When Smoothness is Not Enough: Toward Exact Quantification and Optimization of the Price-of-Anarchy

Rahul Chandan Dario Paccagnan Jason R. Marden

***Abstract*— Today's multiagent systems have grown too complex to rely on centralized controllers, prompting increasing interest in the design of distributed algorithms. In this respect, game theory has emerged as a valuable tool to complement more traditional techniques. The fundamental idea behind this approach is the assignment of agents' local cost functions, such that their selfish minimization attains, or is provably close to, the global objective. Any algorithm capable of computing an equilibrium of the corresponding game will inherit an approximation ratio that is, in the worst case, equal to the price-of-anarchy of the considered class of equilibria. Therefore, a successful application of the game design approach hinges on the possibility to quantify and optimize the equilibrium performance.**

Toward this end, we introduce the notion of generalized smoothness, and show that the resulting efficiency bounds are significantly tighter compared to those obtained using the traditional smoothness approach. Leveraging this newly-introduced notion, we quantify the equilibrium performance for the class of local resource allocation games. Finally, we show how the agents' local decision rules can be designed in order to optimize the efficiency of the corresponding equilibria, by means of a tractable linear program.

I. Introduction

Interest in the field of multiagent systems' control has experienced rapid growth in recent years, as a variety of application domains have emerged [1], [2]. The impact of recent advancements in multiagent control has been far-reaching, revolutionizing traditional industries such as transportation and power networks [3], [4], [5], while also driving the development of novel technologies including robotic swarms and self-driving cars [6], [7].

Modern multiagent systems must adhere to imposing constraints with regards to their spatial distribution, overall scale, privacy requirements and communication bandwidth. As a consequence, the coordination of such systems does not allow for centralized decision making, but instead requires the use of distributed protocols. Ideally, a distributed algorithm will meet the system's requirements for scalability, communication bandwidth, and security, while achieving the desired global objective.

A well-established and fruitful approach to tackle this class of problems consists in the design of a centralized maximization algorithm, that is later distributed by leveraging the structure of the problem considered, e.g., [8], [9]. An alternative approach, termed *game design*, has emerged in parallel as a valuable tool to complement the aforementioned design philosophy [10]. Instead of directly specifying the decision-making process, local cost functions are assigned to the system's agents such that their selfish minimization results in the achievement of the system-level objective.

The advantages of using this approach are two-fold: i) we inherit a pool of algorithms that are distributed by nature, asynchronous, and resilient to external disturbances [11]; and, ii) we obtain access to readily-available performance certificates in the form of efficiency bounds. In fact, any (distributed) algorithm capable of driving the system to an equilibrium configuration (e.g. pure Nash equilibrium, mixed Nash equilibrium, correlated equilibrium, etc.) will inherit an approximation ratio matching the corresponding worst-case equilibrium efficiency, called the *price-of-anarchy*. Motivated by the game-theoretic approach, we aim to develop novel techniques to quantify and minimize the price-of-anarchy in distributed systems.

A. Related Works

This work is inspired by [12] and [13] where, the authors quantify the price of anarchy of covering problems using an approach reminiscent of - but different from - that used in the smoothness framework.

While smoothness arguments [14], [15] have proven useful when characterizing the performance of broad classes of equilibria [16], they have also been applied to a variety of problems, including learning [17], and mechanism design [18]. Unfortunately, as observed in [19] and proven later in this manuscript, traditional smoothness arguments find limited applicability in connection to *design problems*. Generalized smoothness is tailored to resolve this weakness, while retaining all the strengths of the traditional smoothness approach. Leveraging this novel notion, we extend the linear programming approach of [19], [20] and show how to compute and optimize the price-of-anarchy of coarse-correlated equilibria, relative to any local resource allocation game.

B. Our Contributions

In this work, we introduce a broader notion of smoothness, referred to as generalized smoothness, allowing us to provide tighter bounds on the performance of coarse-correlated equilibria. To demonstrate the strength of this novel approach, we apply our result to a general design problem, and show that the bounds are tight. Our main contributions are listed below.

This work is supported by ONR grant #N00014-17-1-2060, NSF grant #ECCS-1638214, and SNSF grant #P2EZP2_181618

R. Chandan and J.R. Marden are with the Department of Electrical and Computer Engineering and the Center of Control, Dynamical Systems and Computation, UC Santa Barbara, USA. Email: rchandan@ucsb.edu, jrmarden@ece.ucsb.edu.

D. Paccagnan is with the Department of Mechanical Engineering and the Center of Control, Dynamical Systems and Computation, UC Santa Barbara, USA. Email: dariop@ucsb.edu.

1) We demonstrate that price-of-anarchy bounds obtained via smoothness arguments are not tight if the sum of players' local cost functions is not equal to the system cost (Theorem 1).
2) We introduce the notion of generalized smoothness, and show that, in general, it provides tighter bounds on the price-of-anarchy compared to current smoothness approaches (Theorems 2 and 3).
3) For the class of local resource allocation problems, we show that generalized smoothness provides tight bounds on the price-of-anarchy (Theorem 4). As a consequence, we show how the price-of-anarchy can be characterized (Theorem 5) and optimized (Theorem 6) using tractable linear programs.

Although we consider only cost-minimization games, the results can be applied to welfare-maximization games with minor modification of the generalized smoothness criterion.

II. Problem Statement

Consider a class of resource allocation problems where $N = \{1, \dots, n\}$ denotes the set of agents, and each agent i must select an action a_i from its action set $\mathcal{A}_i$. The system cost induced by allocation $a = (a_1, \dots, a_n) \in \mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$ is $C(a)$, where $C : \mathcal{A} \to \mathbb{R}$. Our objective is to find an optimal allocation, i.e.,

$$a^{\text{opt}} \in \arg\min_{a \in \mathcal{A}} C(a). \tag{1}$$

Motivated by the discussion of Section I, we turn our attention to deriving a distributed solution to (1). Unfortunately, this class of combinatorial problems is inherently intractable (more precisely, NP-hard). Thus, in the remainder of the paper, we aim to obtain an approximate solution to (1) through a distributed and tractable algorithm, ideally with the best possible approximation ratio[1]. We tackle the problem using the game-theoretic approach discussed in the introduction, and hence present the following game. Consider the game where the agent set is N, each agent's action set is $\mathcal{A}_i$, and in which each agent i evaluates its actions using a local cost function $J_i : \mathcal{A} \to \mathbb{R}$. In the forthcoming analysis, we will focus on the solution concept of Nash equilibrium, defined as any allocation $a^{\text{ne}} \in \mathcal{A}$ such that,

$$J_i(a^{\text{ne}}) \le J_i(a_i, a^{\text{ne}}_{-i}) \quad \forall a_i \in \mathcal{A}_i, \forall i \in N, \tag{2}$$

where $a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$. We represent the game as defined above with the tuple $G = (N, \mathcal{A}, \{J_i\}, C)$, where $\{J_i\} = \{J_1, \dots, J_n\}$. We measure the performance of a given algorithm using the notion of price-of-anarchy,

$$\text{PoA}(G) := \frac{\max_{a \in \text{NE}(G)} C(a)}{\min_{a \in \mathcal{A}} C(a)}, \tag{3}$$

where $\text{NE}(G)$ is the set of all pure Nash equilibria of the game G. Informally, the price-of-anarchy describes the ratio between the worst performing equilibrium and the optimal allocation. A lower price-of-anarchy is indicative of higher overall equilibrium performance. As such, the price-of-anarchy is an upper-bound on the efficiency of *any* equilibrium in the game. In cases where we have a family of games $\mathcal{G}$, the price-of-anarchy is further defined as,

$$\text{PoA}(\mathcal{G}) := \sup_{G \in \mathcal{G}} \text{PoA}(G). \tag{4}$$

Our work centers around the following two questions:

1) Given a class of cost-minimization games, how do we quantify the price-of-anarchy?
2) How can agents' local cost functions be designed in order to minimize the price-of-anarchy?

