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

Vittorio Bilò

Department of Mathematics and Physics "Ennio De Giorgi", University of Salento, Provinciale Lecce-Arnesano, P.O. Box 193, 73100 Lecce, Italy
vittorio.bilo@unisalento.it

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 an 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 prices of anarchy and stability of exact and approximate pure Nash equilibria, as well as the approximation ratio of the solutions achieved after a one-round walk starting from the empty strategy profile, 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 solutions achieved after a one-round walk starting from the empty strategy profile for unweighted players in the cases of quadratic and cubic latency functions.

1 Introduction

Characterizing the quality of self-emerging solutions in non-cooperative systems is one of the leading research direction in Algorithmic Game Theory. Given a game $\mathcal{G}$, a social function $\mathcal{F}$ measuring the quality of any solution which can be realized in $\mathcal{G}$, and the definition of a set $\mathcal{E}$ of certain self-emerging solutions,

* This work was partially supported by the PRIN 2008 research project COGENT "Computational and game-theoretic aspects of uncoordinated networks" funded by the Italian Ministry of University and Research and by the "Progetto 5 per mille per la ricerca": "Collisioni fra vortici puntiformi e fra filamenti di vorticità: singolarità, trasporto e caos" at the University of Salento.

T. Erlebach and G. Persiano (Eds.): WAOA 2012, LNCS 7846, pp. 215–228, 2013.

we are asked to bound the ratio $\mathcal{Q}(\mathcal{G},\mathcal{E},\mathcal{F}) := \mathcal{F}(K)/\mathcal{F}(O)$, 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 solution optimizing $\mathcal{F}$.

In the last ten years, there has been a flourishing of contribution in this topic and, after a first flood of unrelated results, coming as a direct consequence of the fresh intellectual excitement caused by the affirmation of this new research direction, a novel approach, aimed at developing a more mature understanding of which is the big picture standing behind these problems and their solutions, is now arising.

In such a setting, Roughgarden [18] proposes the so-called "smoothness argument" as a unifying technique for proving tight upper bounds on $\mathcal{Q}(\mathcal{G},\mathcal{E},\mathcal{F})$ 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 defined as 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 [16], 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 $\mathcal{Q}(\mathcal{G},\mathcal{E},\mathcal{F})$ 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 [16] 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 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 [16], the Strong Duality Theorem assures that each primal-dual framework can always be used to derive the exact value of $\mathcal{Q}(\mathcal{G},\mathcal{E},\mathcal{F})$ provided that, for any solution S which can be realized in $\mathcal{G}$, $\mathcal{F}(S)$ can be expressed though linear programming and (i) the polyhedron defining $\mathcal{E}(\mathcal{G})$ can be expressed though 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 though 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 solutions (K, O) yield the exact value of $\mathcal{Q}(\mathcal{G},\mathcal{E},\mathcal{F})$, thus being able to construct quite systematically matching lower bounding instances.

We consider three sets of solutions $\mathcal{E}$, namely, (i) ϵ-approximate pure Nash equilibria (ϵ-PNE), that is, outcomes in which no player can improve her situation of more than an additive factor ϵ by unilaterally changing the adopted strategy (in this case, $\mathcal{Q}(\mathcal{G},\mathcal{E},\mathcal{F})$ is called the approximate price of anarchy of $\mathcal{G}$ (ϵ-PoA($\mathcal{G}$)) when K is the worst solution in $\mathcal{E}(\mathcal{G})$, while it is called the approximate price of stability of $\mathcal{G}$ (ϵ-PoS($\mathcal{G}$)) when K is the best solution 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 (by definition, each 0-PNE is a PNE and the terms price of anarchy (PoA($\mathcal{G}$)) and price of stability (PoS($\mathcal{G}$)) are used in this case); (*iii*) solutions achieved after a one-round walk starting from the empty strategy profile [15], 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 solution in $\mathcal{E}(\mathcal{G})$ and $\mathcal{Q}(\mathcal{G}, \mathcal{E}, \mathcal{F})$ is denoted by $\text{Apx}_{\emptyset}^{1}(\mathcal{G})$).

Our Contribution. Our method reveals to be particularly powerful when applied to the class of weighted congestion games. In these games there are n players competing for a set of resources. These games have a particular appeal since, from 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 [19] proved through a potential function argument that PNE are always guaranteed to exist, while general weighted congestion games are guaranteed to possess PNE if and only if the latency functions are either affine or exponential [11–13, 17].

