Taxes for Linear Atomic Congestion Games*

Ioannis Caragiannis, Christos Kaklamanis, and Panagiotis Kanellopoulos

Research Academic Computer Technology Institute and
Dept. of Computer Engineering and Informatics
University of Patras, 26500 Rio, Greece

Abstract. We study congestion games where players aim to access a set of resources. Each player has a set of possible strategies and each resource has a function associating the latency it incurs to the players using it. Players are non–cooperative and each wishes to follow strategies that minimize her own latency with no regard to the global optimum. Previous work has studied the impact of this selfish behavior to system performance. In this paper, we study the question of how much the performance can be improved if players are forced to pay taxes for using resources. Our objective is to extend the original game so that selfish behavior does not deteriorate performance. We consider atomic congestion games with linear latency functions and present both negative and positive results. Our negative results show that optimal system performance cannot be achieved even in very simple games. On the positive side, we show that there are ways to assign taxes that can improve the performance of linear congestion games by forcing players to follow strategies where the total latency suffered is within a factor of 2 of the minimum possible; this result is shown to be tight. Furthermore, even in cases where in the absence of taxes the system behavior may be very poor, we show that the total disutility of players (latency plus taxes) is not much larger than the optimal total latency. Besides existential results, we show how to compute taxes in time polynomial in the size of the game by solving convex quadratic programs. Similar questions have been extensively studied in the model of non-atomic congestion games. To the best of our knowledge, this is the first study of the efficiency of taxes in atomic congestion games.

1 Introduction

We study the well-known *congestion games* introduced by Rosenthal [22]. In a congestion game Π there is a set E of resources and a set N of n players. Each player i has a positive unsplittable demand (or weight) w_i and a set of actions $\mathcal{P}_i \subseteq 2^E$ (each action of player i is a set of resources). Each resource e has a non-negative and non-decreasing latency function f_e defined over non-negative numbers. A resource e used by players with total demand w causes a latency of $f_e(w)$ to each of them. Players are non–cooperative and each wishes

* This work was partially supported by the European Union under IST FET Integrated Project 015964 AEOLUS.

Y. Azar and T. Erlebach (Eds.): ESA 2006, LNCS 4168, pp. 184–195, 2006.

to minimize her own cost (the cumulative latency experienced at the resource used) with no regard to the global optimum. *Network congestion games* can be used to model non-cooperative users in a communication network, where each user i aims to communicate an amount of traffic w_i through a least congested single path connecting two particular nodes s_i and t_i. In this setting, resources correspond to network links and the actions of user i are all the paths connecting node s_i to t_i.

In general, players follow *mixed strategies*, i.e., player i selects a probability distribution $y_i = \{y_{ip}|p \in \mathcal{P}_i\}$ over her actions. Mixed strategies where $y_{ip} \in \{0, 1\}$ are called *pure strategies.* Each player is aware of the strategies selected by all other players. We denote by y_{ie} the probability that player i uses resource e. Clearly, $y_{ie} = \sum_{p \in \mathcal{P}_i : e \in p} y_{ip}$. We use the term *assignment* to refer to the vector of players' strategies. In a pure assignment, all players follow pure strategies. Given an assignment y, we denote by $L_{ip}(y; \Pi)$ the expected latency of player i when selecting action p. Then the expected latency of player i is $L_i(y; \Pi) = \sum_{p \in \mathcal{P}_i} y_{ip} L_{ip}(y; \Pi)$. An assignment y is a (mixed or pure) *Nash equilibrium* if no player has an incentive to unilaterally change her strategy, i.e., $L_i(y; \Pi) \leq L_i(y_{-i}, x_i; \Pi)$ for any player i and for any probability distribution x_i over the actions in $\mathcal{P}_i$, where y_{-i}, x_i denotes the assignment obtained by y when player i deviates from y_i to x_i. The *weighted total latency* defined as $W(y; \Pi) = \sum_i w_i L_i(y; \Pi)$ has been used as a measure of performance of assignment y in game Π. Another natural measure of performance is the *total latency* defined as $T(y; \Pi) = \sum_i L_i(y; \Pi)$. The *price of anarchy* [17, 21] (with respect to the weighted total latency) of a game Π is the maximum of the ratio of $W(y; \Pi)/W(x; \Pi)$ where y is a Nash equilibrium and x is any assignment for Π. Similarly, we may define the price of anarchy with respect to the total latency. We use the terms *unweighted* and *weighted* for congestion games in order to denote whether players have equal weights or not. Clearly, in unweighted congestion games, the weighted total latency equals the total latency.

