

Non-Atomic One-Round Walks in Polynomial Congestion Games

Cosimo Vinci

Gran Sasso Science Institute, L'Aquila - Italy
cosimo.vinci@gssi.infn.it

Abstract. In this paper we study the approximation ratio of the solutions achieved after an ϵ -approximate one-round walk in non-atomic congestion games. Prior to this work, the solution concept of one-round walks had been studied for atomic congestion games with linear latency functions only [Christodoulou et al. 2006, Bilò et al. 2011]. We focus on polynomial latency functions, and, by exploiting the primal-dual technique [Bilò 2012], we prove that the approximation ratio is exactly $((1 + \epsilon)(p + 1))^{p+1}$ for every polynomial of degree p . Then, we show that, by resorting to static (resp. dynamic) resource taxation, the approximation ratio can be lowered to $(1 + \epsilon)^{p+1}(p + 1)^p$ (resp. $(1 + \epsilon)^{p+1}(p + 1)!$).

Keywords: computational social choice, algorithmic game theory, congestion games

1 Introduction

Since the end of the Twentieth Century, the computer science community has been interested in the study of complex systems populated by (numerous) selfish agents interacting with each other, and in how their selfish behavior impacts on the social welfare [24]. As examples, one may think to web users greedily sharing limited resources, or to drivers who want to move as fast as possible from a location to another along a street network.

These systems are often modeled by *congestion games* [25]. In these games, there is a set of non-cooperative selfish players sharing a set of *resources* and each resource incurs a certain *latency* to the players using it. Each player has an available set of strategies, where each strategy is a non-empty subset of resources, and aims at choosing a strategy minimizing her cost which is defined as the sum of the latencies experienced on all the selected resources. A congestion game is called *atomic* when the set of players is finite, and it is called *non-atomic* when

Copyright © by the paper's authors. Copying permitted for private and academic purposes.

V. Biló, A. Caruso (Eds.): ICTCS 2016, Proceedings of the 17th Italian Conference on Theoretical Computer Science, 73100 Lecce, Italy, September 7–9 2016, pp. 11–22 published in CEUR Workshop Proceedings Vol-1720 at <http://ceur-ws.org/Vol-1720>

the set of players is infinite and the contribution of each player to the social welfare is infinitesimally small.

In congestion games, selfish behavior always leads to stable outcomes, called *Pure Nash Equilibrium* [23], in which each player cannot improve her utility by unilaterally deviating from her strategy. In general, Pure Nash Equilibria do not yield an optimal social welfare, hence, to measure the quality of these outcomes, two metrics have been successfully proposed: the *price of anarchy* [21] and the *price of stability* [2]. They are defined as the highest and the lowest ratio between the social welfare at any Nash equilibrium and the social optimum, respectively.

Besides Pure Nash Equilibria, in the setting of atomic congestion games, the notion of *one-round walks starting from the empty state* has been widely investigated due to its simplicity and effectiveness. This concept assumes that, starting from the situation in which no strategy has been specified yet, the players are processed sequentially and, at each iteration, the selected player irrevocably chooses her strategy so as to minimize her cost based on the choices of the previous ones. The *approximation ratio of one-round walks starting from the empty state*, an analogous of the price of anarchy instantiated to one-round walks, measures the quality of the outcomes produced at the end of this process [22,14].

Our Contribution. In this work, we translate the solution concept of one-round walks from atomic congestion games to non-atomic ones. We define the solution concept of ϵ -approximate non-atomic one-round walk starting from the empty state, in which there is a continuous flow of selfish players (instead of a discrete number of players) greedily selecting their strategies with the aim of approximately (up to a factor of $1 + \epsilon$) minimizing their costs, given the choices of their predecessors. The ϵ -approximation ratio of a non-atomic one-round walk starting from the empty state is the highest ratio of the the social value achieved by the final outcome of an ϵ -approximate non-atomic one-round walk and the optimal social value. In particular, we study this metric for non-atomic congestion games with polynomial latencies, by proving a tight bound of $((1 + \epsilon)(p + 1))^{p+1}$, where p is the degree of the polynomial functions. Given that the (exact and approximate) price of anarchy for these games has been proven to be equal to $\Theta(\max\{p/\log(p), (1 + \epsilon)^{p+1}\})$, our result shows that outcomes generated after one-round walks are tremendously worse (even asymptotically) than Pure Nash Equilibria in terms of social welfare.

For such a reason, we also focus on (resource) *taxation* [4]: an approach that has been intensively studied in the literature in order to improve the quality of the outcomes resulting from selfish behavior in congestion games. We prove that, by resorting to static taxation (taxes are constant with respect to resource congestion), the ϵ -approximation ratio drops to $(1 + \epsilon)^{p+1}(p + 1)^p$, thus having a good asymptotic reduction with respect to the case without taxes. By resorting to dynamic taxation (in which taxes can vary as a function of resource congestion), we lower the ϵ -approximation ratio to $(1 + \epsilon)^{p+1}(p + 1)! \in \Theta((1 + \epsilon)^{p+1}(p + 1)^{p+3/2}e^{-p})$, thus having a further improvement.

