The Korean Institute of Communications and Information Sciences

Available online at www.sciencedirect.com

ScienceDirect

ICT Express 3 (2017) 155–159

www.elsevier.com/locate/icte

The social cost of congestion games by imposing variable delays[☆]

Josep Díaz[a], Ioannis Giotis[a,b], Lefteris Kirousis[c,d,*], Ioannis Mourtos[b], Maria Serna[a]

[a] *Departament de Ciències de la Computació, Universitat Politècnica de Catalunya, Barcelona, Spain*
[b] *Department of Management Science and Technology, Athens University of Economics and Business, Greece*
[c] *Department of Mathematics, National and Kapodistrian University of Athens, Greece*
[d] *Computer Technology Institute Press "Diophantus", Patras, Greece*

Received 17 July 2017; accepted 15 November 2017
Available online 13 December 2017

Abstract

In this study, we describe a new coordination mechanism for non-atomic congestion games that leads to a (selfish) social cost which is arbitrarily close to the non-selfish optimal. This mechanism incurs no additional cost, in contrast to tolls that typically differ from the social cost as expressed in terms of delays.

Keywords: Congestion games; Price of anarchy; Coordination mechanisms

1. Preliminaries

A selfish behavior is one of the primary reasons why many systems with multiple agents deviate from desirable outcomes. Allowing players to solely prioritize their own benefit can lead to social inefficiency, even in outcomes where no one is better off compared to an optimal solution. A typical example is transportation and network routing where a selfish selection among possible routes can lead to congestion with accompanying economical and environmental issues. Various approaches have been proposed to steer the selfishly constructed outcome toward optimal social welfare. Typically, the main idea is to incentivize the users to alter their selections to ones that lead to socially better outcomes, usually using tolls or similar measures.

We propose herein an alternative approach that alters the way users experience latency and can offer significant improvements on the social cost. However drivers still get to pick their own route. In more detail, instead of all users experiencing the same latency, we propose the implementation of variable latencies through a prioritization scheme. In other words, we allow for some users to experience latencies smaller than before, while others to experience longer ones. We employ known results to show that our system achieves the optimal social welfare, if users behave selfishly, as they are expected to. We present a discretization of the theoretically continuous functions, which approximates the optimal social welfare, to make such a system practical.

We also wish to emphasize the distributed and decentralized nature of our system. As explained in the next sections, each resource (road or highway in the transportation setting) individually and independently implements the desired changes. Note that our system's average latency on each road, as experienced by the users, is at least equal, and closely matches, the road's average latency without the system in place. Hence our system falls under the notion of coordination mechanisms, in other words no "cheating" in the form of network improvements, which typically carry significant cost, is introduced. Moreover, no imposing of tolls (transfer of social cost to a different type) is conducted. We simply distribute the resource differently. This

---

[*] Corresponding author at: Department of Mathematics, National and Kapodistrian University of Athens, Greece.

E-mail addresses: diaz@cs.upc.edu (J. Díaz), igiotis@cs.upc.edu (I. Giotis), lkirousis@math.uoa.gr (L. Kirousis), mourtos@aueb.gr (I. Mourtos), mjserna@cs.upc.edu (M. Serna).

Peer review under responsibility of The Korean Institute of Communications Information Sciences.

[☆] This paper is part of a special section titled "Special Issue on Intelligent Transportation Communication Systems" guest edited by Prof. Zonghua Gu, Prof. Shigang Yue, Prof. Sanghyun Ahn, Prof. Juan-Carlos Cano, Prof. Joel Rodrigues, Prof. Houbing Song, Prof. Hai L. Vu, Prof. Danda Rawat, Prof. Dongkyun Kim.

https://doi.org/10.1016/j.icte.2017.11.014

holds on any instance, and not just in equilibrium settings, which means that we do not obtain a worse performance even in non-stable situations. Furthermore, we do not need to know the *demand* in advance, i.e. our system delivers close to the social optimum for all possible total amounts of traffic. Our only requirement is that the latency induced on each road is a non-negative, non-decreasing, continuously differentiable and convex function of the traffic.

We believe that our system has a strong applicability potential. For example, some countries have already implemented metered highway entrance ramps which can vary the latency of incoming drivers. Traffic lights may also be used in an urban environment to implement this aspect of our mechanism. We deliberately leave the prioritization scheme generic to allow for different such approaches with our only requirement being that users choosing to alter their current selection are forced to experience maximal latency in their new selection. This is a reasonable requirement as typically someone that alters her selection in a running system ends up at the end of the queue.

We examine our system in the generic scheme of congestion games to emphasize that it admits applications beyond traffic routing. One interesting application could be in the context of job scheduling on computing resources. Again, in a typical model, each user choosing to use a particular resource experiences the same latency (e.g, computing jobs running in parallel on a computer). We can achieve optimal average job completion times under selfish behavior by prioritizing jobs according to our proposed mechanism, so that some jobs complete faster and some slower than before. We note that this can easily be implemented by an administrator (human or computerized) using system priorities.

2. Related work

The fact that selfish behavior can lead to inefficiency has long been studied in the context of transportation theory [1,2]. More recently, Koutsoupias and Papadimitriou introduced the *Price of Anarchy* as a measure of this inefficiency [3,4]. The exploration of this metric in the context of selfish routing was then greatly progressed by Roughgarden and Tardos [5,6] who bounded the Price of Anarchy for different classes of latency functions.

Ways to improve inefficient outcomes have naturally been investigated, with the imposition of *tolls* being a prime example [7–9]. While this approach achieves optimal social welfare regarding latencies, it introduces a cost separation to the players because the tolls' cost is affecting behavior but is not accounted for in the objective function.

