Operations Research Letters 44 (2016) 645–648

Contents lists available at ScienceDirect

Operations Research Letters

journal homepage: www.elsevier.com/locate/orl

Tighter bounds on the inefficiency ratio of stable equilibria in load balancing games

Akaki Mamageishvili *, Paolo Penna

Department of Computer Science, ETH Zurich, Switzerland

ARTICLE INFO

Article history:
Received 15 December 2015
Received in revised form
19 July 2016
Accepted 19 July 2016
Available online 25 July 2016

Keywords:
Nash equilibria
Load balancing games
Potential games
Price of anarchy

ABSTRACT

In this paper we study the inefficiency ratio of stable equilibria in load balancing games introduced by Asadpour and Saberi (2009). We prove tighter lower and upper bounds of 7/6 and 4/3, respectively. This improves over the best known bounds for the problem (19/18 and 3/2, respectively). Equivalently, the results apply to the question of how well the optimum for the L_2-norm can approximate the L_∞-norm (makespan) in identical machines scheduling.

1. Introduction

Load balancing problems are classical optimization problems which are also actively studied in the context of games, where jobs are owned by selfish but rational players. These games are prototypical of resource allocation problems in which users (players) do not act altruistically, therefore leading the system to suboptimal configurations. Naturally, one can consider the *social optimum* as the allocation minimizing the *makespan*, that is, the maximum load over all machines (a classical measure of efficiency). In contrast, in the game-theoretic setting, each player strives to optimize the cost of her own job only. This will typically result in a so-called *Nash equilibrium*, that is, an allocation in which no player can benefit by moving her job to another machine. In this work we consider *pure* Nash equilibria, that is, configurations in which each player chooses one strategy and unilateral deviations are not beneficial. In general, games may also possess mixed Nash equilibria in which players choose strategies according to a probability distribution. The inefficiency of Nash equilibria is a central topic in algorithmic game theory and it measures how much selfishness can impede optimization.

Asadpour and Saberi [3] introduced and studied the *inefficiency ratio of stable equilibria (IRSE)* in several games, including load balancing ones (see Section 2). This notion quantifies the efficiency loss in games when players play certain *noisy* best-response dynamics (see Section 1.2). For load balancing games, the IRSE has another very simple and natural interpretation (which is also of independent interest and studied earlier):

> *Are the allocations minimizing the L_2-norm (sum of the squares of the machine loads) also sufficiently good for minimizing the L_∞-norm (makespan)?*

Intuitively, while the social cost is measured by the L_∞-norm (makespan), the players collectively minimize the L_2-norm (the potential of the game). Therefore, the IRSE on load balancing games is equal to some value c if *every* allocation minimizing the L_2-norm is automatically a c-approximation for the L_∞-norm, i.e., the makespan of this allocation is at most c-times the optimal makespan. An exact bound on the IRSE is not known, as opposed to other measures related to the inefficiency of equilibria (see Section 1.2). Asadpour and Saberi [3] proved an upper bound of 3/2 on IRSE and observed that an example in Alon et al. [1] implies a lower bound on IRSE of 19/18.

1.1. Our contribution

In this work we improve both the previous upper and lower bounds on the IRSE on load balancing games:

- In Section 3 we show an improved lower bound of 7/6. While the previous lower bound from Alon et al. [1] is obtained by a simple instance with 3 machines and 6 jobs, our result is based on a family of instances that depend on the number of machines. Notably, our construction improves the previous lower bound already for three machines to 13/12. As the number of machines grows, the lower bound tends to 7/6.

* Corresponding author.
E-mail address: akaki@inf.ethz.ch (A. Mamageishvili).

http://dx.doi.org/10.1016/j.orl.2016.07.014

- In Section 4 we improve the upper bound to 4/3. Intuitively speaking, the proof consists in showing that in every allocation whose cost is more than 4/3 the optimal makespan, we can redistribute the jobs in a way that reduces the potential, i.e., the L_2-norm. This is the same argument of Asadpour and Saberi [3], and our technical contribution is to show that there are *two* ways to "reshuffle" the jobs, so that one of them decreases the potential.

As mentioned above, these results can be restated by saying that every job allocation minimizing the L_2-norm guarantees a 4/3 approximation for the makespan, while in some instances the optimum for the L_2-norm has makespan at least 7/6 the optimal makespan. Closing the gap between the upper and the lower bound remains an open question, which we discuss at the end of Section 4. In the next section, we discuss further relation with prior work, including studies on the quality of equilibria in games.

1.2. Significance of the results and related work

The inefficiency of Nash equilibria is often measured through two classical notions: the *price of anarchy (PoA)* introduced by Koutsoupias and Papadimitriou [9] for load balancing games on related machines, and the *price of stability (PoS)* introduced in Anshelevich et al. [2], which compare respectively the *worst* and the *best* Nash equilibrium to the social optimum. In some cases, these notions can be considered too extreme as they may include "unrealistic" equilibria.