Related Work. The inefficiency of equilibria for polynomial congestion games has been largely studied for both atomic and non-atomic congestion games. Relatively to the non-atomic setting, Roughgarden [27] proves that the price of anarchy for congestion games with polynomial latency functions of degree p is $[1 - p(p + 1)^{-(p+1)/p}]^{-1}$ and characterizes the price of anarchy for other latency functions.

Christodoulou et al. [12] and Awerbuch et al. [3] prove that the price of anarchy for linear atomic congestion games is $5/2$ for unweighted games and 2.168 for weighted games. Aland. et al [1] generalize the previous results to polynomial congestion games, and provide tight bounds asymptotically equal to $\Theta(p/\log(p))^{p+1}$ on the price of anarchy for both unweighted and weighted atomic congestion games with polynomial latency functions of degree p . Christodoulou et al. [13] provide bounds on the approximate price of anarchy and stability for both non-atomic and atomic polynomial congestion games, which are tight for the former.

Among the mechanisms used to improve the quality of outcomes in congestion games, the use of taxation has been extensively studied for several decades. The marginal cost taxation [4] has been proven to enforce an optimal solution in non-atomic congestion games with very general latency functions. In subsequent works, the existence and the computation of efficient taxes in many variants of non-atomic congestion games has been intensively studied [15,16,19,29].

Caragiannis et al. [11] study the efficiency of taxation for linear atomic congestion games, and, among the obtained results, they prove that the price of anarchy drops from 2.168 to at least 2 for weighted congestion games. Bilò and Vinci [8] extend the previous result to polynomial congestion games, and prove that, by resorting to taxation, the ϵ -approximate price of anarchy of unweighted and weighted atomic congestion games with polynomial latency functions drops to $\mathcal{T}_{p+1}(1 + \epsilon)$, where \mathcal{T}_{p+1} is the $(p + 1)$ -th Touchard polynomial.

Other mechanisms used to reduce the price of anarchy in congestion games are Stackelberg Strategies [26,29,20,7,17], in which the strategies of a fraction of players can be controlled with the aim of reducing the price of anarchy.

Mirroknì and Vetta [22] initiate the study of the social welfare achieved after multiple rounds of best-responses in a Nash-dynamics for a particular class of games, and they also consider the case of a one round of best responses. Christodoulou et al. [14] study for the first time the quality of outcomes obtained after a one round of best responses in linear atomic congestion games. They prove that, if players start from an empty state, the approximation ratio is at most $2 + \sqrt{5}$ for unweighted games and $4 + 2\sqrt{3}$ for weighted games. For the unweighted setting, Bilò et al. [6] prove that the previous upper bound is tight, improving a lower bound of 4 obtained for load balancing games by Caragiannis et al. [10]. They also consider as a social function the maximum utility among all players, and provide asymptotically matching bounds for the approximation ratio of one-round walks starting from the empty state and from an arbitrary state. Bilò [5] studies the approximation ratio for atomic congestion games with quadratic and cubic latency functions. The quality achieved by multiple rounds of best responses in linear congestion games has been studied in [14,18]. Bilò and Vinci

[8] improve via taxation the performances of ϵ -approximated one-round walks starting from the empty state in unweighted and weighted polynomial atomic congestion games.

2 Preliminaries

For any integer n , set $[n] := \{1, 2, \dots, n\}$ and $[n]_0 := [n] \cup \{0\}$.

Non-Atomic Congestion Games. A *non-atomic congestion game* is a tuple $\text{CG} = (\mathbf{N}, E, (\ell_e)_{e \in E}, (r_i)_{i \in \mathbf{N}}, (\Sigma_i)_{i \in \mathbf{N}})$, where \mathbf{N} is a totally ordered set of $n \geq 2$ types of players, E is a set of resources, $\ell_e : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is the latency function of resource $e \in E$, and, given $i \in \mathbf{N}$, $r_i \in \mathbb{R}_{\geq 0}$ is the amount of players of type i and $\Sigma_i = \{S_{i,1}, \dots, S_{i,m_i}\} \subseteq 2^E$ is the set of strategies of a player of type i , that is, each player of type i has $m_i \geq 1$ possible choices. A congestion game has polynomial latencies of degree $p \in \mathbb{N}$ when, for each $e \in E$, $\ell_e(x) := \sum_{d \in [p]_0} \alpha_{e,d} x^d$, with $\alpha_{e,d} \geq 0$ for each $d \in [p]_0$; when $p = 1$, we speak of affine latencies.

