



ELSEVIER

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs



Improving the H_k -bound on the price of stability in undirected Shapley network design games [☆]



Yann Disser ^{a,*}, Andreas Emil Feldmann ^{b,*}, Max Klimm ^{a,*}, Matúš Mihalák ^{c,*}

^a Institut für Mathematik, Technische Universität Berlin, Germany

^b Combinatorics & Optimization, University of Waterloo, Canada

^c Institute of Theoretical Computer Science, ETH Zurich, Switzerland

ARTICLE INFO

Article history:

Received 1 November 2013

Received in revised form 26 September 2014

Accepted 21 October 2014

Available online 24 October 2014

Communicated by G. Ausiello

Keywords:

Undirected Shapley network design game

Price of stability

Potential-optimal price of stability

Potential-optimal price of anarchy

ABSTRACT

In this article we show that the price of stability of Shapley network design games on undirected graphs with k players is at most $\frac{k^3(k+1)/2-k^2}{1+k^3(k+1)/2-k^2} H_k = (1 - \Theta(1/k^4)) H_k$, where H_k denotes the k -th harmonic number. This improves on the known upper bound of H_k , which is also valid for directed graphs but for these, in contrast, is tight. Hence, we give the first non-trivial upper bound on the price of stability for undirected Shapley network design games that is valid for an arbitrary number of players. Our bound is proved by analyzing the price of stability restricted to Nash equilibria that minimize the potential function of the game. We also present a game with $k = 3$ players in which such a restricted price of stability is 1.634. This shows that the analysis of Bilò and Bove (2011) [3] is tight. In addition, we give an example for three players that improves the lower bound on the (unrestricted) price of stability to 1.571.

Crown Copyright © 2014 Published by Elsevier B.V. All rights reserved.

1. Introduction

Infrastructure networks are the lifelines of our civilization. Through generations a tremendous effort has been undertaken to cover the Earth's surface with irrigation canal systems, sewage lines, roads, railways, and – more recently – data networks. Some of these infrastructures are initiated and planned by a central authority that designs the network and decides on its topology and dimension. Many networks, however, arise as an outcome of actions of selfish individuals who are motivated by their own connectivity requirements rather than by optimizing the overall network design. A prominent example of the latter phenomenon is the rise of the Internet. In order to quantify the efficiency of networks, it is crucial to understand the processes that govern their formation. Anshelevich et al. [1] proposed a particularly elegant model for such processes, which is now known as the *Shapley network design game* or the *network design game with fair cost allocation* (for an overview of other models for network formation, see [14]).

The Shapley network design game is played by k players on a graph $G = (V, E)$ with positive edge-costs $c_e \in \mathbb{N}$. Each player $i \in \{1, 2, \dots, k\}$ has a source-target pair $s_i, t_i \in V$ of vertices that she needs to connect with a simple path in G . The choice of such a path is called a *strategy* of the player, and a collection consisting of one strategy for each player is called a *strategy profile*. The cost c_e of every edge e is shared equally among the players using it. Each player i aims at choosing

[☆] A preliminary version of this article appeared at the 8th International Conference on Algorithms and Complexity (CIAC).

* Corresponding authors.

E-mail addresses: disser@math.tu-berlin.de (Y. Disser), andreas.feldmann@uwaterloo.ca (A.E. Feldmann), klimm@math.tu-berlin.de (M. Klimm), matus.mihalak@inf.ethz.ch (M. Mihalák).

<http://dx.doi.org/10.1016/j.tcs.2014.10.037>

0304-3975/Crown Copyright © 2014 Published by Elsevier B.V. All rights reserved.

a path of the smallest possible (individual) cost to herself. This cost is defined as the sum of the cost shares for player i along the path. Players are selfish in that they only care about their own costs. In particular, they do not care about the *social cost*, defined as the sum of all players' individual costs and denoted by $\text{cost}(P)$ for a strategy profile P .

A *Nash equilibrium* of a Shapley network design game is a strategy profile in which no player i can switch to an s_i - t_i path that yields her a smaller individual cost. To quantify the effect of the selfish behavior, it is natural to compare the social cost of a Nash equilibrium of the game with the smallest social cost among all possible strategy profiles [1,14]. Several quantifications of selfish behavior have been studied, based on whether we restrict ourselves to a specific set of Nash equilibria, and whether we compare the worst or best such equilibrium in terms of social cost. In this article, we adopt the notion of the *price of stability*, introduced by Anshelevich et al. [1]. Denoting by \mathcal{N} the set of all Nash equilibria and by O an (optimal) strategy profile that minimizes the social cost of a game, the price of stability of the game is defined as the ratio $\min_{N \in \mathcal{N}} \text{cost}(N) / \text{cost}(O)$.

