settings.

Let us first point out that, in a ring design game, any improvement dynamics starting from the optimal state leads to an equilibrium at most 2 times the cost of the optimal state. In fact, either the optimal state is a Nash equilibrium, or there is a player j wishing to switch from his optimal strategy to the alternative path. At the optimum, the cost of player j is at most $C(\textsc{Opt})$, and thus the cost of the alternative path cannot be more than this quantity. Since the alternative path of j contains edges of the ring not belonging to $E(\textsc{Opt})$, it implies that $C(\textsc{Opt})$ is also an upper bound to $c(E \setminus E(\textsc{Opt}))$. Consequently, the cost of the entire ring, and thus the cost of any state, is at most $2C(\textsc{Opt})$. As we show here, by doing a more careful analysis, we are actually able to prove a tight bound of $3/2 \cdot C(\textsc{Opt})$ on the cost of any equilibrium reachable from the optimum.

Our results. In this paper we show that in a ring design game, differently from what the classical bound of n on the price of anarchy suggests, there always exist good performing Nash equilibria. In particular, we show that there always exists a Nash equilibrium of cost at most $3/2$ times the cost of an optimal state, thus giving a bound on the PoS. We show that such equilibrium can be reached by a dynamics having as initial state an optimal configuration. Such result also gives some insight on the problem of computing an equilibrium in a ring design game. In fact, it reveals that if the cost of the entire ring is larger than $3/2$ times the cost of an optimal state, then the dynamics starting from an optimal state converges quickly, within at most 3 steps, to an equilibrium. We also show that such bound on the PoS is tight, by showing an instance for which $\text{PoS} = 3/2 - \epsilon$.

2 Upper and lower bounds on the price of stability

We start by upper bounding the price of stability. Our technique to prove the bound on the PoS is different from the previously used ones. Previous techniques used potential function arguments and proved that any equilibrium reached by any dynamics starting by an optimal state has potential value at most $H_n \cdot C(\textsc{Opt})$. Here we also bound the cost of a Nash equilibrium reachable by a dynamics from the optimal state but without using potential function arguments. In particular, the analysis is made by cases on the number of moves, and for each such case we write a linear program that captures the most important inequalities. The most important observation we use is that one needs to consider the cases when at most 4 players move. We prove that for higher number of moves the PoS can only be smaller.

Our notation includes the number m representing the amount of steps in which some fixed dynamics reaches a Nash equilibrium starting from an optimal state $\textsc{Opt}$, the Nash equilibrium N obtained after m steps, as well as players making a move in the dynamics, meaning that π_j denotes the player that made the move at step $j = 1, \ldots, m$ during the dynamics. Note that a player could make a move at many different steps of the dynamics. Let $\sigma^0, \ldots, \sigma^j, \ldots, \sigma^m$ be the states corresponding to the considered dynamics, where $\sigma^0 = \textsc{Opt}$ and $\sigma^m = \text{N}$. Also, let f be a set of players of interest. The set f will be composed by a subset of the players moving in the dynamics. The usage of f will be clear in the proof of Theorem 1. For any $A \subseteq f$ the set D_A will denote the edges in $\textsc{Opt}$ which are used by exactly the players in A, and R_A will denote the edges used in $\textsc{Opt}$ which are used by exactly the players in A and at least one player from outside of f, formally:

$$D_A^f = \{e \in E \mid (\forall i \in f.\ e \in \textsc{Opt}_i \iff i \in A) \wedge \neg \exists i \notin f.\ e \in \textsc{Opt}_i\},$$
$$R_A^f = \{e \in E \mid (\forall i \in f. e \in \textsc{Opt}_i \iff i \in A) \wedge \exists i \notin f.\ e \in \textsc{Opt}_i\}.$$

For the sake of simplicity in the sequel we will omit the superscript f when it is clear from the context. Notice that D_A andR_A naturally define a partition of the edges of the ring, and that for any f we have $E(\text{OPT}) = \Big(\bigcup_{A \subseteq f} D_A \cup R_A\Big) \backslash D_\emptyset$. Moreover, let $\lambda > 0$ be such that $c\big(D_\varnothing\big) \leq \lambda C(\text{OPT})$. Since OPT and $D_\emptyset$ is a partition of E and the cost of any equilibrium N can be at most $c(E)$, then:

$$PoS \leq \frac{C(\text{N})}{C(\text{OPT})} \leq \frac{C(\text{OPT}) + c\big(D_\varnothing\big)}{C(\text{OPT})} \leq \frac{C(\text{OPT}) + \lambda \cdot C(\text{OPT})}{C(\text{OPT})} \leq 1 + \lambda. \quad (1)$$

