

A Unifying Tool for Bounding the Quality of Non-Cooperative Solutions in Weighted Congestion Games

Vittorio Bilò¹

Published online: 2 December 2017

© Springer Science+Business Media, LLC, part of Springer Nature 2017

Abstract We present a general technique, based on a primal-dual formulation, for analyzing the quality of self-emerging solutions in weighted congestion games. With respect to traditional combinatorial approaches, the primal-dual schema has at least three advantages: first, it provides an analytic tool which can always be used to prove tight upper bounds for all the cases in which we are able to characterize exactly the polyhedron of the solutions under analysis; secondly, in each such a case, the complementary slackness conditions give us a hint on how to construct matching lower bounding instances; thirdly, proofs become simpler and easy to check. For the sake of exposition, we first apply our technique to the problems of bounding the price of anarchy and stability of exact and approximate pure Nash equilibria, as well as the approximation ratio of the strategy profiles achieved after a one-round walk starting from the empty state, in the case of affine latency functions and we show how all the known upper bounds for these measures (and some of their generalizations) can be easily reobtained under a unified approach. Then, we use the technique to attack the more challenging setting of polynomial latency functions. In particular, we obtain the first known upper bounds on the price of stability of pure Nash equilibria and on the approximation ratio of the strategy profiles achieved after a one-round walk starting from the empty state for unweighted players in the cases of quadratic and cubic latency functions.

A preliminary version of this paper appeared in the Proceedings of the 10th Workshop on Approximation and Online Algorithms (WAOA 2012) [7]. This work was partially supported by the PRIN 2010–2011 research project ARS TechnoMedia: “Algorithmics for Social Technological Networks” funded by the Italian Ministry of University.

✉ Vittorio Bilò
vittorio.bilo@unisalento.it

¹ Department of Mathematics and Physics “Ennio De Giorgi”, University of Salento, Provinciale Lecce-Amesano, P.O. Box 193, 73100, Lecce, Italy

Keywords Price of anarchy and stability · Performance of one-round walks · (Approximate) Nash equilibria · Congestion games · Primal-dual analysis

1 Introduction

Characterizing the quality of self-emerging solutions in non-cooperative systems is one of the leading research directions in Algorithmic Game Theory. Given a game \mathcal{G} , a social function \mathcal{F} measuring the quality of any strategy profile which can be realized in \mathcal{G} , and the definition of a set \mathcal{E} of certain self-emerging solutions, we are asked to bound the ratio $Q(\mathcal{G}, \mathcal{E}, \mathcal{F}) := \mathcal{F}(K)/\mathcal{F}(O)$, where K is some solution in $\mathcal{E}(\mathcal{G})$ (usually either the worst or the best one with respect to \mathcal{F}) and O is the strategy profile optimizing \mathcal{F} .

In such a setting, Roughgarden [40] proposes the so-called “smoothness argument” as a unifying technique for proving tight upper bounds on $Q(\mathcal{G}, \mathcal{E}, \mathcal{F})$ for several notions of self-emerging solutions \mathcal{E} , when \mathcal{G} satisfies some general properties, K is the worst solution in $\mathcal{E}(\mathcal{G})$ and \mathcal{F} is sum-bounded, that is, upper bounded by the sum of the players’ payoffs. He also gives a more refined interpretation of this argument and stresses also its intrinsic limitations, in a subsequent work with Nadav [37], by means of a primal-dual characterization which shares lot of similarities with the primal-dual framework we provide in this paper. Anyway, there is a subtle, yet substantial, difference between the two approaches and we believe that the one we propose is more general and powerful. Both techniques formulate the problem of bounding $Q(\mathcal{G}, \mathcal{E}, \mathcal{F})$ via a (primal) linear program and, then, an upper bound is achieved by providing a feasible solution for the related dual program. But, while in [37] the variables defining the primal formulation are yielded by the strategic choices of the players in both K and O (as one would expect), in our technique the variables are the parameters defining the players’ payoffs in \mathcal{G} , while K and O play the role of fixed, but arbitrary, constants.

As it will be clarified later, such an approach, although preserving the same degree of generality, applies to a broader class of games and allows for a simple analysis facilitating the proof of tight results. In fact, as already pointed out in [37], the Strong Duality Theorem assures that each primal-dual framework can always be used to derive the exact value of $Q(\mathcal{G}, \mathcal{E}, \mathcal{F})$ provided that, for any strategy profile S which can be realized in \mathcal{G} , $\mathcal{F}(S)$ can be expressed through linear programming and

- (i) the polyhedron defining $\mathcal{E}(\mathcal{G})$ can be expressed through linear programming, when K is the worst solution in $\mathcal{E}(\mathcal{G})$ with respect to \mathcal{F} ,
- (ii) the polyhedron defining K can be expressed through linear programming, when K is the best solution in $\mathcal{E}(\mathcal{G})$ with respect to \mathcal{F} .

Moreover, in all such cases, by applying the “complementary slackness conditions”, we can figure out which pairs of strategy profiles (K, O) yield the exact value of $Q(\mathcal{G}, \mathcal{E}, \mathcal{F})$, thus being able to construct quite systematically matching lower bounding instances.

In this work, we consider three sets of solutions \mathcal{E} (see Section 2 for their formal definitions), namely,

- (i) $(1 + \epsilon)$ -approximate pure Nash equilibria (PNE_ϵ), that is, outcomes in which no player can improve her situation by a factor of more than ϵ by unilaterally changing the adopted strategy (i.e., no player possesses an ϵ -approximate improving deviation). In this case, $Q(\mathcal{G}, \mathcal{E}, \mathcal{F})$ is called the ϵ -approximate price of anarchy of \mathcal{G} , denoted as $\text{PoA}_\epsilon(\mathcal{G})$, when K is the worst strategy profile in $\mathcal{E}(\mathcal{G})$, while it is called the ϵ -approximate price of stability of \mathcal{G} , denoted as $\text{PoS}_\epsilon(\mathcal{G})$, when K is the best strategy profile in $\mathcal{E}(\mathcal{G})$;
- (ii) pure Nash equilibria (PNE), that is, the set of outcomes in which no player can improve her situation by unilaterally changing the adopted strategy (i.e., no player possesses an improving deviation). By definition, each PNE_0 is a PNE and the terms price of anarchy ($\text{PoA}(\mathcal{G})$) and price of stability ($\text{PoS}(\mathcal{G})$) are used in this case;
- (iii) strategy profiles achieved after a one-round walk starting from the empty state [36], that is, the set of outcomes which arise when, starting from an initial configuration in which no player has done any strategic choice yet, each player is asked to select, sequentially and according to a given ordering, her best possible current strategy. In this case, K is always defined as the worst strategy profile in $\mathcal{E}(\mathcal{G})$ and $Q(\mathcal{G}, \mathcal{E}, \mathcal{F})$ is denoted as $\text{Apx}_\emptyset^1(\mathcal{G})$.

We observe that approximate pure Nash equilibria (resp. pure Nash equilibria) can also be intended as the output of an algorithm which repeatedly allows players to perform approximate improving deviations (resp. improving deviations) until a steady state is finally reached. Similarly, strategy profiles achieved after a one-round walk starting from the empty state can be seen as the output of a greedy-like online algorithm which aims at minimizing \mathcal{F} by assigning to each player, upon her arrival, the strategy currently minimizing her cost. These relationships are at the basis of our approach: to exploit linear programming, and in particular the duality theory, to bound the efficiency of self-emerging solutions, in the same spirit as for the approximation ratio of polynomial time algorithms and for the competitive ratio of online algorithms.

1.1 Our Contribution

Our method reveals to be particularly powerful when applied to the class of weighted congestion games. In these games, there are n weighted players competing for a set of resources. Each resource has a latency function which only depends on its congestion, i.e., the sum of the weights of its users. These games have a particular appeal since, from the one hand, they are general enough to model a variety of situations arising in real life applications and, from the other one, they are structured enough to allow a systematic theoretical study. For example, for the case in which all players have the same weight (unweighted players), Rosenthal [41] proved through a potential function argument that PNE are always guaranteed to exist, while, in general,

weighted congestion games are guaranteed to possess PNE if and only if the latency functions are either affine or exponential [29–31, 39].

In order to illustrate the versatility and usefulness of our technique, we first consider the well-known and studied case in which the latency functions are affine and \mathcal{F} is the sum of the players' payoffs and show how all the known results (as well as some of their generalizations) can be easily reobtained under a unifying approach. For the PoA_ϵ and the PoS_ϵ in the unweighted case and for the Apx_θ^1 in the weighted case, we reobtain the known upper bounds given in [4, 23, 25, 26] with significantly shorter and simpler proofs (where, by simple, we mean that only basic notions of calculus are needed in the arguments), while for the generalizations of the PoA_ϵ and the PoS_ϵ in the weighted case, we give the first upper bounds known in the literature.

After having introduced the technique, we show how it can be used to attack the more challenging case of polynomial latency functions. In such a scenario, the PoA and the PoA_ϵ were already studied and characterized in [1] and [25], respectively, and both papers pose the achievement of upper bounds on the PoS and the PoS_ϵ as a major open problem in the area. For unweighted players, we show that, for any congestion game \mathcal{G} with quadratic latency functions, $2.1859 \leq \text{PoS}(\mathcal{G}) \leq 2.362$ and $\text{Apx}_\theta^1(\mathcal{G}) \leq 37.5888$ and that, for any congestion game \mathcal{G} with cubic latency functions, $2.7558 \leq \text{PoS}(\mathcal{G}) \leq 3.322$ and $\text{Apx}_\theta^1(\mathcal{G}) \leq 527.323$. Recently, and after the publication of the conference version of this paper, Christodoulou and Gairing [22] gave an exact characterization of the PoS for polynomial latency functions of maximum degree equal to d . In particular, they showed that the two upper bounds we give in this paper for the cases of $d = 2, 3$ are tight.

What we would like to stress here is that, rather than the novelty of the results achieved in this paper, what makes our method interesting is its capability of being easily adapted to a variety of particular situations and we are more than sure of the fact that it will prove to be a powerful tool to be exploited in the analysis of the efficiency achieved by different classes of self-emerging solutions in other contexts as well (see, Section 1.3 for a discussion on some recent application of our technique). To this aim, we show how the method applies also to other social functions, such as the maximum of the players' payoffs, where the smoothness argument cannot even be applied. Moreover, in the case of polynomial latency functions, the primal-dual technique proposed in [37] cannot be used, since the players' costs are not linear in the variables of the problem.

Although most of the times our method may look like a reinterpretation of known techniques (as the use of Nash inequalities and potential inequalities), there is a feature that makes it preferable to classical combinatorial approaches: once the properties possessed by the target solution are suitably modeled as linear constraints and embedded within the formulation, then, by computing its optimal solution, we are guaranteed to determine the best possible upper bound on the worst-case efficiency of the target solution which can be obtained by making use of these properties. Put in other words, we do not need to “optimize” the proof of the upper bound: the method does it all for us, all we need to do is to find the optimal solution of the formulation. Not by chance, in fact, after the publication of the conference version of this paper [7], the primal-dual method has been exploited to provide the exact solution to some open problems in this research field (see, Section 1.3 for a detailed discussion).

1.2 Related Work

The study of the quality of self-emerging solutions in non-cooperative systems initiated with the seminal papers of Koutsoupias and Papadimitriou [32] and Anshelevich et al. [2] which introduced, respectively, the notions of price of anarchy and price of stability.

A lot of results have been achieved since then and we recall here only the ones which are closely related to our scenario of application, that is, weighted congestion games with polynomial latency functions.

For affine latency functions and \mathcal{F} defined as the sum of the players’ payoffs, Christodoulou and Koutsoupias [23] show that the PoA is exactly $5/2$ for unweighted players, while Awerbuch, Azar and Epstein [4] show that it rises to exactly $(3 + \sqrt{5})/2$ in the weighted case. These bounds keep holding also when considering the price of anarchy of generalizations of PNE such as mixed Nash equilibria and correlated equilibria, as shown by Christodoulou and Koutsoupias in [24]. Similarly, for polynomial latency functions with maximum degree equal to d , Aland et al. [1] prove that the price of anarchy of all these equilibria is exactly Φ_d^{d+1} in the weighted case and exactly $\frac{(k+1)^{2d+1} - k^{d+1}(k+2)^d}{(k+1)^{d+1} - (k+2)^d + (k+1)^d - k^{d+1}}$ in the unweighted case, where Φ_d is the unique non-negative real solution to $(x + 1)^d = x^{d+1}$ and $k = \lfloor \Phi_d \rfloor$. These interdependencies have been analyzed by Roughgarden [40], who proves that unweighted congestion games with non-negative and non-decreasing latency functions belong to the class of games for which a so-called “smoothness argument” applies and that such a smoothness argument directly implies that, when \mathcal{F} is the sum of the players’ payoffs, the price of anarchy stays the same, independently of which solution concept among PNE, mixed Nash equilibria, correlated equilibria and coarse correlated equilibria is adopted. Such a result has been extended also to the weighted case by Bhawalkar, Gairing and Roughgarden in [6]. For the alternative model in which \mathcal{F} is defined as the maximum of the players’ payoffs, Christodoulou and Koutsoupias [23] show a PoA of $\Theta(\sqrt{n})$ in the case of affine latency functions.