Coordination mechanisms were recently introduced by Christodoulou et al. [10] as a way to "shape" latency functions and steer the selfishly dictated outcome toward greater social welfare. Two main restrictions are considered in the type of coordination mechanisms defined in [10], namely that the latency per resource is not decreased and that the benchmark optimal social welfare, against which the mechanism is measured, is still the original one without any additional latencies possibly imposed by the mechanism. It has recently been shown that indeed such mechanisms can positively affect social welfare [11]. Our approach sustains the non-decreasing latency on average but not on every user, as a prerequisite for achieving a significantly lower Price of Anarchy than the mechanism of [10]. In fact, the average latency per user within our system can be made arbitrarily close to the unique latency per user without the system in place.

The approach of differentiating the latency per resource is also explored from an algorithmic perspective by Harks et al. [12] but not with the same scheme. The results of Farzad et al. [13] are closer to our work. However, in the later, it is the *strategic* equilibrium without the mechanism in place that matches the optimal under that mechanism (i.e., the strategic equilibrium under that mechanism may in general differ substantially from the non-selfish optimal).

3. Model

For convenience, we define a congestion game $(E, l, \mathcal{S}, P, d)$ in the generic sense using the network routing (or alternatively, flow) terminology. E is a set of edges with an associated non-negative, non-decreasing, continuously differentiable and convex $l_e()$ *latency* function for each edge. P is a set of players, partitioned into n sets P_i, $i = 1, \ldots, n$. A player in P_i is said to be a player of player type i. For each player type i we have a source–sink pair (s_i, t_i). The set $\mathcal{S}$ is partitioned into sets $\mathcal{S}_i$, $i = 1, \ldots, n$ so that each $\mathcal{S}_i$, which is called the strategy set of player type i, is a set of finite sequences of elements of E. The elements $S \in \mathcal{S}_i$ are the strategies of players of player type i (also referred to as *paths* from s_i to t_i). Finally, d is a sequence d_i, $i = 1, \ldots, n$, where each d_i is non-negative number d_i, the flow (or traffic) demand for player type i.

We assume that each player type corresponds to a continuum of nonatomic players, each with a negligible flow. An infinitesimal part of the flow (or traffic) will often referred to as a *user*. Let x_i^S denote a nonnegative real representing the part of demand d_i that uses strategy (path) S and x_i the vector for the strategy set $\mathcal{S}_i$, i.e. $x_i = (x_i^S)_{S \in \mathcal{S}_i}$. The vector x for all x_i's is called a *flow* if for all player types i, $\sum_{S \in \mathcal{S}_i} x_i^S = d_i$. We define the part of the demand of a player type i that uses edge e as follows

$$x_e^i = \sum_{\{S : S \in \mathcal{S}_i, e \in S\}} x_i^S.$$

The total flow through an edge e is defined as follows

$$x_e = \sum_{i=1\ldots n} x_e^i.$$

In the related literature, the *cost* induced to each player type i by a flow x is defined as $c_i(x) = \sum_{e \in E} l_e(x_e) \cdot x_e^i$. The cost of the total flow through an edge e is defined as

$$c_e(x_e) = l_e(x_e) \cdot x_e,$$

whereas the social cost is defined as

$$C(x) = \sum_{e \in E} l_e(x_e) \cdot x_e.$$

We now provide the notion of Wardrop equilibrium in our setting for reference.

جستجو کنید ...

جستجو

در اجرای درخواست شما مشکلی رخ داده است

با سلام□ متأسفانه مشکلی در فرایند اجرای درخواست شما رخ داده است□

همکاران ما در حال تلاش برای رفع این مشکل هستند□

لطفاً درخواست خود را در ساعات دیگری مجدداً تکرار فرمایید و اگر باز هم با این مشکل رو به رو شدید، از طریق فرم تماس با ما به واحد پشتیبانی اطلاع دهید□

برای یافتن مطلب مورد نظر خود می توانید از روش های جستجوی زیر استفاده فرمایید :

جستجو در میان موضوعات

برای جستجو در میان موضوعات، به محض این که عبارت خود را در فیلد زیر بنویسید، موضوع های مرتبط در درخت سمت چپ با رنگ متمایزی مشخص می شوند.

جستجو ...

جستجو در میان مقالات

اگر موضوع مورد نظر شما در لیست موضوعات اصلی وجود نداشت، با استفاده از فیلد زیر می توانید آن را در بین کل مقاله های سایت جستجو فرمایید.

جستجو ...

جستجو

لیست درختی موضوعات

- الگوریتم های هیوریستیک
- بازاریابی و مدیریت بازار
- حسابداری و حسابرسی
- روش های آماری
- سازمان و مدیریت
- سیستم های اطلاعاتی
- علوم اقتصادی
- مدیریت استراتژیک
- مدیریت امور فرهنگی
- مدیریت تولید
- مدیریت دولتی
- مدیریت رفتار سازمانی
- مدیریت مالی
- مدیریت منابع انسانی
- مطالعات اسلامی
- کارآفرینی و مدیریت کسب و کار

تماس با واحد پشتیبانی

همکاران ما در واحد پشتیبانی آمادگی دارند تمامی درخواست های شما عزیزان را بررسی نموده و در اسرع وقت رسیدگی نمایند.

پیگیری خرید مقاله

پس از خرید هر مقاله، یک کد رهگیری منحصر به فرد به شما تقدیم خواهد شد که با استفاده از آن می توانید وضعیت خرید خود را پیگیری فرمایید.

کد رهگیری

ارسال

پیگیری سفارش ترجمه

با ثبت کد رهگیری پرداخت، می توانید سفارش خود را پیگیری نموده و به محض اتمام ترجمه، فایل ترجمه مقاله خود را دانلود نمایید.

کد رهگیری

ارسال

g+