A *strategy profile* is an n -tuple $\Delta = (\Delta_1, \dots, \Delta_n)$, where $\Delta_i : \Sigma_i \rightarrow \mathbb{R}_{\geq 0}$ is a function denoting, for each strategy $S_i \in \Sigma_i$, the amount $\Delta_i(S_i)$ of players of type i selecting strategy S_i , so that $\sum_{S_i \in \Sigma_i} \Delta_i(S_i) = r_i$. For a strategy profile Δ , the congestion of resource $e \in E$ in Δ , denoted as $k_e(\Delta) := \sum_{i \in \mathbf{N}, S_i \in \Sigma_i: e \in S_i} \Delta_i(S_i)$, is the total amount of players using resource e in Δ . The cost of a player selecting a strategy $S_i \in \Sigma_i$ is defined as $c_{S_i}(\Delta) = \sum_{e \in S_i} \ell_e(k_e(\Delta))$ and each player aims at minimizing it.

Non-Atomic One-Round Walks. For any type $i \in \mathbf{N}$ of players, we extend the set Σ_i of strategies with the empty strategy \emptyset_i , so as to include also the cases in which some players have not chosen their strategies yet; in particular, we denote with \emptyset the empty state, that is, the strategy profile in which none of the players has chosen a strategy.

For any $\epsilon \geq 0$, a strategy $S'_i \in \Sigma_i \setminus \{\emptyset_i\}$ is an ϵ -*approximate best-response* (ϵ -best-response, for brevity) for players of type i in Δ if, for each $S'_i \in \Sigma_i \setminus \{\emptyset_i\}$, $c_{S'_i}(\Delta) \leq (1 + \epsilon)c_{S_i}(\Delta)$. Let $M := \sum_{i \in \mathbf{N}} r_i$ and let $f : [0, M] \rightarrow \mathbf{N}$ be a right-continuous function such that $\int_{f^{-1}[\{i\}]} dx = r_i$ for each $i \in \mathbf{N}$. We call f an *ordering function*. Let $(\Delta^{f,t})_{t \in [0, M]}$ be a family of strategy profiles such that: (1) $\sum_{S \in \Sigma_i \setminus \{\emptyset_i\}} \Delta_i^{f,t}(S) = \int_{[0,t] \cap f^{-1}[\{i\}]} dx$ for each $i \in \mathbf{N}$ and $t \in [0, M]$ (then, $\Delta_i^{f,t}(\emptyset_i) = r_i - \int_{[0,t] \cap f^{-1}[\{i\}]} dx$), (2) $\Delta_i^{f,t}(S)$ is non-decreasing in t . Such a family of strategy profiles is called a *weak one-round walk starting from the empty state*.

Observe that $\Delta_i^{f,t}(S)$ has to be necessarily a non-decreasing and Lipschitzian function with respect to t . Informally, a weak one-round walk models a family of strategy profiles generated by a flow of players sequentially selecting their strategies, in such a way, for any $t \in [0, M]$, there is an amount $\Delta_i^{f,t}(S)$ of players of type i which have already selected strategy S (Point 1), and these players cannot change their strategy (Point 2). Moreover, observe that f defines the ordering in which the players appear in the game.

A weak one-round walk is *simple* if $t \leq t' \Rightarrow f(t) \leq f(t')$, and, for each $i \in \mathbf{N}$, there exists a strategy $S_i \in \Sigma_i$ such that $\Delta_i^{f,M}(S_i) = \int_{f^{-1}[\{i\}]} dx$. Informally, a weak one-round walk is simple if all players select their strategy according to the ordering by which their types are defined in \mathbf{N} , and if players of the same type play the same strategy. A simple weak one-round walk can be univocally represented as a sequence of strategies $(S_i)_{i \in \mathbf{N}}$ such that S_i is the strategy played by a player of type i . The strategy profile generated by a weak one-round walk is $\Delta^{f,M}$.

A weak one-round walk $(\Delta^{f,t})_{t \in [0,M]}$ is an ϵ -approximate non-atomic one-round walk starting from the empty state (ϵ -one-round walk, for brevity) if, for any $t \in [0, M]$, $\Delta_{f(t)}^{f,t}(S)$ is right-increasing at t only if strategy S is an ϵ -best-response in $\Delta^{f,t}$ for player $f(t)$. Informally, an ϵ -one-round walk is a weak one-round walk in which all players sequentially select an ϵ -best-response.

Efficiency Metrics. A social function that is usually used as a measure of the quality of a strategy profile in non-atomic congestion games, is the *total latency*, defined as $\text{TL}(\Delta) = \sum_{i \in \mathbf{N}, S_i \in \Sigma} \Delta_i(S_i) \cdot c_{S_i}(\Delta) = \sum_{e \in E} k_e(\Delta) \cdot \ell_e(k_e(\Delta))$. A *social optimum* is a strategy profile Δ^* minimizing TL . The *approximation ratio of ϵ -approximate non-atomic one-round walks starting from the empty state* (ϵ -approximation ratio, for brevity) of a congestion game CG (resp. a class of congestion games) with respect to the total latency is the supremum of the ratio $\text{TL}(\Delta)/\text{TL}(\Delta^*)$, where Δ is the strategy profile generated by an ϵ -one-round walk for CG (resp. for some game CG in the class) and Δ^* is a social optimum for CG.

