Distributed Welfare Games with Applications to Sensor Coverage

Submission Supplement:

Proofs that were excluded from the submission due to space constraints are included in the appendix.

Jason R. Marden[*] Adam Wierman[†]

Abstract

We consider a variation of the resource allocation problem. In the traditional problem, there is a global planner who would like to assign a set of players to a set of resources so as to maximize social welfare. 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 social 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 social welfare to the players? We measure the efficacy of a distribution rule in two ways: (i) Does a pure Nash equilibrium exist? (ii) How does the social welfare of a pure Nash equilibrium compare to the social welfare of the optimal allocation? This comparison is referred to as the price of anarchy. To answer these questions, 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 a price of anarchy for distributed welfare games in a variety of settings. Lastly, we highlight the implications of these results using the sensor coverage problem.

1 Introduction

Resource allocation problems have garnered significant research attention across many disciplines There are a number of practical examples of resource allocation problems that have fueled this research focus, see [9, 2, 8, 19] 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 [13]. A second example is task assignment where the objective is to assign people to tasks so as to maximize a 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 some form of a global objective, i.e., the social welfare.

Traditionally, researchers have aimed at developing algorithms to determine near optimal allocations for resource allocation problems [6, 1, 3]. In this paper, we will approach the problem from a different,

[*]Jason R. Marden is postdoctoral scholar in the Information Science and Technology center at California Institute of Technology, marden@caltech.edu.

[†]Adam Wierman is an Assistant Professor in the Department of Computer Science at California Institute of Technology, adamw@caltech.edu.

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 sensors 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 [4, 14]. The primary challenge is: how can a global planner entice the players to settle on something desirable with regards to the social welfare? Equivalently, in the case of engineered systems, such as sensor coverage, how can a global planner *design local utility functions* for the players so that they will settle on a desirable allocation?

Motivated by this question, in this paper, we analyze a class of games we refer to as *distributed welfare games*. We will define the class formally in Section 3, but, at a high level, the general theme of distributed welfare games is that the social welfare generated from a particular allocation is distributed to the participating players at the global planner's discretion. This class is formulated in order to isolate the following design question: "how should the global planner distribute the social welfare to the players?".

To answer this question we need to measure the quality of a distribution rule. We use two common measures for this purpose. The first measure that we consider is whether a given distribution rule guarantees the existence of a pure Nash equilibrium. In a non-cooperative setting where players are self-interested, a pure Nash equilibrium represents an individually agreeable allocation. We seek to identify characteristics of the distribution rule that guarantee the existence of a pure Nash equilibrium in distributed welfare games. We identify three properties of the distribution rule, see Conditions 4.1–4.3, that guarantee the existence of a Nash equilibrium in any distributed welfare games where players can select only one resource. Further, we show that a simple distribution rule guarantees the existence of a Nash equilibrium in any anonymous distributed welfare game.

The second measure that we consider is the relationship between the social welfare of a pure Nash equilibrium and the social welfare of the optimal (centralized) allocation. We will measure the suboptimality of a Nash equilibrium using the *price of anarchy*. The price of anarchy is defined as the worst case ratio between the social welfare evaluated at any pure Nash equilibrium and the optimal social welfare and been studied extensively, [22, 11, 20]. In general, the price anarchy can be arbitrarily bad in resource allocation problems; however, when we restrict our attention to certain classes of social welfare functions, such as submodular, we can develop distribution rules that obtain social welfare within 1/2 of that of the optimal assignment. Further, we derive specific bounds for sensor coverage by exploiting the structure of the social welfare function.

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" [24, 17, 16, 15, 14] as a local control mechanism for coordinating group behavior.

The remainder of the paper will be organized as follows. We will begin in Section 2 by highlighting one resource allocation problem of particular interest: the sensor coverage problem. Furthermore, all relevant game theoretic concepts used in this paper will be introduced in this section. We will define the class of distributed welfare games in Section 3. Next, in Section 4, we will discuss the task of designing a distribution rule that guarantees the existence of a Nash equilibrium. Then, in Section 5, we will discuss the efficiency of Nash equilibria in distributed welfare games. To highlight the impact of our results for distributed welfare games, we will discuss their implications for sensor coverage in Section 6. Finally, we will conclude in Section 7.

