

Multiagent Maximum Coverage Problems: The Trade-off Between Anarchy and Stability

Vinod Ramaswamy, Dario Paccagnan, and Jason R. Marden

Abstract—The price of anarchy and price of stability are two well-studied performance metrics that seek to characterize the inefficiency of equilibria in distributed systems. The distinction between these two performance metrics centers on the equilibria that they focus on: the price of anarchy characterizes the quality of the worst-performing equilibria, while the price of stability characterizes the quality of the best-performing equilibria. While much of the literature focuses on these metrics from an analysis perspective, in this work we consider these performance metrics from a design perspective. Specifically, we focus on the setting where a system operator is tasked with designing local agent utility functions to optimize these performance metrics in a class of games termed covering games. Our main result characterizes a fundamental trade-off between the price of anarchy and price of stability in the form of a fully explicit Pareto frontier. Within this setup, we observe that optimizing the price of anarchy comes directly at the expense of the price of stability (and vice versa). Our second result demonstrates how a system-operator could incorporate an additional piece of system-level information into the design of the agents’ utility functions to breach these limitations and improve the efficiency guarantees associated with the resulting equilibria. Informally, this valuable piece of system-level information pertains to the value of the largest uncovered resource in our covering game.

I. INTRODUCTION

A multiagent system can be characterized by a collection of individual subsystems, each making independent decisions in response to locally available information. Such a decision-making architecture can either emerge naturally as the result of self-interested behavior, e.g., drivers in a transportation network, or be the result of a design choice in engineered system. In the latter case, the need for distributed decision-making stems from the scale, spatial distribution, and sheer quantity of information associated with various problem domains that exclude the possibility for centralized decision making and control. Examples include the use of robotic networks in post-disaster environments [1], [15], [20], task scheduling and management [8], water conservative food production [18], fleets of autonomous vehicles [36], and micro-scale medical treatments [14], [35].

Regardless of the specific problem domain, the central goal associated with the design of a networked control system is

This work is supported by SNSF Grant #P2EZP2-181618, ONR Grant #N00014-20-1-2359, and AFOSR Grant #FA9550-20-1-0054 and NSF Grant #ECCS-1351866. The conference version of this work is given in [34].

V. Ramaswamy was with the Department of Electrical, Computer, and Energy Engineering, University of Colorado at Boulder, mailto:vinodr@gmail.com.

D. Paccagnan is with the Department of Mechanical Engineering and with the Center of Control, Dynamical Systems and Computation, University of California, Santa Barbara, dariop@ucsb.edu.

J. R. Marden is with the Department of Electrical and Computer Engineering, University of California, Santa Barbara, jrmarden@ece.ucsb.edu.

to derive *admissible* control policies for the decision-making entities that ensure the emergent collective behavior is desirable with regards to a given system-level objective. At a high level, this design process entails specifying two key elements: (i) the information available to each subsystem, attained either through sensing or communication, and (ii) a decision-making mechanism that prescribes how each subsystem processes available information to formulate a decision. The quality of the resulting networked control system is ultimately gauged by several dimensions including the stability and efficacy of the emergent collective behavior, characteristics of the transient behavior, in addition to the communication costs associated with propagating information throughout the system. Within this setting, two fundamental questions arise:

- (i) What are the decision-making rules that optimize the performance of the emergent collective behavior?
- (ii) How does informational availability translate to attainable performance guarantees?

This paper seeks to shed light on the answer to these two questions for the class of multiagent maximum coverage problems introduced in [12]. In a multiagent maximum coverage problem we are given a ground set of resources \mathcal{R} , n collections of subsets of the ground set $\mathcal{A}_1, \dots, \mathcal{A}_n$ where $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$ for each $i \in \{1, \dots, n\}$, and a valuation $v_r \geq 0$ for each resource $r \in \mathcal{R}$. Given a specific covering problem, the system-level objective is to select one set from each collection, i.e., one choice $a_i \in \mathcal{A}_i$ for each $i = 1, \dots, n$, so as to maximize the total value of the covered elements. It is important to highlight that there are well-established centralized algorithms that can derive an admissible allocation of agents to resources that is within a factor of $1 - 1/e$ of the optimal allocation’s value in polynomial time for special classes of multiagent maximum coverage problems [9], [19], [30]. Further, no polynomial time algorithm can provide a better approximation, unless $\mathcal{P} = \mathcal{NP}$. Unfortunately, the applicability of such centralized algorithms for the control of multiagent systems is limited given the concerns highlighted above.

This paper focuses on distributed approaches for reaching a near-optimal allocation where the individual agents make their covering selections in response to a *designed* decision-making policy that makes use of locally available information. The central goal here is to design agent decision-making rules that optimize the quality of the emergent collective behavior for a given level of informational availability. Of specific interest will be identifying how the level of information available to the individual agents impacts the attainable performance guarantees associated with the corresponding optimal networked control system.

In the spirit of [12], [26], we approach this problem through a game theoretic lens where we model the individual agents as players in a game and assign each agent a local objective function that guides its decision-making process. Here, we impose a locality constraints that each agent’s local objective can *only* depend on (i) the resources that the agent selected and (ii) the number of other agents that selected these resources. We treat these local objective functions as our design parameter and focus our analysis on characterizing the performance guarantees associated with the resulting equilibria of the designed game. We concentrate our analysis on two well-studied performance metrics in the game theoretic literature termed the *price of anarchy* and *price of stability* [2], [17]. Informally, the price of anarchy provides performance guarantees associated with the worst performing equilibrium relative to the optimal allocation. The price of stability, on the other hand, provides similar performance guarantees when restricting attention to the best performing equilibrium. The lack of uniqueness of equilibria implies that these bounds are often quite different.¹

The work of [12] was one of the first to view price of anarchy as a design objective rather than its more traditional analytical counterpart. Focusing on the same class of multi-agent maximum coverage problems, [12] identifies the agent objective functions that optimize the price of anarchy given the informational limitations discussed above. Note that in this setting, any agent i is unaware of the covering options of any other agents $j \neq i$, as well as any resource values that the agent is unable to cover. Interestingly, [12] demonstrates that this optimal price of anarchy attains the same $1 - 1/e$ guarantees of the best centralized algorithms, meaning that there is no degradation in terms of the worst-case efficiency guarantees when transitioning from the best centralized algorithm to the presented distributed algorithm that adheres to the prescribed informational limitations.

Our Contribution. The first main result of this manuscript addresses the achievable performance guarantees, in terms of the price of anarchy and price of stability, attainable through the design of agent objective functions of the above form in [12]. In Theorem 4.1, we characterize the price of anarchy and price of stability frontier that is achievable through the design of agent objective functions in these multiagent covering problems. This characterization demonstrates a fundamental trade-off between the price of anarchy and price of stability as design objectives in such multiagent covering problems. That is, designing agent objective functions to improve the worst-case performance guarantees necessarily degrades the best-case performance guarantees. As corner cases, we demonstrate that any objective functions that ensure a price of anarchy of $1 - 1/e$ also inherits a price of stability of $1 - 1/e$. Note that having a price of stability smaller than 1 implies that the optimal allocation is not necessarily an equilibrium. On the other hand, any objective functions that ensure a price of stability of 1 also inherit a price of anarchy of at most $1/2$.

The second main result of this manuscript demonstrates that one can move beyond this frontier through the design

¹The justification for focusing purely on equilibria, as opposed to dynamics, derives from the fact that there is a rich body of literature in distributed learning that could be employed to lead the collective behavior to an equilibrium of the derived game, c.f., [11].

of agent objective functions encompassing additional system-wide information. In Theorem 5.1, we identify a minimal (and easily attainable) piece of system-level information that permits the realization of decision-making rules with performance guarantees beyond the price of anarchy / price of stability frontier provided in Theorem 4.1. When agents are provided with this additional information, which can be informally interpreted as the largest value of an uncovered resource in the system, one can derive agent objective functions that yield a price of anarchy of $1 - 1/e$ and a price of stability of 1, which was unattainable without this piece of information.

The importance of this result centers on the fact that specific system-level information, if propagated to the agents, could be exploited in networked control algorithms to improve system performance. Understanding the trade-off between this improvement and the communication costs necessary to propagate this information throughout the system is clearly an important question that warrants future attention.

Related Work. The results contained in this manuscript add to the growing literature of utility design, which can be interpreted as a subfield of mechanism design [7] where the goal is to design admissible agent objective functions to optimize various performance metrics, such as the price of anarchy and price of stability, [16], [21], [27]. While recent work in [13] has identified all design approaches that ensure equilibrium existence in local utility designs, the question of optimizing the worst-case efficiency of the resulting equilibria, i.e., optimizing the price of anarchy, is far less understood. Nonetheless, there are a few positive results in this domain worth reviewing. Beyond [12], alternative problem domains where optimizing the price of anarchy has been explored include concave cost sharing games [24], optimal tolling in congestion games [32] and reverse carpooling games [23]. More recently, the authors in [31], [33] characterize and optimize the price of anarchy relative to a broader class of submodular and supermodular combinatorial optimization problems, rediscovering [12] as a special case. Lastly, a recent result in [10] characterizes a similar trade-off between the price of anarchy and price of stability in a mechanism design setting.

Much of the research regarding optimal utility design has concentrated on a specific class of objectives, termed *budget-balanced* objectives, which imposes the constraint that the sum of the agents’ objectives is equal to the system welfare for every allocation. Within the confines of budget-balanced agent objective functions, several works have identified the optimality of the Shapley value objective design with regards to the price of anarchy guarantees [6], [37], [38]. However, the imposition of the budget-balanced constraint is unwarranted in the context of multiagent system design and its removal allows for improved performance, as shown in [12] and this manuscript.

Paper Organization. In section II we introduce the multiagent coverage problem and corresponding performance metrics. In section III we provide some preliminary analysis on the performance of specific distribution rules. In section IV we present the trade-off between price of anarchy and stability. Lastly, in section V we show how to breach this trade-off by leveraging an additional piece of system-level information.

Notation. We use \mathbb{N} , $\mathbb{R}_{>0}$ and $\mathbb{R}_{\geq 0}$ to denote the set of natural,

positive and non negative real numbers; e is Euler's number.

II. MODEL AND PERFORMANCE METRICS

In this section we introduce the multiagent maximum coverage problem and our game theoretic model for the design of local decision-making mechanisms [12]. Further, we define the objectives and performance metrics of interest, as well as provide a review of the relevant literature.

A. Covering problems

Let $\mathcal{R} = \{r_1, r_2, \dots, r_m\}$ be a finite set of resources where each resource $r \in \mathcal{R}$ is associated with a value $v_r \geq 0$ defining its importance. We consider a covering problem where the goal is to allocate a collection of agents $N = \{1, \dots, n\}$ to resources in \mathcal{R} in order to maximize the cumulative value of the covered resources. The set of possible assignments for each agent $i \in N$ is given by $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$ and we express an admissible allocation by the tuple $a = (a_1, a_2, \dots, a_n) \in \mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$. The agents' assignment sets $\mathcal{A}_1, \dots, \mathcal{A}_n$ are fixed, and need not conform to any specific structure or symmetry. The total value, or welfare, associated with an allocation $a \in \mathcal{A}$ is given by

$$W(a) = \sum_{r \in \cup_{i \in N} a_i} v_r. \quad (1)$$

The goal of the covering problem is to find an optimal allocation, i.e., an allocation $a^{\text{opt}} \in \mathcal{A}$ that satisfies

$$a^{\text{opt}} \in \arg \max_{a \in \mathcal{A}} W(a), \quad (2)$$

where we restrict attention to admissible allocations $a \in \mathcal{A}$.

We will often express an allocation a as (a_i, a_{-i}) with the understanding that $a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$ denotes the collection of choices of the agents other than agent i . Lastly, we will periodically restrict our attention to the special case of single-selection covering problems where each allocation $a_i \in \mathcal{A}_i$ consists of a single resource, i.e., $\mathcal{A}_i \subseteq \mathcal{R}$ as opposed to $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$. Figure 1 provides some illustrative examples of single-selection covering problems.