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 associated with the resources 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 ϵ-PoA and ϵ-PoS in the unweighted case and for $\text{Apx}_{\emptyset}^{1}$, we reobtain known upper bounds with significatively 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 ϵ-PoA and the ϵ-PoS 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 case, the PoA and ϵ-PoA was already studied and characterized in [1] and [9], respectively, and both papers pose the achievement of upper bounds on the PoS and ϵ-PoS as a major open problem in the area. For unweighted players, we show that PoS ≤ 2.362 and $\text{Apx}_{\emptyset}^{1} \leq 37.5888$ for quadratic latency functions and that PoS ≤ 3.322 and $\text{Apx}_{\emptyset}^{1} \leq 527.323$ for cubic latency functions. Extensions to ϵ-PoS and weighted players are left to future work.

What we would like to stress here is that, more than the novelty of the results achieved in this paper, what makes our method significative 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. To this aim, in the full version of this paper, we show how the method applies also to other social functions, such as the maximum of the players' payoffs, and to other subclasses of congestion games such as resource allocation games with fair cost sharing. Note that, in the latter case, as well as in the case of polynomial latency functions, the primal-dual technique proposed in [16] cannot be used, since the players' costs are not linear in the variables of the problem.

Related Works. The study of the quality of self-emerging solutions in non-cooperative systems initiated with the seminal papers of Koutsoupias and Papadimitriou [14] and Anshelevich et al. [2] which introduced, respectively, the notions of price of anarchy and price of stability.

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 [7] show that the PoA is exactly $5/2$ for unweighted players, while Awerbuch et al. [3] 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 [8]. 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 [18] who proves that unweighted congestion games with latency functions constrained in a given set belong to the class of games for which a so-called "smoothness argument" applies and that such a smoothness argument directly implies the fact that the prices of anarchy of PNE, mixed Nash equilibria, correlated equilibria and coarse correlated equilibria are always the same when $\mathcal{F}$ is the sum of the players' payoffs. Such a result has been extended also to the weighted case by Bhawalkar et al. in [4]. For the alternative model in which $\mathcal{F}$ is defined as the maximum of the players' payoffs, Christodoulou and Koutsoupias [7] 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. [6] and by Christodoulou and Koutsoupias [8], respectively, set the PoS to exactly $1+1/\sqrt{3}$. Clearly, being the PoS a best-case measure and being the set of PNE properly contained in the set of all the other equilibrium concepts, again we have a unique bound for PNE and all of its generalizations.

As to ϵ-PNE, in the case of unweighted players, polynomial latency functions and $\mathcal{F}$ defined as the sum of the players' payoffs, Christodoulou et al. [9] show that the ϵ-PoA 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

ϵ-PoS is at least $\frac{2(3+\epsilon+\theta\epsilon^2+3\epsilon^3+2\epsilon^4+\theta+\theta\epsilon)}{6+2\epsilon+5\theta\epsilon+6\epsilon^3+4\epsilon^4-\theta\epsilon^3+2\theta\epsilon^2}$, 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, $\mathrm{Apx}^1_\emptyset$ 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 et al. [10] and by Bilò et al. [5], while, for weighted players, Caragiannis et al. [6] give a lower bound of $3+2\sqrt{2}$ and Christodoulou et al. [10] give an upper bound of $4+2\sqrt{3}$. For $\mathcal{F}$ being the maximum of the players' payoffs, Bilò et al. [5] show that $\mathrm{Apx}^1_\emptyset$ is $\Theta(\sqrt[4]{n^3})$ in the unweighted case and affine latency functions.

Paper Organization. In 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, finally, in Section 5 we use it to address the case of quadratic and cubic latency functions. Due to lack of space, some proofs are omitted and can be found in the full version of the paper.

2 Definitions

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

A *weighted congestion game* $\mathcal{G} = \big([n], E, (\Sigma_i)_{i\in[n]}, (\ell_e)_{e\in E}, (w_i)_{i\in[n]}\big)$ is a non-cooperative strategic game in which there is a set E of m *resources* to be shared among the players in $[n]$. Each player i has an associated weight $w_i \geq 1$ 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 *solutions*) which can be realized in $\mathcal{G}$. Given a strategy profile $S = (s_1, s_2, \ldots, s_n) \in \Sigma$ 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 $0 \leq i \leq d$.

