Theoretical Computer Science 521 (2014) 107–122

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

On the hardness of network design for bottleneck routing games ☆

Dimitris Fotakis [a], Alexis C. Kaporis [b], Thanasis Lianeas [a,*], Paul G. Spirakis [c,d]

[a] *School of Electrical and Computer Engineering, National Technical University of Athens, 15780 Athens, Greece*
[b] *Department of Information and Communication Systems Engineering, University of the Aegean, 83200 Karlovasi, Samos, Greece*
[c] *Department of Computer Science, University of Liverpool, UK*
[d] *Computer Technology Institute and Press – Diophantus, University Campus, 26500 Patras, Greece*

ARTICLE INFO

Article history:
Received 28 February 2013
Received in revised form 24 September 2013
Accepted 26 November 2013
Communicated by X. Deng

Keywords:
Bottleneck routing games
Braess's paradox
Selfish network design

ABSTRACT

In routing games, the selfish behavior of the players may lead to a degradation of the network performance at equilibrium. In more than a few cases however, the equilibrium performance can be significantly improved if we remove some edges from the network. This counterintuitive fact, widely known as Braess's paradox, gives rise to the (selfish) network design problem, where we seek to recognize routing games suffering from the paradox, and to improve their equilibrium performance by edge removal. In this work, we investigate the computational complexity and the approximability of the network design problem for non-atomic bottleneck routing games, where the individual cost of each player is the bottleneck cost of her path, and the social cost is the bottleneck cost of the network, i.e. the maximum latency of a used edge. We first show that bottleneck routing games do not suffer from Braess's paradox either if the network is series-parallel, or if we consider only subpath-optimal Nash flows. On the negative side, we prove that even for games with strictly increasing linear latencies, it is NP-hard not only to recognize instances suffering from the paradox, but also to distinguish between instances for which the Price of Anarchy (PoA) can decrease to 1 and instances for which the PoA cannot be improved by edge removal, even if their PoA is as large as $\Omega(n^{0.121})$. This implies that the network design problem for linear bottleneck routing games is NP-hard to approximate within a factor of $O(n^{0.121-\varepsilon})$, for any constant $\varepsilon > 0$. The proof is based on a recursive construction of hard instances that carefully exploits the properties of bottleneck routing games, and may be of independent interest. On the positive side, we present an algorithm for finding a subnetwork that is almost optimal with respect to the bottleneck cost of its worst Nash flow, when the worst Nash flow in the best subnetwork routes a non-negligible amount of flow on all used edges. We show that the running time is essentially determined by the total number of paths in the network, and is quasipolynomial when the number of paths is quasipolynomial.

☆ This work was supported by the project Algorithmic Game Theory, co-financed by the European Union (European Social Fund – ESF) and Greek national funds, through the Operational Program "Education and Lifelong Learning", under the research funding program Thales, by an NTUA Basic Research Grant (PEBE 2009), by the ERC project RIMACO, and by the EU FP7/2007-13 (DG INFSO G4-ICT for Transport) under Grant Agreement No. 288094 (Project eCompass). An extended abstract of this work appeared in the Proceedings of the 5th International Symposium on Algorithmic Game Theory, SAGT 2012, Maria Serna (Ed.), Lect. Notes Comput. Sci., vol. 7615, Springer, 2012, pp. 156–167.

* Corresponding author.
E-mail addresses: fotakis@cs.ntua.gr (D. Fotakis), kaporisa@gmail.com (A.C. Kaporis), tlianeas@mail.ntua.gr (T. Lianeas), spirakis@cti.gr, P.Spirakis@liverpool.ac.uk (P.G. Spirakis).

0304-3975/$ – see front matter
http://dx.doi.org/10.1016/j.tcs.2013.11.035

1. Introduction

A typical instance of a non-atomic *bottleneck routing game* consists of a directed network, with an origin s and a destination t, where each edge is associated with a non-decreasing function that determines the edge's latency as a function of its traffic. A rate of traffic is controlled by an infinite population of players, each willing to route a negligible amount of traffic through an $s-t$ path. The players are non-cooperative and selfish, and seek to minimize the maximum edge latency, a.k.a. the *bottleneck cost* of their path. Thus, the players reach a *Nash equilibrium flow*, or simply a *Nash flow*, where they all use paths with a common locally minimum bottleneck cost. Bottleneck routing games and their variants have received considerable attention due to their practical applications in communication networks (see e.g., [6,3] and the references therein).