B. A game theoretic model

This paper focuses on deriving distributed mechanisms for attaining near optimal solutions to covering problem where the individual agents make independent choices in response to local available information. Specifically, in this section we assume that each agent i has information only regarding the resources that the agent can select. Rather than directly specifying a decision-making process, here we focus on the design of local agent objective functions that adhere to these informational dependencies and will ultimately be used to guide the agents' selection process. To that end, we consider the framework proposed in [12] where each agent is associated with a local utility or objective function $U_i : \mathcal{A} \rightarrow \mathbb{R}$, and for any allocation $a = (a_i, a_{-i}) \in \mathcal{A}$, the utility of agent i is

$$U_i(a_i, a_{-i}) = \sum_{r \in a_i} v_r \cdot f(|a|_r), \quad (3)$$

where $|a|_r$ captures the number of agents that choose resource r in the allocation a , i.e., the cardinality of the set $\{i \in N :$

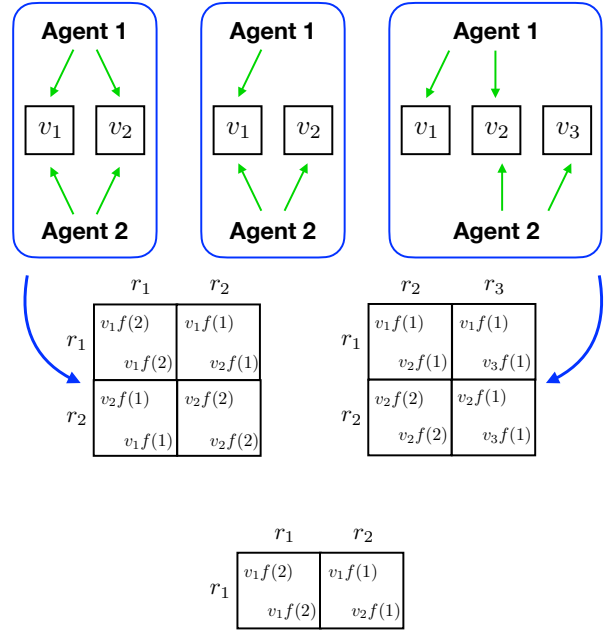


Fig. 1: The figure provides an illustration of three different single-selection covering problems with $N = \{1, 2\}$ where the potential assignments (or actions) of each agent are highlighted by the green arrows. For the covering problem on the left, there are two resources r_1 and r_2 with valuations v_1 and v_2 respectively. The action sets of both agents are symmetric and of the form $\mathcal{A}_1 = \mathcal{A}_2 = \{\{r_1\}, \{r_2\}\}$. The covering problem in the middle is identical to that on the left, with the sole exception being that the action set of agent 1 is now of the form $\mathcal{A}_1 = \{\{r_1\}\}$, meaning that agent 1 is now only capable of choosing resource r_1 . For the covering problem on the right, there are three resources r_1, r_2 , and r_3 with valuations v_1, v_2 , and v_3 respectively. The action set for each agent is now of the form $\mathcal{A}_1 = \{\{r_1\}, \{r_2\}\}$ and $\mathcal{A}_2 = \{\{r_2\}, \{r_3\}\}$. Here, the agents' action sets are no longer symmetric as agent 1 is not able to cover resource $\{r_3\}$ and similarly agent 2 is not able to cover resource $\{r_1\}$. It is important to highlight that the structure of the agents' action sets is part of the problem definition, i.e., not a design choice. Accordingly, for any problem the goal is to find the optimal allocation within the given admissible assignment set. In general, the allocation choices of an agent need not be restricted to singletons as highlighted above. Lastly, the two payoff matrices highlight the utility function given in (3) of both agents for the left and right scenarios respectively. Agent 1 is the Row player, Agent 2 is the Column player, and the numbers in each cell correspond to the agents' payoffs for the given joint action. Here, the top left entry in each cell is agent 1's payoff and the bottom right entry is agent 2's payoff for that particular joint action. Note that an optimal allocation is a feasible allocation that optimizes the maximum coverage objective, i.e., $a^{\text{opt}} \in \arg \max_{a \in \mathcal{A}} W(a)$. On the other hand, a Nash equilibrium is an admissible allocation $a^{\text{ne}} \in \mathcal{A}$ that satisfies $U_i(a_i^{\text{ne}}, a_{-i}^{\text{ne}}) \geq U_i(a_i, a_{-i}^{\text{ne}})$, for all $i \in N$ and $a_i \in \mathcal{A}_i$. The design of $f(\cdot)$ impacts the resulting set of Nash equilibria (as it impacts what function $U_i(\cdot)$ each agent utilizes), but has no impact on the optimal allocation. Hence, our design goal is to design $f(\cdot)$ to make the resulting Nash equilibria close in performance to the optimal allocation.

$r \in a_i\}$, and $f : \mathbb{N} \rightarrow \mathbb{R}$ defines the payoff awarded to the agent at each resource in the allocation a . In order to evaluate this utility, it is important to highlight that each agent i is required to observe (i) the resources in its action set and (ii) the number of other agents covering each of these resources.

We will refer to f as the *distribution rule* throughout.² We will express such a n -agent covering game by the tuple $G^n = \{N, \mathcal{R}, \{\mathcal{A}_i\}_{i \in N}, f, \{v_r\}_{r \in \mathcal{R}}\}$ and drop the subscripts on the above sets, e.g., denote $\{v_r\}_{r \in \mathcal{R}}$ as $\{v_r\}$, for brevity.

The goal of this paper is to derive the distribution rule f that optimizes the performance of the emergent collective behavior. Here, we focus on the concept of pure Nash equilibrium as a model for this emergent collective behavior [29]. A pure Nash equilibrium, which we will henceforth refer to as just an equilibrium, is defined as an allocation $a^{\text{ne}} \in \mathcal{A}$ such that for all $i \in N$ and for all $a_i \in \mathcal{A}_i$, we have

$$U_i(a_i^{\text{ne}}, a_{-i}^{\text{ne}}) \geq U_i(a_i, a_{-i}^{\text{ne}}).$$

In essence, an equilibrium represents an allocation for which no single agent has a unilateral incentive to alter its covering choice given the choices of the other agents. It is important to highlight that an equilibrium might not exist in a general game; however, when restricting attention to a covering game as defined above, an equilibrium is guaranteed to exist as the resulting game is known to be a congestion game [28].³

We will measure the efficiency of an equilibrium allocation in a game G^n through two commonly studied measures, termed *price of anarchy* and *price of stability*, defined as

$$\text{PoA}(G^n) = \min_{a^{\text{ne}} \in \text{NE}(G^n)} \frac{W(a^{\text{ne}})}{W(a^{\text{opt}})} \leq 1,$$

$$\text{PoS}(G^n) = \max_{a^{\text{ne}} \in \text{NE}(G^n)} \frac{W(a^{\text{ne}})}{W(a^{\text{opt}})} \leq 1,$$

where we use $\text{NE}(G^n)$ to denote the set of Nash equilibria of G^n . In words, the price of anarchy characterizes the performance of the worst equilibrium of G^n relative to the performance of the optimal allocation, while the price of stability focuses on the best equilibrium in the game G^n . Such distinction is required as equilibria are guaranteed to exist for the class of utilities considered in (3), but in general they are not unique. By definition $0 \leq \text{PoA}(G^n) \leq \text{PoS}(G^n) \leq 1$.

Throughout, we require that a system designer commits to a distribution rule *without* explicit knowledge of particular covering problem. To that end, the system operator will be aware of the maximum number of agents in the game n , but will have no specific knowledge of the actual number of agents, the resource set \mathcal{R} and valuations $\{v_r\}$, as well as the agents' action sets $\{\mathcal{A}_i\}$. For example, all of the three examples highlighted in Figure 1 represent instances of covering problems with at most two agents, and when designing f a system designer has no a priori knowledge about which specific instance will be realized. Note that once a particular distribution rule f has been chosen, this distribution rule defines a game for *any* realization of the parameters. The objective of the system designer is to provide desirable

²The structure of utility functions considered in (3) is common in the literature and falls under the framework of networked cost-sharing games or distributed welfare games, c.f., [26]. This paper focuses on characterizing the optimal performance guarantees attainable with this utility structure. Once such guarantees are known, one can evaluate whether or not it is desirable to raise or lower the informational dependence on the utility structure.

³There is a rich body of literature that provides distributed algorithms that can coordinate the agents to an equilibrium for the class of games considered in this paper [11], [22], [27], [39], [40]. However, we will not discuss such results due to space considerations.

performance guarantees irrespective of the realization of these parameter, even if they were chosen by an adversary. To that end, let \mathcal{G}_f^n denote the family of covering games with at most n agents that utilize a given distribution rule f , i.e., any game $G \in \mathcal{G}_f^n$ is of the above form. We will measure the quality of a distribution rule f by a worst-case analysis over the set of induced games \mathcal{G}_f^n , which is the natural extension of the price of anarchy and price of stability defined above, i.e.,

$$\text{PoA}(\mathcal{G}_f^n) = \min_{G^n \in \mathcal{G}_f^n} \text{PoA}(G^n), \quad (4)$$

$$\text{PoS}(\mathcal{G}_f^n) = \min_{G^n \in \mathcal{G}_f^n} \text{PoS}(G^n). \quad (5)$$

The price of anarchy $\text{PoA}(\mathcal{G}_f^n)$ for a given distribution rule f provides a bound on the quality of any equilibrium irrespective of the agent set N , resource set \mathcal{R} , action sets $\{\mathcal{A}_i\}$, and resource valuations $\{v_r\}$. The price of stability, on the other hand, provides similar performance guarantees when restricting attention to the best equilibrium.⁴

III. PRELIMINARY RESULTS

In this section we provide a preliminary study on the price of anarchy and price of stability for three well-studied distribution rules. Before doing so, we begin by presenting two general results for characterizing the price of anarchy and price of stability in such games where the distribution rule f is non-increasing, i.e., $f \in \mathcal{F}^n$ where

$$\mathcal{F}^n = \{f \in \mathbb{R}^n : f(k) \geq f(k+1), \forall k \in \{1, \dots, n-1\}\}.$$

This result strengthens the results presented in [12], [31], which focus purely on the set of positive non-increasing distribution rules that satisfies the constraint $f(1) = 1$, i.e.,

$$\mathcal{F}_1^n = \{f \in \mathbb{R}_{\geq 0}^n : f \in \mathcal{F}^n, f(1) = 1\}.$$

A. A Supporting Result for Price of Anarchy

Our first lemma provides a tight characterization of the price of anarchy for any distribution rule $f \in \mathcal{F}^n$. The results demonstrate that the relevant distribution rules are precisely the set of distribution rules given by \mathcal{F}_1^n .

Lemma 3.1: Let $n \geq 2$. The following properties hold:

(i) The price of anarchy associated with any distribution rule $f \in \mathcal{F}_1^n$ is

$$\text{PoA}(\mathcal{G}_f^n) = \frac{1}{1 + \chi_f^n}, \quad (6)$$

where

$$\chi_f^n = \max_{j \leq n-1} \{(j+1)f(j+1) - 1, jf(j) - f(j+1), jf(j+1)\}. \quad (7)$$

(ii) If $f(1) \leq 0$ or $f(k) < 0$ for any $k \geq 2$, the price of anarchy satisfies $\text{PoA}(\mathcal{G}_f^n) = 0$.

⁴One motivation for studying the price of stability is the availability of distributed learning rules that lead the collective behavior to the best equilibrium, e.g., [3], [4], [25]. However, it is important to note that the time complexity, e.g., mixing times, associated with such learning rules can be undesirable depending on the structure of the underlying game [5].

(iii) The optimal price of anarchy over the set of distribution rules \mathcal{F}^n is given by

$$\max_{f \in \mathcal{F}^n} \text{PoA}(\mathcal{G}_f^n) = 1 - \frac{1}{\frac{1}{(n-1)(n-1)!} + \sum_{i=0}^{n-1} \frac{1}{i!}}, \quad (8)$$

which satisfies

$$\lim_{n \rightarrow \infty} \max_{f \in \mathcal{F}^n} \text{PoA}(\mathcal{G}_f^n) = 1 - 1/e. \quad (9)$$

(iv) Lastly, a distribution rule $f^* \in \mathcal{F}^n$ that achieves the optimal price of anarchy in (8) is unique and of the form

$$f^*(j) = (j-1)! \left(\frac{1}{\frac{1}{(n-1)(n-1)!} + \sum_{i=j}^{n-1} \frac{1}{i!}} \right), \quad j \leq n. \quad (10)$$

We will refer to the optimal distribution rule given in (10) as the Gairing distribution rule.