Given a strategy profile $S \in \Sigma$ and a strategy $t \in \Sigma_i$ for player i, we denote with $(S_{-i} \diamond t) = (s_1, \ldots, s_{i-1}, t, s_{i+1}, \ldots, s_n)$ the strategy profile obtained from S when player i changes unilaterally her strategy from s_i to t. We say that $S' = (S_{-i} \diamond t)$ is an *improving deviation* for player i in S if $c_i(S') < c_i(S)$. Given an $\epsilon \geq 0$, a strategy profile S is an *ϵ-approximate pure Nash equilibrium* (ϵ-PNE) if, for any $i \in [n]$ and for any $t \in \Sigma_i$, it holds $c_i(S) \leq (1+\epsilon)c_i(S_{-i} \diamond t)$. For $\epsilon = 0$, the set of ϵ-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 $\textsc{Sum} : \Sigma \mapsto \mathbb{R}_{\geq 0}$ defined as the sum of the players' costs, that is, $\textsc{Sum}(S) = \sum_{i\in[n]} c_i(S)$ and let S^* be the strategy profile minimizing it. Given an $\epsilon \geq 0$ and a weighted congestion game $\mathcal{G}$, let $\mathcal{E}(\mathcal{G})$ be the set of ϵ-approximate Nash equilibria of $\mathcal{G}$. The ϵ-approximate price of anarchy of $\mathcal{G}$ is defined as $\epsilon\text{-PoA}(\mathcal{G}) = \max_{S\in\mathcal{E}(\mathcal{G})} \left\{ \frac{\textsc{Sum}(S)}{\textsc{Sum}(S^*)} \right\}$, while the ϵ-approximate price of stability of $\mathcal{G}$ is defined as $\epsilon\text{-PoS}(\mathcal{G}) = \min_{S\in\mathcal{E}(\mathcal{G})} \left\{ \frac{\textsc{Sum}(S)}{\textsc{Sum}(S^*)} \right\}$.

Given a strategy profile S and a player $i \in [n]$, a strategy profile $t^* \in \Sigma_i$ is a *best-response* for player i in S if it holds $c_i(S_{-i} \diamond t^*) \leq c_i(S_{-i} \diamond t)$ for any $t \in \Sigma_i$. Let $S^{\emptyset}$ be the *empty strategy profile*, i.e., the profile in which no player has performed any strategic choice yet. A one-round walk starting from the empty strategy profile is an $(n+1)$-tuple of strategy profiles $W = (S_0^W, S_1^W, \ldots, S_n^W)$ 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 solution 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 $\mathcal{G}$. The approximation ratio of the solutions achieved after a one-round walk starting from the empty strategy profile in $\mathcal{G}$ is defined as $\text{Apx}_{\emptyset}^1(\mathcal{G}) = \max_{W\in\mathcal{W}(\mathcal{G})} \left\{ \frac{\textsc{Sum}(S_n^W)}{\textsc{Sum}(S^*)} \right\}$.

3 The Primal-Dual Technique

Fix a weighted congestion game $\mathcal{G}$, a social function $\mathcal{F}$ and a class of self-emerging solutions $\mathcal{E}$. Let $O = (s_1^O, \ldots, s_n^O)$ be the strategy profile optimizing $\mathcal{F}$ and $K = (s_1^K, \ldots, s_n^K) \in \mathcal{E}(\mathcal{G})$ be the worst-case solution in $\mathcal{E}(\mathcal{G})$ with respect to $\mathcal{F}$. For any $e \in E$, we set, for the sake of brevity, $O_e = \mathcal{L}_e(O)$ and $K_e = \mathcal{L}_e(K)$.

Since O and K are fixed, we can maximize the inefficiency yielded by the pair (K, O) by suitably choosing the coefficients $\alpha_{e,i}$, for each $e \in E$ and $0 \leq i \leq d$, so that $\mathcal{F}(K)$ is maximized, $\mathcal{F}(O)$ is normalized to one and K 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 $LP(K, O)$.