Taxes. A *dynamic tax-function* $T := (T_e)_{e \in E}$ is a class of functions $T_e : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ increasing the latency function perceived by each player. The presence of T determines a new congestion game $\text{CG}(T) = (\mathbf{N}, E, (\ell'_e)_{e \in E}, (r_i)_{i \in \mathbf{N}}, (\Sigma_i)_{i \in \mathbf{N}})$, equal to CG except the latency functions, such that $\ell'_e(x) := \ell_e(x) + T_e(x)$ for each $e \in E$. However, the total latency function is still evaluated with respect to the initial latency functions ℓ_e s. T is a *static tax-function* if each T_e is a constant function (with respect to the resource congestion).

3 Approximation Ratio of One-Round Walks

3.1 Upper Bound

We prove our upper bound on the ϵ -approximation ratio by using the primal-dual method: a technique introduced by Bilò in [5] to prove bounds on the performance guarantee of self-emerging solutions (such as approximate pure Nash equilibria and their generalizations, approximate one-round walks, and so on) in atomic and non-atomic congestion games.

In our setting, we want to establish an upper bound on the worst-case performance guarantee of ϵ -one-round walks with respect to the total latency in the class of polynomial congestion games. For a general, but fixed, congestion game

CG, the method requires the construction of a linear program $\text{LP}(\text{CG})$ based on the following steps: (1) the objective function is $\text{TL}(\Delta)$, where Δ is the strategy profile generated by an arbitrary fixed ϵ -one-round walk for CG, (2) add the constraint $\text{TL}(\Delta^*) = 1$, where Δ^* is an arbitrary fixed social optimal for CG, (3) relate Δ and Δ^* by adding, for each $i \in \mathbb{N}$, suitable constraints that need to be satisfied by the choices performed by all players of type i , (4) the coefficients of the latency functions are treated as variables, while all the other quantities (e.g. resources congestions) are treated as fixed parameters.

By the generality of CG, it follows that the optimal solution of $\text{LP}(\text{CG})$ is an upper bound on the ϵ -approximation ratio in polynomial congestion games. By the weak-duality theorem [9], any feasible solution to the dual formulation of $\text{LP}(\text{CG})$ yields an upper bound on the ϵ -approximation ratio in polynomial congestion games. The more the constraints defined during step (3) provide an accurate characterization of the properties of ϵ -one-round walks, the more the achieved upper bound will be significant (and possibly tight).

Theorem 1. *The approximation ratio of ϵ -approximate one-round walks starting from the empty state in congestion games with polynomial latency functions of degree p is at most $((1 + \epsilon)(p + 1))^{p+1}$.*

Proof. For an integer $p \geq 1$, fix a congestion game CG having polynomial latencies of degree p . Let f be an ordering function such that $(\Delta^{f,t})_{t \in [0, M]}$ is an ϵ -approximate one-round walk starting from the empty state, and let $\Delta := \Delta^{f, M}$. Let Δ^* be an optimal strategy profile, and let $(o(t))_{t \in [0, M]}$ be a family of strategies such that $o(t)$ is right-continuous with respect to t , $o(t)$ is a strategy of player $f(t)$ and $\int_{o^{-1}[\{S\}] \cap f^{-1}[\{i\}]} dt = \Delta_i^*(S)$. Observe that such a family of strategies always exists. For the sake of conciseness, we set $k_e := k_e(\Delta)$ and $o_e := k_e(\Delta^*)$. By applying the primal-dual method, we get the following linear program $\text{LP}(\text{CG})$:

$$\max \underbrace{\sum_{e \in E} \sum_{d=0}^p \alpha_{e,d} k_e^{d+1}}_{\text{TL}(\Delta)} \quad (1)$$

$$s.t. \quad c^\epsilon(S^*, o(t), \Delta^{f,t}) \leq 0, \quad \forall S^* \in \Sigma_{f(t)}^\epsilon(\Delta^{f,t}), \quad \forall t \in [0, M] \quad (2)$$

$$\underbrace{\sum_{e \in E} \sum_{d=0}^p \alpha_{e,d} o_e^{d+1}}_{\text{TL}(\Delta^*)} = 1 \quad (3)$$

$$\alpha_{e,d} \geq 0, \quad \forall e \in E, \forall d \in [p]_0 \quad (4)$$

where $\Sigma_i^\epsilon(\Delta')$ is the set of all ϵ -best-responses of players of type i in a strategy profile Δ' , and $c^\epsilon(S, S', \Delta')$ is defined as follows:

$$c^\epsilon(S, S', \Delta') := \underbrace{\sum_{e \in S} \sum_{d=0}^p \alpha_{e,d} k_e^d(\Delta')}_{c_S(\Delta')} - (1 + \epsilon) \underbrace{\sum_{e \in S'} \sum_{d=0}^p \alpha_{e,d} k_e^d(\Delta')}_{c_{S'}(\Delta')} \quad (5)$$