Proof: Part (i) merely restates the results given [12], [31], c.f., Theorem 2 in [31]. We will prove Part (ii) of this lemma by constructing a specific game instance where the price of anarchy is 0, hence implying that the price of anarchy over the class of games \mathcal{G}_f^n is also 0. Suppose $f(1) = 0$ and consider a single-selection maximum coverage problem with three resources $\{r_1, r_2, r_3\}$ and valuations v_1, v_2 , and $v_3 = 0$. Furthermore, suppose $\mathcal{A}_1 = \{\{r_1\}, \{r_2\}\}$ and $\mathcal{A}_j = \{\{r_3\}\}$ for all agents $j \in \{2, \dots, n\}$. For any resource valuations $v_1 > v_2$, the action profile (r_2, r_3, \dots, r_3) is an equilibrium while the action profile (r_1, r_3, \dots, r_3) is the optimal allocation. Since v_1 can be arbitrarily large relative to v_2 , this completes the proof. A similar class of examples can be constructed for the case where $f(k) < 0$ for some $k \geq 2$. For Part (iii), consider the family of distribution rules $\mathcal{F}_+^n = \{f \in \mathbb{R}_{\geq 0}^n : f \in \mathcal{F}^n, f(1) > 0\}$. Building upon Part (ii), we know that

$$\max_{f \in \mathcal{F}^n} \text{PoA}(\mathcal{G}_f^n) = \max_{f \in \mathcal{F}_+^n} \text{PoA}(\mathcal{G}_f^n) = \max_{f \in \mathcal{F}_1^n} \text{PoA}(\mathcal{G}_f^n), \quad (11)$$

where the second equality comes from the fact that the price of anarchy is invariant to scaling, i.e., $\text{PoA}(\mathcal{G}_f^n) = \text{PoA}(\mathcal{G}_{\alpha f}^n)$ for any $f \in \mathcal{F}^n$ and $\alpha > 0$. Part (iv) follows from the results given [12], [31] coupled with the results in Part (iii), i.e., there is no loss in restricting attention to distribution rules to \mathcal{F}_1^n .⁵ \square

B. A Supporting Result for Price of Stability

In this section we provide a characterization of the price of stability for any distribution rule $f \in \mathcal{F}_1^n$. Our result provides an upper-bound on the price of anarchy for arbitrary settings. However, this bound is tight if we restrict attention to single-selection games.

Theorem 3.1: Let $n \geq 2$ and consider any distribution rule $f \in \mathcal{F}_1^n$. The price of stability associated with the induced family of games \mathcal{G}_f^n satisfies

$$\text{PoS}(\mathcal{G}_f^n) \leq \min_{1 \leq j \leq n} \left\{ \frac{1}{1 + (j-1)f(j)} \right\}. \quad (12)$$

⁵Interestingly, none of the results associated with Lemma 3.1 require f to be non-increasing. However, we assume that f is non-increasing to be consistent with the forthcoming results on price of stability that require this property.

Further, the optimal price of stability satisfies

$$\max_{f \in \mathcal{F}} \text{PoS}(\mathcal{G}_f^n) = 1, \quad (13)$$

and a distribution rule that achieves this price of stability over the induced games \mathcal{G}_f^n is the marginal contribution distribution rule defined as

$$f^{\text{mc}}(j) = \begin{cases} 1 & \text{if } j = 1, \\ 0 & \text{otherwise.} \end{cases} \quad (14)$$

Lastly, (12) is satisfied with equality for single-selection covering games with n agents.

Note that proving a bound on the price of stability is often technically challenging, as we need to perform a worst-case analysis over a reduced set of equilibria, i.e., over the best-performing equilibria. As a result of this, the proof of Theorem 3.1 is non-trivial, and therefore deferred to the Appendix.

C. Analysis of Existing Distribution Rules

We will now proceed with our analysis of three well-known distribution rules. The first distribution rule that we focus on is the *equal share distribution rule*, which is defined as

$$f^{\text{es}}(j) = \frac{1}{j}, \quad \forall j \geq 1. \quad (15)$$

Corollary 3.1: Consider the class of maximum coverage games $\mathcal{G}_{f^{\text{es}}}^n$ with the equal share distribution rule f^{es} and at most n agents. The price of anarchy and price of stability over the set of games $\mathcal{G}_{f^{\text{es}}}^n$ satisfies

$$\begin{aligned} \lim_{n \rightarrow \infty} \text{PoS}(\mathcal{G}_{f^{\text{es}}}^n) &= 1/2, \\ \lim_{n \rightarrow \infty} \text{PoA}(\mathcal{G}_{f^{\text{es}}}^n) &= 1/2. \end{aligned}$$

Proof of Corollary 3.1: We will start with the price of anarchy result. First, note that $f^{\text{es}} \in \mathcal{F}_1^n$. Using Lemma 3.1, the price of anarchy is $\text{PoA}(\mathcal{G}_{f^{\text{es}}}^n) = \frac{1}{1 + \chi_{f^{\text{es}}}^n}$ where

$$\chi_{f^{\text{es}}}^n = \max_{j \leq n-1} \left\{ 0, 1 - \frac{1}{j+1}, \frac{j}{j+1} \right\} = \frac{n-1}{n}. \quad (16)$$

Accordingly, we have that

$$\lim_{n \rightarrow \infty} \text{PoA}(\mathcal{G}_{f^{\text{es}}}^n) = \lim_{n \rightarrow \infty} \frac{1}{1 + \frac{n-1}{n}} = 1/2.$$

Moving to the price of stability, we know that $\text{PoS}(\mathcal{G}_{f^{\text{es}}}^n) \geq \text{PoA}(\mathcal{G}_{f^{\text{es}}}^n)$ by definition. Invoking Theorem 3.1, we have that

$$\begin{aligned} \text{PoS}(\mathcal{G}_f^n) &\leq \min_{1 \leq j \leq n} \left\{ \frac{1}{1 + (j-1)f(j)} \right\} = \\ &\leq \frac{1}{1 - (n-1)\frac{1}{n}}, \\ &= \text{PoA}(\mathcal{G}_f^n). \end{aligned}$$

Since $\text{PoA}(\mathcal{G}_{f^{\text{es}}}^n) \leq \text{PoS}(\mathcal{G}_{f^{\text{es}}}^n)$ by definition, we have that $\text{PoA}(\mathcal{G}_{f^{\text{es}}}^n) = \text{PoS}(\mathcal{G}_{f^{\text{es}}}^n)$ which completes the result. \square

The second distribution rule that we focus on is the *marginal contribution distribution rule* as defined in (14).

Corollary 3.2: Consider the class of maximum coverage games $\mathcal{G}_{f^{\text{mc}}}^n$ with the marginal cost distribution rule f^{mc} and

at most n agents. The price of anarchy and price of stability over the set of games $\mathcal{G}_{f^{mc}}^n$ satisfies

$$\begin{aligned} \lim_{n \rightarrow \infty} \text{PoS}(\mathcal{G}_{f^{mc}}^n) &= 1, \\ \lim_{n \rightarrow \infty} \text{PoA}(\mathcal{G}_{f^{mc}}^n) &= 1/2. \end{aligned}$$

Proof of Corollary 3.2: We will start with the price of anarchy result. Using Lemma 3.1, the price of anarchy is $\text{PoA}(\mathcal{G}_{f^{mc}}^n) = \frac{1}{1 + \chi_{f^{mc}}^n}$ where $\chi_{f^{mc}}^n = 1$, i.e., the maximum of (7) is attained for $j = 1$. Accordingly, the price of anarchy is $\text{PoA}(\mathcal{G}_{f^{mc}}^n) = 1/2$. With regards to the price of stability, it is well-known that the price of stability satisfies $\text{PoS}(\mathcal{G}_{f^{mc}}^n) = 1$. This follows from [27] as the resulting game is a potential game with potential function coinciding with the system welfare W . The fact that the potential function coincides with W can be easily seen upon substituting the expression for f^{mc} in the potential in (37). This completes the proof. \square

The last distribution rule that we focus on is the Gairing distribution rule given in (10).

Corollary 3.3: Consider the class of maximum coverage games $\mathcal{G}_{f^*}^n$ with the Gairing distribution rule f^* defined in (10) and at most n agents. The price of anarchy and price of stability over the set of games $\mathcal{G}_{f^*}^n$ satisfies

$$\begin{aligned} \lim_{n \rightarrow \infty} \text{PoS}(\mathcal{G}_{f^*}^n) &= 1 - 1/e, \\ \lim_{n \rightarrow \infty} \text{PoA}(\mathcal{G}_{f^*}^n) &= 1 - 1/e. \end{aligned}$$

Proof of Corollary 3.3: The price of anarchy result comes directly from Lemma 3.1. With regards to the price of stability, invoking Theorem 3.1 gives us

$$\begin{aligned} \text{PoS}(\mathcal{G}_{f^*}^n) &\leq \min_{1 \leq j \leq n} \left\{ \frac{1}{1 + (j-1)f^*(j)} \right\} \\ &\leq \frac{1}{1 + (n-1)f^*(n)} \\ &= 1 - \frac{1}{\frac{1}{(n-1)(n-1)!} + \sum_{i=0}^{n-1} \frac{1}{i!}} \\ &= \text{PoA}(\mathcal{G}_{f^*}^n). \end{aligned}$$

Hence, $\text{PoS}(\mathcal{G}_{f^*}^n) = \text{PoA}(\mathcal{G}_{f^*}^n)$ which completes the result. \square

The performance guarantees associated with these three distribution rules is summarized in Figure 2. First, note that there is never an incentive to utilize the equal share distribution rule as the marginal contribution distribution rule achieves the same price of anarchy guarantees while also ensuring a better price of stability (1 as opposed to 1/2). The comparison between the marginal contribution distribution rule and the Gairing distribution rule is not as straightforward. If the goal is to have a better price of anarchy, then the Gairing distribution rule is the better choice and one inherits a suboptimal price of stability ($1 - 1/e$ as opposed to 1). If the goal is to have a better price of stability, then the marginal contribution distribution rule is the better choice and one inherits a suboptimal price of anarchy (1/2 as opposed to $1 - 1/e$). The focus of the ensuing sections is on the existence of alternative distribution rules f that provide a more desirable balance between these two efficiency guarantees.




Fig. 2: On the left, distribution rule values for the equal share distribution rule in (15), marginal cost distribution rule in (14), and the Gairing distribution rule in (10) for the case when $n = 10$. On the right, corresponding prices of anarchy and prices of stability as $n \rightarrow \infty$.

IV. THE TRADE-OFF BETWEEN THE PRICE OF STABILITY AND PRICE OF ANARCHY

The goal of this paper is to investigate the design of distribution rules that optimize the metrics introduced in (4) and (5). The distribution rules highlighted in the previous section hint at the fact that there may be a fundamental trade-off between the price of anarchy and price of stability. That is, optimizing the price of anarchy necessarily comes at the expense of the price of stability (and vice versa). This section provides a characterization of the precise trade-off between these two important performance measures.

A. Characterizing the Trade-off

In this section we provide our first main result that characterizes the inherent tension between the price of anarchy and price of stability as design objectives in multiagent maximum coverage problems. For the remainder of the manuscript, let $\mathcal{F} = \lim_{n \rightarrow \infty} \mathcal{F}_1^n$ denote the set of all relevant distribution rules, i.e., any distribution rule $f \in \mathcal{F}$ satisfies $f: \mathbb{N} \rightarrow \mathbb{R}$, $f(1) = 1$, and f is non-increasing. Further, let $\mathcal{G}_f = \lim_{n \rightarrow \infty} \mathcal{G}_f^n$ capture all maximum coverage games with an arbitrary number of players and a given distribution rule $f \in \mathcal{F}$. First, note that for any $f \in \mathcal{F}$ and any $n \geq 1$, the set of games \mathcal{G}_f^n is well defined, even though certain elements of f might not be utilized. Second, for any $f \in \mathcal{F}$ and $n \geq 1$, the price of anarchy (and similarly the price of stability) satisfies $\text{PoA}(\mathcal{G}_f) \leq \text{PoA}(\mathcal{G}_f^n)$ since $\mathcal{G}_f^n \subset \mathcal{G}_f$.