2 Motivational example

To ground the discussion in this paper, we will highlight many of the issues we discuss using a specific resource allocation problem: the sensor coverage problem. This problem is of particular interest given the growing deployment of sensor networks in a wide range of applications including surveillance, military, environmental monitoring, and beyond.

2.1 The sensor coverage problem

One of the fundamental problems in sensor networks 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. For a more detailed introduction to the problem, refer to [10, 13].

In modeling this problem, we will divide the mission space into a finite set of sectors denoted as X and define an events density function, or relative reward function, $R(x)$, over X, where $R(x) \geq 0, \forall x \in X$. This formulation of the problem is common [5, 10, 26]. Note that $R(x)$ often has a very intuitive meaning, e.g., in the case of enemy submarine tracking, $R(x)$ represents the a priori probability that an enemy submarine is situated in sector x.

There are a finite number of autonomous sensors (or players) denoted as $\mathcal{P} = \{\mathcal{P}_1, ..., \mathcal{P}_n\}$ allocated to the mission space. Each sensor $\mathcal{P}_i$ is capable of sensing activity in (monitoring) possibly multiple sectors simultaneously based on its chosen location. The set of possible monitoring choices for sensor $\mathcal{P}_i$ is denoted as $\mathcal{A}_i \subseteq 2^X$. We will refer to $\mathcal{A}_i$ as the action set of player $\mathcal{P}_i$. Similarly, let $\mathcal{A} := \prod_{\mathcal{P}_i} \mathcal{A}_i$ be the set of joint actions, or monitoring choices, for all players. The probability of sensor $\mathcal{P}_i$ detecting an event in sector x given his current monitoring choice a_i is denoted as $p_i(x, a_i)$. We will assume that the detection probabilities satisfy:

$$\begin{aligned} x \in a_i \quad &\Leftrightarrow \quad p_i(x, a_i) > 0, \\ x \notin a_i \quad &\Leftrightarrow \quad p_i(x, a_i) = 0, \end{aligned}$$

This sensor model is quite general and can accommodate a variety of settings, such as a sensing capability that degrades exponentially in distance, obstacles, etc.

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

$$P(x, a) = 1 - \prod_{\mathcal{P}_i \in \mathcal{P}} [1 - p_i(x, a_i)].$$

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 social welfare function

$$\phi(a) = \sum_{x \in X} R(x) P(x, a). \tag{1}$$

Computing the optimal sensor allocation is an NP-hard combinatorial optimization problem – [19] shows this for a structurally equivalent version of the weapon targeting problem. Resultantly, research has traditionally centered around developing heuristic methods to quickly obtain near optimal allocations, where the degree of suboptimality is dependent on the structure of the global objective, e.g., [3].

2.2 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. To model this situation, we will view the interactions of the sensors as a *non-cooperative* game where each sensor $\mathcal{P}_i$ is assigned a utility function $U_i : \mathcal{A} \rightarrow \mathbb{R}$ that defines his payoff (utility) for each monitoring profile.[1] Note that a sensor's utility could be adversely affected by the monitoring choices of other sensors.

We will refer to the non-cooperative game theoretic formulation of the sensor coverage problem as the *sensor coverage game*. The engineering question in this setting is how to design the sensor utility functions, or equivalently, how should a global planner distribute the social welfare to the sensors. A first thought would be to distribute the global social welfare equally to each agent, which would lead to the utility function

$$U_i(a_i, a_{-i}) = \frac{1}{n} \sum_{x \in X} R(x) P(x, a). \tag{2}$$

One potential problem with this design is that each player must be aware of the complete action profile. In the case of a large sensors networks, this informational demand is quite restrictive. Consequently, *local* utility functions are desirable. In a local design, each player's utility depends only on players monitoring similar sectors. Two examples of local utility functions are

$$U_i(a_i, a_{-i}) = \sum_{x \in a_i} \left(\frac{1}{\sum_j I\{x \in a_j\}} \right) P(x, a) R(x), \tag{3}$$

where $I\{\cdot\}$ is the usual indicator function, and

$$U_i(a_i, a_{-i}) = \sum_{x \in a_i} \left(\frac{p_i(x, a)}{\sum_j p_j(x, a)} \right) P(x, a) R(x). \tag{4}$$

In both cases, $(\cdot)$ captures how the effective reward of cell x, $P(x, a)R(x)$, is distributed amongst the sensors that are monitoring the sector. We will refer to (3) as *equally shared utility*, since the effective reward is distributed evenly across all the sensors monitoring the sector, and (4) as *proportional share utility design*, since the effective reward is distributed in proportion to the detection probability.