The constraints (2) come from the definition of ϵ -best-responses and ϵ -one-round walks. By replacing constraint (2) with

$$\hat{c}^\epsilon(S^*, o(t), \Delta^{f,t}) := \sum_{e \in S^*} \sum_{d=0}^p \alpha_{e,d} k_e^d(\Delta^{f,t}) - (1 + \epsilon) \sum_{e \in o(t)} \sum_{d=0}^p \alpha_{e,d} k_e^d(\Delta) \leq 0 \quad (6)$$

we obtain a relaxation of LP(CG). Given $i \in \mathbf{N}$ and $S \in \Sigma_i \setminus \emptyset_i$, $\hat{c}^\epsilon(S, o(t), \Delta^{f,t})$ can be defined as a function of $\Delta_i^{f,t}(S)$ except for the intervals on which $\Delta_i^{f,t}(S)$ is constant. By exploiting (6), we obtain

$$0 \geq \sum_{i \in \mathbf{N}, S \in \Sigma_i} \int_0^{\Delta_i(S)} \hat{c}^\epsilon(S, o(t), \Delta^{f,t}) d(\Delta_i^{f,t}(S)) \quad (7)$$

$$= \sum_{e \in E} \sum_{d=0}^p \alpha_{e,d} \int_0^{k_e} k_e^d(\Delta^{f,t}) d \left(\sum_{i \in \mathbf{N}, S \in \Sigma_i: e \in S} (\Delta_i^{f,t}(S)) \right) \quad (8)$$

$$- (1 + \epsilon) \int_0^M \left(\sum_{e \in o(t)} \sum_{d=0}^p \alpha_{e,d} k_e^d \right) dt \quad (9)$$

$$= \sum_{e \in E} \sum_{d=0}^p \alpha_{e,d} \int_0^{k_e} k_e^d(\Delta^{f,t}) d(k_e(\Delta^{f,t})) - (1 + \epsilon) \sum_{e \in E} \int_{\{t: e \in o(t)\}} \left(\sum_{d=0}^p \alpha_{e,d} k_e^d \right) dt \quad (10)$$

$$= \sum_{e \in E} \sum_{d=0}^p \alpha_{e,d} \int_0^{k_e} k_e^d dk - (1 + \epsilon) \sum_{e \in E} \int_0^{o_e} \left(\sum_{d=0}^p \alpha_{e,d} k_e^d \right) dk \quad (11)$$

$$= \sum_{e \in E} \sum_{d=0}^p \alpha_{e,d} \frac{k_e^{d+1}}{d+1} - (1 + \epsilon) \sum_{e \in E} o_e \sum_{d=0}^p \alpha_{e,d} k_e^d \quad (12)$$

The equivalences between all pairs of integrals in the above derivation (along with the fact that they are well-defined) come from the properties of the $\Delta_i^{f,t}(S)$ s as functions of t , and, more generally, from the Lebesgue theory of integration [28]. In conclusion, by replacing the left-hand side of (2) with (12), we obtain the following relaxation of LP(CG):

$$\max \sum_{e \in E} \sum_{d=0}^p \alpha_{e,d} k_e^{d+1} \quad (13)$$

$$s.t. \sum_{e \in E} \sum_{d=0}^p \alpha_{e,d} \frac{k_e^{d+1}}{d+1} - (1 + \epsilon) \sum_{e \in E} o_e \sum_{d=0}^p \alpha_{e,d} k_e^d \leq 0 \quad (14)$$

$$\sum_{e \in E} \sum_{d=0}^p \alpha_{e,d} o_e^{d+1} = 1 \quad (15)$$

$$\alpha_{e,d} \geq 0, \quad \forall e \in E, \forall d \in [p]_0 \quad (16)$$

The dual of this problem is the following linear problem $\text{LP}'(\text{CG})$:

$$\min \quad \gamma \tag{17}$$

$$s.t. \quad x \left(\frac{k_e^{d+1}}{d+1} - (1+\epsilon)o_e k_e^d \right) + \gamma o_e^{d+1} \geq k_e^{d+1} \quad \forall e \in E, d \in [p]_0 \tag{18}$$

$$x, \gamma \geq 0 \tag{19}$$

Clearly, any feasible solution for $\text{LP}'(\text{CG})$ is an upper bound on the ϵ -approximation ratio. We show that, by setting $\gamma := ((1+\epsilon)(p+1))^{p+1}$ and $x := (p+1)^2$, the constraints (18) are always satisfied. Indeed, if $o_e = 0$, the constraints (18) are satisfied. Conversely, if $o_e > 0$, by rewriting the constraints (18) in terms of the quantity $t_e := k_e/o_e$, we get the following inequality:

$$g(t_e, d) := -t_e^{d+1} + (p+1)^2 \left(\frac{t_e^{d+1}}{d+1} - (1+\epsilon)t_e^d \right) + ((1+\epsilon)(p+1))^{p+1} \geq 0 \tag{20}$$

for each $e \in E$, $d \in [p]_0$ and $t_e \geq 0$. This inequality is always satisfied for each $d \in [p]_0$ and $t_e \geq 0$ (see the full version), and this fact completes the proof of the theorem. \square

