

Balancing Asymptotic and Transient Efficiency Guarantees in Set Covering Games

Rohit Konda, Rahul Chandan, David Grimsman and Jason R. Marden

Abstract—Game theoretic approaches have gained traction as robust methodologies for designing distributed local algorithms that induce a desired overall system configuration in multi-agent settings. However, much of the emphasis in these approaches is on providing asymptotic guarantees on the performance of a network of agents, and there is a gap in the study of efficiency guarantees along transients of these distributed algorithms. Therefore, in this paper, we study the transient efficiency guarantees of a natural game-theoretic algorithm in the class of set covering games, which have been used to model a variety of applications. Our main results characterize the optimal utility design that maximizes the guaranteed efficiency along the transient of the natural dynamics. Furthermore, we characterize the Pareto-optimal frontier with regards to guaranteed efficiency in the transient and the asymptote under a class of game-theoretic designs. Surprisingly, we show that there exists an extreme trade-off between the long-term and short-term guarantees in that an asymptotically optimal game-theoretic design can perform arbitrarily bad in the transient.

I. INTRODUCTION

Recently, game theory has emerged as a significant framework in the study of multi-agent systems: the agents in the system are modeled as agents in a game, each with its own actions and individual rewards that are functions of the joint action of the other agents. Some prominent examples where these models have found success include wireless communication networks [1], UAV swarm task allocation [2], news subscription services [3], vaccinations during an epidemic [4], security systems [5], facility location [6], coordinating the charging of electric vehicles [7], and national defense [8], among others. Indeed, game theory has played a significant role in many applications and across various disciplines.

A central focus in using game-theoretic models to control multi-agent systems is the design of local objective (or reward) functions for the agents that result in a desirable emergent system-level outcome, as measured by a system welfare function [9]. Generally the emergent outcome is assumed to be a coordinated decision which is a Nash equilibrium of the game, wherein each agent cannot improve its own objective function by switching actions. A significant research effort has been devoted to designing local objective functions that ensure good performance of Nash equilibria. For example, [10] shows that for the class of set covering games, there exists a set of local objective functions which

guarantee that the performance of any Nash equilibrium is within a multiplicative factor of $1 - \frac{1}{e}$ of the optimal decision set. Furthermore, no other polynomial time algorithm can perform better [11].

Rather than focusing solely on the system welfare at the Nash equilibria, this work studies the welfare of decision sets along the transient path toward equilibrium. This may be important when the relevant system parameters evolve on faster time scales, agent modalities vary with time etc. The class of *potential games* [12] is particularly relevant to this goal as the class of *best response/reply processes* and its many variants (see, e.g., [13]) are known to converge to a Nash equilibrium asymptotically from any initial configuration. Best reply processes are intimately tied with the development of game theory, present in Cournot processes, Nash's seminal paper [14], as well as stable and evolutionary equilibrium [15]. Efficiency guarantees along these best response processes are much less studied than the asymptotic guarantees that come with Nash equilibrium. In this work, we restrict our attention to an important subclass of potential games known as *set covering games*, in which agents with limited capabilities try to coordinate to cover as large of an area as possible. For set covering games, we characterize bounds on the efficiency for best response processes, and compare utility designs that maximize either the short- or long-efficiency.

While few works study the guarantees that can arise through best response processes, we highlight a subset of important works that deal with analyzing the emergent behavior of best response dynamics. The authors in [16], [17] analyze lower bounds for the efficiency guarantees resulting after a single best response round in linear congestion and cut games. These guarantees were verified through a linear programming approach in [18] and extended to quadratic and cubic congestion games. Best response dynamics are also studied in [19], where it was found that the efficiency resulting from best response dynamics can be arbitrarily bad for non-potential games. The speed of convergence to efficient states is studied in [20]. While similar in spirit, our paper extends the previous literature by moving from characterization results of prespecified utility designs to instead identifying optimal game-theoretic designs that maximize the efficiency guarantees along a best response process.

The contributions of this paper are all with respect to the class of set covering games and are as follows:

- 1) We identify the local objective that optimizes the transient efficiency guarantee along a best response dynamic;
- 2) We identify a trade-off between transient and asymp-

R. Konda (rkonda@ucsb.edu), R. Chandan, and J. R. Marden are with the Department of Electrical and Computer Engineering at the University of California Santa Barbara and D. Grimsman is with the Department of Computer Science at Brigham Young University. This work is supported by ONR Grant #N00014-20-1-2359 and AFOSR Grant #FA9550-20-1-0054.

otic guarantees, and provide explicit characterizations of the Pareto frontier and Pareto-optimal agent objective functions.

