Information Processing Letters ••• (••••) •••–•••

Contents lists available at ScienceDirect

Information Processing Letters

www.elsevier.com/locate/ipl

An $O(\frac{\log n}{\log\log n})$ upper bound on the price of stability for undirected Shapley network design games

Jian Li [1]

University of Maryland, Department of Computer Science, A.V. Williams Building, College park, MD, United States

ARTICLE INFO

Article history:
Received 27 January 2009
Received in revised form 14 March 2009
Accepted 22 April 2009
Available online xxxx
Communicated by F.Y.L. Chin

Keywords:
Algorithms
Price of stability
Shapley network design game

ABSTRACT

In this paper, we consider the Shapley network design game on undirected networks. In this game, we have an edge weighted undirected network $G(V,E)$ and n selfish players where player i wants to choose a low cost path from source vertex s_i to destination vertex t_i. The cost of each edge is equally split among players who pass it. The price of stability is defined as the ratio of the cost of the best Nash equilibrium to that of the optimal solution. We present an $O(\log n/\log\log n)$ upper bound on price of stability for the single sink case, i.e., $t_i = t$ for all i.

1. Introduction

We consider the Shapley network design game, which is also called network design games with fair cost allocation, introduced in [2]. In this game, we are given a network and n selfish players, where player i wants to go from source vertex s_i to destination vertex t_i. The cost of each edge is shared in a fair manner among players who pass through it. Each player tries to minimize the cost of the path it chooses. We are interested in stable status of the network where no player has an incentive to deviate from its current strategy, which can be modeled by Nash equilibria. The price of stability, defined as the ratio of the cost of the best Nash equilibrium and that of an optimal solution, is used to measure the inefficiency of Nash equilibria. We imagine a network where the traffic will be initially designed by a central network coordinator. However, the coordinator is unable prevent the network users from selfishly deviating from the designated paths. Therefore, in this scenario, the best Nash equilibrium is an obvious solution to propose. In this sense, we can think the price of stability as the degree of degradation of the solution quality for the outcome being stable.

The price of stability was first studied in Schulzan and Moses [1] and was so-called in Anshelevich et al. [2] where the Shapley network design game was also first explored. They showed that a pure-strategy Nash equilibrium always exists and the price of stability of this game is at most the nth harmonic number $H(n)$ and also provide an example showing that this upper bound is the best possible in directed networks. For undirected networks, Anshelevich et al. [2] presented a tight bound on price of stability of 4/3 for single source and two players case. However, whether there is a tighter bound for arbitrarily many players in undirected networks was left as an open question. Fiat et al. [3] improved the upper bound to $O(\log\log n)$ for a special case where each node of the network has a player and they are required to connect to a common destination. Chen and Roughgarden [4] considered the weighted version of the game where each player has a weight and the cost of an edge is shared among the players who pass it in proportion to their weights. As opposed to the ordinary Nash equilibrium considered before, Albers [5] investigated the situation where coordination among players is allowed and showed nearly matching upper and lower bounds on the price of stability with respect to the notion of *strong Nash equilibrium*.

E-mail address: lijian@cs.umd.edu.

[1] Research supported by NSF grants CCF-0430650.

0020-0190/$ – see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.ipl.2009.04.015

Our results. We prove that for undirected graphs with a distinguished destination to which all players must connect, the price of stability of the Shapley network design game is $O(\frac{\log n}{\log\log n})$ where n is the number of players.

2. Preliminaries

We first introduce notations and formally state the problem. We are given a undirected network $G(V, E)$ and n selfish players. Player i has to choose a path from source vertex s_i to destination vertex t_i. Let $\mathcal{P}_i$ denote the set of simple $s_i - t_i$ paths. The cost of an edge e, $c(e)$, is shared equally by all players who pass e. An *outcome* of the game is specified by a set of n path, each chosen by one player. For an outcome $(P_1, P_2, \ldots, P_n)$ for $P_i \in \mathcal{P}_i$, the cost assigned to player i is $c_i(P_1, P_2, \ldots, P_n) = \sum_{e\in P_i} \frac{c_e}{f_e}$ where f_e is the number of paths that include e. We define the cost of the outcome as

$$c(P_1, P_2, \ldots, P_n) = \sum_i c_i(P_1, P_2, \ldots, P_n) = \sum_{e\in\bigcup_i P_i} c_e.$$

Let P_{-i} denote the vector of paths chosen by the players other than i. An outcome $(P_1, P_2, \ldots, P_n)$ is a *Nash equilibrium* if for every player i, $c_i(P_i, P_{-i}) = \min_{\tilde{P}_i\in\mathcal{P}_i} c_i(\tilde{P}_i, P_{-i})$.