By providing a feasible solution for the dual program $DLP(K, O)$, we can obtain an upper bound on the optimal solution of $LP(K, O)$. Our task is to uncover, among all possibilities, the pair (K^*, O^*) yielding the highest possible optimal solution for $LP(K, O)$. To this aim, the study of the dual formulation plays a crucial role: 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 (K^*, O^*) maximizing the inefficiency of the class of solutions $\mathcal{E}$. 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, because, once established the values of the congestions K_e^* and O_e^* for any $e \in E$, there are still infinite many ways to split them among the players using resource e in both K and O.

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 ϵ-PoA, ϵ-PoS and $\text{Apx}_{\emptyset}^{1}$ already known in the literature can be reobtained in a unified manner for both weighted and unweighted players.

We say that the game $\mathcal{G}' = ([n], E', \Sigma', \ell', w)$ is equivalent to the game $\mathcal{G} = ([n], E, \Sigma, \ell, w)$ if there exists a one-to-one mapping $\varphi_i : \Sigma_i \mapsto \Sigma'_i$ for any $i \in [n]$ such that for any strategy profile $S = (s_1, \ldots, s_n) \in \Sigma$, it holds $c_i(S) = c_i(\varphi_1(s_1), \ldots, \varphi_n(s_n))$ for any $i \in [n]$.

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

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

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

By definition, we have that if K is an ϵ-PNE then, for any $i \in [n]$, it holds

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

Thus, the primal formulation $LP(K,O)$ assumes the following form.

$$\begin{aligned} & \textit{maximize} \sum_{e\in E} \left(\alpha_e K_e^2\right) \\ & \textit{subject to} \\ \sum_{e\in s_i^K} (\alpha_e K_e) - (1+\epsilon) \sum_{e\in s_i^O} (\alpha_e(K_e+w_i)) &\le 0, \qquad \forall i \in [n] \\ \sum_{e\in E} \left(\alpha_e O_e^2\right) &= 1, \\ \alpha_e &\ge 0, \qquad \forall e \in E \end{aligned}$$

The dual program $DLP(K,O)$ is

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

Let $\psi = \frac{1+\epsilon+\sqrt{\epsilon^2+6\epsilon+5}}{2}$ and $z = \lfloor \psi \rfloor$. For unweighted players we reobtain the upper bound proved in [9] with a much simpler and shorter proof, while for the weighted case we give the first known upper bound.

Theorem 1. *For any $\epsilon \ge 0$, it holds ϵ-PoA$(\mathcal{G}) \le \frac{(1+\epsilon)(z^2+3z+1)}{2z-\epsilon}$ when $\mathcal{G}$ has unweighted players, while it holds ϵ-PoA$(\mathcal{G}) \le \psi^2$ when $\mathcal{G}$ has weighted players.*

Proof. For the unweighted case, since $w_i = 1$ for each $i \in [n]$, by choosing $y_i = \frac{2z+1}{2z-\epsilon}$ for any $i \in [n]$ and $\gamma = \frac{(1+\epsilon)(z^2+3z+1)}{2z-\epsilon}$, the dual inequalities become of the form

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

which is equivalent to

$$K_e^2 - (2z+1)(K_e O_e + O_e) + (z^2+3z+1)O_e^2 \ge 0.$$

Easy calculations (it suffices solving the inequality for K_e) show that this is always verified for any pair of non-negative *integers* (K_e, O_e). Note that the definition of z guarantees that $2z-\epsilon \ge 0$, so that the proposed dual solution is a feasible one.

For the weighted case, by choosing $y_i = \left(1 + \frac{\sqrt{1+\epsilon}}{\sqrt{5+\epsilon}}\right) w_i$ for any $i \in [n]$ and $\gamma = \psi^2$, each dual inequality is verified when it holds

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

which is equivalent to

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

Easy calculations show that this is always verified for any pair of non-negative *reals* (K_e, O_e) such that $K_e, O_e \in \{0\} \cup [1, \infty)$ for any $e \in E$ (it suffices solving the inequality for K_e and noting that it has no solutions for $O_e \in \{0\} \cup [1, \infty)$ when $\epsilon \geq 0$). □

When $\epsilon = 0$, we reobtain the well-known prices of anarchy of $5/2$ and $(3+\sqrt{5})/2$ which hold for PNE in the unweighted and weighted case, respectively.