3.2 Lower Bound

Theorem 2. *For each $\delta > 0$ there exists a congestion game having polynomial latency functions of degree p such that the approximation ratio of ϵ -approximate one-round walks starting from the empty state is higher than $((1+\epsilon)(p+1))^{p+1} - \delta$.*

Proof. Let $n, m \in \mathbb{N}$ and $q := \lfloor (1+\epsilon)(p+1)n \rfloor - 1$. Let $\text{CG}_{m,n}$ be a congestion game such that the set of players' types is $\mathbf{N} := \{a_1^* \dots, a_{q+1}^*, a'_1, a_1, \dots, a'_m, a_m\}$, the set of resources is $E := \{e_1, \dots, e_{n+m+q-1}\}$, the amount of players of type a_1^* and types a'_i s is 2, while the amount of players of the remaining types is 1. Players of type a_i^* can only select the strategy $s_i^* := \{e_1, e_2, \dots, e_{n+i-2}\}$, players of type a'_i can only select the strategy $s'_i := \{e_{n+i-1}\}$, and players of type a_i can select $s_i := \{e_{n+i}, e_{n+i+1}, \dots, e_{n+i+q-1}\}$ or $o_i := \{e_i, e_{i+1}, \dots, e_{n+i-1}\}$. The latency function is $\ell_e(k_e) := k_e^p/n^{p+1}$ for each resource $e \in E$.

Consider the simple weak one-round walk

$$(\Delta^{f,t})_{t \in [0, M]} := (s_1^* \dots, s_{q+1}^*, s'_1, s_1, \dots, s'_m, s_m).$$

To prove that $(\Delta^{f,t})_{t \in [0, M]}$ is an ϵ -one-round walk it is sufficient to show that s_i is an ϵ -best-response in $\Delta^{f,t}$ for each t such that $f(t) = a_i$. Given such a t , we get

$$c_{s_i}(\Delta^{f,t}) \leq \sum_{j=1}^q \frac{j^p}{n^{p+1}} \leq \int_0^{(1+\epsilon)(p+1)n} \frac{1}{n} \left(\frac{x}{n} \right)^p dx = \int_0^{(1+\epsilon)(p+1)} y^p dy \tag{21}$$

$$= (1+\epsilon)^{p+1} (p+1)^p \leq (1+\epsilon) \left(\frac{q+2}{n} \right)^p \tag{22}$$

$$= (1+\epsilon)n \frac{1}{n} \left(\frac{q+2}{n} \right)^p = (1+\epsilon)c_{o_i}(\Delta^{f,t}) \tag{23}$$

We conclude that s_i is an ϵ -best-response in the strategy profile $\Delta^{f,t}$ for a player of type $a_i = f_1(t)$. Therefore, $(\Delta^{f,t})_{t \in [0, M]}$ is an ϵ -one-round walk, and generates a strategy profile $\Delta_{m,n}$ in which all players of type a_i select the strategy s_i . Let $\Delta_{m,n}^*$ be the strategy profile in which each player of type a_i selects the strategy o_i .




Fig. 1. The figure depicts a lower bounding instance with $n = 4$, $m = 6$, $\ell(k_e) = k_e/4$ and $\epsilon = 0.4$. Observe that $q = \lfloor 1.4 \cdot 2 \cdot 4 \rfloor - 1 = 10$ and the total amount of players is $M = 3m + 2 + q = 30$. The abscissa represents the amount of players which already selected their strategy, and the ordinate represents the resources. The ordinate values of coloured squares having abscissa $x \in [30]$ denote the strategies of all players of type $f(t)$ with $t \in [x-1, x)$. The squares labeled with a cross/a black square/a dark grey square/a tic are related to the resources selected in the the strategies $s_i^*/s'_i/o_i/s_i$. For instance, when an amount $t_1/t_2/t_3$ of players have already entered in the game, a player of type $a_4^*/a_2/a_4'$ can select the strategy $s_4^* = \{e_1, e_2, \dots, e_6\}$ (crosses)/ $s_2 = \{e_6, e_7, \dots, e_{15}\}$ (tics) or $o_2 = \{e_2, e_3, \dots, e_5\}$ (dark grey squares)/ $s'_4 = \{e_7\}$ (black squares).

We get ($o(m)$ is related to the usual asymptotic notation):

$$\lim_{n \rightarrow \infty} \lim_{m \rightarrow \infty} \frac{\text{TL}(\Delta_{m,n})}{\text{TL}(\Delta_{m,n}^*)} \geq \lim_{n \rightarrow \infty} \lim_{m \rightarrow \infty} \frac{q \frac{1}{n} \left(\frac{q}{n}\right)^p (m - o(m)) + o(m)}{\left(\frac{n+2}{n}\right)^{p+1} (m - o(m)) + o(m)} \quad (24)$$

$$= \lim_{n \rightarrow \infty} \left(\frac{q}{n}\right)^{p+1} = \lim_{n \rightarrow \infty} \left(\frac{\lfloor (1+\epsilon)(p+1)n \rfloor - 1}{n}\right)^{p+1} \quad (25)$$

$$= ((1+\epsilon)(p+1))^{p+1}. \quad (26)$$

which completes the proof. \square

4 Efficiency of Taxation

In the following theorems, we prove that by resorting to static or dynamic taxation the ϵ -approximation ratio can be consistently lowered.