III. The smoothness framework

The smoothness framework developed in [14] has proven to be versatile, bringing a number of different price-of-anarchy results under a common analytical language, and producing tight bounds on the price-of-anarchy for different classes of problems [22], [23]. In this section, we revisit the notion of smooth games, and recall how smoothness arguments can be employed to bound the corresponding price-of-anarchy. The following proposition, adapted from [14], makes this statement precise.

Proposition 1 (Smoothness [14])**.** *Let $\mathcal{G}$ denote a class of cost-minimization games where $\sum_{i=1}^n J_i(a) \ge C(a)$. Further, suppose there exist $\lambda > 0$ and $\mu < 1$ such that for every game $G \in \mathcal{G}$, and any two action profiles $a, a^* \in \mathcal{A}$, it holds,*

$$\sum_{i=1}^{n} J_i(a_i^*, a_{-i}) \le \lambda\, C(a^*) + \mu\, C(a). \tag{5}$$

Then the price-of-anarchy satisfies,

$$\text{PoA}(\mathcal{G}) \le \frac{\lambda}{1-\mu}.$$

The above proposition demonstrates that an upper-bound on the price-of-anarchy can be computed by determining parameters $\lambda > 0$, $\mu < 1$ that satisfy (5). Accordingly, the best price-of-anarchy bound that can be derived using the smoothness framework, termed the *robust price-of-anarchy* [14], is given by,

$$\text{RPoA}(\mathcal{G}) := \inf_{\lambda>0, \mu<1} \left\{ \frac{\lambda}{1-\mu} \text{ s.t. (5) holds } \forall \text{G} \in \mathcal{G} \right\}. \tag{6}$$

Note that, in general, $\text{PoA}(\mathcal{G}) \le \text{RPoA}(\mathcal{G})$. Fortunately, $\text{PoA}(\mathcal{G}) = \text{RPoA}(\mathcal{G})$ for the well-studied class of congestion games, in which $\sum_{i=1}^n J_i(a) = C(a)$ for all $a \in \mathcal{A}$, see [14].

However, smoothness arguments are not applicable to games where $\sum_{i=1}^n J_i(a) < C(a)$ for at least one $a \in \mathcal{A}$. Additionally, the robust price-of-anarchy does not provide a tight upper-bound when $\sum_{i=1}^n J_i(a) > C(a)$ for all $a \in \mathcal{A}$, as we demonstrate in the following theorem.

[1]For ease of presentation, most of our analysis will focus on Nash equilibria, which are intractable to find in general. Nevertheless, we will show in Lemma 1 that our results generalize to the much broader set of coarse-correlated equilibria, which can be found in polynomial time [21].

Theorem 1. *For the given game G, assume $\sum_{i=1}^n J_i(a) > C(a)$ holds for all $a \in \mathcal{A}$. Then,*

$$\mathrm{RPoA(G)} > \mathrm{PoA(G)}. \tag{7}$$

Proof. By assumption, there must exist $\gamma > 1$ such that $\sum_{i=1}^n J_i(a) \geq \gamma C(a)$ for all $a \in \mathcal{A}$. Observe that, for $\lambda > 0$ and $\mu < 1$ as in (5),

$$\begin{aligned}\gamma C(a^{\text{ne}}) \leq \sum_{i=1}^n J_i(a^{\text{ne}}) &\leq \sum_{i=1}^n J_i(a^{\text{opt}}, a^{\text{ne}}_{-i}) \\ &\leq \lambda\, C(a^{\text{opt}}) + \mu\, C(a^{\text{ne}}),\end{aligned}$$

where the above inequalities hold by assumption, by (2), and by (5), respectively. As the equilibrium conditions in (2) are scale-invariant, it must be that

$$\mathrm{PoA(G)} \leq \frac{\lambda^*}{\gamma - \mu^*} < \frac{\lambda^*}{1-\mu^*} = \mathrm{RPoA(G)},$$

where $\lambda^* > 0$, $\mu^* < 1$ optimize (6). □

IV. Generalized smoothness

In the previous section, we showed that traditional smoothness arguments are unsuitable for bounding equilibrium performance when the sum of agents' local costs is not equal to the system cost. In the following, we introduce a new notion of smoothness that provides tight bounds for a broader class of games.

Theorem 2 (Generalized Smoothness)**.** *Suppose there exist $\lambda > 0$ and $\mu < 1$ such that for every game $G \in \mathcal{G}$, and any two action profiles $a, a^* \in \mathcal{A}$,*

$$\sum_{i=1}^n J_i(a_i^*, a_{-i}) - \sum_{i=1}^n J_i(a) + C(a) \leq \lambda\, C(a^*) + \mu\, C(a). \tag{8}$$

Then, the price-of-anarchy satisfies,

$$\mathrm{PoA}(\mathcal{G}) \leq \frac{\lambda}{1-\mu}.$$

Proof. For all $G \in \mathcal{G}$, for all $a^{\text{ne}} \in \mathrm{NE}(G)$ and $a^{\text{opt}} \in \mathcal{A}$,

$$\begin{aligned}C(a) &\leq \sum_{i=1}^n J_i(a_i^{\text{opt}}, a^{\text{ne}}_{-i}) - \sum_{i=1}^n J_i(a^{\text{ne}}) + C(a) \\ &\leq \lambda\, C(a^{\text{opt}}) + \mu\, C(a^{\text{ne}}),\end{aligned} \tag{9}$$

where the first inequality holds by (2), and the second, by (8). Rearranging (9), one gets the desired result. □

We use the name *generalized smoothness* as this novel notion of smoothness reduces to traditional smoothness when $\sum_{i=1}^n J_i(a) = C(a)$. Observe that generalized smoothness does not even require $\sum_{i=1}^n J_i(a) \geq C(a)$, and thus applies to a much broader class of games. In parallel to the previous section, we define the *generalized price-of-anarchy* as the best price-of-anarchy bound that can be derived using the generalized smoothness framework,

$$\mathrm{GPoA}(\mathcal{G}) := \inf_{\lambda>0, \mu<1} \left\{ \frac{\lambda}{1-\mu} \text{ s.t. (8) holds } \forall \mathrm{G} \in \mathcal{G} \right\}. \tag{10}$$

In the following theorem, we demonstrate that the proposed generalized smoothness framework provides upper bounds on the price-of-anarchy that are at least as tight as those provided by traditional smoothness.