Theorem 4.1: Consider the class of maximum coverage games introduced in Section II. The following holds:

- (i) The optimal price of anarchy satisfies

$$\max_{f \in \mathcal{F}} \text{PoA}(\mathcal{G}_f) = 1 - 1/e. \quad (17)$$

- (ii) Given a desired price of anarchy $\alpha \in [0, 1/2]$, the best attainable price of stability satisfies

$$\max_{f \in \mathcal{F}: \text{PoA}(\mathcal{G}_f) \geq \alpha} \text{PoS}(\mathcal{G}_f) = 1. \quad (18)$$

- (iii) Given a desired price of anarchy $\alpha \in (1/2, 1 - 1/e]$ and $n \geq 2$, the best attainable price of stability satisfies

$$\max_{f \in \mathcal{F}: \text{PoA}(\mathcal{G}_f^n) \geq \alpha} \text{PoS}(\mathcal{G}_f^n) \leq Z(\alpha, n), \quad (19)$$

where $Z(\alpha, n)$ equals

$$1 + \frac{1}{\max_{1 \leq j \leq n-1} j j! \left(1 - \left(\frac{1}{\alpha} - 1\right) \sum_{i=1}^j \frac{1}{i!}\right)}. \quad (20)$$

The bound in (19) is satisfied with equality if we restrict attention to single-selection maximum coverage games.

- (iv) Given a desired price of anarchy $\alpha = 1 - 1/e$, the best attainable price of stability satisfies

$$\max_{f \in \mathcal{F}: \text{PoA}(\mathcal{G}_f) \geq 1-1/e} \text{PoS}(\mathcal{G}_f) = 1 - 1/e. \quad (21)$$

Theorem 4.1 confirms the intuition gleaned in the previous section that there is indeed a trade-off between the price of anarchy (performance guarantees associated with the worst performing equilibrium) and price of stability (performance guarantees associated with the best performing equilibrium). Hence, there is an inherent tension between these two measures of efficiency as improving the performance of the worst equilibria necessarily comes at the expense of the performance of the best equilibria, and vice versa. The explicit trade-off given in Theorem 4.1 is illustrated in Figure 3. In particular, Theorem 4.1 establishes that there does not exist a distribution rule f that attains joint price of stability and price of anarchy in the red region of the figure. The expression of $Z(\alpha, n)$ defines this trade-off and satisfies $\lim_{n \rightarrow \infty} Z(1 - 1/e, n) = 1 - 1/e$. The last result demonstrates that this trade-off curve is tight in the context of single-selection maximum coverage games, i.e., there are distribution rules f that achieve joint price of anarchy and price of stability guarantees highlighted on the curve $Z(\alpha, n)$. The specific distribution rules will be characterized in the proof of Theorem 4.1.

B. Proof of Theorem 4.1

In this section we complete the Proof of Theorem 4.1. Part (i) follows immediately from Lemma 3.1. Part (ii) follows from Corollary 3.2. We now turn the attention to Part (iii). Towards this goal, let $\mathcal{F}(\alpha)$, $\alpha \in (0, 1]$ be the family of distribution rules that guarantee a price of anarchy of at least α , i.e.,

$$\mathcal{F}(\alpha) = \{f \in \mathcal{F} : \text{PoA}(\mathcal{G}_f^n) \geq \alpha\}.$$

Additionally, let

$$\text{PoS}(\mathcal{G}^n; \alpha) = \max_{f \in \mathcal{F}(\alpha)} \text{PoS}(\mathcal{G}_f^n) \quad (22)$$

be the best achievable price of stability given that the price of anarchy is guaranteed to exceed α , where $\alpha \in (0, 1)$.

We now prove that (19) holds with the equality sign, when restricting to single-selection covering games. It follows immediately that (19) hold with the inequality sign for general (not necessarily single-selection) covering game.

Restricting our attention to single-selection covering games, from Theorem 3.1 and (22) we have

$$\text{PoS}(\mathcal{G}^n; \alpha) = \max_{f \in \mathcal{F}(\alpha)} \min_{1 \leq j \leq n} \left\{ \frac{1}{1 + (j-1)f(j)} \right\}$$

or equivalently

$$\text{PoS}(\mathcal{G}^n; \alpha) = \max_{f \in \mathcal{F}(\alpha)} \min_{1 \leq j \leq n-1} \left\{ \frac{1}{1 + jf(j+1)} \right\}, \quad (23)$$

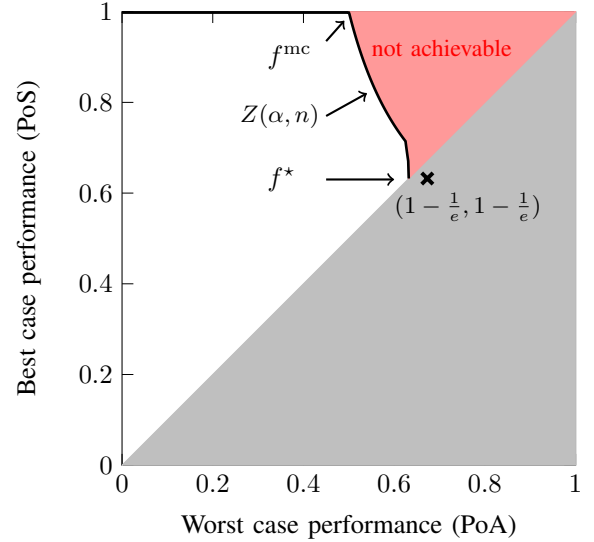


Fig. 3: The figure provides an illustration of the inherent trade-off between the price of anarchy and price of stability. First, note that the gray region is not achievable since $0 \leq \text{PoA}(\mathcal{G}_f) \leq \text{PoS}(\mathcal{G}_f) \leq 1$ by definition. Theorem 4.1 demonstrates that the red region is also *not* achievable. That is, there does not exist a distribution rule with joint price of anarchy and price of stability guarantees in the red region. For example, if the desired price of anarchy is $\alpha \leq 1/2$, then a price of stability of 1 is attainable while meeting this price of anarchy demand. However, if the desired price of anarchy is $\alpha = 1 - 1/e$, then a price of stability of 1 is no longer attainable. In fact, the best attainable price of stability is now also $1 - 1/e$.

since for $j = 1$ it is $(j-1)f(j) = 0$, and $(j-1)f(j) \geq 0$ for all other $2 \leq j \leq n$. Thanks to Lemma 3.1, the constraint $\text{PoA}(\mathcal{G}_f^n) \geq \alpha$ can be written as

$$\begin{aligned} jf(j) - \left(\frac{1}{\alpha} - 1\right) &\leq f(j+1) & 1 \leq j \leq n-1, \\ jf(j+1) &\leq \left(\frac{1}{\alpha} - 1\right) & 1 \leq j \leq n-1, \\ (j+1)f(j+1) - 1 &\leq \left(\frac{1}{\alpha} - 1\right) & 1 \leq j \leq n-1. \end{aligned}$$

Thus, the optimization problem in (23) is equivalent to

$$\begin{aligned} \min_f \max_{1 \leq j \leq n-1} jf(j+1) \\ \text{s.t.} \quad jf(j) - \left(\frac{1}{\alpha} - 1\right) &\leq f(j+1) & 1 \leq j \leq n-1, \\ jf(j+1) &\leq \frac{1}{\alpha} - 1 & 1 \leq j \leq n-1, \\ (j+1)f(j+1) &\leq \frac{1}{\alpha} & 1 \leq j \leq n-1, \\ f(j) &\geq 0 & 1 \leq j \leq n, \\ f(1) &= 1. \end{aligned}$$

Recall that $f(1) = 1$. Recursively applying the first set of inequalities, it follows that for every j with $1 \leq j \leq n-1$

$$jf(j+1) \geq j j! \left(1 - \left(\frac{1}{\alpha} - 1\right) \sum_{i=1}^j \frac{1}{i!}\right).$$

Since our objective is precisely to minimize the quantity $\max_{1 \leq j \leq n-1} j\hat{f}(j+1)$, we find a candidate solution by first solving for the distribution rule that satisfies the $n-1$ linear inequalities with equality, or set it to zero if the term $j\hat{f}(j) - (1/\alpha - 1)$ is negative. Such a distribution rule can be computed recursively, and is of the form

$$\hat{f}(j) = \max \left\{ (j-1)! \left(1 - \left(\frac{1}{\alpha} - 1 \right) \sum_{i=1}^{j-1} \frac{1}{i!} \right), 0 \right\}.$$

By construction, the above distribution satisfies $\hat{f}(1) = 1$ as well as the first and fourth set of constraints in the optimization problem above. In the following we verify that \hat{f} also satisfies the remaining set of constraints, and thus is optimal. Note that the terms $\hat{f}(k)$, $k > n$, are irrelevant in the given analysis. We now show that \hat{f} also satisfies the second and third set of constraints. We begin with the second set of constraints, i.e., we wish to show $j\hat{f}(j+1) \leq (1/\alpha - 1)$ for all α with $\frac{1}{2} < \alpha \leq \max_{f \in \mathcal{F}} \text{PoA}(\mathcal{G}_f^n)$ and for all j with $1 \leq j \leq n-1$. This is equivalent to

$$j j! \left(1 - \left(\frac{1}{\alpha} - 1 \right) \sum_{i=1}^j \frac{1}{i!} \right) \leq \frac{1}{\alpha} - 1,$$

as the result follows readily when \hat{f} takes the value $\hat{f}(j) = 0$. After some manipulation this is equivalent to showing

$$\frac{1}{j j!} + \sum_{i=1}^j \frac{1}{i!} \geq \left(\frac{1}{\alpha} - 1 \right)^{-1}$$

for all α with $\frac{1}{2} < \alpha \leq \max_{f \in \mathcal{F}} \text{PoA}(\mathcal{G}_f^n)$ and for all j with $1 \leq j \leq n-1$. Since $\alpha \leq \max_{f \in \mathcal{F}} \text{PoA}(\mathcal{G}_f^n)$, the result follows if we are able to prove the above inequality for $\alpha = \max_{f \in \mathcal{F}} \text{PoA}(\mathcal{G}_f^n)$, i.e.,

$$\frac{1}{j j!} + \sum_{i=1}^j \frac{1}{i!} \geq \left(\frac{1}{\max_{f \in \mathcal{F}} \text{PoA}(\mathcal{G}_f^n)} - 1 \right)^{-1}. \quad (24)$$

The expression for the optimal price of anarchy was found [12, Corollary 1], and amounts to

$$\max_{f \in \mathcal{F}} \text{PoA}(\mathcal{G}_f^n) = 1 - \frac{1}{\frac{1}{(n-1)(n-1)!} + \sum_{i=0}^{n-1} \frac{1}{i!}}.$$

Replacing this in (24), we are left to prove that for all $1 \leq j \leq n-1$

$$\frac{1}{j j!} + \sum_{i=0}^j \frac{1}{i!} \geq \frac{1}{(n-1)(n-1)!} + \sum_{i=0}^{n-1} \frac{1}{i!}. \quad (25)$$

To do so, let us define the function $g : \mathbb{N} \rightarrow \mathbb{R}$ as

$$g(j) = \frac{1}{j j!} + \sum_{i=0}^j \frac{1}{i!},$$

and observe that $g(j)$ is non-increasing in j . To see this note

that

$$\begin{aligned} g(j) &\geq g(j+1) && \iff \\ \frac{1}{j j!} + \sum_{i=0}^j \frac{1}{i!} &\geq \frac{1}{(j+1)(j+1)!} + \sum_{i=0}^{j+1} \frac{1}{i!} && \iff \\ \frac{1}{j j!} - \frac{1}{(j+1)(j+1)!} &\geq \frac{1}{(j+1)!} && \iff \\ \frac{1}{j} &\geq \frac{1}{j+1}, \end{aligned}$$

which holds for any $j \in \mathbb{N}$. Thus, we conclude that for j with $1 \leq j \leq n-1$ it holds

$$g(j) \geq g(n-1),$$

which is precisely the desired result in (25), thus proving that the third set of inequalities is satisfied.