The rest of the paper is organized as follows: In Section II, we introduce the necessary notation and definitions for best response dynamics, as well as for set covering games. We then outline our main results in Section III. We present a simulation study in Section IV to examine how our findings translate to practice. Finally, we conclude in Section V.

II. MODEL

Notation. We use \mathbb{R} and \mathbb{N} to denote the sets of real and natural numbers respectively. Given a set S , $|S|$ represents its cardinality. $\mathbf{1}_S$ describes the indicator function for set S ($\mathbf{1}_S(e) = 1$ if $e \in S$, 0 otherwise).

A. Set Covering Games and Best Response Dynamics

We present all of our results in the context of *set covering games* [10], a well-studied class of potential games. We first recall some definitions to outline our game-theoretic model, which is used to describe the given class of distributed scenarios. Set covering games are characterized by a finite set of resources \mathcal{R} that can be possibly covered by a set of n agents $\mathcal{I} = \{1, \dots, n\}$. Each resource r has an associated value $v_r \geq 0$ determining its relative importance to the system. Each agent i has a finite action set $\mathcal{A}_i = \{a_i^1, \dots, a_i^c\} \subset 2^{\mathcal{R}}$, representing the sets of resources that each agent can opt to select. The performance of a joint action $a = (a_1, \dots, a_n) \in \mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$ comprised of every agent's actions is evaluated by a system-level objective function $W : \mathcal{A} \rightarrow \mathbb{R}_{\geq 0}$, defined as the total value covered by the agents, i.e.,

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

The goal is to coordinate the agents to a joint action a^{opt} that maximizes this welfare $W(a)$. Thus, each agent $i \in \mathcal{I}$ is endowed with its own utility function $U_i : \mathcal{A} \rightarrow \mathbb{R}$ that dictates its preferences among the joint actions. We use $a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$ to denote the joint action without the action of agent i . A given set covering game G can be summarized as the tuple $G = (\mathcal{I}, \mathcal{A}, W, \{U_i\}_{i \in \mathcal{I}})$. We consider the utility functions of the form

$$U_i(a_i, a_{-i}) = \sum_{r \in a_i} v_r \cdot f(|a|_r), \quad (2)$$

where $|a|_r$ is the number of agents that choose resource r in action profile a , and the *utility rule* $f : \{1, \dots, n\} \rightarrow \mathbb{R}_{\geq 0}$ defines the resource independent agent utility determined by $|a|_r$. This utility rule will be our design choice for defining the agents preferences towards each of their decisions. To motivate our results, we provide an example application of set covering games below, which represents just one among many [21]–[23].

Example 1 (Wireless Sensor Coverage [24]). Consider a group of sensors s_1, s_2, \dots, s_n that can sense a particular

region depending on the orientation, physical placement of the sensor, etc. The choice of these parameters constitute the different possible actions \mathcal{A}_i for each sensor. We partition the entire space into disjoint regions $r \in \mathcal{R}$ and each region r is attributed a weight $v_r \geq 0$ to signify the importance of covering it. As a whole, the set of sensors wish to maximize the total weight of the regions covered. Using a game-theoretic model of this scenario allows us to abstract away much of the complexity that arises from having to consider specific sensor types and their respective capabilities, in contrast to [24]. In this scenario, the results in the paper allow us to study guarantees of performance of local coordination algorithms that the sensors may employ. We further note that short-term performance guarantees may also be desirable especially when there is some drift in the parameters of the game (e.g., in v_r or \mathcal{A}_i).

Next, we introduce some notation relevant to the study of best reply processes. For a given joint action $\bar{a} \in \mathcal{A}$, we say the action \hat{a}_i is a *best response* for agent i if

$$\hat{a}_i \in \text{BR}_i(\bar{a}_{-i}) = \arg \max_{a_i \in \mathcal{A}_i} U_i(a_i, \bar{a}_{-i}). \quad (3)$$

Note that $\text{BR}_i(\bar{a}_{-i})$ may be set valued. In this fashion, the underlying discrete-time dynamics of the best response process is

$$\begin{aligned} a_i[m+1] &\in \text{BR}_i(a_{-i}[m]) \text{ for some agent } i \\ a_{-i}[m+1] &= a_{-i}[m], \end{aligned} \quad (4)$$

where at each step, a single agent acts in a best response to the current joint action profile. It is important to stress that the fixed points of this best response process are Nash equilibria of the game, i.e., a joint action $a^{\text{ne}} \in \mathcal{A}$ such that

$$a_i^{\text{ne}} \in \text{BR}_i(a_{-i}^{\text{ne}}) \text{ for all agents } i \in \mathcal{I}. \quad (5)$$

In the case of potential games [12], the set of Nash equilibrium NE is also globally attractive with respect to the best response dynamics. A potential game is defined as any game for which there exists a potential function ϕ where for all agents i and pairs of actions a, a' ,

$$\phi(a_i, a_{-i}) - \phi(a'_i, a_{-i}) = U_i(a_i, a_{-i}) - U_i(a'_i, a_{-i}). \quad (6)$$

In these games, the potential function decreases monotonically along the best response process in a similar manner to a Lyapunov function. For non-potential games, there is a possibility that the best response process ends up in a cycle [25], but this falls outside the scope of this paper.