[9, 11, 12, 13, 16, 17, 19] study various games which can be thought of as special cases of congestion games with respect to the complexity of computing equilibria of best/worst social cost and the price of anarchy when the social cost is defined as the maximum latency experienced by any player. These include *linear congestion games*, i.e., games with latency functions of the form $f_e(w) = \alpha_e w + b_e$ with non-negative constants α_e and b_e, and *load balancing games*, i.e., linear congestion games where the actions of players are singleton sets. In load balancing terminology, we refer to the resources of a load balancing game as machines. The performance measure of the weighted total latency has been studied in [1, 5, 6, 18, 24]. Awerbuch et al. [1] and, independently, Christodoulou and Koutsoupias [6] prove tight bounds on the price of anarchy of congestion games. Among other results concerning polynomial latency functions, they show that the price of anarchy of pure Nash equilibria in unweighted linear congestion games is 5/2 while for mixed Nash equilibria or pure Nash equilibria of weighted players it is 2.618. Bounds on the price of anarchy of pure Nash equilibria were recently proved to be tight even for load balancing games [5] while better bounds exist only for load balancing games on machines with identical latency functions [5, 24]. The authors of [18] study symmetric load balancing games where all machines are actions for all players.

In order to mitigate the impact of selfish behavior on system performance, we introduce *taxes* to the resources. We use a *tax function* $\delta : E \times Q^+ \rightarrow Q^+$ that assigns a tax $\delta_e(w)$ to each player of weight w that wishes to use e. Assuming selfish behavior of the players, we obtain a new *extended game* (Π, δ) where each player now aims to minimize the expected latency she suffers plus the taxes she pays. The tax paid by player i when selecting action p is $\Delta_{ip}(\Pi, \delta) = \sum_{e \in p} \delta_e(w_i)$. Given an assignment y, the expected tax paid by player i is $\Delta_i(y; \Pi, \delta) = \sum_{p \in \mathcal{P}_i} y_{ip} \Delta_{ip}(\Pi, \delta)$. Again, y is a Nash equilibrium for the extended game if no player has an incentive to unilaterally change her strategy, i.e., $L_i(y; \Pi) + \Delta_i(y; \Pi, \delta) \leq L_i(y_{-i}, x_i; \Pi) + \Delta_i(y_{-i}, x_i; \Pi, \delta)$. We use two measures of performance in the extended game (Π, δ) extending the measures of total latency and weighted total latency in congestion games without taxes. The *total cost* of an assignment y is $T(y; \Pi, \delta) = \sum_i (L_i(y; \Pi) + \Delta_i(y; \Pi, \delta))$, while the *weighted total cost* of an assignment y is $W(y; \Pi, \delta) = \sum_i w_i (L_i(y; \Pi) + \Delta_i(y; \Pi, \delta))$.

Motivated by [8], we distinguish between refundable and non-refundable taxes. In the former case, we assume that the collected taxes can be feasibly returned (directly or indirectly) to the players (e.g., as a "lump-sum refund") and therefore do not contribute to the overall system disutility. However, refunding the collected taxes could be logistically or economically infeasible; the latter case models this scenario.

Definition 1. *A function $\delta : E \times Q^+ \rightarrow Q^+$ is a ρ-mixed-efficient refundable tax for the congestion game Π with respect to the total latency (resp. weighted total latency) if $T(y; \Pi, 0) \leq \rho \cdot T(x; \Pi, 0)$ (resp. $W(y; \Pi, 0) \leq \rho \cdot W(x; \Pi, 0)$) for any mixed Nash equilibrium y in the extended game (Π, δ) and any assignment x. A function $\delta : E \times Q^+ \rightarrow Q^+$ is a ρ-mixed-efficient non-refundable tax for the congestion game Π with respect to the total cost (resp. weighted total cost) if $T(y; \Pi, \delta) \leq \rho \cdot T(x; \Pi, 0)$ (resp. $W(y; \Pi, \delta) \leq \rho \cdot W(x; \Pi, 0)$) for any mixed Nash equilibrium y in the extended game (Π, δ) and any assignment x.*

Similarly, we define ρ-pure-efficient refundable and non-refundable taxes by constraining y to be a pure Nash equilibrium. We use the terms pure-optimal and mixed-optimal to refer to 1-pure-efficient and 1-mixed-efficient taxes, respectively.