Theorem 3. *The approximation ratio of ϵ -approximate one-round walks starting from the empty state in congestion games with polynomial latency functions of degree p is at most $(1 + \epsilon)^{p+1}(p + 1)^p$ by resorting to the static tax-function:*

$$T_e := \alpha_{e,p} o_e^p \xi, \quad \text{with} \quad \xi := \frac{(1 + \epsilon)^p (p + 1)^{p-1}}{p}. \quad (27)$$

Proof. For an integer $p \geq 1$, fix a congestion game $\text{CG}(T)$, where T is the tax-function defined in (27). Define $(\Delta^{f,t})_{t \in [0,M]}$, $o(t)$, Δ , Δ^* , k_e and o_e as in the proof of Theorem 1. As done in the proof of Theorem 1, by applying the primal-dual method and by relaxing the constraints modelling ϵ -best-responses (which take into account taxation in this case), we get the following linear program:

$$\max \sum_{e \in E} \sum_{d=0}^p \alpha_{e,d} k_e^{d+1} \quad (28)$$

$$\text{s.t.} \quad \sum_{e \in E} \left(\sum_{d=0}^p \alpha_{e,d} \left(\frac{k_e^{d+1}}{d+1} - (1 + \epsilon) o_e k_e^d \right) + \alpha_{e,p} (o_e^p \xi k_e - o_e^{p+1} \xi) \right) \leq 0 \quad (29)$$

$$\sum_{e \in E} \sum_{d=0}^p \alpha_{e,d} o_e^{d+1} = 1 \quad (30)$$

$$\alpha_{e,d} \geq 0, \quad \forall e \in E, \forall d \in [p]_0 \quad (31)$$

The dual program is

$$\min \quad \gamma \quad (32)$$

$$\text{s.t.} \quad x \left(\frac{k_e^{d+1}}{d+1} - (1 + \epsilon) o_e k_e^d \right) + \gamma o_e^{d+1} \geq k_e^{d+1} \quad \forall e \in E, d \in [p-1]_0 \quad (33)$$

$$x \left(\frac{k_e^{p+1}}{p+1} + \xi o_e^p k_e - (1 + \epsilon) o_e (k_e^p + \xi o_e^p) \right) + \gamma o_e^{p+1} \geq k_e^{p+1} \quad \forall e \in E \quad (34)$$

$$x, \gamma \geq 0 \quad (35)$$

By setting $\gamma := (1 + \epsilon)^{p+1}(p + 1)^p$ and $x := (p + 1)p$, the constraints (33) and (34) are satisfied (see the full version), and the claim follows. \square

Theorem 4. *The approximation ratio of ϵ -approximate one-round walks starting from the empty state in congestion games with polynomial latency functions of degree p is at most $(1 + \epsilon)^{p+1}(p + 1)!$ by resorting to the dynamic tax-function (set $(d)_j := d!/(d-j)!$):*

$$T_e(k_e) := \sum_{d=0}^p \alpha_{e,d} T_{e,d}(k_e) \quad \text{with} \quad T_{e,d}(k_e) := \sum_{j=1}^d (1 + \epsilon)^j (d)_j k_e^{d-j} o_e^j \quad (36)$$

Proof. By reconsidering the proof of Theorem 3, but with the tax (36) instead of the tax (27), we get a dual program having the following constraints:

$$-k_e^{d+1} + x \left(\frac{k_e^{d+1}}{d+1} + \int_0^{k_e} T_{e,d}(u) du - (1+\epsilon)o_e(k_e^d + T_{e,d}(k)) \right) + \gamma o_e^{d+1} \geq 0 \quad (37)$$

for each $e \in E$ and $d \in [p]_0$. These constraints are always satisfied if $x := p+1$ and $\gamma := (1+\epsilon)^{p+1}(p+1)!$ (see the full version), and the claim follows. \square

Remark 1. Observe that finding an optimal strategy profile requires to solve a convex minimization problem. Therefore, for each $\delta > 0$, one can compute in polynomial time a strategy profile whose social value is at most $1+\delta$ times the social optimum [9]. Then, if \tilde{o}_e is the congestion of resource e in that strategy profile, by using the tax (27) (resp. (36)) with \tilde{o}_e in place of o_e , one can easily prove that the resulting ϵ -approximation ratio is at most $(1+\delta)$ times that of Theorem 3 (resp. Theorem 4).