In this work we assume that the best response process is such that the agents are ordered from 1 to n and perform successive best responses in that order. While other best response processes have been previously studied (see, e.g., k covering walks [16]), we limit our analysis to this specific process for tractability. We also assume that each agent has a null action a_i^{\emptyset} available to them representing their non-participation in the game G .

Definition 1. Given a game G with n agents, a *k-best response round* is a $n \cdot k$ step best reply process with an

initial state $a[0] = a^\emptyset$ and the corresponding dynamics

$$\begin{aligned} a_i[m+1] &\in \text{BR}_i(a_{-i}[m]) \text{ for agent } i = (m+1 \bmod n) \\ a_{-i}[m+1] &= a_{-i}[m]. \end{aligned} \quad (7)$$

This process results in a set of possible action trajectories $\Sigma^k \ni \sigma^k = (a[0], a[1], \dots, a[kn-1], a[kn])$ with a set of possible end action states $\text{end}(k) = \{\sigma^k(kn) : \sigma^k \in \Sigma^k\}$ resulting after running for $k \cdot n$ time steps.

B. Price of Best Response and Price of Anarchy

The primary focus of this paper is to analyze guarantees on the efficiency of possible end actions $\text{end}(k)$ after a k -best response round with respect to the system welfare W . In this regard, we introduce the *price of best response* metric to evaluate the worst case ratio between the welfare of the outcome of the k -best response round and the optimal welfare in a game.

$$\text{PoB}(G, k) = \frac{\min_{a \in \text{end}(k)} W(a)}{\max_{a \in \mathcal{A}} W(a)}. \quad (8)$$

Our study of the price of best response will be in direct comparison with the standard game-theoretic metric of *price of anarchy*, which is defined as

$$\text{PoA}(G) = \frac{\min_{a \in \text{NE}} W(a)}{\max_{a \in \mathcal{A}} W(a)}, \quad (9)$$

where NE is the set of Nash equilibrium. The price of anarchy is a well established metric, with a number of results on its characterization, complexity, and design [26]. We frame price of anarchy as a backdrop to many of our results, where Nash equilibria are the set of end states that are possible, following a k -best response round in the limit as $k \rightarrow \infty$. Under this notation, we can analyze the transient guarantees, which have been relatively unexplored, in these distributed frameworks in comparison to the asymptotic setting.

To isolate the efficiency analysis over the utility rule f and the number of agents n , we consider the price of best response with respect to the class of all set covering games $\mathcal{G}_{f,n}$ with a fixed utility rule f and n number of agents, and any resource set \mathcal{R} and corresponding values v_r and set of agent actions \mathcal{A} :

$$\text{PoB}(f, n, k) = \inf_{G \in \mathcal{G}_{f,n}} \text{PoB}(G, k). \quad (10)$$

Similarly, we define the price of anarchy over $\mathcal{G}_{f,n}$ as

$$\text{PoA}(f, n) = \inf_{G \in \mathcal{G}_{f,n}} \text{PoA}(G). \quad (11)$$

When appropriate, we may also remove the dependence on the number of agents by denoting $\text{PoB}(f, k) = \inf_{n \in \mathbb{N}} \text{PoB}(f, n, k)$ (and likewise for PoA) to take the worst case guarantees across any number of agents.

We are interested in the transient and asymptotic system guarantees that can result from a given utility rule. Therefore we bring to attention two well-motivated utility rules to highlight the spectrum of guarantees that are possible. The first utility rule $f_{\text{PoA}} \in \arg \max_f \text{PoA}(f)$, introduced in

[10], attains the optimal price of anarchy guarantee for the class of set covering games and is defined as

$$f_{\text{PoA}}(j) = \sum_{\ell=j}^{\infty} \frac{(j-1)!}{\ell!(e-1)}. \quad (12)$$

The second rule is the *marginal contribution* (MC) rule. The MC rule is a well studied [27] utility rule and is defined as

$$f_{\text{MC}}(j) = \begin{cases} 1, & \text{for } j = 1 \\ 0, & \text{otherwise} \end{cases}. \quad (13)$$