The bounds on the price of anarchy of congestion games with respect to the weighted total latency can be also expressed using the above definition. Any tight bound of ρ on the price of anarchy over mixed (resp. pure) Nash equilibria implies that the *trivial tax function* that assigns no tax to the resources is ρ-mixed-efficient (resp. ρ-pure-efficient) and no better in general. Another issue which is related to our study is that of *network design* for selfish players (or *resource removal*). In this setting, the question is whether the performance of the game can be improved by removing some of the resources; this is equivalent to a tax function which assigns to each resource a tax of either 0 or ∞ for all players. [3] proves that deciding whether resource removal for a weighted linear congestion game Π can yield price of anarchy better than 2.618 is NP-complete. Furthermore, there are games where this is not feasible at all, implying that taxes of this type are not better than 2.618-pure-efficient.

The problem of computing optimal taxes has been extensively considered in the model of non-atomic congestion games [23]. The main difference of these games from the atomic ones we study in this paper is that each player controls an infinitesimally small demand related to the total demand on the system, thus, the actions of a single player have negligible effect on the system performance. This difference is substantial enough so that the related results (see [8, 10, 15] and the references therein) do not carry over to our model. In fact, even nearly-optimal taxes do not always exist in our model.

In this paper we show the following results. We first study symmetric load balancing games, where we show how to compute pure-optimal taxes for unweighted players. We present lower bounds stating that optimal taxes may not be feasible even in very simple games. In particular, there are unweighted load balancing games on identical machines that do not admit $(11/10-\epsilon)$-pure-efficient taxes, weighted load balancing games on identical machines that do not admit $(9/8-\epsilon)$-pure-efficient taxes (note that this bound matches the upper bound on the price of anarchy for these games [18]), and unweighted load balancing games on identical machines that do not admit $(2-\epsilon)$-mixed-efficient taxes. Even simple non-load-balancing congestion games with unweighted players may not admit $(6/5-\epsilon)$-pure-efficient taxes either. For unweighted congestion games, we present a *universal tax function* by showing that, for a particular value of the parameter τ which is shown to be best possible, the function $\delta_e = \alpha_e \tau$ is $(1+2/\sqrt{3})$-pure-efficient, thus beating the lower bound of 5/2 on the price of anarchy of pure Nash equilibria. This is an interesting result since the tax function does not depend at all on the game played on the resources; it depends only on the resources themselves. Next we exploit solutions of convex quadratic programs to compute 2-mixed-efficient taxes for congestion games with respect to both the weighted total latency and the total latency. Note that the first result beats the lower bound of 2.618 on the price of anarchy while when the total latency is of concern, the price of anarchy is unbounded. Both bounds are tight.

We also consider the case of non-refundable taxes. When considering the weighted total cost, it seems that there is not much room for beating the lower bounds on the price of anarchy. However, we show that weighted load balancing games on identical machines admit $(1+\sqrt{2})$-mixed-efficient non-refundable taxes. This is an existential result since the tax defined uses an *optimal assignment* (i.e., the pure assignment minimizing the weighted total latency). It can be made algorithmic and yield a marginally worse $1+\sqrt{2}+\epsilon$ bound when the number of machines is constant exploiting a PTAS from [4] for approximating the optimal assignment. This result should be compared to the lower bound of 5/2 on the price of anarchy over pure Nash equilibria proved recently in [5]. Recall that the price of anarchy of weighted congestion games is unbounded when the total latency is of concern. Somehow surprisingly, we show that any congestion game admits 4-mixed-efficient non-refundable taxes with respect to the total latency. Furthermore, 6-mixed-efficient non-refundable taxes for these games can be computed in polynomial time. Here, we exploit semi-pure assignments with particular properties which are obtained by rounding the fractional solutions of a convex quadratic program to half-integral ones. The use of convex quadratic programming is motivated by [2] where integral solutions of such programs have been used to approximate scheduling on unrelated machines. However, we rarely need integrality; even fractional or half-integral solutions suffice in order to compute taxes. For the analysis of the upper bounds we develop and use two inequalities that characterize Nash equilibria of congestion games with taxes.