Theorem 3. *For all games $G \in \mathcal{G}$ s.t. $\sum_{i=1}^n J_i(a) \geq C(a)$,*

$$\mathrm{PoA(G)} \leq \mathrm{GPoA(G)} \leq \mathrm{RPoA(G)}.$$

Additionally, if for all $a \in \mathcal{A}$, $\sum_{i=1}^n J_i(a) > C(a)$. Then,

$$\mathrm{GPoA(G)} < \mathrm{RPoA(G)}.$$

Proof. The proof can be found in the Appendix. □

Since the result in Theorem 3 holds for every game in the class $\mathcal{G}$, the inequalities hold with $\mathcal{G}$ in the place of G, i.e. for the whole class.

V. Local resource allocation games

In this section, we show how generalized smoothness can be used in order to provide concrete and tight bounds on the price-of-anarchy for the class of local resource allocation games. This analysis will extend the applicability of the linear programming approach presented in [20] to all coarse-correlated equilibria and to multiple resource types.

Consider a finite set of resources $\mathcal{R} = \{r_1, \ldots, r_m\}$ and let each agent's action set $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$. Every resource $r \in \mathcal{R}$ is associated with a pair of functions (c_r, f_r), where c_r and f_r are mappings from $N = \{1, \ldots, n\} \to \mathbb{R}$. We define the system cost and local cost functions as,

$$\begin{aligned}C(a) &= \sum_{r \in \mathcal{R}} c_r(|a|_r), \\ J_i(a_i, a_{-i}) &= \sum_{r \in a_i} f_r(|a|_r),\end{aligned}$$

where $|a|_r$ is the number of agents covering resource r in allocation a. We identify the aforementioned game with the tuple $G = (n, \mathcal{R}, \mathcal{A}, \{(c_r, f_r)\})$, and, for ease of notation, we remove the subscripts of the above sets, e.g. we write $\{(c_r, f_r)\}$ in place of $\{(c_r, f_r)\}_{r\in\mathcal{R}}$. Given a set of resource types $T = \{(c^1, f^1), \ldots, (c^{|T|}, f^{|T|})\}$, we induce the class of games $\mathcal{G}_T^n$ in which, for every resource $r \in \mathcal{R}$, there exists $v_r \geq 0$ such that,

$$\begin{aligned}c_r(\cdot) &= v_r\, c^t(\cdot) \\ f_r(\cdot) &= v_r\, f^t(\cdot)\end{aligned}$$

for some $(c^t, f^t) \in T$. Informally, $\mathcal{G}_T^n$ contains all possible games G where the number of agents is fixed to n, and each pair (c^t, f^t) is selected from T and premultiplied by v_r. A possible interpretation of $v_r \geq 0$ is that it represents the value of the resource r. We will refer to the functions $\{f_r\}_{r\in\mathcal{R}}$ as *distribution rules*, since each f_r describes how the value of covered resources v_r is split among the agents. We observe that many classes of problems studied in the literature can be analyzed using this model. Important examples include set covering problems, vehicle-target assignment problems, and congestion games [13], [24], [11]. In Section V-A, we exemplify one of many problems that can be represented in this form.

With slight abuse of notation, in the following, we redefine the price-of-anarchy for the class $\mathcal{G}_T^n$ as,

$$\mathrm{PoA}(\mathcal{G}_T^n) := \sup_{G \in \mathcal{G}_T^n} \mathrm{PoA}(G), \tag{11}$$

Before presenting our results, we demonstrate the generality of the local resource allocation problem formulation, using the example of atomic congestion games.

A. An Illustrative Example: Atomic Congestion Games

To demonstrate the generality of the local resource allocation problem presented above, we analyze congestion games, a classical cost-minimization problem [25]. A congestion game can be reformulated as a local resource allocation game in which the system cost and local costs depend on the edges' occupancy-dependent latency functions l_e such that

$$C(a) = \sum_{e \in \mathcal{R}} c_e(|a|_e) = \sum_{e \in \mathcal{R}} l_e(|a|_e)\,|a|_e,$$
$$J_i(a_i, a_{-i}) = \sum_{e \in a_i} l_e(|a|_e).$$

Observe that, in this context, the edges of the congestion game correspond to resources in a local resource allocation game. The latency functions l_e play the role of distribution rules f_r, and c_r is substituted with $l_e(|a|_e)|a|_e$.

Congestion games with affine costs [26]. As a particular example, the class of congestion games with affine cost functions can be represented as a class of local resource allocation problems with two resource types; $T = \{(x^2, x), (x, 1)\}$.[2] As an elementary example, consider the game G with resources $\mathcal{R} = \{e_i\}_{i=1}^4$, each of which has nonnegative value $v_e \geq 0$. The resource e_1 and e_3 are associated with type (x^2, x), whereas e_2 and e_4 have type $(x, 1)\}$. If we define all n agents' action sets as $\{(e_1, e_2), (e_3, e_4)\}$, the game can be represented by a two link network as shown in Fig. 1, where $J_i(a) = v_1|a|_1 + v_2$ for an agent i that selects (e_1, e_2), and $J_i(a) = v_3|a|_3 + v_4$ for agents selecting (e_3, e_4).

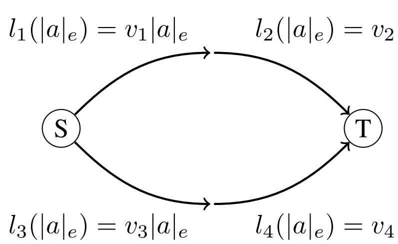

Fig. 1: A simple congestion game with affine cost functions that can be represented as a local resource allocation problem with two types. The system's n agents must select either the top or bottom edge, and experience the corresponding cost.

While we consider the simplistic example of a two-link network here, we note that, in general, any congestion game with affine costs can be represented as a local resource allocation game with T as above. Additionally, given a basis set of all possible edge costs, *any* atomic congestion game can be formulated using our model.

[2]Informally, this means that there are two edge types in the congestion game, those that impose a latency proportional to the number of agents selecting them, and those that have constant latency.

B. Computing the price-of-anarchy

The next theorem shows how the price-of-anarchy in (11) can be recovered by means of the notion of generalized smoothness previously introduced in (8). Before proceeding, we state a few assumptions and some notation.

Standing Assumption. *We assume, with slight abuse of notation and without loss of generality, that $c^t(0) = f^t(0) = 0$, and $f^t(n+1) = f^t(n)$, for all $(c^t, f^t) \in T$.*

Although c^t, f^t were previously defined as mappings from $N \to \mathbb{R}$, we extend their definition to ease the notation. Let

$$\mathcal{I} := \{(x,y,z) \in \mathbb{N}^3 \mid 1 \leq x + y - z \leq n\},$$
$$\mathcal{I}_{\mathcal{R}} := \{(x,y,z) \in \mathcal{I} \mid x + y - z = n \text{ or } (x-z)(y-z)z = 0\}.$$

Theorem 4. *For every class of local resource allocation games $\mathcal{G}_T^n$, it holds*

$$\mathrm{PoA}(\mathcal{G}_{\mathrm{T}}^{\mathrm{n}}) = \mathrm{GPoA}(\mathcal{G}_{\mathrm{T}}^{\mathrm{n}}). \tag{12}$$

Proof. The proof can be found in the Appendix. □

The above theorem shows that generalized smoothness arguments provide tight upper-bounds on the price-of-anarchy in local resource allocation games, and proposes a methodology for constructing worst-case instances. We now exploit this result to obtain easily computable and concrete bounds on the price-of-anarchy.