Note that the induced utility under this mechanism is $U_i(a_i, a_{-i}) = W(a_i, a_{-i}) - W(a_i^\emptyset, a_{-i})$, where the utility for agent i is the added welfare from playing its null action a_i^\emptyset to playing a_i . Using f_{MC} ensures that all Nash equilibria of the game are local optima of the welfare function. However, it is known that the MC rule has sub-optimal asymptotic guarantees as $\text{PoA}(f_{\text{MC}}) = \frac{1}{2} < 1 - \frac{1}{e} = \text{PoA}(f_{\text{PoA}})$ [28]. Thus, under the standard tools of asymptotic analysis, f_{PoA} has been considered to be ‘optimal’ in inducing desirable system behavior. Utilizing the notion of PoB, we inquire if the favorable characteristics of f_{PoA} translate to good short-term behavior of the system. Intuitively, since f_{PoA} guarantees optimal long-term performance, we expect the same in the short-term as well. This discussion motivates the following three questions:

- For a given $k \geq 1$, is $\text{PoB}(f_{\text{PoA}}, n, k)$ always greater than $\text{PoB}(f_{\text{MC}}, n, k)$?
- If not, then for how many rounds does the best response algorithm under f_{PoA} have to be run for the guarantees on the end state performance to surpass the asymptotic performance of marginal contribution? In other words, is there $k \geq 1$ for which $\text{PoB}(f_{\text{PoA}}, n, k) \geq \text{PoA}(f_{\text{MC}}, n) = \frac{1}{2}$?
- For a given $k \geq 1$, what is the utility mechanism that maximizes the end state performance after running the best response algorithm for k rounds, i.e. what is $\arg \max_f \text{PoB}(f, n, k)$?

In the next section, we answer these questions in detail.

III. MAIN RESULTS

In this section, we outline our main results for price of best response guarantees. In this setting, utility functions are designed by carefully selecting the utility rule f and improvement in transient system level behavior corresponds with a higher price of best response. Hence, we characterize the optimal f that achieves the highest possible PoB. We formally state the optimality of the MC rule for the class of set covering games below.

Theorem 1. *Consider the class of set covering games with $n \geq 2$ users. The following statements are true:*

- For any number of rounds $k \geq 1$, the transient performance guarantees associated with the utility mechanisms f_{MC} and f_{PoA} satisfy

$$\text{PoB}(f_{\text{MC}}, n, k) = 1/2, \quad (14)$$

$$\text{PoB}(f_{\text{PoA}}, n, k) \leq 1/2. \quad (15)$$

Furthermore, the marginal contribution rule achieves its asymptotic guarantees after a single round, i.e.,

$$\text{PoB}(f_{MC}, n, k) = \text{PoB}(f_{MC}, n, 1). \quad (16)$$

(b) For any number of rounds $k \geq 1$, the marginal contribution rule f_{MC} optimizes the PoB, i.e.,

$$\max_f \text{PoB}(f, n, k) = \text{PoB}(f_{MC}, n, k) = 1/2. \quad (17)$$

Alternatively, there is no utility rule f with $\text{PoB}(f, n, k) > 1/2$ for any k .

Proof. To prove the statements outlined in items (a) and (b) of the theorem statement, we show, more generally, that for any f and any k , the following series of inequalities is true:

$$\frac{1}{2} = \text{PoB}(f_{MC}, n, 1) = \text{PoB}(f_{MC}, n, k) \geq \text{PoB}(f, n, k).$$

The fact that $\text{PoB}(f_{MC}, n, 1) = 1/2$ comes from classical results in greedy algorithms, such as in [27]. For proving $\text{PoB}(f_{MC}, n, k) = \frac{1}{2}$, we first show that

$$\text{PoB}(f_{MC}, n, k) \geq \text{PoB}(f_{MC}, n, 1). \quad (18)$$

Let $a, a' \in \mathcal{A}$ be joint actions such that a' is a best response to a for agent i . If the MC rule is used as in Equation (13), then

$$\begin{aligned} W(a') - W(a) &= W(a'_i, a_{-i}) - W(a_i^\emptyset, a_{-i}) \\ &\quad + W(a_i^\emptyset, a_{-i}) - W(a_i, a_{-i}) \\ &= U_i(a'_i, a_{-i}) - U_i(a_i, a_{-i}) \\ &\geq 0. \end{aligned}$$