Now let us write the necessary conditions for the fact that player π_j can move in step j of the dynamics, for any $j = 1, \ldots, m$. Such conditions will be expressed by using the above defined variables D_A and R_A. Unfortunately, we do not know the exact usage of edges in sets R_A. Let us define functions $\text{left}_k, \text{right}_k : \Sigma \to \mathbb{R}$ for any players $k \in f$. Set $d_\sigma(k)$ and $r_\sigma(k)$ to be the collection of subsets A of f such that player k is using (all) edges of D_A and R_A respectively in state σ, i.e., $d_\sigma(k) = \{A \in 2^f \mid k \text{ is using edges of } D_A \text{ in } \sigma\}$, $r_\sigma(k) = \{A \in 2^f \mid k \text{ is using edges of } R_A \text{ in } \sigma\}$. Also, define the edges' usage by the players' of interest (i.e., players belonging to f) $\hat{n}_\sigma : 2^E \to \mathbb{N}$ as $\hat{n}_\sigma(H) = \#\{i \in f \mid H \subseteq \sigma_i\}$. Let us define:

$$\text{left}_\sigma(k) = \sum_{A \in d_\sigma(k)} \sum_{e \in D_A} \frac{c(e)}{n_\sigma(e)} + \sum_{A \in r_\sigma(k)} \sum_{e \in R_A} \frac{c(e)}{n} = \sum_{A \in d_\sigma(k)} \frac{c(D_A)}{\hat{n}_\sigma(D_A)} + \sum_{A \in r_\sigma(k)} \frac{c(R_A)}{n}.$$

In the following the function left will be used as a lower bound on the cost of a player. Then wlog we can consider $\frac{c(R_A)}{n}$ to be 0 for any R_A. Therefore in the following we will omit such terms. Moreover, let us define:

$$\begin{aligned}\text{right}_\sigma(k) &= \sum_{A \in d_\sigma(k)} \sum_{e \in D_A} \frac{c(e)}{n_\sigma(e)} + \sum_{A \in r_\sigma(k)} \sum_{e \in R_A} \frac{c(e)}{\hat{n}_\sigma(R_A) + 1} \\ &= \sum_{A \in d_\sigma(k)} \frac{c(D_A)}{\hat{n}_\sigma(D_A)} + \sum_{A \in r_\sigma(k)} \frac{c(R_A)}{\hat{n}_\sigma(R_A) + 1}.\end{aligned}$$

Then the following inequalities hold for any state $\sigma \in \Sigma$:

$$\text{left}_\sigma(k) \leq c_\sigma(k) \leq \text{right}_\sigma(k).$$

The role of functions left_k and right_k is to weaken the inequalities between player's utilities in some neighbour states, so that they become manageable. As we do not know the exact usage of edges in sets R_A, it would be hard to derive the precise bounds. This means that on the lower-hand side we introduce the maximum possible number (i.e., n) of players using edges of sets R_A in σ and on the upper-hand side we introduce the minimum number of players using edges of R_A in σ, i.e., $\hat{n}_\sigma(R_A) + 1$.

The proof of the following lemma will be given in the full version of this paper. This lemma will become useful in the proof of the main theorem.

Lemma 1. *In the ring design game, if in state* Opt *there are at least two players able to perform an improving move (both starting from state* Opt*) then the cost of the whole ring is at most $\frac{3}{2}$ times the cost of an optimal solution, that is* $c(E) \leq \frac{3}{2}C(\text{Opt})$.

Theorem 1. *The price of stability for the ring design game is at most $\frac{3}{2}$.*

Proof. The proof is split into five different cases, depending on the amount m of steps in which some fixed dynamics reaches a Nash equilibrium starting from an optimal state Opt. Moreover notice that since in a ring design game the strategy set of each player i is composed by exactly 2 different strategies, i.e., the clockwise and counterclockwise paths connecting s_i and t_i. This implies that $\pi_j \neq \pi_{j+1}$, for any $j = 1, \ldots, m-1$. We remark that in some cases we get the bound by solving a linear program where constraints are naturally defined by using left and right functions, and where objective functions are proper defined in each of the case.

Case $m = 0$. The equality $m = 0$ trivializes the instance into an example where Opt is a Nash equilibrium, thus PoS $= 1$.

Case $m = 1$. In this case the dynamics reaches a Nash Equilibrium N after one step starting from Opt. Since player π_1 can perform an improving move starting by state Opt, the following inequalities hold: $\text{left}_{\text{N}}(\pi_1) \leq c_{\text{N}}(\pi_1) < c_{\text{OPT}}(\pi_1) \leq \text{right}_{\text{OPT}}(\pi_1)$. Therefore, by setting $f = \{\pi_1\}$ we have that: $\frac{c\left(D_\varnothing\right)}{1} < \frac{c\left(D_{\{\pi_1\}}\right)}{1} + \frac{c\left(R_{\{\pi_1\}}\right)}{2}$. The last inequality directly implies that:

$$\frac{C(\text{N})}{C(\text{OPT})} = \frac{c(D_\varnothing) + c(R_\varnothing) + c(R_{\{\pi_1\}})}{c(D_{\{\pi_1\}}) + c(R_\varnothing) + c(R_{\{\pi_1\}})} \leq \frac{c(D_{\{\pi_1\}}) + c(R_\varnothing) + \frac{3}{2}c(R_{\{\pi_1\}})}{c(D_{\{\pi_1\}}) + c\big(R_\varnothing\big) + c\big(R_{\{\pi_1\}}\big)} \leq \frac{3}{2}.$$