Theorem 5. *For the class $\mathcal{G}_T^n$, $\mathrm{PoA}(\mathcal{G}_{\mathrm{T}}^{\mathrm{n}}) = \mathrm{GPoA}(\mathcal{G}_{\mathrm{T}}^{\mathrm{n}}) = 1/C^*$, where C^* is the value of the following linear program,*

$$\begin{aligned} C^* = &\max_{\nu \in \mathbb{R}_{\geq 0}, \rho \in \mathbb{R}} \rho \\ &\text{s.t. } c^t(y) - \rho c^t(x) \\ &\quad + \nu\left[(x-z)f^t(x) - (y-z)f^t(x+1)\right] \geq 0 \\ &\quad \forall (c^t, f^t) \in T,\ \forall (x,y,z) \in \mathcal{I}_{\mathcal{R}}, \end{aligned} \tag{13}$$

Proof. The proof can be found in the Appendix. □

The above linear program extends [20, Thm. 2] to any class of games $\mathcal{G}_T^n$.

C. Optimizing the price-of-anarchy

Whereas the previous subsection was devoted to calculating the price-of-anarchy for given pairs $\{(c^t, f^t)\}_{t=1}^{|T|}$, we now shift our focus to designing the distribution rules such that the price-of-anarchy is minimized.

Theorem 6. *Consider the cost functions $\{c^1, \ldots, c^{|T|}\}$, and positive integer n. Define f_{OPT}^t as the solution to,*

$$\begin{aligned} f_{\mathrm{OPT}}^t \in &\arg\max_{f \in \mathbb{R}^n, \rho \in \mathbb{R}} \rho \\ \text{s.t. } &1_{\{y \geq 1\}} c^t(y) - \rho 1_{\{x \geq 1\}} c^t(x) + (x-z)f(x) \\ &- (y-z)f(x+1) \geq 0, \forall (x,y,z) \in \mathcal{I}_{\mathcal{R}}, \end{aligned} \tag{14}$$

for all $t \in \{1,\dots,|T|\}$. The set of distribution rules $\mathbf{f}_{\rm OPT} = \{f^1_{\rm OPT},\dots,f^{|T|}_{\rm OPT}\}$ satisfies,

$$\mathbf{f}_{\rm OPT} \in \arg\min_{\mathbf{f}\in\mathbb{R}^{n\times|T|}_{\geq 0}} {\rm PoA}(\mathcal{G}^n_T). \tag{15}$$

Proof. The proof can be found in the Appendix. □

Theorem 6 shows that a set of optimal distribution rules can be calculated using the linear program in [20, Thm. 3]. It is worth noting that for a given class of games $\mathcal{G}^n_T$ with an arbitrary basis set T, it is not possible, in general, to compute the price-of-anarchy as the worst price-of-anarchy over each individual pair (c^t, f^t), i.e. the expression ${\rm PoA}(\mathcal{G}^n_T) = \max_{\mathbf{t}\in T}\{{\rm PoA}\,(\mathcal{G}^{\rm n}_{\mathbf{t}})\}$ does *not hold*, as further elaborated in Lemma 5. Nevertheless, this property is recovered for the specific choice of $f^t = f^t_{\rm OPT}$, see Lemma 6. This constitutes the key observation towards proving Theorem 6.

D. Returning to Atomic Congestion Games

Here we apply the results presented in this section to the class of congestion games, as presented in Section V-A.

Congestion games with affine costs. Using a particular analytical structure, and by exploiting a known polynomial inequality, the authors of [24] prove that the price-of-anarchy for the class of congestion games with affine costs is $5/2$. A direct application of the linear program in Theorem 5 recovers the same result for any number of agents greater than 3. Additionally, we determine a worst-case instance construction with only $n = 3$ agents.[3]

Impossibility results in tolling. The idea of improving the price-of-anarchy in congestion games using a local edge toll as a control mechanism originates from [27]. In order to influence agents' decisions, a toll τ_e is added to the agents' local edge costs such that,

$$J_i(a_i, a_{-i}) = \sum_{e\in a_i} l_e(|a|_r) + \tau_e(|a|_r).$$

The system cost remains unchanged. In this context, the strength of Theorem 6 stems from the simple recipe it provides for constructing a set of optimal tolls. Indeed, solving for the optimal tolls is equivalent to finding the optimal set of distribution rules in the corresponding class of local resource allocation games, where $c^t(x) = x l_e(x)$ for all edges e of type t, for all $t \in \{1,\dots,|T|\}$. Thus, for a given class of congestion games, the result in Theorem 6 provides tight lower-bounds on the best achievable price-of-anarchy using local tolling mechanisms. In Table 1, we compile these bounds for three classes of congestion games considered in [28].

[3]Consider the game G with six edges $\{e_i\}_{i=1}^6$ with identical value (i.e. $v_e = v$) and latency function $l_e(x) = x$. We endow the $n = 3$ agents with the action sets, $\mathcal{A}_1 = \{(e_4, e_5, e_6), (e_1, e_2)\}$, $\mathcal{A}_2 = \{(e_1, e_2, e_5), (e_3, e_4)\}$, and $\mathcal{A}_3 = \{(e_1, e_3, e_4), (e_5, e_6)\}$. The Nash equilibrium $a^{\rm ne}$ corresponds to each agent selecting its three-tuple action, and the optimal actions in $a^{\rm opt}$ are the two-tuple actions. It can easily be verified that (2) is met. ${\rm PoA(G)} = 5/2$, since the system costs are $C(a^{\rm ne}) = 45v$ and $C(a^{\rm opt}) = 18v$. Note that, in general, drawing a worst-case instance as a graph requires additional edges e with value $v_e = 0$.

TABLE 1: Prices-of-anarchy of the original and optimally tolled congestion games, when the number of agents is $n = 10$.

Edge Costs	PoA without tolling	PoA with optimal tolls
Affine	2.50	2.01
Quadratic	9.58	5.10
Cubic	41.5	15.53

VI. Extensions

In this section, we demonstrate that the above results extend to i) coarse-correlated equilibria; and, ii) welfare-maximization problems.

A. Coarse-Correlated Equilibria

A significant advantage of using a smoothness argument is that it provides performance bounds for the class of *coarse-correlated equilibria*, a far broader class of equilibria compared to the class of pure Nash equilibria [22], [29]. A coarse-correlated equilibrium is a probability distribution σ over all allocations $a \in \mathcal{A}$ such that for all $i \in [n]$, and $a' \in \mathcal{A}$, it holds that,

$$\mathbb{E}_{a\sim\sigma}[J_i(a)] \leq \mathbb{E}_{a\sim\sigma}[J_i(a'_i, a_{-i})],$$

where $\mathbb{E}_{a\sim\sigma}[J_i(a)]$ is the expected utility for the agent i. In the next lemma, we show that the price-of-anarchy bounds stemming from generalized smoothness arguments extend to all coarse-correlated equilibria.

Lemma 1. *For every game G in the class of games $\mathcal{G}$,*

$${\rm GPoA(G)} \geq \frac{\max_{\sigma\in CCE({\rm G})} \mathbb{E}_{\rm a\sim\sigma}[{\rm C(a)}]}{\min_{\rm a\in\mathcal{A}} {\rm C(a)}},$$

where $CCE(G)$ is the set of all coarse-correlated equilibria of the game G.