Some details related to the extended game as well as convex quadratic programs are presented in Section 2. The results on refundable and non-refundable taxes are presented in Sections 3 and 4, respectively. We conclude with open problems in Section 5. Due to lack of space, many proofs have been omitted from this extended abstract.

2 Preliminaries

Properties of the extended game. For a weighted congestion game Π and a tax function δ, the extended game (Π, δ) can be seen as a congestion game with player-specific latency functions [14, 20]. Although this is not always true for such games, we can show that the extended game always has a pure Nash equilibrium and hence pure-efficient taxes are well defined. In order to prove this fact, we define a potential function over pure assignments of the extended game by slightly modifying the potential function of weighted linear congestion games in [12].

In our proofs, we use the equivalent expressions of the (weighted) total cost of assignments in the extended game given in the next lemma. The proof easily follows by the definitions.

Lemma 1. *For each assignment y in a weighted congestion game Π with linear latency functions of the form $f_e(w) = \alpha_e(w) + b_e$ and a tax function δ, the following equations hold*

$$W(y; \Pi, \delta) = \sum_e \left(\alpha_e \left(\left(\sum_i y_{ie} w_i \right)^2 + \sum_i y_{ie} (1 - y_{ie}) w_i^2 \right) + b_e \sum_i y_{ie} w_i + \sum_i y_{ie} w_i \delta_e(w_i) \right)$$

$$T(y; \Pi, \delta) = \sum_e \left(\alpha_e \left(\left(\sum_i y_{ie} \right) \left(\sum_i y_{ie} w_i \right) + \sum_i y_{ie} (1 - y_{ie}) w_i \right) + b_e \sum_i y_{ie} + \sum_i y_{ie} \delta_e(w_i) \right)$$

where w_i denotes the weight of player i.

In our analysis, we use the two inequalities stated in the following, which characterize Nash equilibria of the extended game. Although complicated at first glance, when examined carefully (and together with the expressions in Lemma 1), these inequalities provide insight about what efficient taxes should look like.

Lemma 2. *Given a weighted congestion game Π and a tax function δ, consider a mixed Nash equilibrium y and any assignment x of (Π, δ). Then*

$$W(y;\Pi,\delta) \leq \sum_e \left(\alpha_e \left(\left(\sum_i x_{ie} w_i\right)\left(\sum_i y_{ie} w_i\right) + \sum_i x_{ie}(1-y_{ie}) w_i^2\right) + b_e \sum_i x_{ie} w_i + \sum_i x_{ie} w_i \delta_e(w_i)\right) \quad (1)$$

$$T(y;\Pi,\delta) \leq \sum_e \left(\alpha_e \left(\left(\sum_i x_{ie}\right)\left(\sum_i y_{ie} w_i\right) + \sum_i x_{ie}(1-y_{ie}) w_i\right) + b_e \sum_i x_{ie} + \sum_i x_{ie} \delta_e(w_i)\right) \quad (2)$$

Computation of taxes. In most cases, in order to compute taxes, we wish to compute assignments that satisfy some property; these correspond to solutions of programs of the form:

$$\begin{array}{lll} \text{(QP1)} & \text{minimize } g(x) & \\ & \text{subject to } x_{ie} \geq \sum_{p \in \mathcal{P}_i : e \in p} x_{ip}, & i \in N, e \in E \\ & \sum_{p \in \mathcal{P}_i} x_{ip} \geq 1, & i \in N \\ & x_{ie}, x_{ip} \geq 0, & i \in N, e \in E, p \in \mathcal{P}_i \end{array}$$

where $g(x)$ is a convex quadratic function. Convex quadratic programs can be solved within any additive error ϵ in time polynomial in the size of the program and $1/\epsilon$. So, programs like (QP1) are solvable in polynomial time when the total number of actions is polynomial. In many interesting cases like in network congestion games, actions may be exponentially many. However, we can overcome this difficulty for these games and efficiently solve (QP1) in time polynomial in the number of resources and the number of players by considering it as a flow problem. Details will appear in the final version of the paper.

One would hope to solve (QP1) with the objective functions $W(x;\Pi,0)$ or $T(x;\Pi,0)$ and obtain optimal assignments, i.e., mixed assignments of minimum (weighted) total latency. Unfortunately, these functions are non-convex and, furthermore, they are always optimized at pure assignments. This is not difficult to see since the (weighted) total latency of a mixed assignment can be seen as the expectation of the (weighted) total latency of the pure assignments implied by the corresponding probability distributions. Hence, optimizing these functions would also contradict hardness results in [4].