Case $m = 2$. When $m = 2$, the player π_1 leads the dynamic from Opt to σ^1 and player π_2 leads the dynamics from σ^1 to N. Therefore the following must hold:

$$\text{left}_{\sigma^1}(\pi_1) \leq c_{\sigma^1}(\pi_1) < c_{\text{OPT}}(\pi_1) \leq \text{right}_{\text{OPT}}(\pi_1),$$
$$\text{left}_{\text{N}}(\pi_2) \leq c_{\text{N}}(\pi_2) < c_{\sigma^1}(\pi_2) \leq \text{right}_{\sigma^1}(\pi_2).$$

Therefore, by setting $f = \{\pi_1, \pi_2\}$ we have that:

$$\frac{c\big(D_\varnothing\big)}{1} + \frac{c\big(D_{\{\pi_2\}}\big)}{2} < \frac{c\big(D_{\{\pi_1\}}\big)}{1} + \frac{c\big(R_{\{\pi_1\}}\big)}{2} + \frac{c\big(D_{\{\pi_1,\pi_2\}}\big)}{2} + \frac{c\big(R_{\{\pi_1,\pi_2\}}\big)}{3}$$

$$\frac{c\big(D_\varnothing\big)}{2} + \frac{c\big(D_{\{\pi_1\}}\big)}{1} < \frac{c\big(D_{\{\pi_2\}}\big)}{2} + \frac{c\big(R_{\{\pi_2\}}\big)}{2} + \frac{c\big(D_{\{\pi_1,\pi_2\}}\big)}{1} + \frac{c\big(R_{\{\pi_1,\pi_2\}}\big)}{2}.$$

Without loss of generality we can add the following constraints:

$$\sum_{e \in OPT} c(e) \leq 1, \qquad \forall e \in E.\ c(e) \geq 0.$$

We need to bound the value of $c\big(D_\varnothing\big) - \frac{1}{2}c\big(D_{\{\pi_1,\pi_2\}}\big)$ with respect to the above inequalities. Such a bound can be obtained by forming a linear program from all the above equations including the appropriate objective function. We have solved this linear program on a computer using a standard LP solver. This way we have obtained the following bound: $c\big(D_\varnothing\big) - \frac{1}{2}c\big(D_{\{\pi_1,\pi_2\}}\big) \leq \frac{5}{11} < \frac{1}{2}$.

In the remainder of the proof similar bounds have been obtained in the same way by using a LP solver. Further, the cost of states N and OPT are:

$$\begin{aligned} C(\mathrm{N}) = c\big(D_\varnothing\big) + c\big(R_\varnothing\big) + c\big(D_{\{\pi_1\}}\big) + c\big(R_{\{\pi_1\}}\big) \\ + c\big(D_{\{\pi_2\}}\big) + c\big(R_{\{\pi_2\}}\big) + c\big(R_{\{\pi_1,\pi_2\}}\big), \end{aligned}$$

and

$$\begin{aligned} C(\textsc{Opt}) = c\big(D_{\{\pi_1,\pi_2\}}\big) + c\big(R_\varnothing\big) + c\big(D_{\{\pi_1\}}\big) + c\big(R_{\{\pi_1\}}\big) \\ + c\big(D_{\{\pi_2\}}\big) + c\big(R_{\{\pi_2\}}\big) + c\big(R_{\{\pi_1,\pi_2\}}\big), \end{aligned}$$

respectively. Therefore, by using the upper bound on $c\big(D_\varnothing\big) - \frac{1}{2}c\big(D_{\{\pi_1,\pi_2\}}\big)$ we obtain that: $\frac{C(\mathrm{N})}{C(\textsc{Opt})} \leq \frac{16}{11} < \frac{3}{2}$.