Proof. The proof follows the same reasoning as [14, Thm. 3.2], and is omitted due to space constraints. □

Since the sets of pure and mixed Nash equilibria of a game are subsets of its coarse-correlated equilibria, the GPoA is an upper-bound on the efficiency of *all* equilibria within these classes. This result is particularly important toward the tractability of the final algorithm. Indeed, although finding a pure Nash equilibrium can be intractable, coarse-correlated equilibria can often be computed in polynomial time [21].

B. Welfare-Maximization Problems

In welfare-maximization problems, we consider games for which the agent set is N, each agent's action set is $\mathcal{A}_i$, and each agent i evaluates its actions using some local utility function $U_i : \mathcal{A} \to \mathbb{R}$. The system-level welfare induced by a given allocation $a \in \mathcal{A}$ is measured by the function $W : \mathcal{A} \to \mathbb{R}$. A Nash equilibrium is defined as any allocation $a^{\rm ne} \in \mathcal{A}$ such that $U_i(a^{\rm ne}) \geq U_i(a_i, a^{\rm ne}_{-i})$ for all $a_i \in \mathcal{A}_i$ and all $i \in N$, and our objective is to find an allocation $a^{\rm opt} \in \mathcal{A}$ that solves the system-level optimization problem,

$$a^{\rm opt} \in \arg\max_{a\in\mathcal{A}} W(a). \tag{16}$$

The results obtained for cost-minimization problems in this work may be derived analogously for welfare-maximization, although the various linear program formulations may differ. Toward this goal, we propose the following modified notion of generalized smoothness:

Consider a class of welfare-maximization games $\mathcal{G}$, and suppose there exist $\lambda > 0$, $\mu > -1$ such that, for all allocations $a, a^* \in \mathcal{A}$, it holds that,

$$\sum_{i\in[n]} U_i(a_i^*, a_{-i}) - \sum_{i\in[n]} U_i(a) \geq \lambda W(a^*) - (1+\mu)W(a).$$

The corresponding lower-bound on equilibrium efficiency for all games $G \in \mathcal{G}$ is $\lambda/(1+\mu)$.

VII. Conclusions

In this manuscript, we aimed at providing a novel methodology for characterizing and optimizing the price of anarchy in connection to a broad class of problems, including congestion games. Toward this goal, we introduced the notion of generalized smoothness. Compared to traditional smoothness arguments, we showed that generalized smoothness is more widely applicable, and provides tighter bounds. We applied generalized smoothness arguments to the class of local resource allocation problems (which include congestion games) and observed that it provides tight bounds on the price-of-anarchy. Relative to this class of problems, we were able to compute and optimize the price-of-anarchy of coarse-correlated equilibria, by means of concrete and tractable linear programs. Along with other possible future research directions, this work paves the way for the design of optimal tolling schemes through the linear programming approach presented in [19], [20].

References

[1] S. Bussmann, N. R. Jennings, and M. Wooldridge, *Multiagent systems for manufacturing control: a design methodology*. Springer Science & Business Media, 2013.

[2] J. Wu, *Distributed system design*. CRC press, 2017.

[3] X. Yu and Y. Xue, "Smart grids: A cyber–physical systems perspective," *Proceedings of the IEEE*, vol. 104, no. 5, 2016.

[4] K. Spieser, K. Treleaven, R. Zhang, E. Frazzoli, D. Morton, and M. Pavone, "Toward a systematic approach to the design and evaluation of automated mobility-on-demand systems: A case study in singapore," in *Road vehicle automation*. Springer, 2014.

[5] T. Le, P. Kovács, N. Walton, H. L. Vu, L. L. Andrew, and S. S. Hoogendoorn, "Decentralized signal control for urban road networks," *Transportation Research Part C: Emerging Technologies*, vol. 58, 2015.

[6] M. I. El-Hawwary and M. Maggiore, "Distributed circular formation stabilization for dynamic unicycles," *IEEE Transactions on Automatic Control*, vol. 58, no. 1, 2013.

[7] H.-T. Zhang, Z. Chen, and M.-C. Fan, "Collaborative control of multivehicle systems in diverse motion patterns," *IEEE Transactions on Control Systems Technology*, vol. 24, no. 4, 2016.

[8] A. Nedic and A. Ozdaglar, "Distributed subgradient methods for multi-agent optimization," *IEEE Transactions on Automatic Control*, vol. 54, no. 1, 2009.

[9] E. Wei, A. Ozdaglar, and A. Jadbabaie, "A distributed newton method for network utility maximization–i: Algorithm," *IEEE Transactions on Automatic Control*, vol. 58, no. 9, 2013.

[10] J. S. Shamma, *Cooperative control of distributed multi-agent systems*. Wiley Online Library, 2007.

[11] G. Arslan, J. R. Marden, and J. S. Shamma, "Autonomous vehicle-target assignment: A game-theoretical formulation," *Journal of Dynamic Systems, Measurement, and Control*, vol. 129, no. 5, 2007.

[12] M. Gairing, "Covering games: Approximation through non-cooperation," in *International Workshop on Internet and Network Economics*. Springer, 2009.

[13] V. Ramaswamy, D. Paccagnan, and J. R. Marden, "The impact of local information on the performance of multiagent systems," *arXiv preprint arXiv:1710.01409*, 2017.

[14] T. Roughgarden, "Intrinsic robustness of the price of anarchy," *Journal of the ACM (JACM)*, vol. 62, no. 5, 2015.

[15] ——, "The price of anarchy in games of incomplete information," in *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 2012.

[16] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, M. Kyropoulou, B. Lucier, R. P. Leme, and É. Tardos, "Bounding the inefficiency of outcomes in generalized second price auctions," *Journal of Economic Theory*, vol. 156, 2015.

[17] D. J. Foster, Z. Li, T. Lykouris, K. Sridharan, and E. Tardos, "Learning in games: Robustness of fast convergence," in *Advances in Neural Information Processing Systems*, 2016.

[18] V. Syrgkanis and E. Tardos, "Composable and efficient mechanisms," in *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*. ACM, 2013.

[19] D. Paccagnan, R. Chandan, and J. R. Marden, "Distributed resource allocation through utility design-part i: optimizing the performance certificates via the price of anarchy," *arXiv preprint arXiv:1807.01333*, 2018.

[20] R. Chandan, D. Paccagnan, and J. R. Marden, "Optimal price of anarchy in cost-sharing games," *2019 Annual American Control Conference (ACC)*, 2019.

[21] C. H. Papadimitriou and T. Roughgarden, "Computing correlated equilibria in multi-player games," *Journal of the ACM (JACM)*, vol. 55, no. 3, 2008.

[22] T. Roughgarden, "Intrinsic robustness of the price of anarchy," in *Proceedings of the forty-first annual ACM symposium on Theory of computing*. ACM, 2009.

[23] T. Roughgarden, V. Syrgkanis, and E. Tardos, "The price of anarchy in auctions," *Journal of Artificial Intelligence Research*, vol. 59, 2017.

[24] G. Christodoulou and E. Koutsoupias, "The price of anarchy of finite congestion games," in *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*. ACM, 2005, pp. 67–73.

[25] R. W. Rosenthal, "A class of games possessing pure-strategy nash equilibria," *International Journal of Game Theory*, vol. 2, no. 1, pp. 65–67, 1973.