3 Refundable Taxes

We start with an encouraging result concerning pure-optimal refundable taxes.

Theorem 1. *Pure-optimal refundable taxes in unweighted symmetric load balancing games always exist and are computable in polynomial time.*

Proof. (*Sketch*) Consider an unweighted symmetric load balancing game Π with latency functions $f_e(w) = \alpha_e w + b_e$. Let e' be the machine with the smallest a_e among all machines e with non-zero α_e. Let $\epsilon = \alpha_{e'}/2$. Also, let o_e be the number of players that select machine e in an optimal assignment. Let e^* be the machine with maximum $\alpha_e o_e + b_e$ among all machines. For each machine e with $\alpha_e > 0$, we define $\delta_e = \alpha_{e^*} o_{e^*} + b_{e^*} - \alpha_e o_e - b_e$. Let e_0 be the machine with minimum b_e among all machines e with $\alpha_e = 0$. We define $\delta_{e_0} = \alpha_{e^*} o_{e^*} + b_{e^*} + \epsilon$ and $\delta_e = \infty$ for all other machines e with $\alpha_e = 0$. We can show that the function δ is a pure-optimal refundable tax for game Π.

Polynomial time computability follows since optimal assignments are easy to compute through a reduction to a minimum cost flow problem. We construct a network F as follows. For each resource e of the game, F has two nodes u_e and v_e connected through n parallel directed edges g_e^i of unit capacity and cost $\alpha_e(2i - 1) + b_e$, for $i = \{1, \cdots, n\}$. s is connected through directed edges to nodes u_e and all nodes v_e are connected through directed edges to t. All edges adjacent to either s or t have zero cost and capacity n. Then, it easily follows that an optimal assignment for the original game can be obtained by computing a minimum cost flow of size n from s to t. □

Unfortunately, the next theorem rules out the possibility of obtaining optimal taxes even in simple congestion games.

Theorem 2.
a) There exists a weighted symmetric load balancing game on identical machines that does not admit ρ–pure–efficient refundable taxes with respect to the weighted total latency for any $\rho < 9/8$.
b) For any $\epsilon > 0$, there exists an unweighted symmetric load balancing game on identical machines that does not admit $(2 - \epsilon)$–mixed–efficient refundable taxes.
c) There exists an unweighted load balancing game on identical machines that does not admit ρ–pure–efficient refundable taxes for any $\rho < 11/10$.
d) There exists an unweighted congestion game that does not admit ρ–pure–efficient refundable taxes for any $\rho < 6/5$.

Next, we present a universal tax function for unweighted congestion games in the sense that it does not depend at all on the congestion game; it depends only on the resources themselves.

Theorem 3. *Let $\tau = \frac{3}{2}\sqrt{3} - 2$. For any unweighted congestion game Π with linear latency functions $f_e(w) = \alpha_e w + b_e$, the function $\delta_e = \alpha_e \tau$ is a $\left(1 + \frac{2}{\sqrt{3}}\right)$–pure–efficient refundable tax for Π.*

Our next result indicates that the selection of parameter τ in Theorem 3 is the best possible.

Theorem 4. *For any $\tau \geq 0$ and $\epsilon > 0$, there exists an unweighted load balancing game for which the function $\delta_e = \alpha_e \tau$ is not $\left(1 + \frac{2}{\sqrt{3}} - \epsilon\right)$–pure–efficient refundable tax.*

In the rest of this section we construct 2-mixed-efficient refundable taxes. Given a congestion game, we use a particular assignment in order to compute the tax function. In the case of the weighted total latency, we use the solution of the quadratic program (QP1) with the convex quadratic objective function

$$g_1(x) = \sum_e \left(\alpha_e \left(\left(\sum_i x_{ie} w_i \right)^2 + \sum_i x_i w_i^2 \right) + b_e \sum_i x_{ie} w_i \right)$$

Lemma 3. *Consider a weighted congestion game Π and let x be an assignment which is the optimal solution of (QP1) with the objective function g_1. Then, the function $\delta_e(w) = \alpha_e \sum_i x_{ie} w_i$ is a 2–mixed–efficient refundable tax for Π with respect to the weighted total latency.*