We now turn our attention to the third set of constraints, and wish to prove that $(j+1)\hat{f}(j+1) \leq 1/\alpha$ for all α with $\frac{1}{2} < \alpha \leq \max_{f \in \mathcal{F}} \text{PoA}(\mathcal{G}_f^n)$ and for all j with $1 \leq j \leq n-1$. This follows directly as a consequence of $j\hat{f}(j+1) \leq (1/\alpha - 1)$ previously shown. Indeed

$$j\hat{f}(j+1) \leq \frac{1}{\alpha} - 1 \iff \hat{f}(j+1) \leq \frac{1}{j} \left(\frac{1}{\alpha} - 1 \right),$$

which implies the desired result

$$\begin{aligned} (j+1)\hat{f}(j+1) &\leq \frac{j+1}{j} \left(\frac{1}{\alpha} - 1 \right) = \left(\frac{1}{\alpha} - 1 \right) + \frac{1}{j} \left(\frac{1}{\alpha} - 1 \right) \\ &\leq \left(\frac{1}{\alpha} - 1 \right) + 1 = \frac{1}{\alpha}, \end{aligned}$$

where the last inequality follows from $j \geq 1$ and $1/\alpha - 1 \leq 1$ (due to $\alpha > 1/2$). This concludes the above reasoning, and shows that \hat{f} is a solution to (23). Using Theorem 3.1 on such distribution rule gives the desired result.

The result in Part (iv) can be proven by replacing α in (20) with the optimal price of anarchy from [12], [31], i.e., set

$$\alpha^* = 1 - \frac{1}{\frac{1}{(n-1)(n-1)!} + \sum_{i=0}^{n-1} \frac{1}{i!}}.$$

After some manipulation, we obtain

$$\text{PoS}(\mathcal{G}^n; \alpha^*) = \frac{\frac{1}{(n-1)(n-1)!} + \sum_{i=1}^{n-1} \frac{1}{i!}}{\frac{1}{(n-1)(n-1)!} + \sum_{i=0}^{n-1} \frac{1}{i!}} = \max_{f \in \mathcal{F}} \text{PoA}(\mathcal{G}_f^n).$$

Taking the limit as $n \rightarrow \infty$ gives the final result. \square

V. USING INFORMATION TO BREACH THE ANARCHY / STABILITY FRONTIER

The previous section highlights a fundamental tension between the price of stability and price of anarchy for the given covering problem when restricted to local agent objective functions of the form (3). In this section, we challenge the role of locality in these fundamental trade-offs. That is, we show how to move beyond the price of anarchy / price of stability frontier given in Theorem 4.1 if we allow the agents to condition their choice on a higher degree of system-level information. To do so, we restrict our attention to single-selection covering games.

With this goal in mind, we introduce a minimal and easily attainable piece of system-level information that can permit the realization of decision-making rules with efficiency guarantees beyond this frontier. To that end, for each allocation $a \in \mathcal{A}$ we define the information flow graph (V, E) where each node of the graph represents an agent and we construct a directed edge $i \rightarrow j$ if $a_i \in \mathcal{A}_j$ for $i \neq j$ (no self loops). This describes the situation in which the resource selected by agent i in allocation a also belongs to the allocation set of agent j . Based on this allocation-dependent graph, we define for each agent i the set $\mathcal{N}_i(a) \subseteq N$ consisting of all the agents that can reach i through a path in the graph (V, E) . Similarly, for each agent i we define

$$Q_i(a) = (\cup_{j \in \mathcal{N}_i(a)} \mathcal{A}_j) \cup \mathcal{A}_i, \quad (26)$$

which consists of the union of \mathcal{A}_i and all the sets of other agents that can reach i through a path in the graph. An example is shown in Figure 4. Building upon this graph we define the following quantities:

$$V_i(a) = \max_{r \in \mathcal{A}_i \setminus a_{-i}} v_r, \quad (27)$$

$$x_i(a) = \max_{r \in Q_i(a) \setminus a} v_r. \quad (28)$$

The term $V_i(a)$ captures the highest valued resource in agent i 's choice set \mathcal{A}_i that is not covered by any agent. If the set $\mathcal{A}_i \setminus a_{-i}$ is empty, we set $V_i(a) = 0$. Similarly, the term $x_i(a)$ captures the highest-valued resource in the enlarged set $Q_i(a)$ not currently covered by any other agent. If the set $Q_i(a) \setminus a$ is empty, we set $x_i(a) = 0$.

We are now ready to specify the information based covering game with a set of agents N , where each agent has an action set $\mathcal{A}_i \subseteq \mathcal{R}$. Here, we consider a state-based distribution rule that toggles between the two extreme optimal distribution rules f^* and f^{mc} . More formally, the distribution rule for agent i is now of the form

$$f_i^{\text{sb}}(a_i = r, a_{-i}) = \begin{cases} f^{\text{mc}}(|a|_r) & \text{if } V_i(a) \geq x_i(a), \\ f^*(|a|_r) & \text{otherwise,} \end{cases} \quad (29)$$

and the corresponding utilities are given by


$$U_i(a_i = r, a_{-i}) = v_r \cdot f_i^{\text{sb}}(a),$$

as we allow the system-level information $x_i(a)$ and $V_i(a)$ to prescribe which distribution rule each agent applies. We denote with $f^{\text{sb}} = \{f_i^{\text{sb}}\}_{i \in N}$ the collection of distribution rules in (29) and informally refer to it as the state-based distribution rule. Throughout, we express the distribution rule as merely $f_i^{\text{sb}}(a)$ instead of $f_i^{\text{sb}}(x_i(a), V_i(a))$ for brevity.

The next theorem demonstrates how f^{sb} attains performance guarantees beyond the price of stability / price of anarchy frontier established in Theorem 4.1.

Theorem 5.1: Consider any single-selection maximum coverage game with a state-based distribution rule f^{sb} as defined above. First, an equilibrium is guaranteed to exist in any game $G \in \mathcal{G}_{f^{\text{sb}}}$. Furthermore, the price of anarchy and price of stability associated with the induced family of games $\mathcal{G}_{f^{\text{sb}}}$ are

$$\begin{aligned} \text{PoS}(\mathcal{G}_{f^{\text{sb}}}) &= 1, \\ \text{PoA}(\mathcal{G}_{f^{\text{sb}}}) &= \max_{f \in \mathcal{F}} \text{PoA}(\mathcal{G}_f) = 1 - \frac{1}{e}. \end{aligned}$$



Agent i	$\mathcal{N}_i(a)$	$Q_i(a)$
1	{2, 3, 4}	$\mathcal{A}_1 \cup \mathcal{A}_2 \cup \mathcal{A}_3 \cup \mathcal{A}_4$
2	{3, 4}	$\mathcal{A}_2 \cup \mathcal{A}_3 \cup \mathcal{A}_4$
3	{2, 4}	$\mathcal{A}_2 \cup \mathcal{A}_3 \cup \mathcal{A}_4$
4	{2, 3}	$\mathcal{A}_2 \cup \mathcal{A}_3 \cup \mathcal{A}_4$
5	\emptyset	\mathcal{A}_5

Fig. 4: Example of graph (V, E) . The sets $\mathcal{N}_i(a)$ and $Q_i(a)$ are enumerated for each agent in the table on the right. Observe that the graph (V, E) as well as the sets $\mathcal{N}_i(a)$ and $Q_i(a)$ are allocation dependent.

Recall from Theorem 4.1 that a consequence of attaining a price of anarchy of $1 - 1/e$ was a price of stability also of $1 - 1/e$ and this was achieved by f^* defined in (10). Using the state-based rule given in (29), a system designer can achieve the optimal price of anarchy without any consequences for the price of stability. Hence, the identified piece of system-level information was crucial for moving beyond the inherited performance limitations by adhering to our notion of local information. Whether alternative forms of system-level information could move us beyond these guarantees, or achieve these guarantees with less information, is an open question.

A. Proof of Theorem 5.1

We will now present the proof of Theorem 5.1. For readability, we will split the proof into the following two subsections focusing on the price of anarchy and price of stability respectively.

1) *Proof of PoS Result:* We begin our proof with a lemma that identifies an important structure regarding the state based distribution rule f^{sb} .

Lemma 5.1: Let a be any allocation. Then for each agent $i \in N$, one of the following two statements is true.

- $U_i(a'_i = r, a_{-i}) = v_r \cdot f^{\text{mc}}(|a_{-i}| + 1)$, $\forall a'_i \in \mathcal{A}_i$; or
- $U_i(a'_i = r, a_{-i}) = v_r \cdot f^*(|a_{-i}| + 1)$, $\forall a'_i \in \mathcal{A}_i$.

Informally, this lemma states that for a given allocation a , the state based distribution rule will either evaluate every resource at f^* or f^{mc} for a given agent i .

Proof of Lemma 5.1 Let a be any allocation and i be any agent. Extend the definition of $x_i(a)$ in (28) as

$$y_i(a) = \max_{r \in (Q_i(a) \setminus a) \cap \mathcal{A}_i} v_r, \quad (30)$$

$$z_i(a) = \max_{r \in (Q_i(a) \setminus a) \setminus \mathcal{A}_i} v_r, \quad (31)$$

and note that $x_i(a) = \max\{y_i(a), z_i(a)\}$. First observe that for any $a'_i \in \mathcal{A}_i$: (i) $y_i(a) \leq V_i(a)$; (ii) $z_i(a) = z_i(a'_i, a_{-i})$; and (iii) $V_i(a) = V_i(a'_i, a_{-i})$. Accordingly, $x_i(a) > V_i(a)$ if and only if $x_i(a) = z_i(a)$. Consequently, $x_i(a) > V_i(a)$ if and only if $x_i(a'_i, a_{-i}) > V_i(a'_i, a_{-i})$ which completes the proof. \square

We will now prove that an equilibrium exists and the price of stability is 1. In particular, we will show that the optimal allocation a^{opt} is in fact an equilibrium.

Lemma 5.2: Consider any single-selection maximum coverage game with a state-based distribution rule f^{sb} . The optimal allocation a^{opt} is an equilibrium.

Proof of Lemma 5.2 Let a^{opt} be an optimal allocation. We begin by showing that $V_i(a^{\text{opt}}) \geq x_i(a^{\text{opt}})$ for all $i \in N$. Suppose this was not the case, and there exists an agent i such that $V_i(a^{\text{opt}}) < x_i(a^{\text{opt}})$. This implies that there exists a resource $r \in Q_i(a^{\text{opt}}) \setminus a^{\text{opt}}$ such that $v_r > v_{\bar{r}} = \max_{r_0 \in \mathcal{A}_i: |a_{-i}^{\text{opt}}|=0} v_{r_0}$. By definition of $Q_i(a)$ and $\mathcal{N}_i(a)$ there exists a sequence of players $\{i_0, i_1, \dots, i_m\}$ such that $a_{i_k}^{\text{opt}} \in \mathcal{A}_{i_{k-1}}$ for all $k \geq 1$, and $r \in \mathcal{A}_{i_m}$. Hence, consider a new allocation where $a_{i_k} = a_{i_{k+1}}^{\text{opt}}$ for all $k \in \{0, \dots, m-1\}$ and $a_{i_m} = r$, where $a_j = a_j^{\text{opt}}$ for all other agents $j \notin \{i_0, i_1, \dots, i_m\}$. The welfare of this allocation is $W(a) \geq W(a^{\text{opt}}) + v_r - v_{\bar{r}} > W(a^{\text{opt}})$, which contradicts the optimality of a^{opt} . This means that every agent will be using the marginal cost distribution rule to evaluate its utility in the allocation a^{opt} . Now, suppose a^{opt} is not an equilibrium for sake of contradiction. This means, that there exists an agent i with an action $a_i \in \mathcal{A}$ such that $U_i(a_i, a_{-i}^{\text{opt}}) > U_i(a^{\text{opt}})$. Since agents are using the marginal contribution distribution rule, which follows from $V_i(a^{\text{opt}}) \geq x_i(a^{\text{opt}})$, we have

$$W(a_i, a_{-i}^{\text{opt}}) = W(a^{\text{opt}}) - U_i(a^{\text{opt}}) + U_i(a_i, a_{-i}^{\text{opt}}) > W(a^{\text{opt}}),$$

which contradicts the optimality of a^{opt} . This completes the proof. \square

2) *Informal Discussion of PoA Result:* In the following we give an informal discussion for the price of anarchy result. This will be followed by a formal proof.