[26] B. Awerbuch, Y. Azar, and A. Epstein, "The price of routing unsplittable flow," in *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*. ACM, 2005, pp. 57–66.

[27] H. Yang and H.-J. Huang, *Mathematical and economic theory of road pricing*, 2005.

[28] T. Roughgarden, *Selfish routing and the price of anarchy*. MIT press Cambridge, 2005, vol. 174.

[29] H. Moulin and J.-P. Vial, "Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon," *International Journal of Game Theory*, vol. 7, no. 3-4, pp. 201–221, 1978.

[30] J. R. Marden and T. Roughgarden, "Generalized efficiency bounds in distributed resource allocation," *IEEE Transactions on Automatic Control*, vol. 59, no. 3, 2014.

Appendix

Proof of Theorem 3

Proof. When $\sum_{i\in N} J_i(a) = C(a)$, (8) is equivalent to (5). When $\sum_{i\in N} J_i(a) > C(a)$, for all λ, μ satisfying (5), the following must hold for all $a, a^* \in \mathcal{A}$,

$$\sum_{i\in[n]} J_i(a_i^*, a_{-i}) - \sum_{i\in[n]} J_i(a) + C(a) < \lambda C(a^*) + \mu C(a).$$

Thus, there must exist some $\epsilon > 0$ such that (8) holds for $\bar{\lambda} = \lambda^* - \epsilon$ or $\bar{\mu} = \mu^* - \epsilon$, where λ^*, μ^* optimize (6). Since $\lambda/(1-\mu)$ is increasing in both λ and μ, GPoA $< \lambda^*/(1-\mu^*) =$ RPoA. □

Preliminaries to the proof of Theorem 4

Definition 1. *$\mathcal{S}(\mathcal{G}_T^n)$ is the set of parameters $\lambda > 0$, $\mu < 1$ such that, for all $(c^t, f^t) \in T$ and all $(x, y, z) \in \mathcal{I}_\mathcal{R}$,*

$$(z-x)f^t(x) + (y-z)f^t(x+1) + c^t(x) \leq \lambda c^t(y) + \mu c^t(x). \tag{17}$$

Definition 2. *$\gamma(\mathcal{G}_T^n)$ is defined as,*

$$\gamma(\mathcal{G}_T^n) := \inf_{\lambda,\mu} \left\{ \frac{\lambda}{1-\mu} : (\lambda, \mu) \in \mathcal{S}(\mathcal{G}_T^n) \right\} \tag{18}$$

Lemma 2. *For the given class of games $\mathcal{G}_T^n$,*

$$\gamma(\mathcal{G}_T^n) \geq \text{GPoA}(\mathcal{G}_T^n).$$

Proof. Let $|a^{\text{ne}}| = \{x_1, \ldots, x_m\}$, and $|a^{\text{opt}}| = \{y_1, \ldots, y_m\}$. We define z_r as the number of agents that select resource r in both a^{ne} and a^{opt},

$$z_r := |\{i \in N : r \in a_i^{\text{ne}}\} \cap \{i \in N : r \in a_i^{\text{opt}}\}|$$

where $z_r \leq \min\{x_r, y_r\}$, and $1 \leq x_r + y_r - z_r \leq n$.

The following simplification, adapted from [30], is instrumental in our proof of tightness,

$$\begin{aligned} &\sum_{i\in[n]} J_i(a_i^{\text{opt}}, a_{-i}^{\text{ne}}) - \sum_{i\in[n]} J_i(a^{\text{ne}}) + C(a^{\text{ne}}) \qquad (19)\\ &= \sum_{r\in\mathcal{R}} [z_r f_r(x_r) + (y_r - z_r) f_r(x_r+1)] \\ &\quad - \sum_{r\in\mathcal{R}} x_r f_r(x_r) + \sum_{r\in\mathcal{R}} c_r(x_r) \\ &= \sum_{r\in\mathcal{R}} [(z_r - x_r) f_r(x_r) + (y_r - z_r) f_r(x_r+1) + c_r(x_r)]. \end{aligned}$$

We have shown that (19) can be represented as a sum over a subset of the left-hand side expressions in (17) corresponding to the resources in $\mathcal{R}$ weighted by their values. For the proof that it is sufficient to consider $(x, y, z) \in \mathcal{I}_\mathcal{R}$, see the second part of the proof of [20, Thm. 2], and note that (x, y, z) in this paper are equivalent to (j, l, x) in their notation. Thus, the parameters $(\lambda, \mu) \in \mathcal{S}(\mathcal{G}_T^n)$ are sure to satisfy the constraint in (8). This is because $C(a^{\text{ne}})$ is guaranteed to be less than or equal to (19). This implies that $\gamma(\mathcal{G}_T^n) \geq \text{GPoA}(\mathrm{G}_\mathrm{T}^\mathrm{n})$. □

Lemma 3. *Consider the class of games $\mathcal{G}_T^n$. Suppose there exist $(\hat{\lambda}, \hat{\mu}) \in \mathcal{S}(\mathcal{G}_T^n)$ such that,*

$$\frac{\hat{\lambda}}{1-\hat{\mu}} = \gamma(\mathcal{G}_T^n).$$

Then, there must be (c^1, f^1), (c^2, f^2) in T, (x_1, y_1, z_1), (x_2, y_2, z_2) in $\mathcal{I}_\mathcal{R}$, and $\eta \in [0, 1]$ such that,

$$\begin{aligned} &(z_j - x_j) f^j(x_j) + (y_j - z_j) f^j(x_j+1) + c^j(x_j) \\ &= \hat{\lambda}\, c^j(y_j) + \hat{\mu}\, c^j(x_j) \end{aligned} \tag{20}$$

for $j = 1, 2$; and,

$$\begin{aligned} &\eta[z_1 f^1(x_1) + (y_1 - z_1) f^1(x_1+1)] \\ &+ (1-\eta)[z_2 f^2(x_2) + (y_2 - z_2) f^2(x_2+1)] \\ &= \eta\, x_1 f^1(x_1) + (1-\eta)\, x_2 f^2(x_2). \end{aligned} \tag{21}$$

Proof. We define $\mathcal{H}_{c,f,x,y,z}$ as the set of $(\lambda, \mu) \in \mathbb{R}_{>0} \times \mathbb{R}_{<1}$ that satisfy, for given c, f, x, y and z,

$$(z-x)f(x) + (y-z)f(x+1) + c(x) \leq \lambda\, c(y) + \mu\, c(x).$$

We denote by $\delta\mathcal{H}_{c,f,x,y,z}$ the boundary of the set, i.e. the points (λ, μ) that satisfy the above inequality with equality. Some simplifications can be made for the cases when either $x = 0$ or $y = 0$. When $x = 0$ and $y > 0$, then $z = 0 = \min\{x, y\}$, and the set $\mathcal{H}_{c,f,0,y,0}$ contains all $\mu < 1$, and all $\lambda \geq f(1)y/c(y)$. When $x > 0$ and $y = 0$, $z = 0$ once again, the set $\mathcal{H}_{c,f,x,0,0}$ contains all $\lambda > 0$, and all $\mu \geq 1 - xf(x)/c(x)$. For the halfplanes with $x > 0$ and $y > 0$, the boundary is,

$$\mu = -\frac{c(y)}{c(x)}\lambda + \frac{1}{c(x)}\Big[(z-x)f(x) + (y-z)f(x+1) + c(x)\Big].$$

Note that finding $\gamma(\mathcal{G}_T^n)$ is equivalent to finding the point along the boundary that intersects the line with μ-intercept equal to 1 with the most negative slope. Thus, we can find the optimal (λ, μ) by starting on the boundary at $\lambda = \max_y f(1)y/c(y)$, then following the boundary until we reach a line with μ-intercept less than 1. There are three possibilities for the optimal λ, μ: at $\lambda = \max_y f(1)y/c(y)$, at the intersection of two halfplanes with $x_1, x_2 > 0$, $y_1, y_2 > 0$, $z_1, z_2 \geq 0$, or at $\mu = 1 - \min_x xf(x)/c(x)$.

