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

Vittorio Bilò [1]

[1]Dipartimento di Matematica e Fisica "Ennio De Giorgi", Università del Salento, Italy

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, 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 [6] 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 [5], 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 [5] 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. 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 [5], 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, 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 $\mathrm{Apx}^1_\emptyset(\mathcal{G})$).

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 [7] 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 [3,4].

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 $\mathrm{Apx}^1_\emptyset$, 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 were already studied and characterized in [1] and [2], 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, for any congestion game $\mathcal{G}$ with quadratic latency functions, $\mathrm{PoS}(\mathcal{G}) \leq 2.362$ and $\mathrm{Apx}^1_\emptyset(\mathcal{G}) \leq 37.5888$ and that, for any congestion game $\mathcal{G}$ with cubic latency functions, $\mathrm{PoS}(\mathcal{G}) \leq 3.322$ and $\mathrm{Apx}^1_\emptyset(\mathcal{G}) \leq 527.323$.

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, 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 [5] cannot be used, since the players' costs are not linear in the variables of the problem).

REFERENCES

1. S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. Exact Price of Anarchy for Polynomial Congestion Games. *SIAM Journal on Computing*, 40(5): 1211–1233, 2011.
2. G. Christodoulou, E. Koutsoupias, and P. G. Spirakis. On the Performance of Approximate Equilibria in Congestion Games. *Algorithmica*, 61(1): 116–140, 2011.
3. T. Harks and M. Klimm. On the Existence of Pure Nash Equilibria in Weighted Congestion Games. *Mathematics of Operations Research*, 37(3): 419–436, 2012.
4. T. Harks, M. Klim and R.H. Möhring. Characterizing the Existence of Potential Functions in Weighted Congestion Games. *Theory of Computing Systems*, 49(1): 46–70, 2011.
5. U. Nadav and T. Roughgarden. 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, Springer, pp. 319–326, 2010.
6. T. Roughgarden. Intrinsic Robustness of the Price of Anarchy. *Commununications of the ACM*, 55(7): 116–123, 2012.
7. R. W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory*, 2:65–67, 1973.