1.1. Previous work and motivation

Every bottleneck routing game is known to admit a Nash flow that is optimal for the network, in the sense that it minimizes the maximum latency on any used edge, a.k.a. the bottleneck cost of the network (see e.g., [3, Corollary 2]). On the other hand, bottleneck routing games usually admit many different Nash flows, some with a bottleneck cost quite far from the optimum. Hence, there has been a considerable interest in quantifying the performance degradation due to the players' non-cooperative and selfish behavior in (several variants of) bottleneck routing games. This is typically measured by the *Price of Anarchy* (PoA) [13], which is the ratio of the bottleneck cost of the worst Nash flow to the optimal bottleneck cost of the network.

Simple examples (see e.g., [7, Fig. 2]) demonstrate that the PoA of bottleneck routing games with linear latency functions can be as large as $\Omega(n)$, where n is the number of vertices of the network. For atomic splittable bottleneck routing games, where the population of players is finite, and each player controls a non-negligible amount of traffic which can be split among different paths, Banner and Orda [3] observed that the PoA can be unbounded, even for very simple networks, if the players have different origins and destinations and the latency functions are exponential. On the other hand, Banner and Orda proved that if the players use paths that, as a secondary objective, minimize the number of bottleneck edges, then all Nash flows are optimal. For a variant of non-atomic bottleneck routing games, where the social cost is the average (instead of the maximum) bottleneck cost of the players, Cole, Dodis, and Roughgarden [7] proved that the PoA is 4/3, if the latency functions are affine and a subclass of Nash flows, called *subpath-optimal Nash flows*, is only considered. Subsequently, Mazalov et al. [18] studied the inefficiency of the best Nash flow under this notion of social cost.

For atomic unsplittable bottleneck routing games, where each player routes a unit of traffic through a single $s-t$ path, Banner and Orda [3] proved that for polynomial latency functions of degree d, the PoA is $O(m^d)$, where m is the number of edges of the network. On the other hand, Epstein, Feldman, and Mansour [8] proved that for series-parallel networks with arbitrary latency functions, all Nash flows are optimal. Subsequently, Busch and Magdon-Ismail [5] proved that the PoA of atomic unsplittable bottleneck routing games with identity latency functions can be bounded in terms of natural topological properties of the network. In particular, they proved that the PoA of such games is bounded from above by $O(l+\log n)$, where l is the length of the longest $s-t$ path, and by $O(k^2+\log^2 n)$, where k is length of the longest cycle.

With the PoA of bottleneck routing games so high and crucially depending on topological properties of the network, a natural approach to improving the performance at equilibrium is to exploit the essence of Braess's paradox [4], namely that removing some edges may change the network topology (e.g., it may decrease the length of the longest path or cycle), and significantly improve the bottleneck cost of the worst Nash flow (see e.g., Fig. 1). This approach gives rise to the (selfish) *network design problem*, where we seek to recognize bottleneck routing games suffering from the paradox, and to improve the bottleneck cost of the worst Nash flow by edge removal. In particular, given a bottleneck routing game, we seek for the *best subnetwork*, namely, the subnetwork for which the bottleneck cost of the worst Nash flow is best possible. In this setting, one may distinguish two extreme classes of instances: *paradox-free* instances, where edge removal cannot improve the bottleneck cost of the worst Nash flow, and *paradox-ridden* instances, where the bottleneck cost of the worst Nash flow in the best subnetwork is equal to the optimal bottleneck cost of the original network (see also [20,11]).

The approximability of selective network design, a generalization of network design where we cannot remove certain edges, was considered by Hou and Zhang [12]. For atomic unsplittable bottleneck routing games with a different traffic rate and a different origin and destination for each player, they proved that if the latency functions are polynomials of degree d, it is NP-hard to approximate selective network design within a factor of $O(m^{d-\varepsilon})$, for any constant $\varepsilon > 0$. Moreover, for atomic k-splittable bottleneck routing games with multiple origin-destination pairs, they proved that selective network design is NP-hard to approximate within any constant factor.

However, a careful look at the reduction of [12] reveals that their strong inapproximability results crucially depend on both (i) that we can only remove certain edges from the network, so that the subnetwork actually causing a high PoA cannot be destroyed, and (ii) that the players have different origins and destinations (and also are atomic and have different traffic rates). As for the importance of (ii), in a different setting, where the players' individual cost is the sum of edge latencies on their path and the social cost is the bottleneck cost of the network, it is known that Braess's paradox can be dramatically more severe for instances with multiple origin-destination pairs than for instances with a single origin-destination pair. More precisely, Lin et al. [14] proved that if the players have a common origin and destination, the removal of at most k edges from the network cannot improve the equilibrium bottleneck cost by a factor greater than $k+1$. On the other hand, Lin et al.