If the optimal λ, μ occurs at $\lambda = \max_y f(1)y/c(y)$, then $x_1 = 0$, $y_1 > 0$ and $z_1 = 0$, and the other halfplane has μ-intercept less than one, and $x_2 > 0$, $y_2 \geq 0$ and $z_2 \geq 0$. Note that it is also possible for the optimal λ, μ to occur at $\lambda = \max_{y>0} f(1)y/c(y)$ and $\mu = 1 - \min_{x>0} xf(x)/C(x)$.

For all these cases, there exists $\eta \in [0, 1]$ such that,

$$\begin{aligned} &\eta\,[z_1 f^1(x_1) + (y_1 - z_1) f^1(x_1+1)] \\ &+ (1-\eta)\,[z_2 f^2(x_2) + (y_2 - z_2) f^2(x_2+1)] \\ &= \eta x_1 f^1(x_1) + (1-\eta) x_2 f^2(x_2) \end{aligned}$$

If the optimal (λ, μ) occur on a halfplane $\mathcal{H}_{c,f,x,y,z}$ with μ-intercept equal to 1, we select $c_1 = c_2 = c$, $f^1 = f^2 = f$, $x_1 = x_2 = x$, $y_1 = y_2 = y$ and $z_1 = z_2 = z$, where any $\eta \in [0, 1]$ will satisfy the equality. □

Lemma 4. *For the class of games $\mathcal{G}_T^n$, suppose no point $(\lambda,\mu)\in\mathcal{S}(\mathcal{G}_T^n)$ satisfies $\frac{\lambda}{1-\mu}=\gamma(\mathcal{G}_T^n)$. Then, there exists $(f,c)\in T$ and $(x,y,z)\in\mathcal{I}_{\mathcal{R}}$ such that*

$$\gamma(\mathcal{G}_T^n)=\frac{c(x)}{c(y)} \tag{22}$$

$$(y-z)f(x+1)+zf(x)>xf(x) \tag{23}$$

Proof. Borrowing the notation and reasoning of the proof for Lemma 3, we know that the strictest constraint must come from a line corresponding to some $(f,c)\in T$ that for some values of x, y and z has μ-intercept greater than 1, and the least negative slope among all constraints. Since the μ-intercept is greater than 1, $(z-x)f(x)+(y-z)f(x+1)>0$, which implies that $(y-z)f(x+1)+zf(x)>xf(x)$. The least negative slope results from selecting $y=\arg\min_{j\in N}c(j)$ and $x=\arg\max_{j\in N}c(j)$. Much like in [14, Lem. 5.5], we construct a sequence $\{(\lambda_k,\mu_k)\}$ in $\mathcal{S}(\mathcal{G}_T^n)$ such that $\frac{\lambda_k}{1-\mu_k}\downarrow\gamma(\mathcal{G}_t^n)$. Since $\frac{\lambda}{1-\mu}$ is increasing in both λ and μ, it can be assumed that every point (λ_k,μ_k) lies on the boundary of $\mathcal{S}(\mathcal{G}_T^n)$. The values λ_k are bounded from below by the constraints (17) where $x=z=0$, and for finite $\gamma(\mathcal{G}_T^n)$, $\mu_k\le b<1$. Since $\frac{\lambda}{1-\mu}$ is continuous, $\frac{\lambda_k}{1-\mu_k}\downarrow\gamma(\mathcal{G}_t^n)$ and $\gamma(\mathcal{G}_t^n)$ is not attained, the sequence $\{\lambda_k,\mu_k)\}$ has no limit point. Thus, after some rearranging of (17),

$$\begin{aligned}\gamma(\mathcal{G}_T^n)&=\lim_{k\to\infty}\frac{\lambda_k}{1-\mu_k}\\&=\lim_{k\to\infty}\frac{c(x)}{c(y)}+\frac{(z-x)f(x)+(y-z)f(x+1)}{c(y)(1-\mu_k)}=\frac{c(x)}{c(y)},\end{aligned}$$

since $\mu_k\to-\infty$, which completes the proof. □

Proof of Theorem 4

Proof. We first consider the case where the value $\gamma(\mathcal{G}_T^n)$ is not attained for any point $(\lambda,\mu)\in\mathcal{S}(\mathcal{G}_T^n)$ as in Lemma 4. We recover the pair $(f,c)\in T$ that result in the strictest constraint at $\lambda\to\infty$, $\mu\to-\infty$, as well as the values x, y and z that give the least negative slope. We setup a game with $l=\min\{x+y,n\}$ resources organized in a cycle and l agents, i.e. $\mathcal{R}=\{r_1,\dots,r_l\}$ and $N=[l]$, where every resource has type corresponding to the pair (f,c). Each agent $i\in[l]$ is endowed with two actions, the first is to select x consecutive resources starting with r_i and ending with $r_{i+x-1 \bmod l}$, while the second is to select y consecutive resources ending with $r_{i+z-1 \bmod l}$. Condition (23) implies that the former strategy is a Nash equilibrium, and by (22), the price-of-anarchy is at least $\frac{c(x)}{c(y)}=\gamma(\mathcal{G}_T^n)$, as required.

We retrieve (c^1,f^1), (c^2,f^2), (x_1,y_1,z_1), (x_2,y_2,z_2) and η; the optimality parameters as in Lemma 3, where $\gamma(\mathcal{G}_T^n)$ is an upper-bound on $\mathrm{GPoA}(\mathcal{G}_T^n)$ and is guaranteed to be attained, by Lemma 2 and Definition 1. The worst-case game G is constructed in the following way, define two disjoint cycles E_1 and E_2 each with $l=\min\{\max\{x_1+y_1,x_2+y_2\},n\}$ resources enumerated from 1 to l. The resources in E_1 are assigned cost function c_1, distribution rule f^1 and value η, whereas the resources in E_2 are assigned c_2, f^2 and $(1-\eta)$. There are also $l\le n$ players, enumerated 1 through l and we restrict the action set $\mathcal{A}$ to two strategies, $a^{\rm ne}$ and $a^{\rm opt}$. In the first strategy, each player $i\in[l]$ selects x_1 consecutive resources in E_1, $[i,i+1,\dots,i+x_1-1] \mod l$, and x_2 consecutive resources in E_2 starting with resource i. In the second strategy, player i selects y_1 consecutive resources in E_1 ending with resource $i-1$, and y_2 consecutive resources in E_2 ending with resource $i-1$.