Consider a game $G = (N, \mathcal{R}, \{\mathcal{A}_i\}, \{f_i\}, \{v_r\})$. Let a^{ne} be any equilibrium of the game G . A crucial part of the forthcoming analysis will center on a new game G' derived from the original game G and the equilibrium a^{ne} , i.e.,

$$(G, a^{\text{ne}}) \longrightarrow G'.$$

This new game G' possesses the identical player set, resource set, and valuations of the resources as the game G . The difference between the games are (i) the action sets and (ii) the new game G' employs the Gairing distribution rule, f^* , as opposed to the state-based distribution rule f^{sb} . Informally, the proof proceeds in the following two steps:

– **Step 1:** We prove that the equilibrium a^{ne} of G is also an equilibrium of G' . Since the player set, resource set, and valuations of the resources are unchanged we have that

$$W(a^{\text{ne}}; G') = W(a^{\text{ne}}; G),$$

where we write the notation $W(a^{\text{ne}}; G')$ to mean the welfare accrued at the allocation a^{ne} in the game G' .

– **Step 2:** We show that the optimal allocation $a^{\text{opt}'}$ in the new game G' is at least as good as the optimal allocation in the original game G , i.e.,

$$W(a^{\text{opt}'}; G') \geq W(a^{\text{opt}}; G).$$

Combining the results from Step 1 and Step 2 give us

$$\frac{W(a^{\text{ne}}; G)}{W(a^{\text{opt}}; G)} \geq \frac{W(a^{\text{ne}}; G')}{W(a^{\text{opt}'}; G')} \geq \text{PoA}(\mathcal{G}_{f^*}),$$

where the last inequality follows from Lemma 3.1 since G' employs f^* .

3) *Proof of PoA Result:* In this section we present the formal proof for the price of anarchy result pertaining to the state based design. We begin by providing the formal details on the game construction highlighted above.

Construction of game G' : We will now provide the construction of the new game G' from the game G and the equilibrium a^{ne} . We begin with some notation that we will use to construct the new agents' action sets in the game G' . For each $i \in N$, let $r_i = a_i^{\text{ne}}$ and define

$$H_i = \{r \in \mathcal{A}_i : v_r f^*(|a^{\text{ne}}|_{r_i}) < v_r f^*(1 + |a^{\text{ne}}|_r)\}, \quad (32)$$

to be the set of resources that will give a strictly better payoff to agent i if the agent was to use f^* instead of f^{sb} . Further, denote

$$\mathcal{I} = \{i \in N : H_i \neq \emptyset\}, \quad (33)$$

as the set of agents that would move from the equilibrium a^{ne} if they were to use f^* instead of f^{sb} . Finally, for each agent $i \in \mathcal{I}$, let

$$B_i = \{r' \in Q_i(a^{\text{ne}}) : |a^{\text{ne}}|_{r'} = 0\}, \quad (34)$$

i.e. the enlarged set of resources of (26) that are not chosen by anyone at equilibrium.

We are now ready to construct the new game G' . As noted above, the agent set, resource set, and resource valuations are identical to those in G . The new action set of each agent i in game G' is defined as

$$\mathcal{A}'_i = \begin{cases} (\mathcal{A}_i \setminus H_i) \cup B_i \cup \emptyset & \text{if } i \in \mathcal{I}, \\ \mathcal{A}_i & \text{otherwise.} \end{cases}$$

Lastly, the utility functions of the agents are derived using f^* . We denote this game G' by the tuple $G' = (N, \mathcal{R}, \{\mathcal{A}'_i\}, f^*, \{v_r\})$.

Formal Proof of Step 1: In the first part of the proof, we will establish that the equilibrium allocation a^{ne} of game G is also an equilibrium allocation of game G' . By definition of the action sets $\{\mathcal{A}'_i\}$, we know that $a_i^{\text{ne}} \in \mathcal{A}'_i$ for all agents $i \in N$. It remains to show that for any agent $i \in N$

$$U_i(a_i^{\text{ne}}, a_{-i}^{\text{ne}}; G') \geq U_i(a'_i, a_{-i}^{\text{ne}}; G'), \quad \forall a'_i \in \mathcal{A}'_i, \quad (35)$$

where we use the notation $U_i(a; G')$ to denote the utility of agent i for allocation a in the game G' .

Based on the above definition of the action sets $\{\mathcal{A}'_i\}$, we only need to concentrate our attention on agents $i \in \mathcal{I}$ with choices $a'_i \in B_i$, as (35) follows immediately in the other cases. Since a^{ne} is an equilibrium of game G , then by Lemma 5.1 we know that each agent $i \in \mathcal{I}$ employs the marginal contribution distribution rule everywhere. Accordingly, we have

$$U_i(a_i^{\text{ne}}, a_{-i}^{\text{ne}}; G) = \max_{r \in \mathcal{A}_i \setminus a_i^{\text{ne}}} v_r = V_i(a^{\text{ne}}) \geq x_i(a^{\text{ne}}), \quad (36)$$

where the first equality follows from the equilibrium con-

ditions coupled with the use of the marginal contribution distribution rule and the inequality follows from the use of the marginal contribution distribution rule and (29).

We will conclude the proof by a case study on the potential values of $U_i(a_i^{\text{ne}}, a_{-i}^{\text{ne}}; G)$. For brevity in the forthcoming arguments we let $r = a_i^{\text{ne}}$.

– Case (i): Suppose $U_i(a_i^{\text{ne}}, a_{-i}^{\text{ne}}; G) = 0$. In this case, we have that $x_i(a^{\text{ne}}) = 0$ from (36) which implies that any resource $r' \in B_i$ has value $v_{r'} = 0$. Hence, $U_i(a_i^{\text{ne}}, a_{-i}^{\text{ne}}; G') = U_i(r', a_{-i}^{\text{ne}}; G') = 0$ and we are done.

– Case (ii): Suppose $U_i(a_i^{\text{ne}}, a_{-i}^{\text{ne}}; G) > 0$. Based on the definition of the marginal cost distribution rule, this implies that $|a^{\text{ne}}|_r = 1$, and hence

$$U_i(a^{\text{ne}}; G) = v_r f^{\text{mc}}(|a^{\text{ne}}|_r) = v_r f^*(|a^{\text{ne}}|_r) = U_i(a^{\text{ne}}; G'),$$

which follows from the fact that $f^{\text{mc}}(1) = f^*(1) = 1$. Hence, $U_i(a^{\text{ne}}; G') = v_r$. For any $r' \in B_i$, we have

$$U_i(r', a_{-i}^{\text{ne}}; G') = v_{r'} f^*(1 + |a_{-i}^{\text{ne}}|_{r'}) = v_{r'} \leq x_i(a^{\text{ne}}),$$

where the last inequality follows for the definition of $x_i(a^{\text{ne}})$. Combining with (36) gives us

$$U_i(a_i^{\text{ne}}, a_{-i}^{\text{ne}}; G') \geq U_i(r', a_{-i}^{\text{ne}}; G'),$$

which completes the proof. \square

Formal Proof of Step 2: We begin with a lemma that highlights a structure associated with the action sets $\{\mathcal{A}'_i\}$ in the new constructed game G' .

Lemma 5.3: If $r \in \mathcal{A}'_i \setminus \mathcal{A}_i$ for some agent $i \in N$ and resource $r \in \mathcal{R}$, then there exists an agent $j \neq i$ such that $a_j^{\text{ne}} = r$ and consequently $r \in \mathcal{A}'_j$.

Proof of Lemma 5.3 Suppose $r \in \mathcal{A}'_i \setminus \mathcal{A}_i$ for some agent $i \in N$ and resource $r \in \mathcal{R}$. Then, $r \in H_i$ by definition of the set \mathcal{A}'_i . By Lemma 5.1, each agent must either be a marginal contribution agent, i.e., uses f^{mc} at all resources, or a Gairing agent, i.e., uses f^* at all resources. Since, $i \in \mathcal{I}$ and a^{ne} is an equilibrium, agent i must be a marginal contribution agent, i.e.,

$$U_i(a_i^{\text{ne}} = \tilde{r}, a_{-i}^{\text{ne}}) = v_{\tilde{r}} \cdot f^{\text{mc}}(|a^{\text{ne}}|_{\tilde{r}}) \geq v_r \cdot f^{\text{mc}}(|a_{-i}^{\text{ne}}|_r + 1),$$

where the inequality follows from the equilibrium conditions. Suppose by contradiction that $|a_{-i}^{\text{ne}}|_r = 0$. In this case we have

$$v_r \cdot f^{\text{mc}}(|a_{-i}^{\text{ne}}|_r + 1) = v_r \cdot f^*(|a_{-i}^{\text{ne}}|_r + 1).$$

which follows from the fact that $f^{\text{mc}}(1) = f^*(1) = 1$. Since $f^*(k) \geq f^{\text{mc}}(k)$ for all $k \geq 1$, this implies

$$v_{\tilde{r}} \cdot f^*(|a^{\text{ne}}|_{\tilde{r}}) \geq v_{\tilde{r}} \cdot f^{\text{mc}}(|a^{\text{ne}}|_{\tilde{r}}) \geq v_r \cdot f^*(|a_{-i}^{\text{ne}}|_r + 1).$$

Hence, $r \notin H_i$ leading to the contradiction. This completes the proof. \square

We exploit the result of Lemma 5.3 to prove Step 2. Towards this goal, we construct an allocation $a \in \mathcal{A}'$ that satisfies $W(a; G') = W(a^{\text{opt}}; G)$ where $a^{\text{opt}} \in \arg \max_{a \in \mathcal{A}} W(a; G)$. We begin with an initial allocation a where for each agent $i \in N$

$$a = \begin{cases} a_i^{\text{opt}} & \text{if } a_i^{\text{opt}} \in \mathcal{A}'_i, \\ \emptyset & \text{else,} \end{cases}$$

That is, we assign each agent the agent's optimal allocation choice if it is available to them in the new action set \mathcal{A}'_i . If all agents received their optimal choice, then the proof is complete.

If this is not the case, then there will be a set of uncovered resources $\mathcal{U}(a) = \{r \in a^{\text{opt}} : |a|_r = 0\}$ which we denote by $\mathcal{U}(a) = \{r_1, \dots, r_m\}$. We will now argue that we can construct a new allocation a' that covers one additional resource from the set $\mathcal{U}(a)$, i.e., $|\mathcal{U}(a')| = |\mathcal{U}(a)| - 1$ and $a \subseteq a'$, where we denote with $|\mathcal{U}(a)|$ the cardinality of $\mathcal{U}(a)$.

To that end, consider any uncovered resource $r_0 \in \mathcal{U}(a)$. By definition, there exists an agent $i_0 \in N$ such that $a_{i_0}^{\text{opt}} = r_0$ but $r_0 \notin \mathcal{A}'_{i_0}$. Consequently, we have that $r_0 \in H_{i_0}$ and by Lemma 5.3 we know that there exists an agent $i_1 \neq i_0$ such that $a_{i_1}^{\text{ne}} = r_0$. Since $a_{i_1}^{\text{ne}} = r_0$ we also have that $r_0 \in \mathcal{A}'_{i_1}$ by definition. We now analyze the following three cases:

– Case 1: Suppose $a_{i_1} = \emptyset$. Then define a new allocation $a'_{i_1} = r_0$ and $a'_j = a_j$ for all $j \neq i_1$, and we are done.

– Case 2: Suppose $a_{i_1} = r_1$ and $|a^{\text{ne}}|_{r_1} = 0$, meaning that there are no agents at the resource r_1 in the equilibrium allocation. Then by definition $r_1 \in B_{i_0}$ and $r_1 \in \mathcal{A}'_{i_0}$. Define the allocation $a'_{i_0} = r_1$, $a'_{i_1} = r_0$, and $a'_j = a_j$ for all $j \neq i_0, i_1$ and we are done.

– Case 3: Suppose $a_{i_1} = r_1$ and $|a^{\text{ne}}|_{r_1} > 0$, meaning that there are agents at the resource r_1 in the equilibrium allocation. Select any agent i_2 such that $a_{i_2}^{\text{ne}} = r_1$.