Then, after any $m \geq 1$ best response rounds, $W(\text{end}(m)) \geq W(\text{end}(1))$, and therefore the claim is shown. For the other direction as well as the inequality $\frac{1}{2} \geq \text{PoB}(f, n, k)$, we construct a game construction (also shown in Figure 1) such that for any utility rule f , $\text{PoB}(G^f, k) \leq \frac{1}{2}$. The game G^f is constructed as follows with 2 agents. Let the resource set be $\mathcal{R} = \{r_1, r_2, r_3, r_4\}$ with $v_{r_1} = 1$, $v_{r_2} = 1$, $v_{r_3} = f(2)$, and $v_{r_4} = 0$. For agent 1, let $\mathcal{A}_1 = \{a_1^\emptyset, a_1^1, a_1^2\}$ with $a_1^1 = \{r_1\}$ and $a_1^2 = \{r_2\}$. For agent 2, let $\mathcal{A}_2 = \{a_2^\emptyset, a_2^1, a_2^2\}$ with $a_2^1 = \{r_3\}$ and $a_2^2 = \{r_1\}$. It can be seen in this game, if $f(2) \leq 1$, the optimal joint action is (a_1^2, a_2^2) resulting in a welfare of 2, and if $f(2) > 1$, the optimal joint action is (a_1^2, a_2^1) resulting in a welfare of $1 + f(2) \geq 2$. Additionally, starting from $(a_1^\emptyset, a_2^\emptyset)$, note the best response path $((a_1^\emptyset, a_2^\emptyset), (a_1^1, a_2^\emptyset)(a_1^1, a_2^1), \dots, (a_1^1, a_2^1), (a_1^1, a_2^1), (a_1^1, a_2^2))$ exists, ending in the action (a_1^1, a_2^2) with a welfare 1. Since (a_1^1, a_2^2) is a Nash equilibrium, both agents stay playing their first action a^1 for $k - 1$ rounds and agent 2 only switches to a_2^2 in the k -round. Therefore $\text{PoB}(G^f, k) \leq \frac{1}{2}$ for this game. This game construction G^f can be extended to $n > 2$ by fixing the actions of agents $i \geq 2$ to be $\mathcal{A}_i = \{a_i^\emptyset, a_i^1\}$, where the only nonempty action available to them is $a_i^1 = \{r_4\}$ that selects the 0 value resource.

Since G^f may not be the worst case game, $\text{PoB}(f, n, k) \leq \text{PoB}(G^f, k) \leq \frac{1}{2}$ and the final claim is shown. \square




Fig. 1. Under this n -agent game construction G^f , outlined in Theorem 1's proof, $\text{PoB}(G^f, k) \leq \frac{1}{2}$ for any mechanism f . There are four resources $\mathcal{R} = \{r_1, r_2, r_3, r_4\}$ with values $\{1, 1, f(2), 0\}$ respectively. Under the allocation (a_1^2, a_2^2) , the resources r_1 , r_2 , and r_4 are selected to produce a welfare of 2. However for any number of rounds $k \geq 1$, there exists a best response path that ends in an allocation that selects only the resources r_1 and r_4 to produce a welfare of 1. Then $\text{PoB}(f, n, k) \leq \text{PoB}(G^f, k) \leq \frac{1}{2}$.

Observe that Theorem 1 answers Questions (a)–(c) posed in the previous section. Specifically, we have shown that for any number of rounds, the performance guarantees of f_{PoA} never overcome the performance guarantees of f_{MC} . Remarkably, for any fixed number of rounds k , $\text{PoB}(f_{MC}, k)$ remains greater than $\text{PoB}(f_{\text{PoA}}, k)$ and the advantages of using f_{PoA} only appear when the system settles to a Nash equilibrium as $k \rightarrow \infty$. Furthermore, the marginal contribution rule is in fact the utility mechanism that optimizes the price of best response and that this optimal efficiency guarantee is achieved after 1 round of best response. Thus, the transient guarantees of f_{MC} are quite strong in comparison to f_{PoA} . Note that this is in stark contrast with their price of anarchy values, where f_{PoA} significantly outperforms f_{MC} [10], [28]. Clearly, there exists a trade-off between short-term and long-term performance guarantees in this setting.

In our next theorem, we characterize the trade-off between short-term and long-term performance guarantees by providing explicit expressions for the Pareto curve between $\text{PoB}(f, 1)$ and PoA and the corresponding Pareto-optimal utility mechanisms. Surprisingly, we show that $\text{PoB}(f_{\text{PoA}}, 1)$ can be arbitrarily bad. A full proof can be found at [29].

Theorem 2. For a fixed $\text{PoA}(f) = C \in [\frac{1}{2}, 1 - \frac{1}{e}]$, the optimal $\text{PoB}(f, 1)$ is

$$\max_f \text{PoB}(f, 1) = \left[\sum_{j=0}^{\infty} \max\{j!(1 - \frac{1-C}{C} \sum_{\tau=1}^j \frac{1}{\tau!}), 0\} + 1 \right]^{-1}. \quad (19)$$

Furthermore, if $\text{PoA}(f) = 1 - \frac{1}{e}$ for a given f , then $\text{PoB}(f, 1)$ must be equal to 0.