We continue by demonstrating that the first strategy satisfies the conditions for a Nash equilibrium,

$$\begin{aligned}J_i(a^{\rm ne})&=\eta\,x_1\,f^1(x_1)+(1-\eta)\,x_2\,f^2(x_2)\\&=\eta\,[z_1\,f^1(x_1)+(y_1-z_1)\,f^1(x_1+1)]\\&\quad+\,(1-\eta)\,[z_2\,f^2(x_2)+(y_2-z_2)\,f^2(x_2+1)]\\&=J_i(a_i^{\rm opt},a_{-i}^{\rm ne}),\end{aligned}\tag{24}$$

where (24) holds due to Lemma 3. Now we show that the price-of-anarchy of the game is lower-bounded by $\gamma(\mathcal{G}_T^n)$, thus implying equality.

$$\begin{aligned}C(a^{\rm ne})&=C(a^{\rm ne})-\sum_{i=1}^k J_i(a^{\rm ne})+\sum_{i=1}^k J_i(a_i^{\rm opt},a_{-i}^{\rm ne})\\&=k\,\eta\left[\hat\lambda\,c_1(y_1)+\hat\mu\,c_1(x_1)\right]\\&\quad+\,k\,(1-\eta)\left[\hat\lambda\,c_2(y_2)+\hat\mu\,c_2(x_2)\right]\\&=\hat\lambda\,C(a^{\rm opt})+\hat\mu\,C(a^{\rm ne})\end{aligned}$$

In the above, $\gamma(\mathcal{G}_T^n)=\mathrm{PoA(G)}\le\mathrm{PoA}(\mathcal{G}_{\rm T}^{\rm n})$. Since $\gamma(\mathcal{G}_T^n)\ge\mathrm{GPoA}(\mathcal{G}_T^n)\ge\mathrm{PoA}(\mathcal{G}_T^n)$ by Lemma 2, $\mathrm{GPoA}(\mathcal{G}_T^n)$ must be tight. □

Proof of Theorem 5

Proof. We begin by noting that, by Definition 1, we need only consider $(x,y,z)\in\mathcal{I}_{\mathcal{R}}$ when calculating the price-of-anarchy in local resource allocation games. Observe that the constraints in the linear program are equivalent to the simplified conditions for (λ,μ)-generalized smoothness in (17). The linear program constraints read as,

$$c(y)-\rho\,c(x)+\nu\,[(x-z)\,f(x)-(y-z)\,f(x+1)]\ge 0,$$

for all $(x,y,z)\in\mathcal{I}_{\mathcal{R}}$, where $\rho=\frac{1-\mu}{\lambda}$, and $\nu=\frac{1}{\lambda}$. Substituting the expressions for ν and ρ into the above, and rearranging, we are left with,

$$\begin{aligned}&(z-x)\,f(x)+(y-z)\,f(x+1)+c(x)\\&\le\lambda\,c(y)+\mu\,c(x),\end{aligned}$$

for all $(x,y,z)\in\mathcal{I}_{\mathcal{R}}$, which is identical to (17) when there is a single type. Next, observe that maximizing ρ is equivalent to minimizing $\lambda/(1-\mu)$, which concludes the proof. □

Lemma 5. *For a given class of local resource allocation games $\mathcal{G}_T^n$, it holds that,*

$$\mathrm{PoA}(\mathcal{G}_T^n)\ge\max_{\mathbf{t}\in T}\{\mathrm{PoA}\left(\mathcal{G}_{\mathbf{t}}^n\right)\}. \tag{25}$$

Proof. We begin by proving that it is impossible to have $\mathrm{PoA}(\mathcal{G}_T^n)<\max_{\mathbf{t}\in T}\{\mathrm{PoA}(\mathcal{G}_{\mathbf{t}}^n)\}$. Simply note that the worst-case game in $\mathcal{G}_{\mathbf{t}}^n$ for each $\mathbf{t}\in T$ is a member of the class of games $\mathcal{G}_T^n$.

Next, consider the class of games with $n = 3$, and $T = \{T_1, T_2\} = \{(x^2, x), (x, x)\}$. By [20, Thm. 2], the prices-of-anarchy for the games with the individual types are $\text{PoA}(\mathcal{G}^3_{T_1}) = 1.857$ and $\text{PoA}(\mathcal{G}^3_{T_2}) = 2.0$. But, $\text{PoA}(\mathcal{G}^{\text{n}}_{\text{T}}) = 2.6$ by (13). □

Lemma 6. *For a given class of local resource allocation games $\mathcal{G}^n_T$, there exist scaling parameters $\alpha_t \in \mathbb{R}_{\geq 0}$, $t \in |T|$ such that,*

$$\text{PoA}(\mathcal{G}^n_\tau, n) = \max_{\mathbf{t} \in T} \{\text{PoA}(\mathcal{G}^n_{\mathbf{t}})\},$$

where $\tau = \{(c^t, \alpha_t f^t)\}_{t=1}^{|T|}$.

Proof. We denote by (ν_t^*, ρ_t^*) the solution to [20, Thm. 2] for the class of games with one type, $(c^t, f^t) \in T$. First, note that uniform scaling of the distribution rules does not affect the equilibrium conditions, so $\text{PoA}(\{(c^t, \alpha_t f^t)\}, n) = \text{PoA}(\{(c^t, f^t)\}, n)$ for all $\alpha_t > 0$ and all $\mathbf{t} \in T$. Thus, recalling Lemma 5,

$$\text{PoA}(\mathcal{G}^n_\tau) \geq \max_{\mathbf{t} \in T} \{\text{PoA}(\mathcal{G}^n_{\mathbf{t}})\},$$

where $\tau = \{(c^t, \alpha_1 f^t)\}_{t=1}^{|T|}$. Select $\alpha_t = \nu_t^*$ for all $t \in [|T|]$, such that $\tau := \{(c^t, \nu_t^* f^t)\}_{t=1}^{|T|}$. We define $\hat{\rho} := \min_{t \in [|T|]} \rho_t^*$. By construction, $(\hat{\rho}, 1)$ satisfies all the constraints in (13) for types in τ. Thus, $\text{PoA}(\mathcal{G}^n_\tau) \leq 1/\hat{\rho} = \max_{\mathbf{t} \in T}\{\text{PoA}(\mathcal{G}^n_{\mathbf{t}})\}$. □

Proof of Theorem 6

Proof. By Lemma 5, the lowest achievable price-of-anarchy is $\max_{\mathbf{t} \in T^*}\{\text{PoA}(\mathcal{G}^n_{\mathbf{t}})\}$ where $T^* := \{(c^t, f^t_{\text{OPT}})\}_{t=1}^{|T|}$. Additionally, each of the f^t_{OPT} minimizes its corresponding $\text{PoA}(\mathcal{G}^n_{\mathbf{t}})$ by [20, Thm. 3]. Finally, we have that the following statement,

$$\text{PoA}(\mathcal{G}^n_{T^*}) = \max_{\mathbf{t} \in T} \{\text{PoA}(\mathcal{G}^n_{\mathbf{t}})\},$$

holds by the construction of $f_{\text{OPT},t}$ in [20, Thm. 3]; the linear program already multiplies the distribution rule and λ_t^*, and it was shown in the proof of Lemma 6 that $\alpha_t = \lambda_t^*$ for all $t \in T$ is an optimal set of scaling parameters. □