Case $m = 3$. Similarly to the previous case, we will construct a suitable linear program. We know that $\pi_1 \neq \pi_2$ and $\pi_2 \neq \pi_3$. If $\pi_1 = \pi_3$ then by Lemma 1 we have that $PoS \leq \frac{3}{2}$. Therefore we can assume that $\pi_1 \neq \pi_3$. The following must hold:

$$\begin{aligned} \mathrm{left}_{\sigma^1}(\pi_1) \leq c_{\sigma^1}(\pi_1) < c_{\textsc{Opt}}(\pi_1) \leq \mathrm{right}_{\textsc{Opt}}(\pi_1), \\ \mathrm{left}_{\sigma^2}(\pi_2) \leq c_{\sigma^2}(\pi_2) < c_{\sigma^1}(\pi_2) \leq \mathrm{right}_{\sigma^1}(\pi_2), \\ \mathrm{left}_{\mathrm{N}}(\pi_3) \leq c_{\mathrm{N}}(\pi_3) < c_{\sigma^2}(\pi_3) \leq \mathrm{right}_{\sigma^2}(\pi_3). \end{aligned}$$

By setting $f = \{\pi_1, \pi_2, \pi_3\}$ we obtain a set of constraints that along with $C(\textsc{Opt}) \leq 1$ and maximization target $c\big(D_\varnothing\big) - \frac{1}{2}c\big(D_{\{\pi_1,\pi_2,\pi_3\}}\big)$ constitute a linear program with a solution $c\big(D_\varnothing\big) - \frac{1}{2}c\big(D_{\{\pi_1,\pi_2,\pi_3\}}\big) \leq \frac{198}{487} < \frac{1}{2}$. Substituting it into the ratio of the costs of N and OPT we get that: $\frac{C(\mathrm{N})}{C(\textsc{Opt})} \leq \frac{685}{487} < \frac{3}{2}$.

Case $m \geq 4$. Here it is enough to consider the case $m = 4$. This is due to the fact that the inequalities obtained by the dynamics of the first 4 players are strong enough to bound the cost of the whole ring. This gives the bound on the cost of any Nash equilibrium the dynamics will converge to, because even if more players move the cost of the final state will be smaller than the cost of the whole ring. We show that if $m = 4$ then $c\big(D_\varnothing\big) < \frac{1}{2} \cdot C(\textsc{Opt})$. Clearly, adding new constraints for $m > 4$ cannot increase this bound. Then let us consider $m = 4$. As in the previous case we have that $\pi_1 \neq \pi_2$ and $\pi_2 \neq \pi_3$ and $\pi_1 \neq \pi_3$, anyway we are not able to derive any conclusion about π_4. It follows that we have to consider 3 subcases, i.e., $\pi_4 = \pi_1$, $\pi_4 = \pi_2$ and $\pi_4 \neq \pi_z$ for $z = 1, 2, 3$. As usually in this proof, we are going to derive sets of constraints that must hold at every step of the dynamics by using functions left and right.

By summarizing we get three different sets of constraints corresponding to three different linear programs. In each of them it suffices to consider maximization target $c\big(D_{\varnothing}\big)$ assuming that $C(\textsc{Opt}) \leq 1$ wlog. In all cases the maximum value of $c\big(D_{\varnothing}\big)$ turns out to be smaller than $\frac{1}{2}$. Hence, by (1) for all these cases we know that PoS is bounded by $\frac{3}{2}$. □

Corollary 1. *In a ring design game, if the cost of the entire ring is larger than* $3/2$ *times the cost of an optimal state, then the improvement dynamics starting from an optimal state converges quickly, within at most 3 steps, to a Nash equilibrium.*

The following theorem will be given in the full version of this paper and constructs an example (Figure 1) when the above upper bound is reached.

Theorem 2. *Given any* $\epsilon > 0$*, there exists an instance of the ring design game such that the price of stability is at least* $\frac{3}{2} - \epsilon$*.*

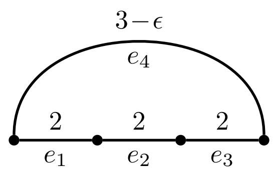

Fig. 1. The lower bound example for PoS on the ring. There is a player associated with each edge. The optimum uses three edges of cost 2 whereas the only Nash equilibrium uses the whole ring.

References

1. E. Anshelevich, A. Dasgupta, J. M. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4), pp. 1602-1623, 2008.
2. V. Bilò, I. Caragiannis, A. Fanelli and G. Monaco. Improved Lower Bounds on the Price of Stability of Undirected Network Design Game. *Theory of Computing Systems*, To appear. DOI: 10.1007/s00224-012-9411-6.
3. H.-L. Chen and T. Roughgarden. Network design with weighted players. *Theory of Computing Systems*, 45, pp. 302-324, 2009.
4. G. Christodoulou, C. Chung, K. Ligett, E. Pyrga, and R. van Stee. On the price of stability for undirected network design. In *(WAOA 2009)*, LNCS 5893, Springer, pp. 86-97.
5. A. Fiat, H. Kaplan, M. Levy, S. Olonetsky, and R. Shabo. On the price of stability for designing undirected networks with fair cost allocations. In *(ICALP 2006)*, LNCS 4051, Springer, Part I, pp. 608-618.
6. R. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2, pp. 65-67, 1973.