Proof Sketch. Given a utility rule f and n number of agents, it is possible to give a closed form expression for both the price of anarchy and price of best response as the following.

$$\begin{aligned} \frac{1}{\text{PoB}(f, n, 1)} &= \left(\sum_{j=1}^n f(j) - \min_j f(j) \right) / f(1) + 1, \\ \frac{1}{\text{PoA}(f, n)} &= 1 + \left(\max_{1 \leq j \leq n-1} \{ (j+1)f(j+1) - f(1), \right. \\ &\quad \left. jf(j) - f(j+1), jf(j+1) \} \right) / f(1). \end{aligned}$$

First, we can show that for a utility rule f to be Pareto-optimal, it must be non-increasing. Under this assumption, we can derive a recursive equation for the Pareto-optimal utility rule f^C with a price of anarchy $\text{PoA}(f^C) = C \in [\frac{1}{2}, 1 - \frac{1}{e}]$ as the following.

$$\begin{aligned} f^C(1) &= 1 \\ f^C(j+1) &= \max\{j f^C(j) - C, 0\}. \end{aligned} \quad (20)$$

Using techniques from time-varying discrete systems, we can rewrite the recursive form of f^C into a closed-form expression and substitute it into the PoB expression to give Eq. (19). \square

The price of anarchy is only taken on the interval $[\frac{1}{2}, 1 - \frac{1}{e}]$, since the greatest achievable $\text{PoA}(f)$ is $1 - \frac{1}{e}$ [10] and the price of anarchy of f_{MC} is $\frac{1}{2}$ (which attains the optimal $\text{PoB}(f, 1)$ as shown in Theorem 1). The corresponding Pareto-optimal utility rules are given in Eq. (20). The optimal trade-off curve between PoA and PoB is plotted in Figure 2.




Fig. 2. All possible pairs $(\text{PoA}(f), \text{PoB}(f, 1)) \in [0, 1] \times [0, 1]$ are depicted in this figure. The attainable set where there exists a mechanism f that can achieve either a greater $\text{PoB}(f, 1)$ or $\text{PoA}(f)$ is depicted in green and the complement is depicted in red. The Pareto curve that results from the optimal trade-off described in Theorem 2 is depicted by the black curve. We observe that there is an extreme drop in transient performance to $\text{PoB}(f, 1) = 0$ when using any utility rule that optimizes $\text{PoA}(f)$.

Theorem 2 highlights the disparity between the efficiency guarantees of the asymptotic and transient values, quantified by the pair of achievable PoA and PoB. While there is a decline in PoA guarantees when using the MC rule, the deterioration in PoB that results from using the mechanism f_{PoA} that optimizes PoA is much more extreme. In fact, the transient performance of f_{PoA} can be arbitrarily bad. This is quite surprising, it appears that a utility rule's asymptotic performance guarantees do not reflect on its transient guarantees. This extreme trade-off has significant consequences on distributed design, especially when the short-term performance of the multi-agent system is critical, and prompts a more careful interpretation of the asymptotic results in such scenarios.

IV. SIMULATIONS

We analyze our theoretical results empirically through a Monte Carlo simulation, presented in Figure 4. We simulate 200 instances where a random game G is generated and the welfare along a 4-round best response path is recorded. Each game G has 10 agents and two sets of 10 resources, where the values v_r are drawn uniformly from the continuous interval $[0, 1]$. Each agent in the game has two singleton actions, where in one action the agent selects a resource in the first set and in the other action the agent selects a resource in the second set, as seen in Figure 3. The two resources in each agent's actions are selected uniformly, at random. These parameters were chosen to induce a natural set of game circumstances that have a rich best response behavior. Additionally, the game is either endowed with the utility rule f_{MC} that optimizes $\text{PoB}(f, 1)$ or f_{PoA} that optimizes $\text{PoA}(f)$. Then the best response round algorithm is performed, where each agent successively performs a best response. The ratio of the welfare achieved at each time step following the two utility rules f_{MC} , f_{PoA} is displayed in Figure 4.

	resources	r_1^1	r_2^1	r_3^1	r_4^1	r_5^1	r_6^1	r_7^1	r_8^1	r_9^1	r_{10}^1	r_1^2	r_2^2	r_3^2	r_4^2	r_5^2	r_6^2	r_7^2	r_8^2	r_9^2	r_{10}^2
Players	values	.3	.2	.5	0	.5	.3	.4	.8	1	.7	.5	.3	.3	1	.5	.3	.3	.6	.2	.5
1		a_1^1																			a_1^2
2			a_2^1			a_2^1						a_2^2									
3			a_3^1												a_3^2						
4								a_4^1				a_4^2									
5							a_5^1												a_5^2		
6					a_6^1							a_6^2									
7					a_7^1										a_7^2						
8					a_8^1											a_8^2					
9			a_9^1																	a_9^2	
10					a_{10}^1														a_{10}^2		

Fig. 3. An example of a game that was generated within the Monte Carlo simulation described in Section IV. The highlighted boxes represent the possible single resource selections that each agent can select. The two sets of resources are divided by a bold line. For each game generated, the welfare ratio between the best response sequences corresponding with the utility rules f_{MC} and f_{PoA} was recorded at each time step.