The *price of stability* is defined as the ratio of the cost of the best Nash equilibrium of the game to that of an optimal solution. We note that the optimal solution is the min-cost steiner forest satisfying all connectivity requirement (s_i, t_i)s.

We consider the following potential function, also used in [2], that maps every outcome into a numeric value

$$\Phi(P_1, \ldots, P_k) = \sum_{e\in E}\sum_{i=1}^{f_e} \frac{c_e}{i} = \sum_{e\in E} c_e \cdot H(f_e), \tag{1}$$

where f_e denotes the number of paths P_i that include edge e and $H(n) = 1 + \frac{1}{2} + \frac{1}{3} + \cdots + \frac{1}{n}$ is the nth Harmonic number.

The most important property of the potential function is that if a single player i changes its strategy then the difference between the potential of the new state and that of the original state is exactly the change in the cost of player i [2].

In a finite game, *better-response dynamics* is the following process: If the current outcome is not a Nash equilibrium, there exists a player who can decrease its cost by switching its strategy. The player updates its strategy to an arbitrary superior one, and repeat until a Nash equilibrium is reached. While better response dynamics need not terminate in general, it must terminate in finite steps in Shapley network design games since the potential Φ strictly decreases during the process and no outcome appears twice in a finite game.

3. An $O(\frac{\log n}{\log\log n})$ upper bound for the single sink case

We assume the network is connected and all players share the same destination t. It is easy to see an optimal solution is a steiner tree with terminals $\{s_i\}_{i=1,\ldots,n} \cup \{t\}$. Suppose the outcome $NASH = (P_1^N, \ldots, P_n^N)$ is a Nash equilibrium which is obtained by better-response dynamics from an optimal solution $OPT = (P_1^O, \ldots, P_n^O)$. The property of the potential function ensures that $\Phi(NASH) \leqslant \Phi(OPT)$. We denote paths of $NASH$ and that of OPT by $\{P_i^N\}_{i=1,\ldots,n}$ and $\{P_i^O\}_{i=\{1,\ldots,n\}}$, respectively. It is proven that the edge set used by $NASH$ forms a tree [3]. We denote the tree of $Nash$ by $T^N = \bigcup_i P_i^N$ and that of OPT by $T^O = \bigcup_i P_i^O$. Let $|NASH|$ and $|OPT|$ be their costs respectively.

Let f_e^N denote the number of paths that include edge e in $NASH$. Let $f^N(i) = \sum_{e:\ f_e^N=i} c_e$ and $g^N(j) = \sum_{e:\ f_e^N\geqslant j} c_e = \sum_{i\geqslant j} f^N(i)$. It is easy to see $|NASH| = \sum_i f^N(i) = g^N(1)$.

For ease of discussion, we create a dummy player 0 residing in $s_0 = t$. We can see this player has no influence on either $NASH$ or OPT. First we consider the tree $T^O = \bigcup_i P_i^O$. Doubling all edges in T^O forms a Eulerian tour. Traversing this tour gives a sequence S of vertices in T^O. Suppose ϕ is a permutation of $\{s_i\}_{i=0,\ldots,n}$ according to their first appearance in S. For simplicity of notation, we let $\phi(n+1) = \phi(0)$. It is easy to see $\sum_{i=0}^n d(\phi(i), \phi(i+1)) \leqslant 2|T^O| = 2|OPT|$ where $d(u, v)$ is the length of the shortest path between u and v.

For any two players i and j, let $LCA(i, j)$ be the least common ancestor of s_i and s_j in tree T^N (taking t as the root). We let P_i^j be the subpath of P_i^N starting from s_i and ending at $LCA(i, j)$. From the definition of Nash equilibrium, we know the cost of player i in $NASH$ is less than that of first reaching s_j and then following the path P_j^N to t. Thus, we have the following:

$$\sum_{e\in P_i^j} \frac{c_e}{f_e^N} \leqslant d(s_i, s_j) + \sum_{e\in P_j^i} \frac{c_e}{f_e^N + 1}.$$

Similarly, we have

$$\sum_{e\in P_j^i} \frac{c_e}{f_e^N} \leqslant d(s_i, s_j) + \sum_{e\in P_i^j} \frac{c_e}{f_e^N + 1}.$$

Adding them together, we get

$$\sum_{e\in P_i^j} \frac{c_e}{f_e^N(f_e^N + 1)} + \sum_{e\in P_j^i} \frac{c_e}{f_e^N(f_e^N + 1)} \leqslant 2d(s_i, s_j).$$

We denote the left-hand side of last equality by $A(i, j)$. We have

$$\sum_{i=0}^n A\big(\phi(i), \phi(i+1)\big) \leqslant 2\sum_{i=0}^n d\big(\phi(i), \phi(i+1)\big) \leqslant 4|OPT|. \tag{2}$$