We will focus on equilibrium behavior in the sensor coverage game. A well-known equilibrium concept that emerges in non-cooperative games, such as the sensor coverage game, is that of a pure Nash equilibrium.

Definition 2.1 (Pure Nash Equilibrium) An action profile $a^* \in \mathcal{A}$ is called a *pure Nash equilibrium* if for all players $\mathcal{P}_i \in \mathcal{P}$,

$$U_i(a_i^*, a_{-i}^*) = \max_{a_i \in \mathcal{A}_i} U_i(a_i, a_{-i}^*). \tag{5}$$

A Nash equilibrium can be thought of as an "agreeable" allocation for the players as it represents a scenario for which no player has an incentive to unilaterally deviate. We will henceforth refer to a pure Nash equilibrium as simply an equilibrium. For a more comprehensive review of the game-theoretic concepts introduced in this section, we refer the readers to [7, 24, 25, 20].

[1]The utility $U_i(a_i, a_{-i})$ defines the payoff sensor $\mathcal{P}_i$ receives for monitoring sectors a_i given that all other sensors are monitoring sectors according to $a_{-i} := \{a_1, ..., a_{i-1}, a_{i+1}, ..., a_n\}$. We will occasionally abuse notation and write $U_i(a_i, a_{-i})$ as just $U_i(a)$.

3 Defining distributed welfare games

The structure and design issues in the sensor coverage game are very similar to what is found in a broad range of resource allocation problems. These similarities motivate us to consider a broad class of games that we term *distributed welfare games*.

A distributed welfare game consists of a set of players $\mathcal{P} := \{\mathcal{P}_1, ..., \mathcal{P}_n\}$ and a finite set of resources X that are to be shared by the players. The action set of player $\mathcal{P}_i$ is denoted as $\mathcal{A}_i \subseteq 2^X$, meaning that a player is potentially capable of selecting multiple resources. Furthermore, there is a social welfare function, $\phi : \mathcal{A} \to \mathbb{R}$, that is *separable*, i.e., $\phi(a) = \sum_{x \in X} \phi_x(a)$, where $\phi_x(a)$ is the effective reward from resource x for the action profile a and $\mathcal{A} := \prod_i \mathcal{A}_i$. Note that in the sensor coverage game we have $\phi_x(a) = P(x, a)R(x)$.

A distributed welfare game enforces the following structure on players' utility functions

$$U_i(a_i, a_{-i}) = \sum_{x \in a_i} g_i(x, a)\, \phi_x(a), \tag{6}$$

where $\{g_1(x, a), ..., g_n(x, a)\}$ defines how the social reward garnered from resource x is distributed across the players. This *distribution rule* is the key component of distributed welfare games, and it must satisfy the following natural properties. For any player $\mathcal{P}_i \in \mathcal{P}$, resource $x \in X$, and action profile $a \in \mathcal{A}$

(i) $g_i(x, a) \geq 0$,

(ii) $x \notin a_i \;\Rightarrow g_i(x, a) = 0$,

(iii) $\sum_i g_i(x, a) \leq 1$.

The focus of this paper is understanding what distribution rules guarantee the existence and efficiency of equilibria. There are two main differentiating features of distribution rules.

A first distinction between distribution rules relates to the informational requirement. A distribution rule can be either *omniscient*, i.e., depend on the structure of the social welfare function and how it is affected by other players, or *trait-based*, i.e., independent of the structure of the social welfare function but (possibly) dependent on characteristics, i.e., traits of the players. Note that from a distributed design perspective it is often required (or desirable) that distribution rules be trait-based rather than omniscient.

A second differentiating feature of distribution rules is whether the social welfare is completely distributed to the participating players. Property (iii) says that the reward distributed to the players at resource x is *at most* the reward gathered at resource x. While this is natural, in many cases it is required (or desirable) to have (iii) be satisfied with equality, so that *all* of the reward from each resource is distributed among the players at the resource. Thus, the global planner does not profit from the players. We will refer to distribution rules that satisfy (iii) with equality as *complete distribution rules*. The equally shared utility in (3) is an example of a complete trait-based distribution rule.

4 Equilibrium existence

Given the definition of distributed welfare games, the natural question is *"does there exist a distribution rule that guarantees the existence of an equilibrium?"* We will begin to answer this question in this section. We start with some motivating examples and then consider the cases of single-selection distributed welfare games and multi-selection anonymous distributed welfare games.

4.1 Exploring distributed welfare games

To elucidate the challenges in designing such a distribution rule, consider the following simple example of a two player sensor coverage game.

Example 4.1 (Equally shared utilities) Consider a sensor coverage game with two sectors, $X = \{x_1, x_2\}$, rewards $R(x_1) = R(x_2) = 10$, and two sensors $\mathcal{P} = (\mathcal{P}_1, \mathcal{P}_2)$ that can monitor either sector, but not both, i.e., $\mathcal{A}_1 = \mathcal{A}_2 = X$, with invariant detection probabilities[2] $p_1 = 1$ and $p_2 = 0.1$. Suppose the sensors are assigned the equally shared utility in (3). This setup yields the following payoff matrix, which illustrates that no equilibrium exists.[3]

		Player 2 x_1	*Player 2* x_2
Player 1	x_1	$5, 5$	$10, 1$
Player 1	x_2	$10, 1$	*5,5*

The problem with this design is that players' utility functions were not always "aligned" with their contribution to the global objective. An alternative rule that remedies this is the *wonderful life utility design* [23]. In this design, each player's utility is set as his marginal contribution to the global objective, i.e.,

$$U_i(a_i, a_{-i}) = \phi(a_i, a_{-i}) - \phi(a_i^0, a_{-i}), \tag{7}$$

where a_i^0 designate the null action for player $\mathcal{P}_i$. This form of the utility design results in the distribution rule

$$g_i(x, a) = \frac{\phi_x(a_i, a_{-i}) - \phi_x(a_i^0, a_{-i})}{\phi_x(a)}. \tag{8}$$

It is almost immediate to see that any action profile that maximizes social welfare is an equilibrium, hence the name "wonderful life."[4] However, there are two limitations of the wonderful life utility design that often prevent it from being applicable.

1. The wonderful life utility may distribute more (or less) reward than is gathered, i.e., it does not always satisfy property (iii) in the definition of distributed welfare games.

2. The wonderful life utility is an *omniscient* distribution rule.

Example 4.2 (Wonderful life utility) Suppose there are two resources $\{x_1, x_2\}$ and two players $\mathcal{P} = (\mathcal{P}_1, \mathcal{P}_2)$ that can select either resource, but not both, i.e., $\mathcal{A}_1 = \mathcal{A}_2 = \{x_1, x_2\}$. Consider the following separable social welfare function:

		Player 2 $\emptyset$	*Player 2* x_1	*Player 2* x_2
Player 1	$\emptyset$	*0*	*4*	*1*
Player 1	x_1	*6*	*6*	*7*
Player 1	x_2	*5*	*9*	*10*

[2] An invariant detection probability means that for each player $\mathcal{P}_i$, $p_i(x, a_i) = p_i I\{x = a_i\}$.

[3] The payoff matrix designates the utility garnered for each player under all possible action profiles. The first entry in a box is the payoff for player $\mathcal{P}_1$ and the second entry in the box in the payoff for player $\mathcal{P}_2$. For example, $U_1(a_1 = x_1, a_2 = x_2) = 10$ and $U_2(a_1 = x_1, a_2 = x_2) = 1$.

[4] It is important to remember that other equilibria might also exist under the wonderful life utility design.

Suppose that the effective reward from a particular resource is distributed according to the wonderful life utility design. The payoff matrix is

		Player 2 $\emptyset$	x_1	x_2
	$\emptyset$	*0,0*	*4,0*	*1,0*
Player 1	x_1	*6,0*	*2,0*	*6,1*
	x_2	*5,0*	*5,4*	*9, 5*

Clearly, the action profile (x_2, x_2) maximizes the social welfare and is also an equilibrium. However, at (x_2, x_2) more reward is distributed than is gathered. Further, at (x_1, x_1) less reward is distributed than is gathered.