For the PoS, only the case of unweighted players, affine latency functions and \mathcal{F} defined as the sum of the players’ payoffs, has been considered so far. The upper and lower bounds achieved by Caragiannis et al. [20] and by Christodoulou and Koutsoupias [24], respectively, set the PoS to exactly $1 + 1/\sqrt{3}$.

As to the PNE_ϵ , in the case of unweighted players, polynomial latency functions and \mathcal{F} defined as the sum of the players’ payoffs, Christodoulou, Koutsoupias and Spirakis [25] show that the PoA_ϵ is exactly $\frac{(1+\epsilon)((z+1)^{2d+1} - z^{d+1}(z+2)^d)}{(z+1)^{d+1} - z^{d+1} - (1+\epsilon)((z+2)^d - (z+1)^d)}$, where z is the maximum integer satisfying $\frac{z^{d+1}}{(z+1)^d} < 1 + \epsilon$, and that, for affine latency functions, the PoS_ϵ is at least $\frac{2(3+\epsilon+\theta\epsilon^2+3\epsilon^3+2\epsilon^4+\theta\epsilon)}{6+2\epsilon+5\theta\epsilon+6\epsilon^3+4\epsilon^4-\theta\epsilon^3+2\theta\epsilon^2}$, where $\theta = \sqrt{3\epsilon^3 + 3 + \epsilon + 2\epsilon^4}$, and at most $(1 + \sqrt{3})/(\epsilon + \sqrt{3})$.

Finally, for affine latency functions and \mathcal{F} defined as the sum of the players’ payoffs, the ApX_\emptyset^1 has been shown to be exactly $2 + \sqrt{5}$ in the unweighted case as a consequence of the upper and lower bounds provided, respectively, by Christodoulou, Mirrokni and Sidiropoulos [26] and by Bilò et al. [10], while, for weighted players, Caragiannis et al. [20] give a lower bound of $3 + 2\sqrt{2}$ and Christodoulou, Mirrokni

and Sidiropoulos [26] give an upper bound of $4 + 2\sqrt{3}$. For \mathcal{F} being the maximum of the players' payoffs, Bilò et al. [10] show that the Apx_\emptyset^1 is $\Theta(\sqrt[4]{n^3})$ in the unweighted case and affine latency functions.

Strategy profiles realized after a one-round walk starting from the empty state can be reinterpreted as the output of an online greedy algorithm for the problem of minimizing the social function \mathcal{F} . Under this perspective, the bounds given by Awerbuch et al. [5], Caragiannis et al. [20] and Suri, Tóth, and Zhou [42] on the competitive ratio of greedy algorithms for the problem of minimizing the total latency in online load balancing can be seen as bounds for the Apx_\emptyset^1 in the special case of affine congestion games in which the strategies of all players are singleton sets.

The use of primal-dual analysis, and in particular of factor-revealing linear programs, for bounding the performance of greedy-like algorithms has been introduced in the seminal paper by Jain et al. [33] and later exploited by Athanassopoulos et al. [3] for variants of set cover, by Mahdian and Yan [34] for online matching, by Feige and Jozeph [27] for maximum directed cut, and by Caragiannis [17] for wavelength routing.

1.3 Subsequent Work

After the publication of the conference version of this paper [7], the primal-dual method has been fruitfully used in a variety of situations.

Bilò, Flammini and Gallotti [12] consider congestion games with affine latency functions under the assumption that the players' knowledge is restricted by the presence of an underlying social knowledge graph. In particular, each player is only aware of her choice and of the choices of all the players representing her neighborhood in the social knowledge graph, whereas in the traditional model each player has full knowledge of the strategic choices of all the other players in the game. By exploiting the primal-dual method, tight bounds on the price of anarchy and almost tight bounds on the price of stability are determined for the case in which the social knowledge graph has a special topology allowing for a “bidimensional” reinterpretation of congestion games with restricted social knowledge.

Bilò et al. [13] exploit the primal-dual method to extrapolate a matching (and unexpected) lower bounding instance for the sequential price of anarchy, that is the price of anarchy of subgame perfect equilibria in sequential games [38], of cut games. This contribution, in particular, illustrates an application of our method to a class of self-emerging solutions not considered in this paper.

The primal-dual method has been used by Bilò [8] to derive tight bounds on the worst-case price of stability of pure Nash equilibria in congestion games with affine latency functions and altruistic players, thus solving an open problem risen by Caragiannis et al. [19]; by Bilò and Paladini [14] to derive tight bounds on the approximation ratio of the strategy profiles achieved after a one-round walk of $(1+\epsilon)$ -approximate best-responses starting from any initial strategy profile in cut games, for any $\epsilon \geq 0$, thus solving an open problem left by Christodoulou, Mirrokni and Sidiropoulos [26]; by Bilò, Fanelli and Moscardelli [11] to derive significant upper

bounds on the price of anarchy of lookahead equilibria in congestion games with affine latency functions, thus improving on a previous result by Mirrokni, Thain and Vetta [35]; and by Bilò and Vinci [15] to obtain exact bounds on the price of anarchy of Stackelberg strategies in congestion games with affine latency functions, thus solving an open problem left by Fotakis [28].

Recently, building on the primal-dual method, Bilò and Vinci [16] design efficient tax functions for weighted congestion games with polynomial latency functions, thus generalizing and extending previous results achieved by Caragiannis et al. [18].

Moreover, very recently, Bilò [9] used the primal-dual method to show that, for a variety of (even non-sum-bounded) social functions and for a broad generalization of the class of weighted congestion games with non-negative (and possibly decreasing) latency functions, the worst-case price of anarchy of $(1 + \epsilon)$ -approximate coarse correlated equilibria still coincides with that of $(1 + \epsilon)$ -approximate pure Nash equilibria, for any $\epsilon \geq 0$, thus significantly generalizing the tightness result obtained by Bhawalkar, Gairing and Roughgarden in [6] by making use of the smoothness argument (delving into a detailed discussion about differences and similarities between the smoothness argument and the primal-dual method is, in our opinion, beyond the scope of this paper; however, the interested reader can fully satisfy his/her curiosity by looking at [9]).

1.4 Paper Organization

In the next section, we give all the necessary definitions and notation, while, in Section 3, we briefly outline the primal-dual method. Then, in Section 4, we illustrate how it applies to affine latency functions and, in Section 5, we use it to address the case of quadratic and cubic latency functions. In Section 6, we show how the method applies to the social function defined as the maximum of the players' payoffs. Finally, in the last section, we discuss further research and open problems.

2 Definitions

For a given integer $n > 0$, we denote as $[n]$ the set $\{1, \dots, n\}$.

A *weighted congestion game* $\mathcal{G} = ([n], E, (\Sigma_i)_{i \in [n]}, (\ell_e)_{e \in E}, (w_i)_{i \in [n]})$ is a non-cooperative strategic game in which there is a set E of m resources to be shared among the players in $[n]$, where $n \geq 2$. Each player i has an associated *weight* $w_i \geq 0$ and the special case in which $w_i = 1$ for any $i \in [n]$ is called the *unweighted case*. The strategy set Σ_i , for any player $i \in [n]$, is a non-empty subset of resources, i.e., $\Sigma_i \subseteq 2^E \setminus \{\emptyset\}$. The set $\Sigma = \times_{i \in [n]} \Sigma_i$ is called the set of *strategy profiles* (or *strategy profiles*) which can be realized in \mathcal{G} . Given a strategy profile $s = (s_1, s_2, \dots, s_n) \in \Sigma$ and a resource $e \in E$, the sum of the weights of all the players using e in s , called the *congestion* of e in s , is denoted by $\mathcal{L}_e(s) = \sum_{i \in [n]: e \in s_i} w_i$. A *latency function* $\ell_e : \mathbb{R}_{\geq 0} \mapsto \mathbb{R}_{\geq 0}$ associates each resource $e \in E$ with a *latency* depending on the congestion of e in s . The *cost* of player i in the strategy profile s is given by $c_i(s) = \sum_{e \in s_i} \ell_e(\mathcal{L}_e(s))$. This work is concerned only with *polynomial latency functions* of

maximum degree d , i.e., the case in which $\ell_e(x) = \sum_{i=0}^d \alpha_{e,i} x^i$ with $\alpha_{e,i} \in \mathbb{R}_{\geq 0}$, for any $e \in E$ and $0 \leq i \leq d$. We denote as \mathcal{WCG} and \mathcal{UCG} the class of weighted and unweighted congestion games, respectively.

Given a strategy profile $s \in \Sigma$ and a strategy $t \in \Sigma_i$ for player i , we denote with $s_{-i} \diamond t = (s_1, \dots, s_{i-1}, t, s_{i+1}, \dots, s_n)$ the strategy profile obtained from s when player i changes unilaterally her strategy from s_i to t . A strategy $t \in \Sigma_i$ is an *improving deviation* for player i in s , if $c_i(s_{-i} \diamond t) < c_i(s)$.

Definition 1 Given an $\epsilon \geq 0$, a strategy profile s is a $(1 + \epsilon)$ -*approximate pure Nash equilibrium* (NE_ϵ) if, for any $i \in [n]$ and $t \in \Sigma_i$, $c_i(s) \leq (1 + \epsilon)c_i(s_{-i} \diamond t)$.

Note that the set of 1-approximate pure Nash equilibria collapses to that of *pure Nash equilibria* (PNE), that is, the set of strategy profiles in which no player possesses an improving deviation.

Consider the social function $\text{SUM} : \Sigma \mapsto \mathbb{R}_{\geq 0}$ defined as the sum of the players' costs, that is, $\text{SUM}(s) = \sum_{i \in [n]} c_i(s)$ and let s^* be the strategy profile minimizing it. Given an $\epsilon \geq 0$ and a weighted congestion game \mathcal{G} , let $\mathcal{E}_\epsilon(\mathcal{G})$ be the set of $(1 + \epsilon)$ -approximate Nash equilibria of \mathcal{G} .

Definition 2 The ϵ -*approximate price of anarchy* of \mathcal{G} is defined as $\text{PoA}_\epsilon(\mathcal{G}) = \max_{s \in \mathcal{E}_\epsilon(\mathcal{G})} \left\{ \frac{\text{SUM}(s)}{\text{SUM}(s^*)} \right\}$, while the ϵ -*approximate price of stability* of \mathcal{G} is defined as $\text{PoS}_\epsilon(\mathcal{G}) = \min_{s \in \mathcal{E}_\epsilon(\mathcal{G})} \left\{ \frac{\text{SUM}(s)}{\text{SUM}(s^*)} \right\}$.

Definition 3 Given a strategy profile s and a player $i \in [n]$, a strategy $t^* \in \Sigma_i$ is a *best-response* for player i in s if $c_i(s_{-i} \diamond t^*) \leq c_i(s_{-i} \diamond t)$ for any $t \in \Sigma_i$.

Let s^\emptyset be the *empty state*, i.e., the profile in which no player has performed any strategic choice yet.

Definition 4 A *one-round walk starting from the empty state* is an $(n + 1)$ -tuple of strategy profiles $W = (s_0^W, s_1^W, \dots, s_n^W)$ such that $s_0^W = s^\emptyset$ and, for any $i \in [n]$, $s_i^W = s_{i-1}^W \diamond t^*$, where t^* is a best-response for player i in s_{i-1}^W . The profile s_n^W is called the strategy profile achieved after the one-round walk W .

Clearly, depending on how the players are ordered from 1 to n and on which best-response is selected at step i when more than one best-response is available to player i in s_{i-1}^W , different one-round walks can be generated. Let $\mathcal{W}(\mathcal{G})$ denote the set of all possible one-round walks which can be generated in game \mathcal{G} .

Definition 5 The *approximation ratio of the strategy profiles achieved after a one-round walk starting from the empty state* in \mathcal{G} is defined as $\text{Apx}_\emptyset^1(\mathcal{G}) = \max_{W \in \mathcal{W}(\mathcal{G})} \left\{ \frac{\text{SUM}(s_n^W)}{\text{SUM}(s^*)} \right\}$.

3 The Primal-Dual Technique

Fix a weighted congestion game \mathcal{G} , a social function \mathcal{F} and a class of self-emerging strategy profiles \mathcal{E} . Let $s^* = (s_1^*, \dots, s_n^*)$ be the strategy profile optimizing \mathcal{F} and $s = (s_1, \dots, s_n) \in \mathcal{E}(\mathcal{G})$ be the worst-case strategy profile in $\mathcal{E}(\mathcal{G})$ with respect to \mathcal{F} . For any $e \in E$, we set, for the sake of brevity, $O_e = \mathcal{L}_e(s^*)$ and $K_e = \mathcal{L}_e(s)$. Note that both O_e and K_e belong to the set of non-negative integers when \mathcal{G} is unweighted, while, otherwise, they belong to the set of non-negative reals.