4.2 Bounding the Approximate Price of Stability

Recall that, since the ϵ-PoS 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 ϵ-PNE. It is not known how to do this at the moment, thus all the approaches used so far in the literature approximate the best ϵ-PNE with an ϵ-PNE minimizing a certain potential function. In [9], it is shown that, for unweighted players, any strategy profile S which is a local minimum of the function

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

called ϵ-approximate potential, is an ϵ-PNE. Thus, it is possible to get an upper bound on the ϵ-PoS by measuring the ϵ-PoA 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 in [9]. Assume that K is the global minimum of Φ_ϵ. We can use the inequality $\Phi_\epsilon(K) \leq \Phi_\epsilon(O)$ which results in the constraint

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

Then, we also have $\sum_{i\in[n]}\left(\Phi_\epsilon(K)-\Phi_\epsilon(K_{-i}\diamond s_i^O)\right) \leq 0$ which results in the constraint

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

Thanks to these two constraints, the dual formulation becomes the following one.

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

Thus, for unweighted players, we obtain the following result for any $\epsilon \in [0,1)$ (this is the only interesting case, since [9] shows that, for any $\epsilon \geq 1$, ϵ-PoS$(\mathcal{G}) = 1$ for any $\mathcal{G}$).

Theorem 2. *For any $\epsilon \in [0,1)$ and $\mathcal{G}$ with unweighted players, it holds ϵ-PoS$(\mathcal{G}) \leq \frac{1+\sqrt{3}}{\epsilon+\sqrt{3}}$.*

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 dual inequalities become

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

Easy calculations (it suffices solving the inequality for K_e) show that this is always verified for any pair of non-negative *integers* (K_e, O_e). □

4.3 Bounding the Approximation Ratio of One-Round Walks

For a one-round walk W, we set $K = S_n^W$. Define $K_e(i)$ as the sum of the weights of the players using resource e in K before player i performs her choice. $LP(K,O)$ in this case has the following form, where the first constraint comes from the fact that when player i enters the game and solution S_{i-1}^K is already constructed, this player picks s_i^K instead of s_i^O.

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

$DLP(K,O)$ is as follows.

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

For both unweighted and weighted players we easily reobtain the upper bounds on $\text{Apx}_{\emptyset}^1$ given in [10].

Theorem 3. *For any $\mathcal{G}$ with unweighted players, it holds $Apx_{\emptyset}^1(\mathcal{G}) \leq 2 + \sqrt{5}$, while for any $\mathcal{G}$ with weighted players, it holds $Apx_{\emptyset}^1(\mathcal{G}) \leq 4 + 2\sqrt{3}$.*

Proof. In the unweighted case, by choosing $y_i = 1 + \sqrt{5}$ for any $i \in [n]$ and $\gamma = 2 + \sqrt{5}$, since for any i such that $e \in s_i^O$ it holds $K_e(i) \leq K_e$, each dual inequality is verified when it holds

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

which is equivalent to

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

Easy calculations (it suffices solving the inequality for K_e) show that this is always verified for any pair of non-negative *integers* (K_e, O_e).

In the weighted case, 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 $e \in s_i^O$ it holds $K_e(i) \leq K_e$, each dual inequality is verified when it holds

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

which is true if it holds

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

Easy calculations (it suffices solving the inequality for K_e) show that this is always verified for any pair of non-negative *reals* (K_e, O_e). □

5 Quadratic and Cubic Latency Functions

In this section, we show how to use the primal-dual method to bound PoS and $\mathrm{Apx}_{\emptyset}^1$ 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. 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. We also leave the extension to ϵ-PNE and weighted players for future works. We restrict to the cases in which the latency functions are of the form $\ell_e(x) = \alpha_e x^d$, since it is possible to show that this can be supposed without loss of generality. In such a setting, [19] shows that, for unweighted players, any strategy profile S which is a local minimum of the potential function

$$\Phi(S) = \sum_{e \in E} \sum_{i=1}^{\mathcal{L}_e(S)} (\alpha_e x^d)$$

is a PNE.

5.1 Bounding the Price of Stability

For the case $d = 2$, it holds

$$\Phi(S) = \frac{1}{6}\sum_{e\in E}\left(\alpha_e \mathcal{L}_e(S)(\mathcal{L}_e(S)+1)(2\mathcal{L}_e(S)+1)\right).$$

Thus, the constraint $\Phi(K) \leq \Phi(O)$ becomes

$$\sum_{e\in E}\left(\alpha_e\left(K_e(K_e+1)(2K_e+1) - O_e(O_e+1)(2O_e+1)\right)\right) \leq 0$$

and the constraint $\sum_{i\in[n]}\left(\Phi(K) - \Phi(K_{-i}\diamond s_i^O)\right) \leq 0$ becomes

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

Hence, $DLP(K,O)$ is the following.

$$\begin{array}{c} \textit{minimize}\ \gamma \\ \textit{subject to} \\ \left(y(K_e(K_e+1)(2K_e+1) - O_e(O_e+1)(2O_e+1))\right) \\ +\left(z(K_e^3 - O_e(K_e+1)^2)\right) + \gamma O_e^3 \geq K_e^3,\ \forall e\in E \\ y,z\geq 0 \end{array}$$

Theorem 4. *For any $\mathcal{G}$ with quadratic latency functions and unweighted players, it holds $PoS(\mathcal{G}) \leq 2.362$.*

Proof. The claim follows by setting $y = 0.318$, $z = 0.453$ and $\gamma = 2.362$. □

For the case $d = 3$, it holds

$$\Phi(S) = \frac{1}{4}\sum_{e\in E}\left(\alpha_e\left(\mathcal{L}_e(S)(\mathcal{L}_e(S)+1)\right)^2\right).$$

Thus, the constraint $\Phi(K) \leq \Phi(O)$ becomes

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

and the constraint $\sum_{i\in[n]}\left(\Phi(K) - \Phi(K_{-i}\diamond s_i^O)\right) \leq 0$ becomes

$$\sum_{e\in E}\left(\alpha_e\left(K_e^4 - O_e(K_e+1)^3\right)\right) \leq 0.$$

Hence, $DLP(K,O)$ is defined as follows.

$$\begin{array}{c} \textit{minimize}\ \gamma \\ \textit{subject to} \\ \left(y(K_e^2(K_e+1)^2 - O_e^2(O_e+1)^2)\right) \\ +\left(z(K_e^4 - O_e(K_e+1)^3)\right) + \gamma O_e^4 \geq K_e^4,\ \forall e\in E \\ y,z\geq 0 \end{array}$$

Theorem 5. *For any $\mathcal{G}$ with cubic latency functions and unweighted players, it holds $PoS(\mathcal{G}) \leq 3.322$.*

Proof. The claim follows by setting $y = 0.747$, $z = 0.331$ and $\gamma = 3.322$. □

By extending the instance given in [9] for lower bounding the PoS in the linear case, the following lower bounds can be easily achieved.

Theorem 6. *For any $\delta > 0$, there exist an unweighted congestion game with quadratic latency functions $\mathcal{G}_1$ and an unweighted congestion game with cubic latency functions $\mathcal{G}_2$ such that $PoS(\mathcal{G}_1) \geq 2.1859 - \delta$ and $PoS(\mathcal{G}_2) \geq 2.7558 - \delta$.*

5.2 Bounding the Approximation Ratio of One-Round Walks

For the case $d = 2$, $DLP(K, O)$ is defined as follows.

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

Theorem 7. *For any $\mathcal{G}$ with quadratic latency functions and unweighted players, it holds $Apx_\emptyset^1(\mathcal{G}) \leq 37.5888$.*

For the case $d = 3$, $DLP(K, O)$ is defined as follows.

$$\begin{aligned} &\textit{minimize } \gamma \\ &\textit{subject to} \\ &\sum_{i:e \in s_i^K} \left(y_i(K_e(i)+1)^3\right) - \sum_{i:e \in s_i^O} \left(y_i(K_e(i)+1)^3\right) + \gamma O_e^4 \geq K_e^4, \ \forall e \in E \\ &y_i \geq 0, \quad \forall i \in [n] \end{aligned}$$

Theorem 8. *For any $\mathcal{G}$ with cubic latency functions and unweighted players, it holds $Apx_\emptyset^1(\mathcal{G}) \leq \frac{17929}{34} \approx 527.323$.*

References

1. Aland, S., Dumrauf, D., Gairing, M., Monien, B., Schoppmann, F.: Exact Price of Anarchy for Polynomial Congestion Games. In: Durand, B., Thomas, W. (eds.) STACS 2006. LNCS, vol. 3884, pp. 218–229. Springer, Heidelberg (2006)
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 Journal on Computing 38(4), 1602–1623 (2008)
3. Awerbuch, B., Azar, Y., Epstein, A.: The Price of Routing Unsplittable Flow. In: Proceedings of the 37th Annual ACM Symposium on Theory of Computing, pp. 57–66. ACM Press (2005)