The IRSE [3] studies the quality of equilibria selected by certain *noisy* best-response dynamics [4]. Intuitively, these dynamics will most likely rest on pure Nash equilibria *minimizing the potential* of the game, and the IRSE can be seen as the price of anarchy restricted to these selected equilibria. The IRSE is also known in the literature under the name of *potential optimal PoA* by Kawase and Makino [8], who also considered the analogous *potential optimal PoS*. Correa et al. [5] studied earlier potential optimal PoA in the capacitated network routing model, though under a different name. In load balancing games (and several others) we have $PoS < IRSE < PoA$, which in a sense tells that the PoA and the PoS are either too pessimistic or too optimistic. Specifically, on m machines $PoA = \left(2 - \frac{1}{m+1}\right)$ [6,10], $PoS = 1$ [10], while $IRSE$ is between 19/18 and 3/2 [3]. The latter bounds are strengthened in the present paper to $7/6 \leq IRSE \leq 4/3$. This means that players can easily compute a 4/3 (or better) approximation of the optimum via these simple dynamics, but also that in some instances the dynamics is unlikely to choose optimal or nearly optimal solutions either (namely, within a factor smaller than 7/6).

The upper bound 4/3 also suggests an intriguing comparison with the study of *sequential PoA* by Hassin and Yovel [7] for these games: there, players are far from myopic and the equilibrium is meant on an extensive form game in which players are able to reason about future moves of following players. It is an interesting question which of these two dynamics give a better makespan in the end.

2. Preliminaries

In load balancing there are n jobs with weights $w_1, \ldots, w_n$ that we want to put on m identical machines (each job is allocated to one machine). The job allocation determines the load l_j of each machine j, that is the sum of the jobs weights that are allocated to this machine. The goal is to find an allocation that has the lowest possible *makespan*, that is, the maximum load over all machines.

In *load balancing games* each job is a *player* who can choose any of the m possible machines. The cost for player i is simply the load of the machine chosen by this player, and naturally each player aims at minimizing *her own cost*. The strategies of all players specify a job allocation (in the game theoretic terminology this is the strategy profile).

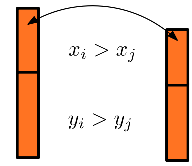

Fig. 1. When a swap of bundles of jobs reduces the potential.

An allocation minimizing the makespan is called a *social optimum*, and its makespan is called a *social optimum makespan*. The *potential* associated to an allocation is the sum of the *squares* of the corresponding machine loads, ${l_1}^2 + \cdots + {l_m}^2$ where l_j is the load of machine j at this allocation. An allocation minimizing the potential function is called a potential optimal allocation or simply *potential optimum*.

It is well known that load balancing games are weighted potential games with the above potential function. This means that all pure Nash equilibria (allocations where no player can unilaterally improve moving to another machine) are actually 'locally optimal' for the potential (a single job move cannot reduce the potential). Asadpour and Saberi [3] introduced and studied the *inefficiency ratio of stable equilibria (IRSE)*, which is the largest (among all the instances of a game) ratio between the *worst* makespan of a potential optimal allocation and the social optimum makespan.

Potential optimal allocations satisfy the following condition (see Fig. 1). Split the total load of each machine into two bundles of jobs, that is, $l_j = x_j + y_j$ where x_j is the sum of the weights of a (possibly empty) subset of jobs allocated to j. If two machines i and j satisfy $x_i < x_j$, then $y_i \geq y_j$ for otherwise swapping x_i with x_j reduces the potential. Pure Nash equilibria satisfy the weaker condition that a single job k in machine i does not improve if moving to another machine j, that is, $l_i - w_k \leq l_j$.

3. Improved lower bound

In this section we strengthen the 19/18 lower bound on IRSE in [3,1]. The idea of the proof is to construct an instance in which the potential optimum is obtained when one single machine has "higher" load, and all others have the same load (Fig. 2 left), while the optimal makespan does the opposite: one machine has "lower" load and all others have the same higher load (Fig. 2 right).

Theorem 3.1. *IRSE in load balancing games is at least* 7/6.

Proof. Consider the instance in Fig. 2 where $m = k + 1 \geq 3$, $n = 2k + 2$ and the weights of jobs are $k, k, k, k + 1, k + 1, \ldots, 2k - 1, 2k - 1, (5k - 1)/2$. We prove that the allocation on the left (see figure) minimizes the potential function, while the one on the right has optimal makespan, thus implying

$$IRSE \geq \frac{(7k-1)/2}{3k} = \frac{7m-8}{6m-6}$$

which tends to 7/6 as m goes to infinity. First note that the potential in both allocations is the same. Consider any job allocation and without loss of generality assume that the job with the largest weight is on machine 1. If we fix a load on machine 1, then any job allocation which balances the load on the other machines minimizes the potential (among all allocations with this fixed load on machine 1). Therefore, if the largest job is alone on machine 1, then the potential is minimized in the social optimum case, while if it is located on machine 1 together with the job of smallest weight, then the potential is minimized when the job