References

1. Aland, S., Dumrauf, D., Gairing, M., Monien, B., Schoppmann, F.: Exact price of anarchy for polynomial congestion games. *SIAM J. Comput.* 40(5), 1211–1233 (2011)
2. Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, E., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. In: *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*. pp. 295–304. FOCS, IEEE Computer Society (2004)
3. Awerbuch, B., Azar, Y., Epstein, A.: The Price of Routing Unsplittable Flow. In: *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing*. pp. 57–66. STOC, ACM (2005)
4. Beckmann, M.: *Studies in the economics of transportation*. Yale University Press for the Cowles Commission for Research in Economics New Haven (1959)
5. Bilò, V.: A Unifying tool for Bounding the quality of Non-Cooperative Solutions in Weighted Congestion Games. *Approximation and Online Algorithms: 10th International Workshop, WAOA 2012* pp. 215–228 (2013)
6. Bilò, V., Fanelli, A., Flammini, M., Moscardelli, L.: Performance of one-round walks in linear congestion games. *Theor. Comp. Sys.* 49(1), 24–45 (2011)
7. Bilò, V., Vinci, C.: On stackelberg strategies in affine congestion games. In: *Web and Internet Economics: 11th International Conference, WINE 2015*. pp. 132–145. Springer Berlin Heidelberg (2015)
8. Bilò, V., Vinci, C.: Dynamic taxes for polynomial congestion games. In: *Proceedings of the 17th ACM Conference on Electronic Commerce. EC (2016)*, To appear.
9. Boyd, S., Vandenberghe, L.: *Convex Optimization*. Cambridge University Press, New York, NY, USA (2004)
10. Caragiannis, I., Flammini, M., Kaklamanis, C., Kanellopoulos, P., Moscardelli, L.: Tight bounds for selfish and greedy load balancing. *Algorithmica* 61(3), 606–637 (2011)
11. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P.: Taxes for Linear Atomic Congestion Games. *ACM Trans. Algorithms* 7(1), 13:1–13:31 (2010)

12. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing. pp. 67–73. STOC '05, ACM (2005)
13. Christodoulou, G., Koutsoupias, E., Spirakis, P.G.: On the performance of approximate equilibria in congestion games. *Algorithmica* 61(1), 116–140 (2011)
14. Christodoulou, G., Mirrokni, V.S., Sidiropoulos, A.: Convergence and approximation in potential games. In: Proceedings of the 23th Annual Conference on Theoretical Aspects of Computer Science. pp. 349–360. STACS, Springer Berlin Heidelberg, Berlin, Heidelberg (2006)
15. Cole, R., Dodis, Y., Roughgarden, T.: How much can taxes help selfish routing? In: Proceedings of the 4th ACM Conference on Electronic Commerce. pp. 98–107. EC, ACM (2003)
16. Cole, R., Dodis, Y., Roughgarden, T.: Pricing Network Edges for Heterogeneous Selfish Users. In: Proceedings of the Thirty-fifth Annual ACM Symposium on Theory of Computing. pp. 521–530. STOC, ACM (2003)
17. Fanelli, A., Flammini, M., Moscardelli, L.: Stackelberg Strategies for Network Design Games. In: Proceedings of the 6th International Conference on Internet and Network Economics. pp. 222–233. WINE, Springer-Verlag (2010)
18. Fanelli, A., Flammini, M., Moscardelli, L.: The speed of convergence in congestion games under best-response dynamics. *ACM Trans. Algorithms* 8(3), 25:1–25:15 (2012)
19. Fleischer, L., Jain, K., Mahdian, M.: Tolls for Heterogeneous Selfish Users in Multicommodity Networks and Generalized Congestion Games. In: Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science. pp. 277–285. FOCS, IEEE Computer Society (2004)
20. Fotakis, D.: Stackelberg Strategies for Atomic Congestion Games. *Theor. Comp. Sys.* 47(1), 218–249 (2010)
21. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: Proceedings of the 16th Annual Conference on Theoretical Aspects of Computer Science. pp. 404–413. STACS, Springer-Verlag (1999)
22. Mirrokni, V.S., Vetta, A.: Convergence issues in competitive games. In: Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, pp. 183–194. Springer Berlin Heidelberg (2004)
23. Nash, J.F.: Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences of the United States of America* 36(1), 48–49 (1950)
24. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V. (eds.): *Algorithmic Game Theory* (2007)
25. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2, 65–67 (1973)
26. Roughgarden, T.: Stackelberg scheduling strategies. In: Proceedings of the Thirty-third Annual ACM Symposium on Theory of Computing. pp. 104–113. STOC '01, ACM (2001)
27. Roughgarden, T.: The price of anarchy is independent of the network topology. *J. Comput. Syst. Sci.* 67(2), 341–364 (2003)
28. Rudin, W.: *Real and Complex Analysis*, 3rd Ed. McGraw-Hill, Inc., New York, NY, USA (1987)
29. Swamy, C.: The Effectiveness of Stackelberg Strategies and Tolls for Network Congestion Games. *ACM Trans. Algorithms* 8(4), 36:1–36:19 (2012)