Proof. We will apply inequality (1) for a mixed Nash equilibrium y of the extended game (Π, δ) and assignment x. The last term in the sum at the definition of $W(y; \Pi, \delta)$ in Lemma 1 becomes $\alpha_e \left(\sum_i x_{ie} w_i\right) \left(\sum_i y_{ie} w_i\right)$ and cancels with the first term in the sum of the right part of (1), while the last term in the sum at the right part of (1) becomes $\alpha_e \left(\sum_i x_{ie} w_i\right)^2$. So, (1) yields

$$\begin{aligned} W(y; \Pi, 0) &\leq \sum_e \alpha_e \sum_i x_{ie} \left(1 - y_{ie}\right) w_i^2 + \sum_e \alpha_e \left(\sum_i x_{ie} w_i \right)^2 + \sum_e b_e \sum_i x_{ie} w_i \\ &\leq \sum_e \alpha_e \sum_i x_{ie} w_i^2 + \sum_e \alpha_e \left(\sum_i x_{ie} w_i \right)^2 + \sum_e b_e \sum_i x_{ie} w_i \\ &\leq \sum_e \alpha_e \sum_i x^*_{ie} w_i^2 + \sum_e \alpha_e \left(\sum_i x^*_{ie} w_i \right)^2 + \sum_e b_e \sum_i x^*_{ie} w_i \\ &\leq 2 \left(\sum_e \alpha_e \left(\sum_i x^*_{ie} w_i \right)^2 + \sum_e b_e \sum_i x^*_{ie} w_i \right) \\ &= 2 \cdot W(x^*; \Pi, 0) \end{aligned}$$

where x^* denotes the pure assignment minimizing the weighted total latency. The last inequality follows due to integrality of x^*. □

In the case of total latency, we use the solution of the quadratic program (QP1) with the convex quadratic objective function

$$g_2(x) = \sum_e \left(\alpha_e \left(\left(\sum_i x_{ie} \right) \left(\sum_i x_{ie} w_i \right) + \sum_i x_i w_i \right) + b_e \sum_i x_{ie} \right)$$

We can show the following result; the proof is similar to the proof of Lemma 3.

Lemma 4. *Consider a weighted congestion game Π and let x be an assignment which is the optimal solution of (QP1) with the objective function g_2. Then, the function $\delta_e(w) = \alpha_j w \sum_i x_{ie}$ is a 2–mixed–efficient refundable tax for Π with respect to the total latency.*

In order to make the above two results constructive, there is a subtle point concerning the validity of the third inequality in the proofs of Lemmas 3 and 4, since, in practice, the solution of the quadratic program has not perfect accuracy. As in [2], we can guarantee the validity of this inequality by making the accuracy parameter sufficiently small. As a corollary we obtain the following statement.

Theorem 5. *There exist polynomial time algorithms for computing 2-mixed-efficient refundable taxes with respect to the total latency and the weighted total latency in weighted congestion games.*

4 Non-refundable Taxes

In this section, we consider non-refundable taxes; we first focus on efficient non–refundable taxes with respect to the weighted total cost. A recent lower bound of 5/2 on the price of anarchy of weighted load balancing games on identical machines [5] implies that the trivial tax function is not $(5/2 - \epsilon)$-pure-efficient for any $\epsilon > 0$. This lower bound can be modified so that resource removal cannot improve the price of anarchy either. We show that better non-refundable taxes do exist. Here, the corresponding tax function uses an optimal assignment. Unfortunately, even computing an approximate such assignment is hard [4]. We can use a PTAS from [4] to show a slightly worse constructive result when the number of machines is constant. We note that the lower bound in [5] uses a constant number of machines.

Theorem 6. *Any weighted load balancing game on identical machines admits $\left(1+\sqrt{2}\right)$–mixed–efficient non–refundable taxes with respect to the weighted total cost.*

The proof of Theorem 6 uses the tax function

$$\delta_e(w) = \begin{cases} \sum_i x_{ie} w_i - w, & \text{if } \sum_i x_{ie} w_i \geq w; \\ 0, & \text{otherwise} \end{cases}$$

for a load balancing game Π with latency function of the form $f(w) = x + b$, where x denotes a pure assignment for Π that minimizes the weighted total latency.