- (i) If $a_{i_2} = \emptyset$, then consider the allocation $a'_{i_1} = r_0$, $a'_{i_2} = r_1$ and $a'_j = a_j$ for all $j \neq i_1, i_2$ and we are done.
- (ii) Otherwise, if $a_{i_2} = r_2$, then let $a'_{i_1} = r_0$, $a'_{i_2} = r_1$, and repeat Case 2 or Case 3 depending on whether $|a^{\text{ne}}|_{r_2} = 0$ or $|a^{\text{ne}}|_{r_2} > 0$. Note that Case 3-(ii) can be repeated at most n iterations until an alternative case that terminates is reached. To see this, note that each time an agent is given a new choice in this process, i.e., $a_i \rightarrow a'_i \neq a_i$, the agent's new choice is the agent's equilibrium choice, i.e., $a'_i = a_i^{\text{ne}}$. Therefore, once an agent is assigned a new choice, the agent will never be reassigned in this process.

Starting from a as defined above, the above process results in a new allocation a' that satisfies $|\mathcal{U}(a')| = |\mathcal{U}(a)| - 1$ and $a \subseteq a'$. As with the allocation a , the allocation a' satisfies $\max_r |a'|_r = 1$ and $a' \subseteq a^{\text{opt}}$. If $a' = a^{\text{opt}}$, we are done. Otherwise, we can repeat the process depicted above to generate a new allocation a'' such that $|\mathcal{U}(a'')| = |\mathcal{U}(a')| - 1$ as nowhere in the process did we rely on the fact that $a_i = a_i^{\text{opt}}$. Repeating these arguments recursively provides the result. \square

VI. CONCLUSIONS

How should a system operator design a networked architecture? The answer to this question is non-trivial and involves weighing the advantages and disadvantages associated with different design choices. In this paper we highlight one novel trade-off pertaining to the worst-case and best-case performance guarantees in distributed maximum coverage problems with local information. Further, we demonstrate how a system designer can move beyond these trade-offs by equipping the agents with additional information about the system. Fully realizing the potential of multiagent systems requires the pursuit of a more formal understanding of the

inherent limitations and performance trade-offs associated with networked architectures. While this paper focused purely on two performance measures, other metrics of interest include convergence rates, robustness to adversaries, and fairness. In each of these settings, it is imperative that a system operator fully understands the role of information within these trade-offs. Only then, will a system operator be able to effectively balance the potential performance gains with the communication costs associated with propagating additional information through the system.

VII. APPENDIX

A. Proof of Theorem 3.1

This section is dedicated to the proof of Theorem 3.1. We will prove the result on the price of stability through a series of intermediate lemmas. We begin by observing that any game G in the class \mathcal{G}_f^n is a congestion game, and thus is a potential game as introduced in [28], with a potential function $\phi : \mathcal{A} \rightarrow \mathbb{R}$ of the form⁶

$$\phi(a) = \sum_{r \in \mathcal{R}} \sum_{j=1}^{|a|_r} v_r f(j). \quad (37)$$

Further recall that an equilibrium is guaranteed to exist in any potential game [28], and one such equilibrium is the allocation that maximizes the potential function ϕ , i.e., $a^{\text{ne}} \in \arg \max_{a \in \mathcal{A}} \phi(a)$.

To prove (12), we restrict our attention to the set of single-selection covering games, where the optimal allocation is disjoint, i.e., $a_i^{\text{opt}} \neq a_j^{\text{opt}}$, for any $i \neq j$. We denote such games by the set $\bar{\mathcal{G}}_f^n \subset \mathcal{G}_f^n$. Our first lemma, stated without proof for brevity, demonstrates that restricting attention to single-selection covering games where the optimal allocation is disjoint is sufficient for characterizing the price of stability in single-selection covering games.

Lemma 7.1: Let $\bar{\mathcal{G}}_f^n$ denote the set of n -agent single-selection maximum coverage games with distribution rule $f \in \mathcal{F}$. It holds $\text{PoS}(\bar{\mathcal{G}}_f^n) = \text{PoS}(\mathcal{G}_f^n)$.

We will now proceed with a series of claims to demonstrate that $\text{PoS}(\bar{\mathcal{G}}_f^n)$ satisfies (12) with equality. The central part of the proof involves focusing on the equilibrium which maximizes the potential function in (37) in the considered class of games $\bar{\mathcal{G}}_f^n$. From this specific equilibrium, we consider a sequence of allocations taking the form $a^0 = a^{\text{ne}}$ and $a^k = (a_{i(k)}^{\text{opt}}, a_{-i(k)}^{k-1})$ for all $k \in \{1, \dots, m\}$ where $i(k)$ is the deviating player in the k -th profile. That is, any two adjacent allocations a^{k-1} and a^k differ by at most one agent's choice, namely agent $i(k)$ that switches from $a_{i(k)}^{\text{ne}}$ to $a_{i(k)}^{\text{opt}}$. In the case that $a_{i(k)}^{\text{ne}} = a_{i(k)}^{\text{opt}}$, the allocations are the same, i.e., $a^k = a^{k+1}$. The selection of the deviating players $\mathcal{I} = \{i(1), \dots, i(m)\}$ is chosen according to the following rules:

- (i) Let $i(1) \in N$ be any arbitrary player.
- (ii) For each $k \geq 1$, if $a_{i(k)}^{\text{opt}} = a_{i(k)}^{\text{ne}}$ or $a_{i(k)}^{\text{opt}} \notin a^{\text{ne}}$ then the sequence is terminated.
- (iii) Otherwise, let $i(k+1)$ be any agent in the set $\{j \in N : a_j^k = a_{i(k)}^{\text{opt}}\}$ and repeat.

⁶A proof of this can be found in [28], where the potential function ϕ is also defined.

The sequence of allocations defined above is merely employed as a convenient mathematical formulation to derive a relationship between $W(a^{\text{ne}})$ and $W(a^{\text{opt}})$. The first part of this derivation is provided in the following lemma.

Lemma 7.2: Define $Q = \cup_{i \in \mathcal{I}} a_i^{\text{ne}}$ and $\bar{Q} = \cup_{i \in \mathcal{I}} a_i^{\text{opt}}$. Then

$$\sum_{i \in \mathcal{I}} U_i(a^{\text{ne}}) \geq \sum_{r \in Q \cap \bar{Q}} v_r f(|a^{\text{ne}}|_r) + \sum_{r \in \bar{Q} \setminus Q} v_r. \quad (38)$$

Proof of Lemma 7.2 We begin with two observations on the above sequence of allocations: (a) the sequence of allocations can continue at most n steps due to the disjointness of a^{opt} and (b) if the sequence continues, it must be that $a_{i(k+1)}^{\text{ne}} = a_{i(k)}^{\text{opt}}$. Observation (b) ensures us that

$$\psi = \sum_{k=1}^{m-1} U_{i(k+1)}(a^k) - U_{i(k)}(a^k) = 0. \quad (39)$$

Accordingly, we have that

$$\begin{aligned} \phi(a^0) - \phi(a^m) &= \sum_{k=0}^{m-1} \phi(a^k) - \phi(a^{k+1}) \\ &= \sum_{k=0}^{m-1} U_{i(k+1)}(a^k) - U_{i(k+1)}(a^{k+1}) \\ &= U_{i(1)}(a^0) - U_{i(m)}(a^m) + \psi \\ &= U_{i(1)}(a^0) - U_{i(m)}(a^m) \geq 0. \end{aligned}$$

The first and third equalities follow by rearranging the terms. The second equality can be shown using the definition of ϕ as in (37); the last equality follows by (39). The inequality derives from the fact that $a^0 = a^{\text{ne}}$ optimizes the potential function. Thanks to observation (b), one can show that

$$Q \setminus \bar{Q} = a_{i(1)}^{\text{ne}} \setminus a_{i(m)}^{\text{opt}} \quad \text{and} \quad \bar{Q} \setminus Q = a_{i(m)}^{\text{opt}} \setminus a_{i(1)}^{\text{ne}}.$$

If $Q \setminus \bar{Q} \neq \emptyset$, it must be that $a_{i(1)}^{\text{ne}} \neq a_{i(m)}^{\text{opt}}$ so that $Q \setminus \bar{Q} = a_{i(1)}^{\text{ne}}$ and $\bar{Q} \setminus Q = a_{i(m)}^{\text{opt}}$. Using $a_{i(1)}^{\text{ne}} \neq a_{i(m)}^{\text{opt}}$ in condition (ii), tells us that $a_{i(m)}^{\text{opt}} \notin a^{\text{ne}}$ and thus the resource $a_{i(m)}^{\text{opt}}$ is not chosen by anyone else in the allocation a^m . Thus, when $Q \setminus \bar{Q} \neq \emptyset$,

$$U_{i(1)}(a^0) - U_{i(m)}(a^m) = \sum_{r \in Q \setminus \bar{Q}} v_r f(|a^{\text{ne}}|_r) - \sum_{r \in \bar{Q} \setminus Q} v_r \geq 0. \quad (40)$$

When $Q \setminus \bar{Q} = \emptyset$, also $\bar{Q} \setminus Q = \emptyset$ and thus the previous inequality still holds. Rearranging the terms and adding $\sum_{r \in Q \cap \bar{Q}} v_r f(|a^{\text{ne}}|_r)$ to each side gives us

$$\sum_{r \in Q} v_r f(|a^{\text{ne}}|_r) \geq \sum_{r \in Q \cap \bar{Q}} v_r f(|a^{\text{ne}}|_r) + \sum_{r \in \bar{Q} \setminus Q} v_r. \quad (41)$$

Finally note that

$$\sum_{i \in \mathcal{I}} U_i(a^{\text{ne}}) \geq \sum_{r \in Q} v_r f(|a^{\text{ne}}|_r)$$

which together with (41) completes the proof. \square

Our next lemma shows that there exists a collection of disjoint sequences that covers all players in N . We will express a sequence merely by the deviating player set \mathcal{I} with the understanding that this set uniquely determines the sequence of allocations.

Lemma 7.3: There exists a collection of deviating players $\mathcal{I}^1, \dots, \mathcal{I}^p$ chosen according to the process described above such that $\cup_k \mathcal{I}^k = N$ and $\mathcal{I}^j \cap \mathcal{I}^k = \emptyset$ for any $j \neq k$.

Proof of Lemma 7.3 Suppose $\mathcal{I}^1, \mathcal{I}^2, \dots, \mathcal{I}^k$ represent the first k sequences of deviating players. Further assume that they are all disjoint. Choose some player $i \in N \setminus \cup_k \mathcal{I}^k$ to start the $(k+1)$ -th sequence. If no such player exists, we are done. Otherwise, construct the sequence according to the process depicted above. If the sequence terminates without selecting a player in $\cup_k \mathcal{I}^k$, then repeat this process to generate the $(k+2)$ -th sequence. Otherwise, let $i^{k+1}(j)$, $j \geq 2$, denote the first player in the $(k+1)$ -th sequence contained in the set $\cup_k \mathcal{I}^k$. Since $a_i^{\text{opt}} \neq a_j^{\text{opt}}$ (for $i \neq j$), this player must be the first player in a previous sequence. Suppose this player is $i^\ell(1)$, where $\ell \in \{1, \dots, k\}$. If this is the case, replace the ℓ -th sequence with $\{i^{k+1}(1), \dots, i^{k+1}(j-1), \mathcal{I}^\ell\}$ which is a valid sequence and disjoint from the others. Then repeat the process above to choose the $(k+1)$ -th sequence. Note that this process can continue at most n -steps and will always result in a collection of disjoint sequences that cover all players in N . This completes the proof. \square

In the following we complete the proof of Theorem 3.1, by means of Lemmas 7.2 and 7.3.