Several significant observations can be made from Figure 4. Observe that after one round (after 30 best response steps), the performance of f_{MC} exceeds the performance f_{PoA} on average. This is indicative of the trade-off shown in Theorem 2. Furthermore, running the best-response algorithm for additional steps recovers the performance difference between f_{MC} and f_{PoA} , with f_{PoA} having a slightly better end performance on average. These results support our claim that optimizing for asymptotic results may induce sub-optimal system behavior in the short term, and those results must be carefully interpreted for application in multi-agent scenarios. By analyzing the performance directly along these dynamics, we see that a richer story emerges in both the theoretical and the empirical results.

V. CONCLUSION

There are many works that focus on efficiency guarantees for the asymptotic behavior of multi-agent systems under the game-theoretic notion of Nash equilibrium. However, much less is understood about the efficiency guarantees in the

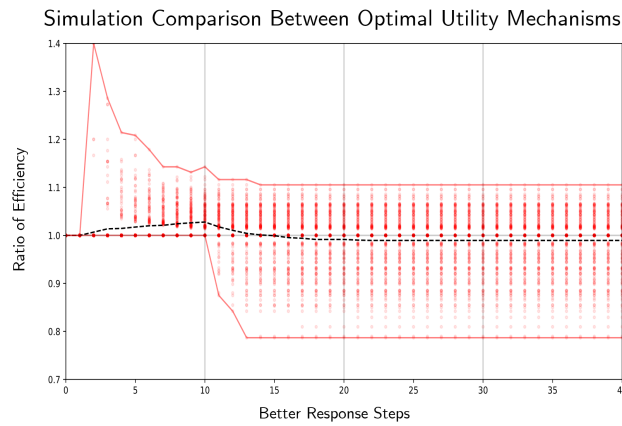


Fig. 4. A Monte Carlo simulation with 10 agents and 2 set of resources was conducted, where the utility rules f_{MC} and f_{POA} were used in separate runs. All the agents are randomly assigned a single resource from each set in their 2 actions. The ratio $W_{MC}(a)/W_{POA}(a)$ (from using either of the utility rules) is plotted at each time step along the two better response paths. The average ratio across all runs is displayed in black and the minimum and maximum ratios are outlined with a red line. While on average, the performance from using f_{MC} is greater than f_{POA} for 1 round, this performance difference is recuperated after running the best response algorithm for more steps.

transient behavior of such systems. In this paper, we examine a natural class of best response dynamics and characterize the efficiency along the resulting transient states. We characterize a utility design that performs optimally in the transient for the class of set covering games. Furthermore, we complement these results by providing an explicit characterization of the efficiency trade-off between 1-round best response and Nash equilibrium that can result from a possible utility design. From these results, we observe that there is a significant misalignment between the optimization of transient and asymptotic behavior for these systems. Interesting future directions include extending our results to a more general class of games and considering other classes of learning dynamics from the literature.

REFERENCES

- [1] Z. Han, D. Niyato, W. Saad, T. Başar, and A. Hjørungnes, *Game theory in wireless and communication networks: theory, models, and applications*. Cambridge university press, 2012.
- [2] J. J. Roldán, J. Del Cerro, and A. Barrientos, “Should we compete or should we cooperate? applying game theory to task allocation in drone swarms,” in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5366–5371, IEEE, 2018.
- [3] C.-C. Hsu, A. Ajorlou, M. Yildiz, and A. Jadbabaie, “Information disclosure and network formation in news subscription services,” in *2020 59th IEEE Conference on Decision and Control (CDC)*, pp. 5586–5591, IEEE, 2020.
- [4] A. R. Hota and S. Sundaram, “Game-theoretic vaccination against networked sis epidemics and impacts of human decision-making,” *IEEE Transactions on Control of Network Systems*, vol. 6, no. 4, pp. 1461–1472, 2019.
- [5] K. Paarporn, R. Chandan, M. Alizadeh, and J. R. Marden, “Characterizing the interplay between information and strength in blotto games,” in *2019 IEEE 58th Conference on Decision and Control (CDC)*, pp. 5977–5982, 2019.
- [6] A. Vetta, “Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions,” in *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, pp. 416–425, 2002.