In order to compute efficient non-refundable taxes with respect to the total cost, we use solutions to the quadratic program (QP1) with the objective function

$$g_3(x) = \sum_e \left(\alpha_e \left(\sum_i x_{ie} \right) \left(\sum_i x_{ie} w_i \right) + b_e \sum_i x_{ie} \right)$$

Ideally, we would like to use optimal pure assignments, i.e., an optimal integral solution x^* of (QP1) with the objective function g_3. However, even approximate semi-pure assignments can be used to obtain efficient non-refundable taxes.

Lemma 5. *Consider a weighted congestion game Π and let x be a semi–pure assignment with $g_3(x) \leq \rho \cdot g_3(x^*)$. Then, the function*

$$\delta_e(w) = \begin{cases} \alpha_e \left(2\sum_i x_{ie} - 1\right) w, & \textit{if } 2\sum_i x_{ie} \geq 1 \\ 0, & \textit{otherwise} \end{cases}$$

is a 4ρ–mixed–efficient non–refundable tax for Π with respect to the total cost.

Hence, by applying Lemma 5 for $\rho = 1$ we obtain the following existential result.

Corollary 1. *Any weighted congestion game admits 4–mixed–efficient non–refundable taxes with respect to the total cost.*

Next, we show how to compute efficient semi–pure assignments to obtain a slightly worse constructive result. We first solve the quadratic program (QP1) with the convex quadratic objective function

$$g_4(x) = \sum_e \left(\alpha_e \left(\frac{1}{2} + \sum_i x_{ie} \right) \left(\sum_i x_{ie} w_i \right) + b_e \sum_i x_{ie} \right)$$

Then, we obtain a half–integral solution $\hat{x}$ by applying randomized rounding to the solution x as follows. For each i, we use a die with one face for each $p \in \mathcal{P}_i$ such that $x_{ip} > 0$ and a probability of x_{ip} associated with the face corresponding to p. We cast the die twice and let p_1 and p_2 be the actions corresponding to the outcomes. If $p_1 = p_2$, we set $\hat{x}_{ip_1} = 1$, while if $p_1 \neq p_2$, we set $\hat{x}_{ip_1} = \hat{x}_{ip_2} = \frac{1}{2}$; we also set $\hat{x}_{ip} = 0$ for each $p \in \mathcal{P}_i \setminus \{p_1, p_2\}$. We also set $\hat{x}_{ie} = \sum_{p \in \mathcal{P}_i} \hat{x}_{ip}$.

Lemma 6. $E\left[g_3(\hat{x})\right] \leq \frac{3}{2} g_3(x^*)$

By using standard probabilistic arguments, we can guarantee that $g_3(\hat{x}) \leq \left(\frac{3}{2} + \epsilon\right) g_3(x^*)$ for any $\epsilon > 0$ by executing the randomized rounding procedure polynomially many times. Hence, Lemma 5 yields the following.

Theorem 7. *There exists a polynomial time algorithm for computing $(6 + \epsilon)$-mixed-efficient non-refundable taxes with respect to the total cost in weighted congestion games.*

5 Open Problems

Our work reveals several interesting open questions. Tightening the bounds for pure-efficient refundable taxes is a challenging task. In particular, extending the results of Theorem 1 and determining the subclass of unweighted congestion games that admit pure-optimal taxes is one of them. The candidate class is that of the unweighted symmetric congestion games which include network congestion games with a single source and a single destination. The existence of efficient non-trivial universal tax functions for weighted congestion games is also open. We conjecture that such taxes do not exist. For non-refundable taxes, the question whether efficient non-trivial taxes for congestion games with respect to the weighted total cost exist is still open. Special cases as simple as unweighted symmetric load balancing are interesting as well. Here, besides the trivial upper bound, we have a preliminary statement that better than 27/23-pure-efficient non-refundable taxes do not exist. We point out that symmetry has not helped so far, since all our lower bounds are in a sense symmetric constructions. The impact of symmetry of games to the existence of efficient taxes needs further investigation. Complexity issues are also very interesting, i.e., given a congestion game Π, how easy is to compute a ρ-mixed/pure-efficient (non)-refundable tax for this particular game? Our results can be thought of as approximation algorithms for this optimization problem. Although we have made no attempt to formally prove this statement, we strongly believe that this problem is computationally hard for some constant $\rho > 1$. Another open problem is to prove bounds on the cost of taxes that force at least one nearly-optimal assignment to become an equilibrium. This is related to the study of the price of stability [5, 7]. Also, having players with different sensitivities to taxes as in the model of [8, 10, 15] is another interesting extension of our model. Finally, it is worth investigating taxes for congestion games with more general (e.g., polynomial) latency functions.