Proof of Theorem 3.1: We begin showing a lower bound on the price of stability. Let $\mathcal{I}^1, \dots, \mathcal{I}^p$ denote a collection of deviating players that satisfies Lemma 7.3. Further, let Q^k and \bar{Q}^k be defined as above for each sequence $k = 1, \dots, p$. Using the result (38) from Lemma 7.2, we have

$$\begin{aligned} \sum_{i \in N} U_i(a^{\text{ne}}) &= \sum_{k=1}^p \sum_{i \in \mathcal{I}^k} U_i(a^{\text{ne}}) \\ &\geq \sum_{k=1}^p \left(\sum_{r \in Q^k \cap \bar{Q}^k} v_r f(|a^{\text{ne}}|_r) + \sum_{r \in \bar{Q}^k \setminus Q^k} v_r \right) \\ &= \sum_{r \in a^{\text{opt}} \cap a^{\text{ne}}} v_r f(|a^{\text{ne}}|_r) + \sum_{r \in a^{\text{opt}} \setminus a^{\text{ne}}} v_r, \end{aligned}$$

where the above equality follows from the fact that $\bar{Q}^i \cap \bar{Q}^j = \emptyset$ for any $i \neq j$ which is due to the disjointness of a^{opt} . Using the definition of $U_i(a^{\text{ne}})$, we have

$$\sum_{r \in a^{\text{ne}} \setminus a^{\text{opt}}} v_r |a^{\text{ne}}|_r f(|a^{\text{ne}}|_r) + \sum_{r \in a^{\text{ne}} \cap a^{\text{opt}}} v_r (|a^{\text{ne}}|_r - 1) f(|a^{\text{ne}}|_r) \geq \sum_{r \in a^{\text{opt}} \setminus a^{\text{ne}}} v_r.$$

Define $\gamma = \max_{j \leq n} (j-1)f(j)$. Working with the above expression we have

$$\sum_{r \in a^{\text{ne}} \setminus a^{\text{opt}}} v_r (\gamma + 1) + \sum_{r \in a^{\text{ne}} \cap a^{\text{opt}}} v_r \gamma \geq \sum_{r \in a^{\text{opt}} \setminus a^{\text{ne}}} v_r,$$

which gives us that

$$(\gamma + 1)W(a^{\text{ne}}) \geq W(a^{\text{opt}})$$

which completes the lower bound.

We will now provide an accompanying upper bound on the price of stability. To that end, consider a family of examples parameterized by a coefficient $j \in \{1, \dots, n\}$. For each j , the game consists of j agents and $(j+1)$ -resources $\mathcal{R} = \{r_0, r_1, \dots, r_j\}$ where the values of the resources are $v_{r_0} = 1$ and $v_{r_1} = \dots = v_{r_j} = f(j) - \varepsilon$ where $\varepsilon > 0$ is an arbitrarily small constant, and the action set of each agent

$i \in \{1, \dots, j\}$ is $\mathcal{A}_i = \{r_0, r_i\}$. The unique equilibrium is of the form $a^{\text{ne}} = (r_0, \dots, r_0)$ as every agent selects resource r_0 and the total welfare is $W(a^{\text{ne}}) = 1$. The optimal allocation is of the form $a^{\text{opt}} = (r_0, r_2, \dots, r_j)$ which generates a total welfare of $W(a^{\text{opt}}) = 1 + (j-1)(f(j) - \varepsilon)$. Performing a worst case analysis over ε and j gives (12). The claim in (13) follows from observing that f^{mc} is always optimal, and returns a price of stability of 1, thanks to Corollary 3.2. \square

REFERENCES

- [1] K. Alexis, G. Nikolakopoulos, A. Tzes, and L. Dritsas. *Coordination of Helicopter UAVs for Aerial Forest-Fire Surveillance*, pages 169–193. Springer, 2009.
- [2] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, pages 295–304, Oct 2004.
- [3] L. Blume. The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5:387–424, 1993.
- [4] L. Blume. Population games. In *The Economy as an evolving complex system II*, pages 425–460. Addison-Wesley, 1997.
- [5] H. Borowski and J. R. Marden. Fast convergence in semianonymous potential games. *IEEE Transactions on Control of Network Systems*, 4(2):246–258, 2017.
- [6] Ho-Lin Chen, Tim Roughgarden, and Gregory Valiant. Designing network protocols for good equilibria. *SIAM Journal on Computing*, 39(5):1799–1832, 2010.
- [7] George Christodoulou, Elias Koutsoupias, and Akash Nanavati. Coordination mechanisms. In *International Colloquium on Automata, Languages, and Programming*, pages 345–357. Springer, 2004.
- [8] Andreas T Ernst, Houyuan Jiang, Mohan Krishnamoorthy, and David Sier. Staff scheduling and rostering: A review of applications, methods and models. *European journal of operational research*, 153(1):3–27, 2004.
- [9] U. Feige and J. Vondrak. Approximation algorithms for allocation problems: Improving the factor of $1 - 1/e$. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, pages 667–676, Oct 2006.
- [10] Aris Filos-Ratsikas, Yiannis Giannakopoulos, and Philip Lazos. The pareto frontier of inefficiency in mechanism design. In *International Conference on Web and Internet Economics*, pages 186–199. Springer, 2019.
- [11] D. Fudenberg and D.K. Levine. *The Theory of Learning in Games*. MIT Press, 1998.
- [12] Martin Gairing. *Covering Games: Approximation through Non-cooperation*, pages 184–195. Springer, 2009.
- [13] Ragavendran Gopalakrishnan, Jason R. Marden, and Adam Wierman. Potential games are necessary to ensure pure nash equilibria in cost sharing games. *Mathematics of Operations Research*, 39(4):1252–1296, 2014.
- [14] K Ishiyama, M Sendoh, and KI Arai. Magnetic micromachines for medical applications. *Journal of Magnetism and Magnetic Materials*, 242:41–46, 2002.
- [15] Hiroaki Kitano, Satoshi Tadokoro, Itsuki Noda, Hitoshi Matsubara, Tomoichi Takahashi, Atsushi Shinjou, and Susumu Shimada. Robocup rescue: Search and rescue in large-scale disasters as a domain for autonomous agents research. In *Systems, Man, and Cybernetics, 1999. IEEE SMC'99 Conference Proceedings. 1999 IEEE International Conference on*, volume 6, pages 739–743. IEEE, 1999.
- [16] Amos Korman and Yoav Rodeh. Multi-round cooperative search games with multiple players. In *ICALP 2019-46th International Colloquium on Automata, Languages and Programming*, 2019.
- [17] Elias Koutsoupias and Christos Papadimitriou. *Worst-Case Equilibria*, pages 404–413. Springer, 1999.
- [18] Toyoki Kozai, Genhua Niu, and Michiko Takagaki. *Plant factory: an indoor vertical farming system for efficient quality food production*. Academic Press, 2015.
- [19] A. Krause and C. Guestrin. Near-optimal observation selection using submodular functions. In *Proceedings of the 22nd Conference on Artificial Intelligence*, 2007.

- [20] Helge-Björn Kuntze, Christian W Frey, Igor Tchouchenkov, Barbara Staehle, Erich Rome, Kai Pfeiffer, Andreas Wenzel, and Jürgen Wöllenstein. Seneka-sensor network with mobile robots for disaster management. In *Homeland Security (HST), 2012 IEEE Conference on Technologies for*, pages 406–410. IEEE, 2012.
- [21] J. R. Marden, G. Arslan, and J. S. Shamma. Connections between cooperative control and potential games. *IEEE Transactions on Systems, Man and Cybernetics. Part B: Cybernetics*, 39:1393–1407, December 2009.
- [22] J. R. Marden, G. Arslan, and J. S. Shamma. Joint strategy fictitious play with inertia for potential games. *IEEE Transactions on Automatic Control*, 54:208–220, February 2009.
- [23] J. R. Marden and M. Effros. The price of selfishness in network coding. *IEEE Transactions on Information Theory*, 58(4):2349–2361, April 2012.
- [24] J. R. Marden and M. Philips. Optimizing the price of anarchy in concave cost sharing games. In *in Proceedings of the 2017 American Control Conference*, pages 5237–5242, May 2017.
- [25] J. R. Marden and J. S. Shamma. Revisiting log-linear learning: Asynchrony, completeness and a payoff-based implementation. *Games and Economic Behavior*, 75:788–808, July 2012.
- [26] Jason R. Marden and Adam Wierman. Distributed welfare games. *Operations Research*, 61(1):155–168, 2013.
- [27] J.R. Marden and J.S. Shamma. Game theory and distributed control. In *Handbook of Game Theory Vol. 4*, pages 861–899. Elsevier, 2014.
- [28] Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124 – 143, 1996.
- [29] John F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950.
- [30] G. Nemhauser, L. Wolsey, and M. Fisher. An analysis of approximations for maximizing submodular set functions – I. *Mathematical Programming*, 14(1):265–294, 1978.
- [31] D. Paccagnan and J. R. Marden. Utility design for distributed resource allocation—Part II: Applications to submodular, covering, and supermodular problems. *arXiv preprint arXiv:1807.01343*, 2018.
- [32] Dario Paccagnan, Rahul Chandan, Bryce L Ferguson, and Jason R Marden. Incentivizing efficient use of shared infrastructure: Optimal tolls in congestion games. *arXiv preprint arXiv:1911.09806*, 2019.
- [33] Dario Paccagnan, Rahul Chandan, and Jason R Marden. Utility design for distributed resource allocation—Part I: Characterizing and optimizing the exact price of anarchy. *IEEE Transactions on Automatic Control (Early Access)*, 2019.
- [34] Vinod Ramaswamy, Dario Paccagnan, and Jason R Marden. Multiagent maximum coverage problems: The trade-off between anarchy and stability. In *2019 18th European Control Conference (ECC)*, pages 1043–1048. IEEE, 2019.
- [35] Ania Servant, Famin Qiu, Mariarosa Mazza, Kostas Kostarelos, and Bradley J Nelson. Controlled in vivo swimming of a swarm of bacteria-like microbotic flagella. *Advanced Materials*, 27(19):2981–2988, 2015.
- [36] Kevin Spieser, Kyle Treleaven, Rick Zhang, Emilio Frazzoli, Daniel Morton, and Marco 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*, pages 229–245. Springer, 2014.
- [37] Philipp von Falkenhausen and Tobias Harks. Optimal cost sharing for resource selection games. *Mathematics of Operations Research*, 38(1):184–208, 2013.
- [38] Philipp von Falkenhausen and Tobias Harks. Optimal cost sharing for capacitated facility location games. *European Journal of Operations Research*, 239(1):187–198, 2014.
- [39] H. P. Young. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton University Press, 1998.
- [40] H. P. Young. *Strategic Learning and its Limits*. Oxford University Press, 2005.



game theory, multi-agent systems and communication networks.

Vinod Ramaswamy received his B.Tech in electronics and communication engineering from Kerala University, India, in 2004 and his M.Tech. degree in electrical engineering, with specialization in communication systems, from Indian Institute of Technology, Madras, India, in 2006. He received his Ph.D. degree in Electrical Engineering from Texas A&M University, College Station in 2013. From 2015 to 2016, he was a Postdoctoral Scholar at University of Colorado, Boulder. He is currently working in Qualcomm. His research interests include



Dario Paccagnan received the Bs.C. (Hons.) and M.Sc. (Hons.) degrees in aerospace engineering from the University of Padova, Padova, Italy, in 2011 and 2014, respectively. He also received the M.Sc. (Hons.) degree in mathematical modeling and computation from the Technical University of Denmark, Kongens Lyngby, Denmark, in 2014, and the PhD degree in optimization and control from ETH Zurich, Zurich, Switzerland, in 2018. He is an Assistant Professor (U.K. Lecturer) with the Department of Computing, Imperial College London, London, U.K., since Fall 2020. Before that, he was a Postdoctoral Fellow with the Center for Control, Dynamical Systems and Computation, University of California, Santa Barbara, Santa Barbara, CA, USA. His research interests are at the interface between game theory and control theory, with a focus on the design of behavior-influencing mechanisms for socio-technical systems. Dr. Paccagnan was recognized with the ETH Medal for his doctoral work, the SNSF Early Postdoc Mobility Fellowship, and the SNSF Doc Mobility Fellowship. He was also a finalist for the 2019 EECI Best PhD. Thesis Award.



Jason R. Marden is a Full Professor in the Department of Electrical and Computer Engineering at the University of California, Santa Barbara. Jason received a BS in Mechanical Engineering in 2001 from UCLA, and a PhD in Mechanical Engineering in 2007, also from UCLA, where he was awarded the Outstanding Graduating PhD Student in Mechanical Engineering. After graduating from UCLA, he served as a junior fellow in the Social and Information Sciences Laboratory at the California Institute of Technology until 2010 when he joined the University of Colorado. In 2015, Jason joined the University of California, Santa Barbara. Jason is a recipient of the NSF Career Award (2014), the ONR Young Investigator Award (2015), the AFOSR Young Investigator Award (2012), the American Automatic Control Council Donald P. Eckman Award (2012), and the SIAG/CST Best SICON Paper Prize (2015). Jason's research interests focus on game theoretic methods for the control of distributed multi-agent systems.