Anshelevich et al. [1] observed that Shapley network design games always have a Nash equilibrium by showing that they belong to the class of congestion games. For these games, the existence of a Nash equilibrium is always guaranteed, as shown by Rosenthal [13]. Rosenthal's existence proof relies on a potential function argument. That is, he showed that there exists a function Φ that maps strategy profiles to real numbers and has the property that if any one player changes her strategy unilaterally, then the value of Φ changes by the exact same value as the cost of the player. This observation, together with the finiteness of the space of all strategy profiles, implies the existence of a Nash equilibrium. In particular, any *potential minimum*, i.e., a strategy profile that globally minimizes the potential function, is a Nash equilibrium. The potential function of a game is unique up to an additive constant (see Monderer and Shapley [12]). Using the special form of the potential function for Shapley network design games, Anshelevich et al. [1] showed that the price of stability of any game is at most $H_k = \sum_{i=1}^k \frac{1}{i}$, the k -th *harmonic number* (which is of order $\log k$). This upper bound is tight for games played on directed graphs. That is, there are Shapley network design games on directed graphs [1] for which the price of stability is arbitrarily close to H_k .

The situation is different for *undirected* Shapley network design games, i.e., games played on undirected graphs. As the same potential arguments remain valid, the price of stability of any game is still at most H_k . Yet, the largest known price of stability (asymptotically) is a constant, more precisely $348/155 \approx 2.245$ (see Bilò et al. [4]). This leaves the question of the worst possible price of stability in undirected Shapley network design games with k players wide open. Remarkably, the largest known price of stability, as provided by Bilò et al., does not come from a simple example, but from a complicated construction. Previously known worst-case games had a price of stability of $4/3 \approx 1.333$ [1], $12/7 \approx 1.714$ [8], and $42/23 \approx 1.826$ [7]. Despite numerous attempts [1,3,4,7,8,11] to narrow the gap of the bounds on the price of stability, there has been little progress in terms of numerical results. It is generally believed that the price of stability is smaller than H_k , and we confirm this belief in this article. For small values of k some smaller upper bounds are known. For $k = 2$ players, the price of stability is at most $4/3 < H_2 = 3/2$ and this is tight [1,7]. Bilò and Bove [3] analyzed the case of $k = 3$ players and showed that the price of stability of any such game is at most $1.634 < H_3 = 1.833$. For this case, however, a considerable gap remains, as the worst example known has a price of stability of $74/48 \approx 1.542$ [7]. Thus, already for $k = 3$ players, the exact worst-case price of stability is unknown.

For several special cases, one can derive better upper bounds on the price of stability. If all players share the same terminal (so called *multicast games*) then the price of stability is at most $O(\log k / \log \log k)$ [11]. For *broadcast games* in which in addition every vertex is the source of at least one player, the price of stability was first shown to be $O(\log \log k)$ [8], then $O(\log \log \log k)$ [10], and very recently the bound was reduced to a constant $O(1)$ [5].

Many of the mentioned upper bounds are not only valid for the best Nash equilibrium of a game, but also for a very specific one – the potential minimum. Potential minima have desirable stability properties. For example, they are reached by certain learning dynamics for players that do not always play rationally (see Blume [6]). This motivates to explicitly study the ratio between the cost of a potential minimum and that of a profile minimizing the social cost – a *social optimum*. To stress the described stability properties of potential minima, Asadpour and Saberi [2] called this ratio the *inefficiency ratio of stable equilibria*. Kawase and Makino [9] called the very same ratio the *potential-optimal price of anarchy*. They also define the *potential-optimal price of stability* of a game in the obvious way as the ratio between the cost of a best potential minimum and that of a social optimum. They prove that the potential-optimal price of anarchy of undirected Shapley network design games is at most $O(\sqrt{\log k})$ for the special case where all players share the same terminal node, and where every vertex is the source of at least one player. They give a construction of a game with potential-optimal price of anarchy $\Omega(\sqrt{\log \log k})$.

1.1. Our contribution

Our main result shows that the price of stability in undirected Shapley network design games is at most $\frac{k^3(k+1)/2-k^2}{1+k^3(k+1)/2-k^2} H_k = (1 - \Theta(1/k^4)) H_k$. Thus, we provide the first general upper bound that shows that the price of stability for k players is strictly smaller than H_k . To prove this upper bound, we generalize the techniques of Christodoulou et al. [7] to any number of players. In short, similar to Christodoulou et al., we obtain a set of inequalities relating the cost of any Nash equilibrium to the cost of a social optimum. We then combine these in a non-trivial way to obtain the claimed upper bound, additionally assuming that the Nash equilibrium has a smaller potential than the social optimum. Interestingly, the resulting upper bound is tight for the case of $k = 2$ players.

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات