Distributed Welfare Games with Applications to Sensor Coverage

Jason R. Marden and Adam Wierman

Abstract— Traditionally resource allocation problems are approached in a centralized manner; however, often centralized control is impossible. We consider a distributed, non-cooperative approach to resource allocation. In particular, we consider the situation where the global planner does not have the authority to assign players to resources; rather, players are self-interested. The question that emerges is how can the global planner entice the players to settle on a desirable allocation with respect to the global welfare? To study this question, we focus on a class of games that we refer to as distributed welfare games. Within this context, we investigate how the global planner should distribute the global welfare to the players. We measure the efficacy of a distribution rule in two ways: (i) Does a pure Nash equilibrium exist? (ii) How efficient are the Nash equilibria as compared with the global optimum? We derive sufficient conditions on the distribution rule that ensures the existence of a pure Nash equilibrium in any single-selection distributed welfare game. Furthermore, we derive bounds on the efficiency of these distribution rules in a variety of settings. Lastly, we highlight the implications of these results in the context of the sensor coverage problem.

I. Introduction

Resource allocation problems have garnered significant research attention across many disciplines, e.g., [2]–[5] and references therein. One example, which we will focus on throughout this paper, is sensor coverage. The goal of the sensor coverage problem is to allocate a fixed number of sensors across a given "mission space" so as to maximize the probability of detecting a particular event [6]. A second example is task assignment where the objective is to assign people to tasks so as to maximize profit or minimize completion time. Regardless of the specific application domain, the underlying goal of any resource allocation problems remains the same; a global planner desires to distribute resources so as to maximize a global objective.

Traditionally, researchers have aimed at developing algorithms to determine near optimal allocations for resource allocation problems [7]–[9]. In this paper, we will approach the problem from a different, game-theoretic, perspective. Suppose that a global planner does not *have*, or possibly *want*, the authority to assign players to resources; rather, players are self-interested. In many environments, such as transportation systems, a global planner does not *have* the authority to assign players to resources. There are also many cases where a global planner may not *want* the authority to assign players to resources. For example, in the case of sensor coverage, one can design sensors as autonomous self-interested entities as opposed to centrally controlled resources.

There are wide-ranging advantages to this game-theoretic form of a distributed architecture, including robustness to agent failures and environmental disturbances, reducing communication requirements, improving scalability, etc. However, several challenges arise when seeking to design and implement such a distributed system [10], [11]. The primary challenge is: how can a global planner entice the players to settle on something desirable with regards to the global welfare? Equivalently, in the case of engineered systems, such as sensor coverage, how can a global planner design local utility functions so that players will settle on a desirable allocation?

The goal of this paper is to establish a general framework, independent of any specific application domain, for (i) investigating the feasibility of non-cooperative resource allocation and (ii) designing desirable utility functions. To that end, in this paper we will consider a class of resource allocation games that we refer to as *distributed welfare games*. A distributed welfare game is a resource allocation game where each player's utility is defined as some fraction of the total welfare garnered (see Section III). Therefore, designing a utility function is equivalent to defining a distribution rule that depicts how the welfare garnered from a specific allocation is distributed to the players. The primary goal is to design distribution rules for distributed welfare games that guarantee the following two properties.

(i) *Existence:* A distribution rule should guarantee that a (pure) Nash equilibrium exists. A Nash equilibrium represents an individually agreeable allocation.

(ii) *Efficiency:* A distribution rule should guarantee that all Nash equilibria are efficient with respect to the global objective.

In addition to the two properties above, which are of primary importance, in many applications there are additional requirements that distribution should satisfy. We will investigate four alternative features of distribution rules including (iii) tractability, (iv) low informational requirement, (v) budget-balance, and (vi) whether the distribution rule results in the formulation of a potential game [12].

With the goal of developing distribution rules that satisfy the above properties, we first focus on developing distribution rules that guarantee the existence of a Nash equilibrium (property i) while satisfying as many of (iii)-(vi) as possible. To accomplish this, we begin by investigating the applicability of cost-sharing methodologies [13]–[15]. Building upon results from [15], we show that one can use cost sharing methodologies as distribution rules in distributed welfare games. However, we demonstrate that these approaches do not satisfy all of the desirable properties above. For example, rules that are budget balanced and guarantee the existence of a Nash equilibrium, often come with significant informational and computational costs.

The natural question that emerges is whether there exists special classes of distributed welfare games where one can establish distribution rules satisfying the above properties. Our second set of results (Section V) investigates this question. We identify three easily verifiable properties of distribution rules, see Conditions 1–3, that guarantee the existence of a Nash equilibrium in any distributed welfare game where players are restricted to selecting a single resource.

In Section VI we explore the efficiency of Nash equilibria (property ii). We measure the efficiency of a Nash equilibrium using the *price of anarchy*. The price of anarchy is defined as the worst-case ratio between the global welfare evaluated at the worst Nash equilibrium and the optimal welfare. In general, the price of anarchy in distributed welfare games can be arbitrarily close to 0; however, when we restrict our attention to submodular welfare functions, which is a common assumption in many resource allocation problems [16], [17], we can develop distribution rules that obtain a welfare within 1/2 of that of the optimal assignment. Furthermore, we tighten this price of anarchy bound in a variety of settings. This compares favorably with the best known results of centralized approximations for resource allocation problems with submodular welfare functions, which guarantee welfare within $1 - 1/e \approx 0.6321$ of the optimal [7]–[9].

To illustrate the impact of the results described above, in Section VII we apply our results to the sensor coverage problem.

It should be noted that this paper predominantly focuses on equilibrium behavior in distributed welfare games; however another natural question is "how do players reach an equilibrium in a distributed fashion?" While not focusing on this question in detail, we do illustrate the potential of the theory of *learning in games* [11], [18]–[21] as a local control mechanism for coordinating group behavior.

II. BACKGROUND

In this paper we consider *resource allocation games* that consist of a set of players $N := \{1, ..., n\}$ and a finite set of resources $\mathcal{R}$ that are to be shared by the players. Each player $i \in N$ is assigned an action set $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$ and a utility function of the form $U_i : \mathcal{A} \to \mathbb{R}$ where $2^{\mathcal{R}}$ denotes the power sets of $\mathcal{R}$ and $\mathcal{A} := \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$ is referred to as the set of joint actions. Therefore, a player may have the option of selecting multiple resources and the player's utility may be influenced by the actions of other players.

For an action profile $a = (a_1, a_2, ..., a_n) \in \mathcal{A}$, let a_{-i} denote the profile of player actions *other than* player i, i.e., $a_{-i} = (a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$. With this notation, we will sometimes write a profile a of actions as (a_i, a_{-i}). Similarly, we may write $U_i(a)$ as $U_i(a_i, a_{-i})$. Furthermore, let $\mathcal{A}_{-i} := \prod_{j \neq i} \mathcal{A}_j$ denote the set of possible collective actions of all players other than player i.

We will focus on analyzing equilibrium behavior in such games. A well-known equilibrium concept that emerges in non-cooperative games is that of a *pure Nash equilibrium*. An action profile $a^* \in \mathcal{A}$ is called a *pure Nash equilibrium* if for all players $i \in N$, $U_i(a_i^*, a_{-i}^*) = \max_{a_i \in \mathcal{A}_i} U_i(a_i, a_{-i}^*)$. A pure Nash equilibrium represents a scenario for which no player has an incentive to unilaterally deviate. We will refer to a pure Nash equilibrium as simply an equilibrium.

One class of games discussed in this paper is *potential games* [12]. In a potential game, the change in a player's utility that results from a unilateral change in strategy equals the change in some global potential function. Specifically, there is a function $\phi : \mathcal{A} \to \mathbb{R}$ such that for every player $i \in N$, for every $a_{-i} \in \mathcal{A}_{-i}$, and for every $a_i', a_i'' \in \mathcal{A}_i$,

$$U_i(a_i', a_{-i}) - U_i(a_i'', a_{-i}) = \phi(a_i', a_{-i}) - \phi(a_i'', a_{-i}). \quad (1)$$

When this condition is satisfied, the game is called a potential game with the potential function ϕ. In potential games, any action profile maximizing the potential function is an equilibrium, hence every potential game has at least one equilibrium.

For a more comprehensive review of the game-theoretic concepts introduced in this section, we refer the readers to [18], [22]–[24].

III. DISTRIBUTED WELFARE GAMES

A *distributed welfare game* is a non-cooperative formalization of a resource allocation game with a specific structure enforced on player utility functions. The formalization includes a wide variety of resource allocation problems including network routing, wireless power management, sensor coverage, and others. To illustrate the applicability of distributed welfare games, we will focus in detail on the sensor coverage problem in Section VII.

To define distributed welfare games, consider a resource allocation game with a global welfare function, $W : \mathcal{A} \to \mathbb{R}_+$. Each player's utility is defined as some portion of the welfare and must satisfy the following properties: for any player $i \in N$ and action profile (allocation) $a \in \mathcal{A}$

(i) $U_i(a) \geq 0$,
(ii) $\sum_{i \in N} U_i(a) \leq W(a)$.

This structure permits the total global welfare to be distributed arbitrarily to the players.

An example of a utility function that satisfies these conditions is equally distributing the global welfare

$$U_i(a) = \frac{1}{n} \, W(a). \quad (2)$$

This utility design guarantees that any allocation that maximizes the welfare is an equilibrium. However, a player needs to know the selections of all other players in addition to the structural form of the global welfare function in order to evaluate her utility, which is typically impractical.

Our focus in this paper is on understanding the degree to which players' utility functions can be *localized* while at the same time guaranteeing both existence and efficiency of equilibria. To this end, we will restrict our attention to *separable* welfare functions of the form

$$W(a) = \sum_{r \in \mathcal{R}} W^r(a^r),$$

where $W^r : 2^N \to \mathbb{R}_+$ is the welfare function for resource r and a^r denotes the subset of players that selected resource r in the joint allocation a, i.e., $a^r := \{i \in N : r \in a_i\}$. To simplify notation, we will commonly write $W^r(a^r)$ as just $W^r(a)$. Because the welfare function is assumed to be separable, we will also restrict player utility functions to be *local* and *separable*, i.e., of the form

$$U_i(a_i, a_{-i}) = \sum_{r \in a_i} f_i(r, a) \, W^r(a), \quad (3)$$

where $\{f_1(r,a), ..., f_n(r,a)\}$ defines how the global reward garnered from resource r is distributed across the players. We will refer to $\{f_1(r,a), ..., f_n(r,a)\}_{r\in\mathcal{R}, a\in\mathcal{A}}$ as the *distribution rule*. A distribution rule must satisfy the following properties: for any player $i \in N$, resource $r \in \mathcal{R}$, and action profile $a \in \mathcal{A}$

(i) $f_i(r,a) \geq 0$,
(ii) $r \notin a_i \Rightarrow f_i(r,a) = 0$,
(iii) $\sum_{i\in N} f_i(r,a) \leq 1$.

We will refer to distribution rules that satisfy (iii) with equality as *budget balanced distribution rules*.

IV. METHODS FOR DISTRIBUTING WELFARE

We will now explore several natural approaches for designing distribution rules in distributed welfare games. These approaches are derived from methodologies in the cost sharing literature [13], such as the Shapley value [14], [15]. It turns out that cost sharing methodologies can be effective as distribution rules in distributed welfare games; however, there are several issues that limit their applicability. Resultantly, we will discuss alternative approaches to designing distribution rules for distributed welfare games in the Section V.

A. Equally Shared Utilities

The utility design in (2) disseminates the total welfare equally to all players and guarantees the existence of efficient equilibria. Unfortunately, it requires players to use global (rather than local) information and thus it is not a local distribution rule of the form (3).

Suppose the welfare from each resource is divided equally amongst the players that selected the resource, i.e.,

$$U_i(a_i, a_{-i}) = \sum_{r\in a_i} \left(\frac{1}{\sum_j I\{r \in a_j\}} \right) W^r(a), \qquad (4)$$

where $I\{\cdot\}$ is the usual indicator function. In general, such a design cannot guarantee the existence of an equilibrium as the following example illustrates.

		Player 2				Player 2	
	$\emptyset$	r^1	r^2		$\emptyset$	r^1	r^2
$\emptyset$	0	4	1	$\emptyset$	0,0	0,4	0,1
Player 1 r^1	6	6	7	r^1	6,0	3,3	6,1
r^2	5	9	10	r^2	5,0	5,4	5,5
		Separable Welfare				Payoffs	

One problem with equally shared utilities is that players' utility functions are not *aligned* with their contribution to the global welfare. However, if players are anonymous with regards to their impact on the global welfare then the equally shared utilities in (4) guarantee the existence of an equilibrium since the resulting game is a congestion game [1].[1]

[1]Players are anonymous with regards to their impact on the global welfare function if for any action profiles $a, a' \in \mathcal{A}$, and any resource $r \in \mathcal{R}$ $\sigma^r(a) = \sigma^r(a') \Rightarrow W^r(a) = W^r(a')$, where $\sigma^r(a) := |\{i \in N : r \in a_i\}|$ denotes the number of players utilizing resource r given the allocation a. Hence, the welfare generated by a particular resource depends only on the number of players utilizing that resource, not the specific players utilizing the resource.

B. Marginal Contribution Utilities

By conditioning a distribution rule not only on the welfare garnered, but also on each player's marginal contribution to the existing welfare, we can guarantee the existence of an equilibrium even when players are not anonymous. Specifically, suppose each player's utility is set as his marginal contribution to the global welfare, i.e.,

$$U_i(a_i, a_{-i}) = W(a_i, a_{-i}) - W(a_i^0, a_{-i}), \qquad (5)$$

where a_i^0 designate the null action for player i, i.e, $a_i^0 = \emptyset$; however, in general a_i^0 could be set as any fixed action in the player's action set. We will refer to the utility design in (5) as the *wonderful life utility* (WLU) [25]. It is well known that assigning each player a utility as in (5) results in a potential game with potential function W; hence any action profile that maximizes the global welfare is an equilibrium. However, other equilibria may also exist under the wonderful life utility design.

There are two limitations of the the marginal contribution utility design. First, each player needs to be able to compute his marginal contribution to the welfare in order to evaluate his utility. Second, the wonderful life utility may distribute more (or less) welfare than is gathered; hence, it may not satisfy condition (iii) of distributed welfare games. While the first limitation cannot be relaxed in this setting, it remains an open question as to whether the second limitation can be addressed utilizing a similar informational requirement.

C. The Shapley Value

While WLU guarantees the existence of an equilibrium in all settings, it may distribute more or less reward than the welfare garnered. It turns out that this problem can be rectified using a common cost sharing methodology known as the Shapley value [14]. For any subset of players $\widetilde{N} \subseteq N$, resource $r \in \mathcal{R}$, and player $i \in \widetilde{N}$, the Shapley value of player i is defined as [14], [15], [26]

$$Sh_i^r(\widetilde{N}) := \sum_{S\subseteq\widetilde{N}:i\in S} \frac{(|\widetilde{N}| - 2)!(|S| - 1)!}{|\widetilde{N}|!} (W^r(S) - W^r(S \setminus \{i\})). \qquad (6)$$

Suppose each player's utility function is defined as

$$U_i(a_i, a_{-i}) := \sum_{r\in a_i} Sh_i^r(a^r), \qquad (7)$$

where $Sh_i^r(a^r)$ is player i's Shapley value at resource r given the allocation of players a^r.

Using the Shapley value as in (7) to distribute welfare leads to a potential game, and thus the existence of an equilibrium.

Proposition 1. *Consider any resource allocation game. If each player is assigned a utility of the form (7), then the ensuing game is a budget balanced distributed welfare game. Furthermore, it is a potential game with the following potential function $\phi : \mathcal{A} \to \mathbb{R}$*

$$\phi(a) := \sum_{r\in\mathcal{R}} \sum_{S\subseteq a^r} \frac{1}{|S|} \left(\sum_{T\subseteq S} (-1)^{|S|-|T|} W^r(T) \right). \qquad (8)$$

Unfortunately, there are two limitations of the Shapley value utility design that may prevent it from being applicable.

First, there is a high informational requirement as each player must be able to compute his marginal contribution to all action profiles in order to evaluate his utility. Second, in general computing a Shapley value is intractable in games with a large number of players. This is highlighted explicitly in (6) where computation of the Shapley value requires a weighted summation over all subsets of players. However, it should be noted that this computational cost is lessened dramatically if there are a limited number of distinct "classes" of players, e.g., if players are anonymous then the Shapley value is equivalent to the equal share distribution rule in (4).

V. SINGLE SELECTION DISTRIBUTED WELFARE GAMES

The previous section provided a number of distribution rule designs; however they all suffered from significant limitations, such as intractability and high informational requirements. In this section, we restrict the setting in the hopes of deriving distribution rules that avoid these limitations. In particular, we focus on a simplified setting where players are only allowed to select a single resource, i.e., $\mathcal{A}_i \subseteq \mathcal{R}$ as opposed to $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$.

To this end, we identify three sufficient conditions that determine whether a distribution rule guarantees the existence of an equilibrium for a particular game. These sufficient conditions translate to pairwise comparisons of players' utility functions. We will illustrate the applicability of these conditions using the sensor coverage problem in Section VII.

Before stating the conditions, we will introduce the following notation. The allocation $(r^{(i)}, r^{(j)}, a_{-ij})$, denotes the situation where player i selects resource $r^{(i)}$, player j selects resource $r^{(j)}$, and all other player select resources according to $a_{-ij} \in \mathcal{A}_{-ij} := \prod_{k \neq i,j} \mathcal{A}_k$. Likewise, let $U_i(r^{(i)}, r^{(j)}, a_{-ij})$ and $U_j(r^{(i)}, r^{(j)}, a_{-ij})$ denote player i and j's respective utilities for the given allocation.

Condition 1. *Let i and j be any two players. If for any resource $r \in \mathcal{R}$ and any allocation $a_{-ij} \in \mathcal{A}_{-ij}$*

$$U_i(r, r, a_{-ij}) > U_j(r, r, a_{-ij}),$$

then for all resources $r' \in \mathcal{R}$ and all allocations $a'_{-ij} \in \mathcal{A}_{-ij}$

$$U_i(r', r', a'_{-ij}) \geq U_j(r', r', a'_{-ij}).$$

In this situation, we will say that player i is stronger than player j. Furthermore, we will require that player strengths are a transitive property, that is if player i is stronger than player j who is stronger than player k, then player i is also stronger than player k.

Condition 2. *Suppose player i is stronger than player j. For any resource $r \in \mathcal{R}$ and any action profile $a_{-ij} \in \mathcal{A}_{-ij}$, the following holds*

$$U_i(r, a_j^0, a_{-ij}) \geq U_i(r, r, a_{-ij}).$$

Condition 3. *Suppose player i is stronger than player j. For any resources $r \in \mathcal{R}$ and action profile $a_{-ij} \in \mathcal{A}_{-ij}$, the following holds*

$$\frac{U_j(a_i^0, r, a_{-ij})}{U_i(r, a_j^0, a_{-ij})} \geq \max_{r^* \in \mathcal{R}} \frac{U_j(r^*, r^*, a_{-ij})}{U_i(r^*, r^*, a_{-ij})}.$$

Theorem 2. *Consider any single selection resource allocation game. If each player's utility function satisfies Conditions 1, 2, and 3 then an equilibrium exists.*

Proof: We begin by renumbering the players in order of strengths with player 1 being the strongest player. This is possible because of Condition 1.

We will construct an equilibrium by letting each player select his action one at a time in order of strength. The general idea of the proof is that once a player selects an action, the player will never seek to deviate regardless of the other player's action. First, player 1 selects the resource $r^{(1)}$ according to

$$r^{(1)} \in \arg\max_{r \in \mathcal{R}} U_1(r, a_{-1}^0) \qquad (9)$$

Next, player 2 selects action $r^{(2)}$ according to

$$r^{(2)} \in \arg\max_{r \in \mathcal{R}} U_2(r^{(1)}, r, a_{-12}^0).$$

If $r^{(1)} \neq r^{(2)}$, then by (9) and Condition 2 we know that

$$U_1(r^{(1)}, a_{-1}^0) \geq U_1(r^{(2)}, a_{-1}^0) \geq U_1(r^{(2)}, r^{(2)}, a_{-12}^0).$$

Therefore, player 1 can not improve his utility by switching his strategy, i.e.,

$$U_1(r^{(1)}, r^{(2)}, a_{-12}^0) \geq U_1(r, r^{(2)}, a_{-12}^0), \ \forall r \in \mathcal{R}.$$

If $r^{(1)} = r^{(2)} = r$, then by Condition 3, we know that for any resource $\widetilde{r} \in \mathcal{R}$, $\widetilde{r} \neq r$,

$$\frac{U_2(\widetilde{r}, r, a_{-12}^0)}{U_1(r, \widetilde{r}, a_{-12}^0)} = \frac{U_2(a_1^0, r, a_{-12}^0)}{U_1(r, a_2^0, a_{-12}^0)} \geq \frac{U_2(r, r, a_{-12}^0)}{U_1(r, r, a_{-ij}^0)}.$$

Using the above inequality, we can conclude that for any resource $\widetilde{r} \in \mathcal{R}$, $U_2(r, r, a_{-12}^0) \geq U_2(r, \widetilde{r}, a_{-12}^0)$ implies $U_1(r, r, a_{-12}^0) \geq U_1(\widetilde{r}, r, a_{-12}^0)$. Therefore, player 1 cannot improve his utility by switching his strategy.

If $n = 2$, then $a = (r^{(1)}, r^{(2)})$ would be an equilibrium. Otherwise this argument could be repeated n times to construct an equilibrium.
$\square$

It remains an open question as to whether Conditions $1-3$ guarantee additional properties pertaining to the structure of the game besides existence of an equilibrium. For example, if each player's utility function satisfies Conditions $1-3$, is the game a potential game or some variant?

VI. EFFICIENCY OF EQUILIBRIA

We will measure the efficiency of equilibria using the *price of anarchy* [24]. The price of anarchy (or more appropriately in this context the "price of localization") is defined as the worst case ratio between the global welfare at any equilibrium and the optimal global welfare. Formally, for any equilibrium a^{ne} and optimal allocation a^{opt}, $W(a^{ne}) \geq \gamma W(a^{opt})$ where γ is the price of anarchy.

Unfortunately, without any assumptions on the global welfare function W, the price of anarchy can be arbitrarily bad in distributed welfare games. However, when the welfare function is submodular it is possible to attain a much better price of anarchy.[2] We can utilize Theorem 3.4 in [16] in the context of distributed welfare games to prove the following proposition.

---

[2]A set valued function $W : 2^{\mathcal{A}} \to \mathbb{R}$ is submodular if $W(X) + W(Y) \geq W(X \cap Y) + W(X \cup Y)$ for all $X, Y \subseteq 2^{\mathcal{A}}$. Submodularity corresponds to the notion of a decreasing marginal contribution and is a common in many resource allocation problems, e.g., [16], [17]. Further, it is a key property underlying the design of many centralized algorithms for these problems.

Proposition 3. *Consider any distributed welfare game with a submodular global welfare function. If*

  (i) players are anonymous and assigned an equally shared utility as in (4), or
  (ii) players are assigned a wonderful life utility as in (5), or
  (iii) players are assigned a Shapley value utility as in (7), or
  (iv) players are assigned a utility function that satisfies Condtions 1–3 (single-selection)

then an equilibrium exists and the price of anarchy is 1/2.

See [1] for the proof of this proposition and examples illustrating that the bound is tight. To provide a basis for comparison, computing the optimal assignment for a general distributed welfare game is NP-complete [5]. Further, the best known approximation algorithms guarantee only to provide a solution that is within $1 - 1/e \approx 0.6321$ of the optimal [7]–[9]. Thus, the 1/2 price of anarchy in this scenario is comparable to the best centralized solution.

VII. THE SENSOR COVERAGE PROBLEM

Consider the sensor coverage problem in [6], [27]. There exists a finite set of sectors denoted as $\mathcal{R}$ and a value function $V : \mathcal{R} \to \mathbb{R}_+$. Note that $V(r)$ often has a very intuitive meaning, e.g., in the case of enemy submarine tracking, $V(r)$ represents the a priori probability that an enemy submarine is situated in sector r.

There are a finite number of autonomous sensors (or players) denoted as $N = \{1, ..., n\}$ allocated to the mission space. Each sensor i is capable of sensing activity in (monitoring) possibly multiple sectors simultaneously based on its chosen location. The set of possible monitoring choices for sensor i is denoted as $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$. Let $\mathcal{A} = \prod_{i \in N} \mathcal{A}_i$ be the set of joint actions, or monitoring choices, for all players. The probability that sensor i detects an event in sector r given his current monitoring choice a_i is denoted as $p_i(r, a_i)$. We will assume that the detection probabilities satisfy:

$$r \in a_i \quad \Leftrightarrow \quad p_i(r, a_i) > 0,$$
$$r \notin a_i \quad \Leftrightarrow \quad p_i(r, a_i) = 0,$$

For a given joint action profile $a := \{a_1, ..., a_n\}$, the joint probability of detecting an event in sector r is

$$P(r, a) = 1 - \prod_{i \in N}[1 - p_i(r, a_i)].$$

The goal of the global planner in this scenario is to allocate the sensors in a way that maximizes the probability of detecting an event, which is characterized by the following global welfare function [6]

$$W(a) = \sum_{r \in \mathcal{R}} V(r) P(r, a). \qquad (10)$$

A. The Sensor Coverage Game

Rather than view the sensor coverage problem as a centralized optimization problem, our focus is on the design of autonomous sensors that are individually capable of making their own independent decision in response to local information. We will model the interactions of the sensors as a non-cooperative resource allocation game where each sensor i is assigned a utility function $U_i : \mathcal{A} \to \mathbb{R}$ that defines his payoff (utility) for each monitoring profile.

We will refer to the non-cooperative game theoretic formulation of the sensor coverage problem as the *sensor coverage game*. The sensor coverage game is a simple example of a distributed welfare game. Thus, we can apply our results in order to design distribution rules that guarantee the existence of an equilibrium. In particular, we can immediately conclude that the wonderful life and the Shapley value utility designs will guarantee the existence of an equilibrium. Further, since the welfare function is submodular, both rules yield a price of anarchy of 1/2. However, both rules come with their respective limitations as previously discussed.

B. Single Selection Sensor Coverage Games

Consider the situation where each sensor is only capable of selecting a single sector, i.e., $\mathcal{A}_i = \mathcal{R}$. Furthermore, suppose each sensor has an invariant detection probability $p_i > 0$ such that $p_i(r, a_i) = p_i$ if $r = a_i$, otherwise $p_i(r, a_i) = 0$.

Utilizing Conditions 1-3, one can prove the following [1].

Corollary 4. *Consider any single selection sensor coverage game where each sensor has an invariant detection probability. An equilibrium is guaranteed to exist under the following distribution rules:*

$$U_i(r, a_{-i}) = \frac{p_i}{\sum_{j \in N : a_j = r} p_j} P(r, a) V(r), \qquad (11)$$

$$U_i(r, a_{-i}) = \left(\frac{\frac{p_i}{1 - p_i}}{\sum_{j \in N : a_j = r} \frac{p_j}{1 - p_j}} \right) P(r, a) V(r). \qquad (12)$$

Furthermore, the price of anarchy in each case is 1/2.

C. Anonymous Sensor Coverage Games

If we specialize further to the case of anonymous sensors, that is all sensors have the same detection probability p, we can obtain a tighter bound on the price of anarchy that illustrates the impact of the number of sensors and the detection probability [1].

Theorem 5. *Consider a single-sector anonymous sensor coverage game with n sensors each having invariant detection probability p. Under the equal share utility design (4) and equilibrium is guaranteed to exist, and the price of anarchy is bounded by*

$$\frac{W(a^{ne})}{W(a^{opt})} \geq \left(\frac{a^*}{n} + \frac{1 - (1-p)^{n-a^*}}{1 - (1-p)^n} \right)^{-1}$$

where $a^* = \begin{cases} n - 1, & p = 1; \\ n - \dfrac{\log\left(n \frac{\log(1/(1-p))}{1-(1-p)^n}\right)}{\log(1/(1-p))}, & p < 1. \end{cases}$

D. Simulation Experiments

Consider a single selection sensor coverage game with 100 sensors with invariant detection probability $p = 0.25$. The mission space is $\mathcal{R} = \{r_1, ..., r_{25}\}$. The value for each sector is randomly assigned from a uniform distribution; two sectors according to $U[0, 6]$, four sectors according to $U[0, 3]$, and the remaining according to $U[0, 1]$. Each sensor is capable of monitoring any of the 25 sectors, i.e., $\mathcal{A}_i = \mathcal{R}$ and uses the equal share utility design (4).

There is a large body of literature analyzing distributed learning algorithms in congestion games, or equivalently potential games [11], [18]–[21]. We will apply *fading memory*
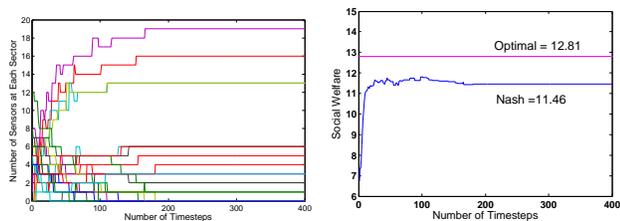
(a) Evolution of Number of Sensors in Each Sector over Mission

(b) Evolution of Welfare Function over Mission

Fig. 1. Simulation results for sensor coverage game.

joint strategy fictitious play with inertia, which guarantees convergence to an equilibrium in any (generic) congestion game while maintaining computational tractability even in large-scale games. We refer the reader to [20] for the details of the learning rule. We use the following discount factor and inertia: $\lambda = 0.5$ and $\epsilon = 0.02$.

Figure 1(a) illustrates the evolution of the number of sensors at each sector. The identity of the sectors is unimportant as the key observation is that behavior settles down at an equilibrium. Figure 1(b) illustrates the evolution of the global welfare in addition to the efficiency gap between the equilibrium and the optimal. The results illustrate that Theorem 5 provides a very conservative estimate of the price of anarchy since the observed efficiency is 0.936.

VIII. CONCLUDING REMARKS

In this paper, we focus on a class of games that we refer to as distributed welfare games. These games are formulated to study how the method used to divide the global welfare among participating players impacts the existence and efficiency of equilibria. We derive three sufficient conditions on distribution rules that guarantee the existence of an equilibrium in the setting where players are only allowed to select a single resource. Further, we illustrate the applicability of these conditions in the case of the sensor coverage problem. In general, designing a distribution rule that guarantees the existence of an equilibrium in distributed welfare games is an open problem. We also derive general bounds on the price of anarchy in distributed welfare games and application specific bounds on the price of anarchy for the sensor coverage problem. Importantly, the structure of the welfare function W for sensor coverage parallels those for many other problems, e.g., weapon targeting and fault detection.

An important open question that remains involves the use of learning rules for distributed welfare games. When players are anonymous, we demonstrated that there are several distributed learning algorithms that guarantee players will reach an equilibrium. However, it remains to design learning algorithms for the distributed welfare games with players that are not anonymous.

REFERENCES

[1] J. R. Marden and A. Wierman, *Distributed Welfare Games*. Under submission, 2008.
[2] T. Ibaraki and N. Katoh, *Resource allocation problems: algorithmic approaches*. Cambridge, MA, USA: MIT Press, 1988.
[3] R. Ahuja, T. Magnanti, and J. Orlin, *Network flows: theory, algorithms, and applications*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1993.
[4] D. Ghose, M. Krichman, J. Speyer, and J. Shamma, "Modeling and analysis of air campaign resource allocation: A spatio-temporal decomposition approach," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. **32**, no. 3, pp. 403–418, 2002.
[5] R. Murphey, "Target-based weapon target assignment problems," in *Nonlinear Assignment Problems: Algorithms and Applications*, P. Pardalos and L. Pitsoulis, Eds. Kluwer Academic Publishers, 1999, pp. 39–53.
[6] W. Li and C. G. Cassandras, "Sensor networks and cooperative control," in *Proc. of Conf. on Decision and Control*, 2005.
[7] U. Feige and J. Vondrak, "Approximation algorithms for allocation problems: Improving the factor of 1 - 1/e," in *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, 2006, pp. 667–676.
[8] A. Ageev and M. Sviridenko, "Pipage rounding: a new method for constructing algorithms with proven performance guarantee," *Journal of Combinatorial Optimization*, vol. 8, pp. 307–328, 2004.
[9] R. K. Ahuja, A. Kumar, K. Jha, and J. B. Orlin, "Exact and heuristic methods for the weapon-target assignment problem," 2003, http://ssrn.com/abstract=489802.
[10] G. Arslan, J. R. Marden, and J. S. Shamma, "Autonomous vehicle-target assignment: a game theoretical formulation," *ASME Journal of Dynamic Systems, Measurement and Control*, vol. 129, pp. 584–596, September 2007.
[11] J. R. Marden, G. Arslan, and J. S. Shamma, "Connections between cooperative control and potential games illustrated on the consensus problem," in *Proceedings of the 2007 European Control Conference (ECC '07)*, July 2007.
[12] D. Monderer and L. Shapley, "Potential games," *Games and Economic Behavior*, vol. 14, pp. 124–143, 1996.
[13] H. P. Young, *Equity*. Princeton, NJ: Princeton University Press, 1994.
[14] L. S. Shapley, "A value for n-person games," in *Contributions to the Theory of Games II (Annals of Mathematics Studies 28)*, H. W. Kuhn and A. W. Tucker, Eds. Princeton, NJ: Princeton University Press, 1953, pp. 307–317.
[15] S. Hart and A. Mas-Colell, "Potential, value, and consistency," *Econometrica*, vol. 57, no. 3, pp. 589–614, May 1989.
[16] A. Vetta, "Nash equilibria in competitive societies with applications to facility location, traffic routing, and auctions," in *Proc. of Symp. on Fdns. of Comp. Sci.*, 2002, pp. 416–425.
[17] A. Krause and C. Guestrin, "Near-optimal obervation selection using submodular functions," in *Proc. of Conf. on Artifical Intelligence*, 2007.
[18] H. P. Young, *Individual Strategy and Social Structure*. Princeton, NJ: Princeton University Press, 1998.
[19] D. Monderer and L. Shapley, "Fictitious play property for games with identical interests," *Games and Economic Theory*, vol. 68, pp. 258–265, 1996.
[20] J. R. Marden, G. Arslan, and J. S. Shamma, "Joint strategy fictitious play with inertia for potential games," *IEEE Transactions on Automatic Control*, 2008, to appear.
[21] ——, "Regret based dynamics: Convergence in weakly acyclic games," in *Proceedings of the 2007 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, Honolulu, Hawaii, May 2007.
[22] D. Fudenberg and J. Tirole, *Game Theory*. Cambridge, MA: MIT Press, 1991.
[23] H. P. Young, *Strategic Learning and its Limits*. Oxford University Press, 2005.
[24] N. Nissan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic game theory*. New York, NY, USA: Cambridge University Press, 2007.
[25] D. Wolpert and K. Tumor, "An overview of collective intelligence," in *Handbook of Agent Technology*, J. M. Bradshaw, Ed. AAAI Press/MIT Press, 1999.
[26] G. Haeringer, "A new weight scheme for the shapley value," *Mathematical Social Sciences*, vol. 52, no. 1, pp. 88–98, July 2006.
[27] S. Iyengar and R. Brooks, *Distributed sensor networks*. Boca Raton, FL, USA: Chapman & Hall, 2005.