Since s^* and s are fixed, we can maximize the inefficiency yielded by the pair of profiles (s, s^*) by suitably choosing the coefficients $\alpha_{e,i}$, for each $e \in E$ and $0 \leq i \leq d$, so that $\mathcal{F}(s)$ is maximized, $\mathcal{F}(s)$ is normalized to one¹ and s meets all the constraints defining the set $\mathcal{E}(\mathcal{G})$. For the sets \mathcal{E} and social functions \mathcal{F} considered in this paper, this task can be easily achieved by creating a suitable linear program $\text{LP}(s, s^*)$, (see, for instance, the linear program $\text{LP}(s, s^*)$ defined at page 9 for bounding PoA_ϵ , or the linear program $\text{LP}(s, s^*)$ defined at page 16 for bounding Apx_{\emptyset}^1).

By providing a feasible solution for the dual program $\text{DLP}(s, s^*)$, we can obtain an upper bound on the optimal solution of $\text{LP}(s, s^*)$. Our task is to uncover, among all possibilities, the pair (\bar{s}, \bar{s}^*) yielding the highest possible optimal solution for $\text{LP}(s, s^*)$. To this aim, the study of the dual formulation plays a crucial role. Since the set of variables of $\text{LP}(s, s^*)$ is given by the set of coefficients $\alpha_{e,i}$, for each $e \in E$ and $0 \leq i \leq d$, we shall have a dual constraint for each $e \in E$ and $0 \leq i \leq d$. Any such constraint will depend on the identities of the players using e in both s and s^* (see, for instance, the linear program $\text{DLP}(s, s^*)$ defined at page 10 for bounding PoA_ϵ , or the linear program $\text{DLP}(s, s^*)$ defined at page 16 for bounding Apx_{\emptyset}^1). Since the power set of a set of n elements has cardinality 2^n , and we have two possible subsets of users (the one in s and the one in s^*), for a game with n players, we can have up to 4^n different dual constraint for each fixed $e \in E$ and $0 \leq i \leq d$, some of which might be “harder” to satisfy than the others. Hence, if we are able to detect the nature of the “worst-case” dual constraints, then we can easily figure out the form of the pair (\bar{s}, \bar{s}^*) maximizing the inefficiency of the class of strategy profiles \mathcal{E} . Put in other words, our task is to determine suitable dual variables, in particular the ones minimizing the dual objective function, which are able to satisfy each of the infinitely many dual constraints (obtained by considering all possible values for the number of players n) that can be generated by our formulation, when varying the pair (s, s^*) . The worst-case dual constraints we are looking for, then, will be the ones which are tight, i.e., satisfied at equality.

Clearly, by the complementary slackness conditions, if we find the optimal dual solution, then we can quite systematically construct the matching primal instance by choosing a suitable set of players and resources so as to implement all the tight dual constraints. This task is much more complicated to be achieved in the weighted case,

¹It is not difficult to see that such an assumption is without loss of generality, since, given any weighted congestion game \mathcal{G} , one can always scale the coefficients $\alpha_{e,i}$ so as to obtain another game \mathcal{G}' possessing the same set of strategy profiles of \mathcal{G} and verifying $\mathcal{F}(s) = 1$ as well as $\mathcal{Q}(\mathcal{G}, E, F) = \mathcal{Q}(\mathcal{G}', E, F)$.

because, once established the congestion values K_e and O_e for any $e \in E$, there are still infinitely many ways to split them among the players using resource e in both s and s^* , whereas, in the unweighted case, such a splitting is univocal.

4 Application to Affine Latency Functions

In order to easily illustrate our primal-dual technique, in this section we consider the well-known and studied case of affine latency functions and social function SUM and show how the results for the PoA_ϵ , the PoS_ϵ and the Apx_θ^1 already known in the literature can be reobtained in a unified manner for both weighted and unweighted players. Before proceeding with the application of the primal-dual method, we show how to simplify the parameters defining the general expression of affine latency functions without losing in generality.

Game $\mathcal{G}' = ([n], E', (\Sigma'_i)_{i \in [n]}, (\ell'_e)_{e \in E}, (w_i)_{i \in [n]})$ is equivalent to game $\mathcal{G} = ([n], E, (\Sigma_i)_{i \in [n]}, (\ell_e)_{e \in E}, (w_i)_{i \in [n]})$ if there exist n bijections $\varphi_i : \Sigma_i \mapsto \Sigma'_i$, one for each $i \in [n]$, such that, for any strategy profile $s = (s_1, \dots, s_n) \in \Sigma$, it holds that $c_i(s) = c_i(\varphi_1(s_1), \dots, \varphi_n(s_n))$ for any $i \in [n]$.

Lemma 1 *For each weighted congestion game with affine latency functions $\mathcal{G} = ([n], E, (\Sigma_i)_{i \in [n]}, (\ell_e)_{e \in E}, (w_i)_{i \in [n]})$ there always exists an equivalent weighted congestion game with linear latency functions $\mathcal{G}' = ([n], E', (\Sigma'_i)_{i \in [n]}, (\ell'_e)_{e \in E}, (w_i)_{i \in [n]})$, i.e., such that $\ell'_e(x) = \alpha_{e,1}x := \alpha_e x$ for any $e \in E'$.*

Proof Consider the weighted congestion game $\mathcal{G} = ([n], E, (\Sigma_i)_{i \in [n]}, (\ell_e)_{e \in E}, (w_i)_{i \in [n]})$ with latency functions $\ell_e(x) = \alpha_e x + \beta_e$ for any $e \in E$. For each $\tilde{e} \in E$ such that $\beta_{\tilde{e}} > 0$, let $N_{\tilde{e}}$ be the set of players who can choose \tilde{e} , that is, $N_{\tilde{e}} = \{i \in [n] : \exists s \in \Sigma_i : \tilde{e} \in s\}$. The set of resources E' is obtained by replicating all the resources in E and adding a new resource $e_{\tilde{e}}^i$ for any $\tilde{e} \in E$ and any $i \in N_{\tilde{e}}$, that is, $E' = E \cup \bigcup_{\tilde{e} \in E, i \in N_{\tilde{e}}} \{e_{\tilde{e}}^i\}$. The latency functions are defined as $\ell'_e(x) = \alpha_e x$ for any $e \in E' \cap E$ and $\ell'_{e_{\tilde{e}}^i}(x) = \frac{\beta_{\tilde{e}}}{w_i} x$ for any $\tilde{e} \in E$ and any $i \in N_{\tilde{e}}$. Finally, for each $i \in [n]$, the bijection φ_i is defined as follows: $\varphi_i(s) = s \cup \bigcup_{\tilde{e} \in s} \{e_{\tilde{e}}^i\}$. It is not difficult to see that for any $s = (s_1, \dots, s_n) \in \Sigma$ and for any $i \in [n]$, it holds that $c_i(s) = c_i(\varphi_1(s_1), \dots, \varphi_n(s_n))$. □

As a consequence of Lemma 1, throughout this section, we restrict to latency functions of the form $\ell_e(x) = \alpha_e x$, for any $e \in E$. In such a setting, we can rewrite the social value of a strategy profile as $\text{SUM}(s) = \sum_{e \in E} (\alpha_e \mathcal{L}_e(s)^2)$.

4.1 Bounding the Approximate Price of Anarchy

Since s is a PNE_ϵ , it follows that, for any $i \in [n]$,

$$c_i(s) = \sum_{e \in s_i} (\alpha_e K_e) \leq (1 + \epsilon) c_i(s_{-i} \diamond s_i^*) \leq (1 + \epsilon) \sum_{e \in s_i^*} (\alpha_e (K_e + w_i)).$$

Thus, the primal formulation $LP(s, s^*)$ assumes the following form.

$$\begin{aligned}
 & \text{maximize } \sum_{e \in E} (\alpha_e K_e^2) \\
 & \text{subject to} \\
 & \sum_{e \in S_i} (\alpha_e K_e) - (1 + \epsilon) \sum_{e \in S_i^*} (\alpha_e (K_e + w_i)) \leq 0, \quad \forall i \in [n] \\
 & \sum_{e \in E} (\alpha_e O_e^2) = 1, \\
 & \alpha_e \geq 0, \quad \forall e \in E
 \end{aligned}$$

The dual program $DLP(s, s^*)$ is

$$\begin{aligned}
 & \text{minimize } \gamma \\
 & \text{subject to} \\
 & \sum_{i: e \in S_i} (y_i K_e) - (1 + \epsilon) \sum_{i: e \in S_i^*} (y_i (K_e + w_i)) + \gamma O_e^2 \geq K_e^2, \quad \forall e \in E \\
 & y_i \geq 0, \quad \forall i \in [n]
 \end{aligned}$$

Define $\psi = \frac{1+\epsilon+\sqrt{\epsilon^2+6\epsilon+5}}{2}$ and $z = \lfloor \psi \rfloor$. For unweighted players, we reobtain the upper bound proved by Christodoulou, Koutsoupias and Spirakis in [25] with a much simpler and shorter proof, while, for the weighted case, we give the first known upper bound.

Theorem 1 For any fixed $\epsilon \geq 0$, $PoA_\epsilon(\mathcal{G}) \leq \frac{(1+\epsilon)(z^2+3z+1)}{2z-\epsilon}$ for any $\mathcal{G} \in \mathcal{UCG}$ and $PoA_\epsilon(\mathcal{G}) \leq \psi^2$ for any $\mathcal{G} \in \mathcal{WCG}$.

Proof For $\mathcal{G} \in \mathcal{UCG}$, since $w_i = 1$ for each $i \in [n]$, by choosing $y_i = \frac{2z+1}{2z-\epsilon}$ for any $i \in [n]$ and $\gamma = \frac{(1+\epsilon)(z^2+3z+1)}{2z-\epsilon}$ (note that the definition of z guarantees that $2z - \epsilon > 0$, so that $y_i \geq 0$ for each $i \in [n]$), the first dual constraint becomes of the form

$$\frac{2z + 1}{2z - \epsilon} \left(K_e^2 - (1 + \epsilon)(K_e + 1)O_e \right) + \frac{(1 + \epsilon)(z^2 + 3z + 1)}{2z - \epsilon} O_e^2 \geq K_e^2$$

which is equivalent to

$$K_e^2 - (2z + 1)(K_e O_e + O_e) + (z^2 + 3z + 1)O_e^2 \geq 0. \tag{1}$$

Easy calculations show that this is always satisfied for any pair of non-negative integers (K_e, O_e) . In fact, for $O_e = 0$, inequality (1) becomes $K_e^2 \geq 0$ which is always true. For $O_e = 1$, inequality (1) becomes $K_e^2 - (2z + 1)K_e + z^2 + z \geq 0$ which is always satisfied for any $K_e \leq z$ and $K_e \geq z + 1$ and, since z is a positive integer, it follows that inequality (1) is satisfied for any integer K_e . Finally, for $O_e \geq 2$, the equation associated with inequality (1) has no real solution when solved for K_e (its discriminant is negative), which implies that inequality (1) is satisfied for any value

of K_e . For an illustration on how the dual variables can be computed, see Section I in the Appendix.

For $\mathcal{G} \in \mathcal{WCG}$, by choosing $y_i = \left(1 + \frac{\sqrt{1+\epsilon}}{\sqrt{5+\epsilon}}\right) w_i$ for any $i \in [n]$ and $\gamma = \psi^2$, the first dual constraint is satisfied when

$$\left(1 + \frac{\sqrt{1+\epsilon}}{\sqrt{5+\epsilon}}\right) \left(K_e^2 - (1+\epsilon)(K_e O_e + O_e^2)\right) + \psi^2 O_e^2 \geq K_e^2$$

which is equivalent to

$$\frac{\sqrt{1+\epsilon}}{\sqrt{5+\epsilon}} K_e^2 - \left(1 + \frac{\sqrt{1+\epsilon}}{\sqrt{5+\epsilon}}\right) (1+\epsilon)(K_e O_e + O_e^2) + \psi^2 O_e^2 \geq 0. \tag{2}$$

Note that, since $\frac{\sqrt{1+\epsilon}}{\sqrt{5+\epsilon}} \left(1 + \frac{\sqrt{1+\epsilon}}{\sqrt{5+\epsilon}}\right) (1+\epsilon) = 2\psi$ and $\frac{\sqrt{1+\epsilon}}{\sqrt{5+\epsilon}} \psi^2 - 2\psi = \psi$, inequality (2) can be rewritten as $K_e^2 - 2\psi K_e O_e + \psi^2 O_e^2 = (K_e - \psi O_e)^2 \geq 0$ which is always satisfied for any pair of reals (K_e, O_e) . \square

When $\epsilon = 0$, we reobtain the well-known price of anarchy of $5/2$ and $(3 + \sqrt{5})/2$ which hold for PNE in the unweighted and weighted case, respectively. Note that, when $\psi = z$, it holds that $\psi^2 = \frac{(1+\epsilon)(z^2+3z+1)}{2z-\epsilon}$. Hence, interestingly enough, the PoA_ϵ in the weighted and unweighted cases coincide for all $\epsilon \geq 0$ such that $\frac{1+\epsilon+\sqrt{\epsilon^2+6\epsilon+5}}{2}$ is a natural number.

4.1.1 Analysis of the Dual Constraints

We now illustrate how the dual formulation can be also used to discover a matching lower bounding instance or to figure out its general structure.

Unweighted case For PNE, that is the case of $\epsilon = 0$, the dual constraints (1) get tight only for pairs (K_e, O_e) of the form $(1, 1)$ and $(2, 1)$. Thus, if the $5/2$ upper bound is tight, the complementary slackness conditions assure us that, in the matching lower bounding instance, only resources implementing the pairs $(1, 1)$ and $(2, 1)$ are needed. This can be easily achieved through a game using 3 players and 3 resources and defined as follows: $\Sigma_1 = \{\{e_1, e_2\}, \{e_3\}\}$, $\Sigma_2 = \{\{e_1\}, \{e_2, e_3\}\}$, $\Sigma_3 = \{\{e_2\}, \{e_3\}\}$ and $\alpha_1 = 5, \alpha_2 = 2, \alpha_3 = 3$. For such an instance, we have $s = (\{e_1, e_2\}, \{e_2, e_3\}, \{e_3\})$ and $s^* = (\{e_3\}, \{e_1\}, \{e_2\})$ for a price of anarchy of $5/2$. Clearly, this instance can be extended to any number of players $n > 3$ by adding a fourth resource e_4 with $\alpha_4 = 0$ and setting $\Sigma_i = \{e_4\}$ for any $i \in [n]$ with $i \geq 4$. Note that these are minimal lower bounding instances (the previous known lower bounding instances presented by Christodoulou and Koutsoupias in [23] used $2n$ resources for any $n \geq 3$).

More generally, for PNE_ϵ , the dual constraints (1) get tight only for pairs of the form $(z, 1)$ and $(z + 1, 1)$. Thus, in order to obtain a matching lower bounding instance, we only need to implement this family of dual constraints, that is, we need an instance with at least $z + 2$ players and a set of resources such that $O_e = 1$ and $K_e \in \{z, z + 1\}$ for any $e \in E$. In fact, the matching lower bounding instances given

by Christodoulou, Koutsoupias and Spirakis in [25] use $z + 2$ players and $2z + 4$ resources, half of which has $K_e = z$ and the other half has $K_e = z + 1$. It is easy to see that these instances are not minimal. In fact, they produce a dual program with $z + 3$ variables and $2z + 4$ constraints, where only $z + 3$ constraints are sufficient to exactly characterize the optimal dual solution. Unfortunately, this set of constraints changes as a function of ϵ and so it is not easy to achieve a general scheme of minimal lower bounding instances.

Weighted case The dual constraints (2) get tight only for pairs (K_e, O_e) such that $K_e = \psi O_e$. For PNE, that is the case $\epsilon = 0$, consider the game with 3 players and 3 resources with $w_1 = 1, w_2 = w_3 = (1 + \sqrt{5})/2, \Sigma_1 = \{\{e_1\}, \{e_2, e_3\}\}, \Sigma_2 = \{\{e_2\}, \{e_1, e_3\}\}, \Sigma_3 = \{\{e_3\}, \{e_2\}\}, \alpha_1 = 2, \alpha_2 = \sqrt{5} - 1$ and $\alpha_3 = 3 - \sqrt{5}$. We have $s = (\{e_2, e_3\}, \{e_1, e_3\}, \{e_2\})$ and $s^* = (\{e_1\}, \{e_2\}, \{e_3\})$ for a price of anarchy equal to $(3 + \sqrt{5})/2$. Again, we have identified a minimal lower bounding instance which is slightly simpler than the previous one given by Awerbuch, Azar and Epstein in [4] which used 4 players and 3 resources.

For general values of ϵ , we are able to provide tight lower bounds only for a subset, although having infinite cardinality, of values of ϵ . Let t and y be two positive integers such that $1 \leq y \leq t + 1$. We set $\epsilon(t, y) = \frac{(t-1)\sqrt{t^2+4y+2y+t^2-t-2}}{\sqrt{t^2+4y+t+2}}$, which is always non-negative since $y \geq 1$ and yields $\psi = \frac{t+\sqrt{t^2+4y}}{2} > t$. We create an instance with $t + 2$ players and $2(t + 1)$ resources, where $w_i = 1$ for any $i \in [t + 1]$ and $w_{t+2} = \psi - t$. The first $t + 1$ resources $(e_j)_{j \in [t+1]}$ have latency $\ell(x) = x$, while the last $t + 1$ resources $(e'_j)_{j \in [t+1]}$ have latency $\ell(x) = x/y$. The set of strategies for each player $i \in [t + 1]$ is $\Sigma_i = \{\{e_i\}, \bigcup_{j \in [t]} \{e_{i+j}\} \cup \bigcup_{j \in [y]} \{e'_{i+j}\}\}$, with the sum of the indices taken circularly, while $\Sigma_{t+2} = \{\bigcup_{j \in [t+1]} \{e'_j\}, \bigcup_{j \in [t+1]} \{e_j\}\}$. The first strategy of each player is the optimal one, while the second strategy is the one played at the $\text{PNE}_{\epsilon(t,y)}$. Note that, for any e , we have $K_e = \psi O_e$, thus implying $\text{PoA}_\epsilon = \psi^2$. It is not difficult to show that s is a $(1 + \epsilon(t, y)) - \text{PNE}$ by exploiting the equality $2\psi = t + \sqrt{t^2 + 4y}$.

Deriving a tight lower bound for any possible value of ϵ remains an interesting open problem.

4.2 Bounding the Approximate Price of Stability

In this case, we have two different approaches for unweighted and weighted games, respectively.

4.2.1 Unweighted Case

Recall that, since the PoS_ϵ is a best-case measure, the primal-dual approach guarantees a tight analysis only if we are able to exactly characterize the polyhedron defining the set of the best quality PNE_ϵ . It is not known how to do this at the moment, thus all the approaches used so far in the literature approximate the best quality PNE_ϵ with a PNE_ϵ minimizing a certain potential function. Christodoulou, Koutsoupias and

Spirakis [25] showed that, for unweighted players, any strategy profile s which is a local minimum of the function

$$\Phi_\epsilon(s) = \frac{1}{2} \sum_{e \in E} \left(\alpha_e \left(\mathcal{L}_e(s)^2 + \frac{1-\epsilon}{1+\epsilon} \mathcal{L}_e(s) \right) \right),$$

called ϵ -approximate potential, is a PNE $_\epsilon$. Thus, it is possible to get an upper bound on the PoS $_\epsilon$ by bounding the PoA $_\epsilon$ of the global minimum of Φ_ϵ .

We now illustrate our approach which yields the same $\frac{1+\sqrt{3}}{\epsilon+\sqrt{3}}$ upper bound achieved by Christodoulou, Koutsoupias and Spirakis in [25]. Assume that s is the global minimum of Φ_ϵ . We can use the inequality $\Phi_\epsilon(s) \leq \Phi_\epsilon(s^*)$ which results in the constraint

$$\sum_{e \in E} \left(\alpha_e \left(K_e^2 + \frac{1-\epsilon}{1+\epsilon} K_e - O_e^2 - \frac{1-\epsilon}{1+\epsilon} O_e \right) \right) \leq 0. \tag{3}$$

Moreover, it also holds that $\sum_{i \in [n]} (\Phi_\epsilon(s) - \Phi_\epsilon(s_{-i} \diamond s_i^*)) \leq 0$. Such an inequality can be expressed as a function of K_e, O_e and ϵ as follows:

$$\begin{aligned} & \Phi_\epsilon(s) - \Phi_\epsilon(s_{-i} \diamond s_i^*) \\ &= \frac{1}{2} \sum_{e \in s_i \setminus s_i^*} \left(\alpha_e \left(K_e^2 + \frac{1-\epsilon}{1+\epsilon} K_e - (K_e - 1)^2 - \frac{1-\epsilon}{1+\epsilon} (K_e - 1) \right) \right) \\ & \quad + \frac{1}{2} \sum_{e \in s_i^* \setminus s_i} \left(\alpha_e \left(K_e^2 + \frac{1-\epsilon}{1+\epsilon} K_e - (K_e + 1)^2 - \frac{1-\epsilon}{1+\epsilon} (K_e + 1) \right) \right) \\ &= \sum_{e \in s_i \setminus s_i^*} \left(\alpha_e \left(K_e - \frac{\epsilon}{1+\epsilon} \right) \right) - \sum_{e \in s_i^* \setminus s_i} \left(\alpha_e \left(K_e + \frac{1}{1+\epsilon} \right) \right). \end{aligned}$$

Now, for each $e \in E$, define $\Delta_e = |\{i \in [n] : e \in s_i \cap s_i^*\}|$. By summing up for each $i \in [n]$, we obtain

$$\begin{aligned} & \sum_{i \in [n]} (\Phi_\epsilon(s) - \Phi_\epsilon(s_{-i} \diamond s_i^*)) \\ &= \sum_{i \in [n]} \left(\sum_{e \in s_i \setminus s_i^*} \left(\alpha_e \left(K_e - \frac{\epsilon}{1+\epsilon} \right) \right) - \sum_{e \in s_i^* \setminus s_i} \left(\alpha_e \left(K_e + \frac{1}{1+\epsilon} \right) \right) \right) \\ &= \sum_{e \in E} \left(\alpha_e \left(\left(K_e - \frac{\epsilon}{1+\epsilon} \right) (K_e - \Delta_e) - \left(K_e + \frac{1}{1+\epsilon} \right) (O_e - \Delta_e) \right) \right) \\ &\geq \sum_{e \in E} \left(\alpha_e \left(K_e^2 - \frac{\epsilon}{1+\epsilon} K_e - K_e O_e - \frac{1}{1+\epsilon} O_e \right) \right), \end{aligned}$$

which implies

$$\sum_{e \in E} \left(\alpha_e \left(K_e^2 - \frac{\epsilon}{1+\epsilon} K_e - K_e O_e - \frac{1}{1+\epsilon} O_e \right) \right) \leq 0. \tag{4}$$

Thanks to (3) and (4), the dual formulation becomes:

$$\begin{aligned}
 & \text{minimize } \gamma \\
 & \text{subject to} \\
 & K_e^2(y+z) + \frac{K_e}{1+\epsilon}(y(1-\epsilon) - z\epsilon) \\
 & - \left(yO_e^2 + zK_eO_e + \frac{O_e}{1+\epsilon}(y(1-\epsilon) + z) \right) + \gamma O_e^2 \geq K_e^2, \quad \forall e \in E \\
 & y, z \geq 0
 \end{aligned}$$

Thus, for unweighted players, we obtain the following result for any $\epsilon \in [0, 1)$ (this is the only interesting case, since Christodoulou, Koutsoupias and Spirakis [25] show that, for any $\epsilon \geq 1$, $\text{PoS}_\epsilon(\mathcal{G}) = 1$ for any \mathcal{G}).

Theorem 2 For any fixed $\epsilon \in [0, 1)$, $\text{PoS}_\epsilon(\mathcal{G}) \leq \frac{1+\sqrt{3}}{\epsilon+\sqrt{3}}$ for any $\mathcal{G} \in \text{UCG}$.

Proof By choosing $y = \frac{2\epsilon+\sqrt{3}(1+\epsilon)}{2(\epsilon+\sqrt{3})}$, $z = \frac{1-\epsilon}{\epsilon+\sqrt{3}}$ and $\gamma = \frac{1+\sqrt{3}}{\epsilon+\sqrt{3}}$, the first dual constraint becomes of the form

$$(\epsilon - 1)((\sqrt{3} - 2)K_e^2 + (2O_e - \sqrt{3})K_e + (2 + \sqrt{3})(O_e - O_e^2)) \geq 0. \tag{5}$$

Easy calculations show that this is always satisfied for any pair of non-negative integers (K_e, O_e) . In fact, for $O_e = 0$, inequality (5) becomes $K_e((2 - \sqrt{3})K_e + \sqrt{3})(1 - \epsilon) \geq 0$ which is always satisfied for any non-negative integer K_e since $\epsilon < 1$. For $O_e = 1$, inequality (5) becomes $K_e(K_e - 1)(1 - \epsilon) \geq 0$ which is always satisfied for any non-negative integer K_e since $\epsilon < 1$. Finally, for any $O_e \geq 2$, the equation associated with inequality (5) has no real solution when solved for K_e (its discriminant is negative), which implies that inequality (5) is satisfied for any value of K_e . \square

Analysis of the Dual Constraints The dual constraint (5) gets tight only for pairs (K_e, O_e) of the form $(0, 1)$ and $(1, 1)$ which are clearly insufficient to achieve a PoS_ϵ greater than 1, since $K_e \leq O_e$ for any $e \in E$. *What is going on here?* The answer is that the lower bound on the PoS_ϵ can be achieved only asymptotically, that is, when n tends to infinity. Thus, we must also check what happens when both K_e and O_e goes to infinity and their ratio remains constant. We obtain that the dual constraint (5) is asymptotically tight for pairs of the form (K_e, O_e) such that $K_e = (2 + \sqrt{3})O_e$ and O_e goes to infinity. The lower bounding instances proposed by Christodoulou, Koutsoupias and Spirakis in [25] have n_1 resources of type $(0, 1)$, $n_1(n_1 - 1)$ resources of type $(1, 1)$ and one resource of type $\left(n_1, \frac{\sqrt{2\epsilon^4+3\epsilon^3+\epsilon+3}+2\epsilon^2+2\epsilon-1}{\sqrt{2\epsilon^4+3\epsilon^3+\epsilon+3+\epsilon^2+\epsilon+1}} n_1 \right)$ with n_1 going to infinity. Thus, such lower bounding instances possess all the combinatorics needed to implement the worst-case dual constraints, but still there is a remarkable gap between upper and lower bounds. Hence, the intuition should suggest that the upper bound is not tight and that additional constraints should be used in the primal formulation so as to better characterize the polyhedron defining the best PNE_ϵ . Note that the inequalities stating the s is a PNE_ϵ is of no use here since they are dominated by the inequality $\sum_{i \in [n]} (\Phi_\epsilon(s) - \Phi_\epsilon(s_{-i} \diamond s_i^*)) \leq 0$.

4.2.2 Weighted Case

In order to deal with the weighted case, it is possible to rephrase the approach of Christodoulou, Koutsoupias and Spirakis [25] to turn the potential given by Fotakis, Kontogiannis and Spirakis in [29] for weighted linear congestion games into an ϵ -potential function for this class of games so as to use the same approach as in the unweighted case.

We define the following ϵ -potential function.

$$\Phi_\epsilon(s) = \frac{1}{2} \sum_{e \in E} (\alpha_e \mathcal{L}_e(s)^2) + \frac{1 - \epsilon}{2(1 + \epsilon)} \sum_{e \in E} \sum_{i: e \in S_i} (\alpha_e w_i^2).$$

Lemma 2 Any profile which is a local minimum of Φ_ϵ is a PNE_ϵ .

Proof Consider a profile $s = (s_1, \dots, s_n)$. We want to compute the change in the ϵ -potential function when player i changes her strategy from s_i to t . The resulting profile $s_{-i} \diamond t$ has

$$\mathcal{L}_e(s_{-i} \diamond t) = \begin{cases} \mathcal{L}_e(s) - w_i, & e \in s_i \setminus t, \\ \mathcal{L}_e(s) + w_i, & e \in t \setminus s_i, \\ \mathcal{L}_e(s), & \text{otherwise.} \end{cases}$$

From this we can compute the difference

$$\begin{aligned} & \Phi_\epsilon(s_{-i} \diamond t) - \Phi_\epsilon(s) \\ &= \sum_{e \in t \setminus s_i} \left(\alpha_e \left(w_i \mathcal{L}_e(s) + \frac{w_i^2}{1 + \epsilon} \right) \right) - \sum_{e \in s_i \setminus t} \left(\alpha_e \left(w_i \mathcal{L}_e(s) - \frac{\epsilon w_i^2}{1 + \epsilon} \right) \right). \end{aligned}$$

We can rewrite this as

$$\begin{aligned} & \Phi_\epsilon(s_{-i} \diamond t) - \Phi_\epsilon(s) \\ &= \sum_{e \in t} \left(\alpha_e \left(w_i \mathcal{L}_e(s) + \frac{w_i^2}{1 + \epsilon} \right) \right) - \sum_{e \in t \cap s_i} (\alpha_e w_i^2) - \sum_{e \in s_i} \left(\alpha_e \left(w_i \mathcal{L}_e(s) - \frac{\epsilon w_i^2}{1 + \epsilon} \right) \right). \end{aligned}$$

Suppose now that s is a local minimum of Φ_ϵ which implies $\Phi_\epsilon(s) \leq \Phi_\epsilon(s_{-i} \diamond t)$ for any $i \in [n]$ and $t \in \Sigma_i$. The cost of player i before the change is $c_i(s) = \sum_{e \in s_i} (\alpha_e \mathcal{L}_e(s))$ and after the change is $c_i(s_{-i} \diamond t) = \sum_{e \in t} (\alpha_e \mathcal{L}_e(s_{-i} \diamond t))$. We show that s is a PNE_ϵ , that is, $c_i(s) \leq (1 + \epsilon)c_i(s_{-i} \diamond t)$.

By exploiting the two different parts defining the ϵ potential function, we obtain

$$c_i(s) = \sum_{e \in s_i} (\alpha_e \mathcal{L}_e(s)) \leq \sum_{e \in s_i} \left(\alpha_e (1 + \epsilon) \left(\mathcal{L}_e(s) - \frac{\epsilon w_i}{1 + \epsilon} \right) \right)$$

which holds because $\mathcal{L}_e(s) \geq w_i$ when $e \in s_i$, and

$$\begin{aligned} c_i(s_{-i} \diamond t) &= \sum_{e \in t} (\alpha_e (\mathcal{L}_e(s_{-i} \diamond t) + w_i)) - \sum_{e \in t \cap s_i} (\alpha_e w_i) \\ &\geq \sum_{e \in t} \left(\alpha_e \left(\mathcal{L}_e(s) + \frac{w_i}{1 + \epsilon} \right) \right) - \sum_{e \in t \cap s_i} (\alpha_e w_i). \end{aligned}$$

which holds for any $\epsilon \geq 0$.

It follows immediately that $c_i(s) \leq (1 + \epsilon)c_i(s_{-i} \diamond t)$, thus s is a PNE_ϵ . □

Using the constraint $\Phi_\epsilon(s) - \Phi_\epsilon(s^*) \leq 0$ in our formulation, we can easily prove the following theorem.

Theorem 3 For any fixed $\epsilon \in [0; 1]$, $\text{PoS}_\epsilon(\mathcal{G}) \leq \frac{2}{1+\epsilon}$ for any $\mathcal{G} \in \mathcal{WCG}$.

Proof In such a case, by choosing $y_i = 0$ for any $i \in [n]$, $z = 1$ and $y_{n+1} = \frac{2}{1+\epsilon}$, the first dual constraint becomes of the form

$$\frac{1 - \epsilon}{1 + \epsilon} K_e \geq 0$$

which is always satisfied for any pair of non-negative reals (K_e, O_e) . □

In this case, no specific lower bounds are known, besides the ones with unweighted players.

4.3 Bounding the Approximation Ratio of One-Round Walks

For a one-round walk W , we set $s = s_n^W$. Define $K_e(i)$ as the sum of the weights of the players using resource e in s before player i performs her choice. $\text{LP}(s, s^*)$ in this case has the following form, where the first constraint comes from the fact that, when player i enters the game and strategy profile s_{i-1}^W is already constructed, this player chooses s_i instead of s_i^* .

$$\begin{aligned} &\text{maximize } \sum_{e \in E} (\alpha_e K_e^2) \\ &\text{subject to} \\ &\sum_{e \in s_i} (\alpha_e (K_e(i) + w_i)) - \sum_{e \in s_i^*} (\alpha_e (K_e(i) + w_i)) \leq 0, \quad \forall i \in [n] \\ &\sum_{e \in E} (\alpha_e O_e^2) = 1 \\ &\alpha_e \geq 0, \quad \forall e \in E \end{aligned}$$

DLP(s, s^*) is as follows.

$$\begin{aligned}
 & \text{minimize } \gamma \\
 & \text{subject to} \\
 & \sum_{i:e \in s_i} (y_i(K_e(i) + w_i)) - \sum_{i:e \in s_i^*} (y_i(K_e(i) + w_i)) + \gamma O_e^2 \geq K_e^2, \quad \forall e \in E \\
 & y_i \geq 0, \quad \forall i \in [n]
 \end{aligned}$$

For both unweighted and weighted players we easily reobtain the upper bounds on the Apx_\emptyset^1 given by Christodoulou, Mirrokni and Sidiropoulos in [26].

Theorem 4 For any $\mathcal{G} \in \text{UCG}$, $\text{Apx}_\emptyset^1(\mathcal{G}) \leq 2 + \sqrt{5}$ and, for any $\mathcal{G} \in \text{WCG}$, $\text{Apx}_\emptyset^1(\mathcal{G}) \leq 4 + 2\sqrt{3}$.

Proof For $\mathcal{G} \in \text{UCG}$, by choosing $y_i = 1 + \sqrt{5}$ for any $i \in [n]$ and $\gamma = 2 + \sqrt{5}$, since for any i such that $e \in s_i^*$ it holds that $K_e(i) \leq K_e$, the first dual constraint is satisfied when

$$(1 + \sqrt{5}) \left(\frac{K_e(K_e + 1)}{2} - (K_e + 1)O_e \right) + (2 + \sqrt{5}) O_e^2 \geq K_e^2$$

which is equivalent to

$$\left(\frac{\sqrt{5} - 1}{2} \right) K_e^2 + (1 + \sqrt{5}) \left(\frac{K_e}{2} - K_e O_e - O_e \right) + (2 + \sqrt{5}) O_e^2 \geq 0. \tag{6}$$

Easy calculations show that this is always satisfied for any pair of non-negative integers (K_e, O_e) . In fact, for $O_e = 0$, inequality (6) becomes $K_e((\sqrt{5} - 1)K_e + \sqrt{5} + 1) \geq 0$ which is always satisfied for any non-negative integer K_e . For $O_e = 1$, inequality (6) becomes $K_e((\sqrt{5} - 1)K_e - \sqrt{5} - 1) \geq -2$ which is always satisfied for any integer K_e . Finally, for any $O_e \geq 2$, the equation associated with inequality (6) has no real solution when solved for K_e (its discriminant is negative), which implies that inequality (6) is satisfied for any non-negative integer value of K_e .

For $\mathcal{G} \in \text{WCG}$, by choosing $y_i = \left(2 + \frac{2}{\sqrt{3}}\right) w_i$ for any $i \in [n]$ and $\gamma = 4 + 2\sqrt{3}$, since for any i such that $e \in s_i^*$ it holds that $K_e(i) \leq K_e$, the first dual constraint is satisfied when

$$\left(2 + \frac{2}{\sqrt{3}}\right) \left(\sum_{i \leq j: e \in s_i^K \cap s_j^K} (w_i w_j) - \sum_{i:e \in s_i^O} (w_i(K_e + w_i)) \right) + (4 + 2\sqrt{3}) O_e^2 \geq K_e^2$$

which is true if

$$\frac{1}{\sqrt{3}} K_e^2 - \left(2 + \frac{2}{\sqrt{3}}\right) K_e O_e + \left(2 + \frac{4}{\sqrt{3}}\right) O_e^2 \geq 0. \tag{7}$$

Note that, after rearranging of the terms, inequality (7) becomes $(K_e - (1 + \sqrt{3})O_e)^2 \geq 0$ which is clearly satisfied for any pair of reals (K_e, O_e) . \square

4.3.1 Analysis of the Dual Constraints

For the unweighted case, the dual constraints get tight for pairs of the form $(1, 1)$, while the asymptotical dual constraints get tight for pairs of the form $\left(\frac{3+\sqrt{5}}{2} O_e, O_e\right)$. These pairs exactly characterize the structure of the lower bounding instance derived by Bilò et al. in [10]. For the weighted case, the worst case dual constraints occur when all players using resource e in the walk have weight 1, while only one player uses e in the optimal strategy profile. Moreover, the asymptotical dual constraints get tight for pairs of the form $((1 + \sqrt{3}) O_e, O_e)$. In this case, the best known lower bound, equal to $3 + 2\sqrt{2}$, has been given by Caragiannis et al. in [20].

5 Application to Quadratic and Cubic Latency Functions

In this section, we show how to use the primal-dual method to bound the PoS and the Apx_{\emptyset}^1 in the case of polynomial latency functions of maximum degree d and unweighted players. We only consider the case $d \leq 3$, that is, quadratic and cubic latency functions. As we will see, it is not difficult to extend the approach to any particular value of d , but it is quite hard to obtain a general result as a function of d because we do not have simple closed formulas expressing some of the summations we need in our analysis for any value of d .

First of all note that, by using the same argument exploited in the proof of Lemma 1, we can assume without loss of generality that, in all games we will consider throughout this section, the latency functions are of the form $\ell_e(x) = \sum_{i=1}^d \alpha_{e,i} x^i$, that is, without the term $\alpha_{e,0}$.

5.1 Bounding the Price of Stability

Recall that, for any $G \in \mathcal{UCG}$, Rosenthal [41] shows that any strategy profile s which is a local minimum of the potential function

$$\Phi(s) = \sum_{e \in E} \sum_{x=1}^{\mathcal{L}_e(s)} \ell_e(x) = \sum_{e \in E} \sum_{x=1}^{\mathcal{L}_e(s)} \sum_{i=1}^d (\alpha_{e,i} x^i)$$

is a PNE.

Let us denote with $\Phi_i(s) = \sum_{e \in E} \sum_{x=1}^{\mathcal{L}_e(s)} (\alpha_{e,i} x^i)$ the contribution of the term $\alpha_{e,i} x^i$ to $\Phi(s)$. We have already seen in the previous section that

$$\Phi_1(s) = \frac{1}{2} \sum_{e \in E} (\alpha_{e,1} (\mathcal{L}_e(s)^2 + \mathcal{L}_e(s))).$$

Moreover,

$$\Phi_2(s) = \frac{1}{6} \sum_{e \in E} (\alpha_e \mathcal{L}_e(s) (\mathcal{L}_e(s) + 1) (2\mathcal{L}_e(s) + 1))$$

and

$$\Phi_3(s) = \frac{1}{4} \sum_{e \in E} (\alpha_e (\mathcal{L}_e(s)(\mathcal{L}_e(s) + 1))^2).$$

According to this decomposition of $\Phi(s)$, the constraint $\Phi(s) \leq \Phi(s^*)$ becomes

$$\sum_{i=1}^d \Phi_i(s) \leq \sum_{i=1}^d \Phi_i(s^*)$$

and the constraint $\sum_{i \in [n]} (\Phi(s) - \Phi(s_{-i} \diamond s_i^*)) \leq 0$ becomes

$$\sum_{i=1}^d \left(\sum_{j \in [n]} (\Phi_i(s) - \Phi_i(s_{-j} \diamond s_j^*)) \right) \leq 0.$$

Again, as shown in the previous section,

$$\sum_{i \in [n]} (\Phi_1(s) - \Phi_1(s_{-i} \diamond s_i^*)) = \sum_{e \in E} (\alpha_{e,1} (K_e^2 - K_e O_e - O_e)).$$

Moreover, using the same derivation as in the previous section, we obtain

$$\sum_{i \in [n]} (\Phi_2(s) - \Phi_2(s_{-i} \diamond s_i^*)) = \sum_{e \in E} (\alpha_{e,2} (K_e^3 - O_e(K_e + 1)^2))$$

and

$$\sum_{i \in [n]} (\Phi_3(s) - \Phi_3(s_{-i} \diamond s_i^*)) = \sum_{e \in E} (\alpha_{e,3} (K_e^4 - O_e(K_e + 1)^3)).$$

Hence, for the case of $d = 2$, $DLP(s, s^*)$ is the following.

$$\begin{aligned} & \text{minimize } \gamma \\ & \text{subject to} \\ & K_e^2 \left(\frac{\gamma}{2} + z\right) + \frac{\gamma}{2} K_e \\ & - \left(\frac{\gamma}{2} O_e^2 + z K_e O_e + O_e \left(\frac{\gamma}{2} + z\right)\right) + \gamma O_e^2 \geq K_e^2, \quad \forall e \in E \\ & \frac{\gamma}{6} (K_e(K_e + 1)(2K_e + 1) - O_e(O_e + 1)(2O_e + 1)) \\ & + z(K_e^3 - O_e(K_e + 1)^2) + \gamma O_e^3 \geq K_e^3, \quad \forall e \in E \\ & \gamma, z \geq 0 \end{aligned}$$

By analyzing the dual formulation, we obtain the following upper bound on the PoS.

Theorem 5 For any $\mathcal{G} \in UCG$ with quadratic latency functions, $\text{PoS}(\mathcal{G}) \leq 2.362$.

Proof The claim follows by setting $\gamma = 1.908$, $z = 0.453$ and $\gamma = 2.362$.

Let us denote as $f_1(K_e, O_e)$ and $f_2(K_e, O_e)$, respectively, the first and second constraint of $DLP(s, s^*)$.

For the first constraint, $f_1(0, O_e) = 1408O_e^2 - 1407O_e \geq 0$ is always satisfied for any integer $O_e \geq 0$, while the equation $f_1(K_e, O_e) = 0$ has no solution (its discriminant is negative) when $K_e \geq 1$. Hence, the first constraint is always satisfied for any pair of non-negative integers (K_e, O_e) .

For the second constraint, note that $\frac{\delta f_2}{\delta O_e}(K_e, O_e) = -3(151K_e^2 + 302K_e - 1726O_e^2 + 636O_e + 257)$ implies that, for any fixed value of K_e , f_2 is increasing in O_e when $O_e \geq 1$. Hence, we just need to show that $f_2(K_e, O_e) \geq 0$ for $O_e \in \{0, 1\}$. It holds that $f_2(K_e, 0) = 89K_e^3 + 954K_e^2 + 318K_e \geq 0$ which is always satisfied for any $K_e \geq 0$ and $f_2(K_e, 1) = 89K_e^3 + 501K_e^2 - 588K_e + 1 \geq 0$ which is always satisfied for any non-negative integer. Hence, also the second constraint is always satisfied for any pair of non-negative integers (K_e, O_e) . \square

For the case $d = 3$, $DLP(s, s^*)$ is the following.

$$\begin{aligned}
 & \text{minimize } \gamma \\
 & \text{subject to} \\
 & K_e^2(\frac{y}{2} + z) + \frac{y}{2}K_e \\
 & - (\frac{y}{2}O_e^2 + zK_eO_e + O_e(\frac{y}{2} + z)) + \gamma O_e^2 \geq K_e^2, \quad \forall e \in E \\
 & \frac{y}{6}(K_e(K_e + 1)(2K_e + 1) - O_e(O_e + 1)(2O_e + 1)) \\
 & + z(K_e^3 - O_e(K_e + 1)^2) + \gamma O_e^3 \geq K_e^3, \quad \forall e \in E \\
 & \frac{y}{4}(K_e^2(K_e + 1)^2 - O_e^2(O_e + 1)^2) \\
 & + z(K_e^4 - O_e(K_e + 1)^3) + \gamma O_e^4 \geq K_e^4, \quad \forall e \in E \\
 & y, z \geq 0
 \end{aligned}$$

By analyzing the dual formulation, we obtain the following upper bound on the PoS.

Theorem 6 For any $\mathcal{G} \in UCG$ with cubic latency functions, $PoS(\mathcal{G}) \leq 3.322$.

Proof The claim follows by setting $y = 2.99$, $z = 0.331$ and $\gamma = 3.322$.

Let us denote as $f_i(K_e, O_e)$ the i th constraint of $DLP(s, s^*)$.

For the first constraint, $f_1(0, O_e) = 1827O_e^2 - 1826O_e \geq 0$ is always satisfied for any integer $O_e \geq 0$, while the equation $f_1(K_e, O_e) = 0$ has no solution when $K_e \geq 1$ (its discriminant is negative). Hence, the first constraint is always satisfied for any pair of non-negative integers (K_e, O_e) .

For the second constraint, note that $\frac{\delta f_2}{\delta O_e}(K_e, O_e) = -933K_e^2 - 1986K_e + 20928O_e^2 - 8970O_e - 2488$ implies that, for any fixed value of K_e , f_2 is increasing in O_e when $O_e \geq 1$. Hence, we just need to show that $f_2(K_e, O_e) \geq 0$ for $O_e \in \{0, 1\}$. It holds that $f_2(K_e, 0) = 983K_e^3 + 4485K_e^2 + 1495K_e \geq 0$ which is always satisfied for any $K_e \geq 0$ and $f_2(K_e, 1) = 983K_e^3 + 3492K_e^2 - 491K_e + 3 \geq 0$ which is always satisfied for any non-negative integer. Hence, also the second constraint is always satisfied for any pair of non-negative integers (K_e, O_e) .

For the third constraint, note that $\frac{\delta f_3}{\delta O_e}(K_e, O_e) = -2(331K_e^3 + 993K_e^2 + 933K_e - 10298O_e^3 + 4485O_e^2 + 1495O_e + 331)$ implies that, for any fixed value of K_e , f_3 is increasing in O_e when $O_e \geq 1$. Hence, we just need to show that $f_3(K_e, O_e) \geq 0$ for $O_e \in \{0, 1\}$. It holds that $f_3(K_e, 0) = 157K_e^2 + 2990K_e + 1495 \geq 0$ which is always satisfied for any $K_e \geq 0$ and $f_3(K_e, 1) = 157K_e^4 + 2328K_e^3 - 491K_e^2 - 1986K_e + 2 \geq 0$ which is always satisfied for any non-negative integer. Hence, also the third constraint is always satisfied for any pair of non-negative integers (K_e, O_e) . \square

By extending the instance given by Christodoulou, Koutsoupias and Spirakis in [25] for lower bounding the PoS in the case of affine latency functions, the following lower bounds can be easily achieved.

Theorem 7 *For any $\delta > 0$, there exist $\mathcal{G}_1 \in UCG$ with quadratic latency functions and $\mathcal{G}_2 \in UCG$ with cubic latency functions such that $\text{PoS}(\mathcal{G}_1) \geq 2.1859 - \delta$ and $\text{PoS}(\mathcal{G}_2) \geq 2.7558 - \delta$.*

Proof Consider a game $\mathcal{G} \in UCG$ with $n = n_1 + n_2$ players divided into two sets P_1 and P_2 with $|P_1| = n_1$ and $|P_2| = n_2$. Each player $i \in P_1$ has two strategies s_i and s_i^* , while all players in P_2 have the same strategy \bar{s} .

There are three types of resources:

- n_1 resources $r_i, i \in [n_1]$, each with latency function $\ell_{r_i}(x) = rx^d$, where the parameter r will be specified later. Resource r_i belongs only to s_i^* ;
- $n_1(n_1 - 1)$ resources $r'_{i,j}, i, j \in [n_1]$ with $i \neq j$, each with latency function $\ell_{r'_{ij}}(x) = r'x^d$, where the parameter r' will be specified later. Resource r'_{ij} belongs only to s_i and to s_j^* ;
- one resource r'' with latency function $\ell_{r''}(x) = x^d$. Resource r'' belongs to s_i for each $i \in [n_1]$ and to \bar{s} ;

The cost of each player $i \in P_1$ adopting strategy s_i when there are exactly k players in P_1 adopting the strategy played in s (and thus there are $n_1 - k$ players in P_1 adopting the strategy played in s^*) is $\text{cost}_s(k) = (4n_1 - 3k - 1)r' + (n_2 + k)^2$ when $d = 2$ and it is $\text{cost}_s(k) = (8n_1 - 7k - 1)r' + (n_2 + k)^3$ when $d = 3$. Similarly, the cost of each player $i \in P_1$ adopting strategy s_i^* when there are exactly k players in P_1 adopting the strategy played in s is $\text{cost}_{s^*}(k) = r + (n_1 + 3k - 1)r'$ when $d = 2$ and it is $\text{cost}_{s^*}(k) = r + (n_1 + 7k - 1)r'$ when $d = 3$.

We now want to select the parameters r and r' so that s is the unique PNE of the game. This is true if, for any $k \in [n_1]$, it holds that $\text{cost}_{s^*}(k - 1) > \text{cost}_s(k)$. Such a condition is always satisfied for the following values of r and r' :

- $r = \frac{2n_2^2 + (n_1 + 1)(n_1 + 2n_2)}{2} + \gamma$ and $r' = \frac{n_1 + 2n_2}{6}$, when $d = 2$,
- $r = \frac{2n_2^3 + (n_1 + 1)(n_1^2 + 3n_1n_2 + 3n_2^2)}{2} + \gamma$ and $r' = \frac{n_1^2 + 3n_1n_2 + 3n_2^2}{14}$, when $d = 3$,

where γ is an arbitrarily small positive value.

Next step is to select n_1 and n_2 so as to maximize the ratio $\frac{\text{SUM}(s)}{\text{SUM}(s^*)} = \frac{r'n_1(n_1 - 1) + (n_1 + n_2)^{d+1}}{rn_1 + r'n_1(n_1 - 1) + n_2^{d+1}}$ for $d = 2, 3$. By choosing $n_1 = 1.5595n_2$ when $d = 2$ and $n_1 = 1.0988n_2$ when $d = 3$ and letting n_2 go to infinity, we obtain the claim. \square

A recent paper by Christodoulou and Gairing [22] shows that $\text{PoS} = 2.362$ when $d = 2$ and $\text{PoS} = 3.322$ when $d = 3$, i.e., our upper bounds are tight, while our lower bounds are not. In particular, while our lower bounds hold for the more restricted notion of dominant strategy equilibria, Christodoulou and Gairing construct instances for which there exists a unique PNE which is not a dominant strategy

equilibria. Hence, an interesting research problem is that of determine the price of anarchy/stability² of dominant strategy equilibria.

5.2 Bounding the Approximation Ratio of One-Round Walks

For the case $d = 2$, $DLP(s, s^*)$ is defined as follows.

$$\begin{aligned}
 & \text{minimize } \gamma \\
 & \text{subject to} \\
 & \sum_{i:e \in s_i} (y_i(K_e(i) + 1)) - \sum_{i:e \in s_i^*} (y_i(K_e(i) + 1)) + \gamma O_e^2 \geq K_e^2, \quad \forall e \in E \\
 & \sum_{i:e \in s_i} (y_i(K_e(i) + 1)^2) - \sum_{i:e \in s_i^*} (y_i(K_e(i) + 1)^2) + \gamma O_e^3 \geq K_e^3, \quad \forall e \in E \\
 & y_i \geq 0, \quad \forall i \in [n]
 \end{aligned}$$

From the above formulation, we obtain the following upper bound.

Theorem 8 For any $\mathcal{G} \in \mathcal{UCG}$ with quadratic latency functions, $\text{Ap}x_{\gamma}^1(\mathcal{G}) \leq 37.5888$.

Proof The worst-case dual constraints occur when each player i using resource e in s^* enters the game after all players using e in s have entered the game yet. The claim follows by choosing $y_i = 5.2944$ for any $i \in [n]$ and $\gamma = 37.5888$.

According to these choices, the dual constraints, namely $f_1(K_e, O_e)$ and $f_2(K_e, O_e)$, become

$$2059K_e^2 + 3309K_e(1 - 2O_e) + 6O_e(7831O_e - 1103) \geq 0$$

and

$$956K_e^3 + 3309K_e^2(1 - 2O_e) + 1103K_e(1 - 12O_e) + 42986O_e^3 - 6618O_e \geq 0.$$

For the first constraint, it holds that $f_1(K_e, 0) = 2059K_e^2 + 3309K_e \geq 0$ which is always satisfied for any $K_e \geq 0$. Moreover, the equation $f_1(K_e, O_e) = 0$ has no solution for any $O_e \geq 1$ (its discriminant is negative). Hence, the first constraint is always satisfied for any pair of non-negative integers (K_e, O_e) .

For the second constraint, note that $\frac{\delta f_2}{\delta O_e}(K_e, O_e) = -6(1103K_e^2 + 2206K_e - 23493O_e^2 + 1103)$ implies that, for any fixed value of K_e , f_2 is increasing in O_e when $O_e \geq 1$. Hence, we just need to show that $f_2(K_e, O_e) \geq 0$ for $O_e \in \{0, 1\}$. It holds that $f_2(K_e, 0) = 956K_e^3 + 3309K_e^2 + 1103K_e \geq 0$ which is always satisfied for any $K_e \geq 0$ and $f_2(K_e, 1) = 956K_e^3 - 3309K_e^2 - 12133K_e + 40368 \geq 0$ which is always satisfied for any non-negative integer. Hence, the second constraint is always satisfied for any pair of non-negative integers (K_e, O_e) . □

²These two metrics coincide in the case of dominant strategy equilibria.

For the case $d = 3$, $DLP(s, s^*)$ is defined as follows.

minimize γ
 subject to

$$\sum_{i:e \in s_i} (y_i(K_e(i) + 1)) - \sum_{i:e \in s_i^*} (y_i(K_e(i) + 1)) + \gamma O_e^2 \geq K_e^2, \quad \forall e \in E$$

$$\sum_{i:e \in s_i} (y_i(K_e(i) + 1)^2) - \sum_{i:e \in s_i^*} (y_i(K_e(i) + 1)^2) + \gamma O_e^3 \geq K_e^3, \quad \forall e \in E$$

$$\sum_{i:e \in s_i} (y_i(K_e(i) + 1)^3) - \sum_{i:e \in s_i^*} (y_i(K_e(i) + 1)^3) + \gamma O_e^4 \geq K_e^4, \quad \forall e \in E$$

$$y_i \geq 0, \quad \forall i \in [n]$$

From the above formulation, we obtain the following upper bound.

Theorem 9 For any $\mathcal{G} \in \mathcal{UCG}$ with cubic latency functions, $\text{Ap}x_{\emptyset}^1(\mathcal{G}) \leq \frac{17929}{34} \approx 527.323$.

Proof The worst-case dual constraints occur when each player i using resource e in s^* enters the game after all players using e in s have entered the game yet. The claim follows by choosing $y_i = \frac{369}{34}$ for any $i \in [n]$ and $\gamma = \frac{17929}{34}$.

According to these choices, the dual constraints, namely $f_1(K_e, O_e)$, $f_2(K_e, O_e)$ and $f_3(K_e, O_e)$, become

$$301K_e^2 + 3699K_e(1 - 2O_e) + 35858O_e^2 - 738O_e \geq 0,$$

$$178K_e^3 + 369K_e^2(1 - 2O_e) + 123K_e(1 - 12O_e) + 35858O_e^3 - 738O_e \geq 0$$

and

$$233K_e^4 + 738K_e^3(1 - 2O_e) + 369K_e^2(1 - 12O_e) - 4428K_eO_e + 71716O_e^4 - 1476O_e \geq 0.$$

For the first constraint, it holds that $f_1(K_e, 0) = 301K_e^2 + 369K_e \geq 0$ which is always satisfied for any $K_e \geq 0$. Moreover, the equation $f_1(K_e, O_e) = 0$ has no solution for any $O_e \geq 1$ (its discriminant is negative). Hence, the first constraint is always satisfied for any pair of non-negative integers (K_e, O_e) .

For the second constraint, note that $\frac{\delta f_2}{\delta O_e}(K_e, O_e) = -6(123K_e^2 + 246K_e - 17929O_e^2 + 123)$ implies that, for any fixed value of K_e , f_2 is increasing in O_e when $O_e \geq 1$. Hence, we just need to show that $f_2(K_e, O_e) \geq 0$ for $O_e \in \{0, 1\}$. It holds that $f_2(K_e, 0) = 178K_e^3 + 369K_e^2 + 123K_e \geq 0$ which is always satisfied for any $K_e \geq 0$ and $f_2(K_e, 1) = 178K_e^3 - 369K_e^2 - 1353K_e + 35120 \geq 0$ which is always satisfied for any non-negative integer. Hence, the second constraint is always satisfied for any pair of non-negative integers (K_e, O_e) .

For the third constraint, note that $\frac{\delta f_3}{\delta O_e}(K_e, O_e) = -4(369K_e^3 + 1107K_e^2 + 1107K_e - 71716O_e^3 + 369)$ implies that, for any fixed value of K_e , f_3 is increasing in O_e when $O_e \geq 1$. Hence, we just need to show that $f_3(K_e, O_e) \geq 0$ for $O_e \in \{0, 1\}$. It holds that $f_3(K_e, 0) = 233K_e^4 + 738K_e^3 + 369K_e^2 \geq 0$ which is always satisfied

for any $K_e \geq 0$ and $f_3(K_e, 1) = 233K_e^4 - 738K_e^3 - 4059K_e^2 - 4428K_e + 70240 \geq 0$ which is always satisfied for any non-negative integer. Hence, the third constraint is always satisfied for any pair of non-negative integers (K_e, O_e) . \square

6 Further Applications: the PoA Under the Social Function Max

In this section, we show how the primal-dual technique can be adapted also to the case in which the social function is the maximum of the players’ payoffs. For the sake of brevity, we consider only the problem of bounding the PoA in unweighted congestion games with affine latency functions. In order to deal with the maximum social function, we assume, without loss of generality, that player n is the one paying the highest cost in s and impose that, in s^* , no player pays more than one.

Thus, $LP(s, s^*)$ has the following form.

$$\begin{aligned}
 & \text{maximize } k \\
 & \text{subject to} \\
 & \sum_{e \in s_i} (\alpha_e K_e) - \sum_{e \in s_i^*} (\alpha_e (K_e + 1)) \leq 0, \quad \forall i \in [n] \\
 & \sum_{e \in s_i} (\alpha_e K_e) \leq k, \quad \forall i \in [n - 1] \\
 & \sum_{e \in s_n} (\alpha_e K_e) = k \\
 & \sum_{e \in s_i^*} (\alpha_e O_e) \leq 1, \quad \forall i \in [n] \\
 & \alpha_e \geq 0, \quad \forall e \in E
 \end{aligned}$$

$DLP(s, s^*)$ is as follows.

$$\begin{aligned}
 & \text{minimize } \sum_{i \in [n]} z_i \\
 & \text{subject to} \\
 & \sum_{i: e \in s_i} (K_e (x_i + y_i)) - \sum_{i: e \in s_i^*} (x_i (K_e + 1) - z_i O_e) \geq 0, \quad \forall e \in E \\
 & \sum_{i \in [n]} y_i \leq -1 \\
 & x_i, z_i \geq 0, \quad \forall i \in [n] \\
 & y_i \geq 0, \quad \forall i \in [n - 1]
 \end{aligned}$$

We easily reobtain the upper bound on the PoA proven by Christodoulou and Koutsoupias in [23].

Theorem 10 For any $\mathcal{G} \in \mathcal{UCG}$, $PoA(\mathcal{G}) = O(\sqrt{n})$.

Proof The claim follows by choosing $x_i = \frac{1}{\sqrt{n}}$, $y_i = 0$ and $z_i = \frac{2}{\sqrt{n}}$ for each $i \in [n - 1]$ and $x_n = 1$, $y_n = -1$ and $z_n = 2\sqrt{n}$. Such a choice clearly satisfies the second dual constraint, thus we just need to show that the first dual constraint, that we denote as $f(K_e, O_e) \geq 0$, is satisfied for any pair of non-negative integers. Note that $f(K_e, O_e)$ may assume different forms depending of which of the following four situations occurs:

- $e \notin s_n \wedge e \notin s_n^*$. In this case, $f(K_e, O_e) = K_e^2 - K_e O_e - O_e + 2O_e^2 \geq 0$ is always satisfied for any pair of (K_e, O_e) since the equation $f(K_e, O_e) = 0$ has no solution (its discriminant is negative).
- $e \notin s_n \wedge e \in s_n^*$. In this case, $f(K_e, O_e) = K_e^2 - K_e(\sqrt{n} + O_e - 1) + 2nO_e - \sqrt{n} + (O_e - 1)(2O_e - 1)$. Note that $\frac{\delta f}{\delta O_e}(K_e, O_e) = 2n - K_e + 4O_e - 3$ which is increasing in O_e . Since $e \in s_n^*$ implies $O_e \geq 1$, we just need to show that $f(K_e, O_e) \geq 0$ for $O_e = 1$. It holds that $f(K_e, 1) = K_e^2 - \sqrt{n}K_e - \sqrt{n} + 2n$ which is always satisfied for any value of K_e (it suffices noting that this inequality coincides with the one considered in the previous case when $O_e = \sqrt{n}$).
- $e \in s_n \wedge e \notin s_n^*$. In this case, $f(K_e, O_e) = K_e^2 - K_e(O_e + 1) - O_e + 2O_e^2$. The equation $f(K_e, O_e) = 0$ has no solution for $O_e \geq 2$ (its discriminant is negative), while it has a unique solution for $O_e = 1$. Hence, it suffices showing the claim for $O_e = 0$. It holds that $f(K_e, O_e) = K_e^2 - K_2 \geq 0$ which is always satisfied for any integer K_e .
- $e \in s_n \wedge e \in s_n^*$. In this case, $f(K_e, O_e) = K_e^2 - K_e(\sqrt{n} + O_e) + 2nO_e - \sqrt{n} + (O_e - 1)(2O_e - 1)$. Note that $\frac{\delta f}{\delta O_e}(K_e, O_e) = 2n - K_e + 4O_e - 3$ which is increasing in O_e . Since $e \in s_n^*$ implies $O_e \geq 1$, we just need to show that $f(K_e, O_e) \geq 0$ for $O_e = 1$. It holds that $f(K_e, 1) = K_e^2 - (\sqrt{n} + 1)K_e - \sqrt{n} + 2n$ which is always satisfied for any value of K_e (it suffices noting that this inequality coincides with the one considered in the previous case when $O_e = \sqrt{n}$).

The claim follows since $\sum_{i=1}^n z_i < 4\sqrt{n}$. □

7 Conclusions and Open Problems

In this paper, we have introduced a primal-dual method for bounding the quality of self-emerging solutions, namely (approximate) pure Nash equilibria and strategy profiles achieved after a one-round walk starting from the empty state, in weighted congestion games. Our technique has revealed itself to be particularly effective in this domain of application, as we could easily exploit it to reobtain, under a unifying approach, all the results already known in the literature for the case of affine latency functions as well as novel results for the cases of quadratic and cubic latency functions.

Our method differs significantly from the primal-dual method proposed by Nadav and Roughgarden in [37]. Roughly speaking, both methods consist in writing down a linear program aiming at measuring the “maximum distance” between a given strategy profile and the social optimum in a given game, but, while in their approach the

strategic choices of the players are treated as variables (whereas the parameters of the game defining the players’ costs/utilities are fixed constant values), in our formulations these choices are reversed. Our key idea is, in fact, that of fixing two “virtual strategy profiles”, representing the self-emerging solution under analysis and the social optimum, which, although fixed, are kept parametric so as to encompass the spectrum of all possible pairs of profiles, while the parameters defining the players’ costs/utilities in the game are now treated as variables in the formulation. The particularly compact succinct representation characterizing weighted congestion games makes them the natural application arena for our technique even in several contexts in which Nadav and Roughgarden’s method cannot be applied since the formulation becomes non-linear.

Several interesting open problems are left open and we discuss in the following some of them which, in our opinion, are worth to be investigated.

The first approach coming into mind is to export our primal-dual method outside the scope of weighted congestion games, by trying to apply it to other succinctly representable games. It would also be nice to understand whether our method can be suitably used to analyze the efficiency of other self-emerging solutions in congestion games. For instance, it is known that the PoA of Strong Nash equilibria in symmetric unweighted congestion games with affine latency functions is strictly smaller than $5/2$ [21] (which corresponds to the PoA of pure Nash equilibria), but no significant bounds have been given up to date.

Moreover, is it possible to strengthen some of the results already known in the literature for the case of unweighted congestion games with linear latencies? For instance, is it possible to narrow the gap between upper and lower bounds on the PoS_{ϵ} ? (To this aim, we believe that adding suitable constraints in the primal formulation should provide better upper bounds.)

Or is it possible to narrow the gap between upper and lower bounds on the Apx_{\emptyset}^1 when all the players’ strategies are made of a single resource? (In such a case, the best-known lower bound is 4 [20], while no improvements on the $2 + \sqrt{5} \approx 4.24$ upper bound are known for this special case.)

Appendix: How to Compute Dual Variables: an Example

In this section, we show how to compute the dual variables exploited within the proof of Theorem 1.

In order to find the best possible upper bound on the PoA_{ϵ} , we need to find the variables γ and y_i , for each $i \in [n]$, such that the dual constraint

$$\sum_{i:e \in s_i} (y_i K_e) - (1 + \epsilon) \sum_{i:e \in s_i^*} (y_i (K_e + w_i)) + \gamma O_e^2 \geq K_e^2$$

is satisfied for each pair of non-negative integers (K_e, O_e) .

Since in our games no player has a special role which may distinguish her from the others, we can assume, without loss of generality, that all players are

indistinguishable, so that we can set $y_i = y$ for each $i \in [n]$. With this simplification, the dual constraint becomes

$$yK_e^2 - (1 + \epsilon)yO_e(K_e + 1) + \gamma O_e^2 \geq K_e^2,$$

by which we obtain

$$\gamma \geq \frac{K_e^2(1 - y) + (1 + \epsilon)yO_e(K_e + 1)}{O_e^2}, \quad (8)$$

for $O_e \geq 1$ and $y \geq 1$ for $O_e = 0$. Moreover, it is also easy to see that, for $K_e \gg O_e$, the dual constraint is satisfied only if $y > 1$. Thus, from now on, we can assume, without loss of generality, that $y > 1$, $O_e \geq 1$ and focus on the satisfiability of (8).

Let us denote with $f(K_e, O_e, y)$ the right-hand side of (8). Our task is to find the value of y minimizing $f(K_e, O_e, y)$ for each pair of non-negative integers (K_e, O_e) with $O_e \geq 1$. We have $\frac{\delta f}{\delta K_e}(K_e, O_e, y) = \frac{y((1+\epsilon)O_e - 2K_e) + 2K_e}{O_e^2}$ which, being a linear function in K_e , is maximized for $K_e = \frac{(1+\epsilon)yO_e}{2(y-1)}$. By using $K_e = \frac{(1+\epsilon)yO_e}{2(y-1)}$ in $f(K_e, O_e, y)$, we obtain a function which is always decreasing in O_e ; hence, we can claim that the maximum value of $f(K_e, O_e, y)$ is attained for $K_e = \frac{(1+\epsilon)y}{2(y-1)}$ and $O_e = 1$.

However, the resulting lower bound on γ might not be tight, as K_e might not be an integer. In order to remain within the realm of the integers, we can assume that the maximum value of $f(K_e, O_e, y)$ is attained for $O_e = 1$ and K_e such that either $K_e = \left\lfloor \frac{(1+\epsilon)y}{2(y-1)} \right\rfloor$ or $K_e = \left\lceil \frac{(1+\epsilon)y}{2(y-1)} \right\rceil$. By setting $\left\lfloor \frac{(1+\epsilon)y}{2(y-1)} \right\rfloor = z$, this is equivalent to saying that the maximum value of $f(K_e, O_e, y)$ is attained for $O_e = 1$ and K_e such that either $K_e = z$ or $K_e = z + 1$. By imposing the condition $f(z, 1, y) = f(z + 1, 1, y)$, we obtain $y = \frac{2z+1}{2z-\epsilon}$ and $\gamma = f(z, 1, \frac{2z+1}{2z-\epsilon}) = (1 + \epsilon) \frac{z^2 + 3z + 1}{2z - \epsilon}$.

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.M., Tardos, E., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. *SIAM J. Comput.* **38**(4), 1602–1623 (2008)
3. Athanassopoulos, S., Caragiannis, I., Kaklamanis, C.: Analysis of Approximation algorithms for k -Set cover using factor-revealing linear programs. *Theory of Computing Systems* **45**(3), 555–576 (2009)
4. Awerbuch, B., Azar, Y., Epstein, A.: The price of routing unsplitable flow. In: Proceedings of the 37th annual ACM symposium on theory of computing (STOC), pp. 57–66. ACM Press, New York (2005)
5. Awerbuch, B., Azar, Y., Grove, E.F., Kao, M.-Y., Krishnan, P., Vitter, J.S.: Load balancing in the l_p norm. In: Proceedings of the 36th annual symposium on foundations of computer science (FOCS). IEEE Computer Society, pp. 383–391 (1995)
6. Bhawalkar, K., Gairing, M., Roughgarden, T.: Weighted congestion Games: price of anarchy, universal worst-Case examples, and tightness. *ACM Trans. Econ. Comput.* **2**(4), 14:1–14:23 (2014)
7. Bilò, V.: A unifying tool for bounding the quality of non-cooperative solutions in weighted congestion games. In: Proceedings of the 10th workshop on approximation and online algorithms (WAOA), LNCS 7846, pp. 215–228. Springer, Berlin (2012)

8. Bilò, V.: On linear congestion games with altruistic social context. In: Proceedings of the 20th international computing and combinatorics conference (COCOON), LNCS 8591, pp. 547–558. Springer, Berlin (2014)
9. Bilò, V.: On the robustness of the approximate price of anarchy in generalized congestion games. In: Proceedings of the 9th international symposium on algorithmic game theory (SAGT), LNCS 9928, pp. 93–1004. Springer, Berlin (2016)
10. Bilò, V., Fanelli, A., Flammini, M., Moscardelli, L.: Performances of One-Round walks in linear congestion games. *Theory of Computing Systems* **49**(1), 24–45 (2011)
11. Bilò, V., Fanelli, A., Moscardelli, L.: On lookahead equilibria in linear congestion games. *Math. Struct. Comput. Sci.* **27**(2), 197–214 (2017)
12. Bilò, V., Flammini, M., Gallotti, V.: On bidimensional congestion games. In: Proceedings of the 19th international colloquium on structural information and communication complexity (SIROCCO), LNCS 7355, pp. 147–158. Springer, Berlin (2012)
13. Bilò, V., Flammini, M., Monaco, G., Moscardelli, L.: Some anomalies of farsighted strategic behavior. *Theory of Computing Systems* **56**(1), 156–180 (2015)
14. Bilò, V., Paladini, M.: On the performance of mildly greedy players in cut games. *J. Comb. Optim.* **32**(4), 1036–1051 (2016)
15. Bilò, V., Vinci, C.: On stackelberg strategies in affine congestion games. In: Proceedings of the 11th international conference on Web and internet economics (WINE), LNCS 9470, pp. 132–145. Springer, Berlin (2015)
16. Bilò, V., Vinci, C.: Dynamic taxes for polynomial congestion games. In: Proceedings of the ACM conference on economics and computation (EC), pp. 839–856. ACM Press, New York (2016)
17. Caragiannis, I.: Wavelength management in WDM rings to maximize the number of connections. *SIAM J. Discret. Math.* **23**(2), 959–978 (2009)
18. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P.: Taxes for linear atomic congestion games. *ACM Trans. Algorithms* **7**(1), 13 (2010)
19. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., Kyropoulou, M., Papaioannou, E.: The impact of altruism on the efficiency of atomic congestion games. In: Proceedings of the 5th international symposium on trustworthy global computing (TGC), LNCS 6084, pp. 172–188. Springer, Berlin (2010)
20. Caragiannis, I., Flammini, M., Kaklamanis, C., Kanellopoulos, P., Moscardelli, L.: Tight bounds for selfish and greedy load balancing. *Algorithmica* **61**(3), 606–637 (2011)
21. Chien, S., Sinclair, A.: Strong and pareto price of anarchy in congestion games. In: Proceedings of the 36th international colloquium on automata, languages and programming (ICALP), LNCS, pp. 279–291. Springer, Berlin (2009)
22. Christodoulou, G., Gairing, M.: Price of stability in polynomial congestion games. *ACM Trans. Econ. Comput.* **42**(2), 10 (2016)
23. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: Proceedings of the 37th annual ACM Symposium on theory of computing (STOC), pp. 67–73. ACM Press, New York (2005)
24. Christodoulou, G., Koutsoupias, E.: On the price of anarchy and stability of correlated equilibria of linear congestion games. In: Proceedings of the 13th annual European symposium on algorithms (ESA), LNCS 3669, pp. 59–70. Springer, Berlin (2005)
25. Christodoulou, G., Koutsoupias, E., Spirakis, P.G.: On the performance of approximate equilibria in congestion games. *Algorithmica* **61**(1), 116–140 (2011)
26. Christodoulou, G., Mirrokni, V.S., Sidiropoulos, A.: Convergence and approximation in potential games. *Theor. Comput. Sci.* **438**, 13–27 (2012)
27. Feige, U., Jozeph, S.: Oblivious algorithms for the maximum directed cut problem. *Algorithmica* **71**(2), 404–428 (2015)
28. Fotakis, D.: Stackelberg strategies for atomic congestion games. *Theory of Computing Systems* **47**(1), 218–249 (2010)
29. Fotakis, D., Kontogiannis, S.C., Spirakis, P.G.: Symmetry in network congestion games pure equilibria and anarchy cost. In: Proceedings of the 3rd workshop on approximation and online algorithms (WAOA), LNCS 3879, pp. 161–175. Springer, Berlin (2005)
30. Harks, T., Klimm, M.: On the existence of pure nash equilibria in weighted congestion games. *Math. Oper. Res.* **37**(3), 419–436 (2012)
31. Harks, T., Klim, M., Möhring, R.H.: Characterizing the existence of potential functions in weighted congestion games. *Theory of Computing Systems* **49**(1), 46–70 (2011)

32. Koutsoupias, E., Papadimitriou, C.: Worst-Case Equilibria. In: Proceedings of the 16th International Symposium on Theoretical Aspects of Computer Science (STACS), LNCS 1563, pp. 404–413. Springer, Berlin (1999)
33. Jain, K., Mahdian, M., Markakis, E., Saberi, A., Vazirani, V.V.: Greedy facility location algorithms analyzed using dual fitting with Factor-Revealing LP. *J. ACM* **50**(6), 795–824 (2003)
34. Mahdian, M., Yan, Q.: Online bipartite matching with random arrivals: an approach based on strongly factor-revealing LPs. In: Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC), pp. 597–606. ACM Press, New York (2011)
35. Mirrokni, V.S., Thain, N., Vetta, A.: A theoretical examination of practical game playing lookahead search. In: Proceedings of the 5th international symposium on algorithmic game theory (SAGT), LNCS 7615, pp. 251–262. Springer, Berlin (2012)
36. Mirrokni, V.S., Vetta, A.: Convergence issues in competitive games. In: Proceedings of the 7th international workshop on approximation algorithms for combinatorial optimization problems (APPROX), LNCS 3122, pp. 183–194. Springer, Berlin (2004)
37. Nadav, U., Roughgarden, T.: The limits of smoothness: a primal-dual framework for price of anarchy bounds. In: Proceedings of the 6th international workshop on internet and network economics (WINE), LNCS 6484, pp. 319–326. Springer, Berlin (2010)
38. Paes Leme, R., Syrgkanis, V., Tardos, É.: The curse of simultaneity. In: Proceedings of innovations in theoretical computer science (ITCS), pp. 60–67. ACM press, New York (2012)
39. Panagopoulou, P.N., Spirakis, P.G.: Algorithms for pure nash equilibria in weighted congestion games. *ACM J. Exp. Algorithmics* **11**(2), 7 (2006)
40. Roughgarden, T.: Intrinsic robustness of the price of anarchy. *Commun. ACM* **55**(7), 116–123 (2012)
41. Rosenthal, R.W.: A class of games possessing pure-Strategy nash equilibria. *Int. J. Game Theory* **2**, 65–67 (1973)
42. Suri, S., Tóth, C., Zhou, Y.: Selfish load balancing and atomic congestion games. *Algorithmica* **47**(1), 79–96 (2007)