4. Bhawalkar, K., Gairing, M., Roughgarden, T.: Weighted Congestion Games: Price of Anarchy, Universal Worst-Case Examples, and Tightness. In: de Berg, M., Meyer, U. (eds.) ESA 2010, Part II. LNCS, vol. 6347, pp. 17–28. Springer, Heidelberg (2010)
5. Bilò, V., Fanelli, A., Flammini, M., Moscardelli, L.: Performances of One-Round Walks in Linear Congestion Games. In: Mavronicolas, M., Papadopoulou, V.G. (eds.) SAGT 2009. LNCS, vol. 5814, pp. 311–322. Springer, Heidelberg (2009)
6. Caragiannis, I., Flammini, M., Kaklamanis, C., Kanellopoulos, P., Moscardelli, L.: Tight Bounds for Selfish and Greedy Load Balancing. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4051, pp. 311–322. Springer, Heidelberg (2006)
7. Christodoulou, G., Koutsoupias, E.: The Price of Anarchy of Finite Congestion Games. In: Proceedings of the 37th Annual ACM Symposium on Theory of Computing, pp. 67–73. ACM Press (2005)
8. Christodoulou, G., Koutsoupias, E.: On the Price of Anarchy and Stability of Correlated Equilibria of Linear Congestion Games. In: Brodal, G.S., Leonardi, S. (eds.) ESA 2005. LNCS, vol. 3669, pp. 59–70. Springer, Heidelberg (2005)
9. Christodoulou, G., Koutsoupias, E., Spirakis, P.G.: On the Performance of Approximate Equilibria in Congestion Games. Algorithmica (to appear)
10. Christodoulou, G., Mirrokni, V.S., Sidiropoulos, A.: Convergence and Approximation in Potential Games. In: Durand, B., Thomas, W. (eds.) STACS 2006. LNCS, vol. 3884, pp. 349–360. Springer, Heidelberg (2006)
11. Fotakis, D., Kontogiannis, S.C., Spirakis, P.G.: Symmetry in Network Congestion Games: Pure Equilibria and Anarchy Cost. In: Erlebach, T., Persinao, G. (eds.) WAOA 2005. LNCS, vol. 3879, pp. 161–175. Springer, Heidelberg (2006)
12. Harks, T., Klimm, M.: On the Existence of Pure Nash Equilibria in Weighted Congestion Games. In: Abramsky, S., Gavoille, C., Kirchner, C., Meyer auf der Heide, F., Spirakis, P.G. (eds.) ICALP 2010. LNCS, vol. 6198, pp. 79–89. Springer, Heidelberg (2010)
13. Harks, T., Klimm, M., Möhring, R.H.: Characterizing the Existence of Potential Functions in Weighted Congestion Games. In: Mavronicolas, M., Papadopoulou, V.G. (eds.) SAGT 2009. LNCS, vol. 5814, pp. 97–108. Springer, Heidelberg (2009)
14. Koutsoupias, E., Papadimitriou, C.: Worst-Case Equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999)
15. Mirrokni, V.S., Vetta, A.: Convergence Issues in Competitive Games. In: Jansen, K., Khanna, S., Rolim, J.D.P., Ron, D. (eds.) APPROX and RANDOM 2004. LNCS, vol. 3122, pp. 183–194. Springer, Heidelberg (2004)
16. Nadav, U., Roughgarden, T.: The Limits of Smoothness: A Primal-Dual Framework for Price of Anarchy Bounds. In: Saberi, A. (ed.) WINE 2010. LNCS, vol. 6484, pp. 319–326. Springer, Heidelberg (2010)
17. Panagopoulou, P.N., Spirakis, P.G.: Algorithms for Pure Nash Equilibria in Weighted Congestion Games. Journal of Experimental Algorithmics 11(2.7) (2006)
18. Roughgarden, T.: Intrinsic Robustness of the Price of Anarchy. In: Proceedings of the 41st Annual ACM Symposium on Theory of Computing, pp. 513–522. ACM Press (2009)
19. Rosenthal, R.W.: A Class of Games Possessing Pure-Strategy Nash Equilibria. International Journal of Game Theory 2, 65–67 (1973)