- [7] J. Martinez-Piauelo, N. Quijano, and C. Ocampo-Martinez, “Decentralized charging coordination of electric vehicles using multi-population games,” in *2020 59th IEEE Conference on Decision and Control (CDC)*, pp. 506–511, IEEE, 2020.
- [8] E. S. Lee, D. Shishika, and V. Kumar, “Perimeter-defense game between aerial defender and ground intruder,” in *2020 59th IEEE Conference on Decision and Control (CDC)*, pp. 1530–1536, IEEE, 2020.
- [9] J. R. Marden and J. S. Shamma, “Game theory and distributed control,” in *Handbook of game theory with economic applications*, vol. 4, pp. 861–899, Elsevier, 2015.
- [10] M. Gairing, “Covering games: Approximation through non-cooperation,” in *International Workshop on Internet and Network Economics*, pp. 184–195, Springer, 2009.
- [11] U. Feige, “A threshold of $\ln n$ for approximating set cover,” *Journal of the ACM (JACM)*, vol. 45, no. 4, pp. 634–652, 1998.
- [12] D. Monderer and L. S. Shapley, “Potential games,” *Games and economic behavior*, vol. 14, no. 1, pp. 124–143, 1996.
- [13] H. P. Young, “The evolution of conventions,” *Econometrica: Journal of the Econometric Society*, pp. 57–84, 1993.
- [14] J. F. Nash et al., “Equilibrium points in n-person games,” *Proceedings of the national academy of sciences*, vol. 36, no. 1, pp. 48–49, 1950.
- [15] W. H. Sandholm, “Evolutionary game theory,” *Complex Social and Behavioral Systems: Game Theory and Agent-Based Models*, pp. 573–608, 2020.
- [16] G. Christodoulou, V. S. Mirrokni, and A. Sidiropoulos, “Convergence and approximation in potential games,” in *Annual Symposium on Theoretical Aspects of Computer Science*, pp. 349–360, Springer, 2006.
- [17] V. Bilò, A. Fanelli, M. Flammini, and L. Moscardelli, “Performances of one-round walks in linear congestion games,” in *International Symposium on Algorithmic Game Theory*, pp. 311–322, Springer, 2009.
- [18] V. Bilò, “A unifying tool for bounding the quality of non-cooperative solutions in weighted congestion games,” *Theory of Computing Systems*, vol. 62, no. 5, pp. 1288–1317, 2018.
- [19] V. S. Mirrokni and A. Vetta, “Convergence issues in competitive games,” in *Approximation, randomization, and combinatorial optimization. algorithms and techniques*, pp. 183–194, Springer, 2004.
- [20] A. Fanelli, M. Flammini, and L. Moscardelli, “The speed of convergence in congestion games under best-response dynamics,” in *International Colloquium on Automata, Languages, and Programming*, pp. 796–807, Springer, 2008.
- [21] D. S. Hochbaum and W. Maass, “Approximation schemes for covering and packing problems in image processing and vlsi,” *Journal of the ACM (JACM)*, vol. 32, no. 1, pp. 130–136, 1985.
- [22] S. Mandal, N. Patra, and M. Pal, “Covering problem on fuzzy graphs and its application in disaster management system,” *Soft Computing*, vol. 25, no. 4, pp. 2545–2557, 2021.
- [23] J. M. Nash, “Optimal allocation of tracking resources,” in *1977 IEEE Conference on Decision and Control including the 16th Symposium on Adaptive Processes and A Special Symposium on Fuzzy Set Theory and Applications*, pp. 1177–1180, IEEE, 1977.
- [24] C.-F. Huang and Y.-C. Tseng, “The coverage problem in a wireless sensor network,” *Mobile networks and Applications*, vol. 10, no. 4, pp. 519–528, 2005.
- [25] M. Goemans, V. Mirrokni, and A. Vetta, “Sink equilibria and convergence,” in *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS’05)*, pp. 142–151, IEEE, 2005.
- [26] 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.
- [27] A. Vetta, “Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions,” in *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, pp. 416–425, IEEE, 2002.
- [28] V. Ramaswamy, D. Paccagnan, and J. R. Marden, “Multiagent maximum coverage problems: The trade-off between anarchy and stability,” in *2019 18th European Control Conference (ECC)*, pp. 1043–1048, IEEE, 2019.
- [29] R. Konda, R. Chandan, D. Grimsman, and J. R. Marden, “Balancing asymptotic and transient efficiency guarantees in set covering games,” *arXiv preprint arXiv:2106.04031*, 2021.