References

1. B. Awerbuch, Y. Azar, and A. Epstein. The price of routing unsplittable flow. In *Proc. of the 37th Annual ACM Symposium on Theory of Computing (STOC '05)*, pp. 57-66, 2005.
2. Y. Azar and A. Epstein. Convex programming for scheduling unrelated parallel machines. In *Proc. of the 37th Annual ACM Symposium on Theory of Computing (STOC '05)*, pp. 331-337, 2005.
3. Y. Azar and A. Epstein. The hardness of network design for unsplittable flow with selfish users. In *Proc. of the 3rd Workshop on Approximation and Online Algorithms (WAOA '05)*, pp. 41-54, 2005.
4. Y. Azar, L. Epstein, Y. Richter, and G. Woeginger. All-norm approximation algorithms. *Journal of Algorithms*, 52(2): 120-133, 2004.
5. I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos, and L. Moscardelli. Tight bounds for selfish and greedy load balancing. In *Proc. of the 33rd International Colloquium on Automata, Languages and Programming (ICALP '06)*, LNCS 4051, Springer, Part I, pp. 311-322, 2006.
6. G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proc. of the 37th Annual ACM Symposium on Theory of Computing (STOC '05)*, pp. 67-73, 2005.
7. G. Christodoulou and E. Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *Proc. of the 13th Annual European Symposium on Algorithms (ESA '05)*, LNCS 3669, Springer, pp. 59-70, 2005.

8. R. Cole, Y. Dodis, and T. Roughgarden. Pricing network edges for heterogeneous selfish users. In *Proc. of the 35th Annual ACM Symposium on Theory of Computing (STOC '03)*, pp. 521-530, 2003.
9. A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria. In *Proc. of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '02)*, pp. 413-420, 2002.
10. L. Fleischer, K. Jain, and M. Mahdian. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In *Proc. of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS '04)*, pp. 277-285, 2004.
11. D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas and P. Spirakis. The structure and complexity of Nash equilibria for a selfish routing game. In *Proc. of the 29th International Colloquium on Automata, Languages and Programming (ICALP '02)*, LNCS 2380, Springer, pp. 123-134, 2002.
12. D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish unsplittable flows. *Theoretical Computer Science*, 340(3): 514-538, 2005.
13. M. Gairing, T. Lücking, M. Mavronicolas and B. Monien. Computing Nash equilibria for scheduling on restricted parallel links. In *Proc. of the 36th Annual ACM Symposium on Theory of Computing (STOC '04)*, pp. 613-622, 2004.
14. M. Gairing, B. Monien, and K. Tiemann. Routing (un-) splittable flow in games with player-specific latency functions. In *Proc. of the 33rd International Colloquium on Automata, Languages and Programming (ICALP '06)*, LNCS, Springer, 2006, to appear.
15. G. Karakostas, and S. Kolliopoulos. Edge pricing of multicommodity networks for heterogeneous selfish users. In *Proc. of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS '04)*, pp. 268-276, 2004.
16. E. Koutsoupias, M. Mavronicolas and P. Spirakis. Approximate equilibria and ball fusion. *Theory of Computing Systems*, 36(6): 683-693, 2003.
17. E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proc. of the 16th International Symposium on Theoretical Aspects of Computer Science (STACS '99)*, LNCS 1563, Springer, pp. 404-413, 1999.
18. T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. A new model for selfish routing. In *Proc. of the 21st International Symposium on Theoretical Aspects of Computer Science (STACS '04)*, LNCS 2996, Springer, pp. 547-558, 2004.
19. M. Mavronicolas and P. Spirakis. The price of selfish routing. In *Proc. of the 33rd Annual ACM Symposium on Theory of Computing (STOC '01)*, pp. 510-519, 2001.
20. I. Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13: 111-124, 1996.
21. C. Papadimitriou. Algorithms, games and the internet. In *Proc. of the 33rd Annual ACM Symposium on Theory of Computing (STOC '01)*, pp. 749-753, 2001.
22. R. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2: 65-67, 1973.
23. T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2): 236-259, 2002.
24. S. Suri, C. Tóth and Y. Zhou. Selfish load balancing and atomic congestion games. In *Proc. of the 16th Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '04)*, pp. 188-195, 2004.