Now we prove

$$\sum_{i=0}^n A\big(\phi(i), \phi(i+1)\big) = \sum_{i=0}^n \Bigg(\sum_{e\in P_{\phi(i+1)}^{\phi(i)} \cup P_{\phi(i)}^{\phi(i+1)}} \frac{c_e}{f_e^N(f_e^N + 1)} \Bigg) \geqslant \sum_{e\in T_N} \frac{c_e}{f_e^N(f_e^N + 1)} = \sum_i \frac{1}{i(i+1)} f^N(i). \tag{3}$$

Actually, we only need to prove every $e \in T_N$ appears in $P_{\phi(i)}^{\phi(i+1)} \cup P_{\phi(i+1)}^{\phi(i)}$ for some $0 \leqslant i \leqslant n$. First it is easy to see $P_i^j \cup P_j^i$ is the unique path from s_i to s_j in T_N. For any $e \in T_N$, let $T_{N,e}^1$ and $T_{N,e}^2$ be two connected components obtained by deleting e from T_N. It is easy to see $T_{N,e}^i \cap \{s_0, \ldots, s_n\} \neq \emptyset$ for $i = 1, 2$ since each leaf of T_N contains at least one player. So, there exists some i such that $\phi(i) \in T_e^1$ and $\phi(i+1) \in T_e^2$ and e must lie in the unique path from $\phi(i)$ to $\phi(i+1)$.

We define $\gamma = \max\{i \mid g^N(i) \geqslant \frac{1}{2} \cdot |NASH|\}$. We can see the following:

$$\begin{aligned}\Phi(NASH) &= \sum_i f^N(i)H(i) \geqslant \sum_{i \geqslant \gamma} f^N(i)H(i) \\ &\geqslant H(\gamma) \sum_{i \geqslant \gamma} f^N(i) = H(\gamma) g^N(\gamma) \\ &\geqslant \frac{1}{2} H(\gamma)|NASH|.\end{aligned}$$

Since $\Phi(NASH) \leqslant \Phi(OPT) \leqslant H(n)|OPT|$, we have

$$|NASH| \leqslant \frac{2H(n)}{H(\gamma)} \cdot |OPT|. \tag{4}$$

From (2) and (3), we can get

$$\begin{aligned}4|OPT| &\geqslant \sum_i \frac{1}{i(i+1)} f^N(i) \geqslant \sum_{i \leqslant \gamma} \frac{1}{i(i+1)} f^N(i) \\ &\geqslant \frac{1}{\gamma(\gamma+1)} \sum_{i \leqslant \gamma} f^N(i) \\ &= \frac{1}{\gamma(\gamma+1)} \big(|NASH| - g^N(\gamma+1)\big) \\ &\geqslant \frac{1}{2\gamma(\gamma+1)} |NASH|,\end{aligned}$$

where the last inequality holds since $g^N(\gamma+1) < \frac{1}{2}|NASH|$ by the definition of γ.

Thus, we have

$$|NASH| \leqslant 8\gamma(\gamma+1) \cdot |OPT|. \tag{5}$$

Combining inequalities (4) and (5), we have $|NASH| \leqslant \min\{\frac{2H(n)}{H(\gamma)}, 8\gamma(\gamma+1)\} \cdot |OPT|$ for any γ. The right-hand side takes maximum value $O(\frac{\log n}{\log\log n}) \cdot |OPT|$ by choosing $\gamma = O(\sqrt{\frac{\log n}{\log\log n}})$. Therefore, we have proved $\frac{|NASH|}{|OPT|} \leqslant O(\frac{\log n}{\log\log n})$.

Acknowledgement

Thanks to Samir Khuller for providing comments on earlier drafts.

References

[1] A.S. Schulzan, N. Stier Moses, On the performance of user equilibria in traffic networks, in: Proceedings of the 14th Annual ACM–SIAM Symposium on Discrete Algorithms (SODA), 2003, pp. 86–87.

[2] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, T. Roughgarden, The price of stability for network design with fair cost allocation, in: Proc of 45th FOCS, 2004, pp. 295–304.

[3] Amos Fiat, Haim Kaplan, Meital Levy, Svetlana Olonetsky, Ronen Shabo, On the price of stability for designing undirected networks with fair cost allocations, in: ICALP06.

[4] Ho-Lin Chen, Tim Roughgarden, Network design with weighted players, in: Proceedings of the Eighteenth Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), 2006, pp. 29–38.

[5] Susanne Albers, On the value of coordination in network design, in: Proceedings of the Nineteenth Annual ACM–SIAM Symposium on Discrete Algorithms, 2008, pp. 294–303.