Given this example, it is natural to wonder whether the wonderful life utility design can be adjusted so that it is a complete distribution rule. A first suggestion is the following distribution rule, which disperses the effective reward proportionally to each player's marginal contribution to the social welfare, i.e.,

$$g_i(x, a) = \frac{\phi_x(a) - \phi_x(a_i^0, a_{-i})}{\sum_{\mathcal{P}_j} \phi_x(a) - \phi_x(a_j^0, a_{-j})}. \tag{9}$$

This form of utility design will be referred to as the *marginal share utility design*. We will demonstrate in Section 6 that the marginal share utility design guarantees the existence of an equilibrium in the single-sector sensor coverage game when players are only allowed to monitor a single sector. However, the following example illustrates that the marginal share utility will not provide this guarantee in general.

Example 4.3 (Marginal share utility design) Consider the setup in Example 4.2 with a distribution rule as in (9). This leads to the following payoff matrix:

		Player 2 $\emptyset$	A	B
	$\emptyset$	*0,0*	*4,0*	*1,0*
Player 1	A	*6,0*	*6,0*	*6,1*
	B	*5,0*	*5,4*	*6.4, 3.6*

4.2 Single-selection distributed welfare games

The previous section illustrates that the design of distribution rules is delicate. In this section, our goal is to characterize the set of distribution rules that guarantee the existence of an equilibrium. Our main result is to provide three sufficient conditions that ensure the existence of an equilibrium.

Before stating the conditions, we will introduce the following notation. The action profile (x, y, a_{-ij}), denotes the situation where player $\mathcal{P}_i$ selects resource x and player $\mathcal{P}_j$ selected resource y. Likewise, $U_i(x, y, a_{-ij})$ and $U_j(x, y, a_{-ij})$ denote the utility player $\mathcal{P}_i$ and $\mathcal{P}_j$ receives when player $\mathcal{P}_i$ selects resource x and player $\mathcal{P}_j$ selects resource y. Let $\mathcal{A}_{-ij} := \prod_{k \neq i,j} \mathcal{A}_k$.

We will now state three sufficient conditions that guarantee the existence of an equilibrium in distributed welfare games. Without loss of generality, we will state the conditions with respect to players' utility functions as opposed to distribution rules. The first condition deals with consistency of players' strengths.

Condition 4.1 (Consistency of Strengths) Let $\mathcal{P}_i$ and $\mathcal{P}_j$ be any two players. If

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

for any resource $x \in X$ and profile $a_{-ij} \in \mathcal{A}_{-ij}$ *then*

$$U_i(\tilde{x}, \tilde{x}, \tilde{a}_{-ij}) \geq U_j(\tilde{x}, \tilde{x}, \tilde{a}_{-ij}), \forall \tilde{x} \in X, a_{-ij} \in \mathcal{A}_{-ij}.$$

If $U_i(x, x, a_{-ij}) \geq U_j(x, x, a_{-ij})$ for all $x \in X$ and $a_{ij} \in \mathcal{A}_{-ij}$, then we will say that player $\mathcal{P}_i$ is "stronger" than player $\mathcal{P}_j$. Similarly, we will say that player $\mathcal{P}_j$ is "weaker" than player $\mathcal{P}_i$.

The second condition states that a stronger player will not exploit a weaker player.

Condition 4.2 (No Exploitation) Suppose $\mathcal{P}_i$ is stronger than player $\mathcal{P}_j$. For any resource $x \in X$ and any action profile $a_{-ij} \in \mathcal{A}_{-ij}$,

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

The third condition relates to the utility relationship between two players utilizing different resources.

Condition 4.3 (Relative Strengths) Suppose $\mathcal{P}_i$ is stronger than player $\mathcal{P}_j$. Let

$$R^x_{i \to j} := \max_{x \in X, a_{-ij} \in \mathcal{A}_{-ij}} \frac{U_j(x, x, a_{-ij})}{U_i(x, x, a_{-ij})},$$

For any resources $x \in X$ and action profile $a_{-12} \in \mathcal{A}_{-12}$, the following holds

$$U_j(a_i^0, x, a_{-ij}) \geq R^x_{i \to j}\, U_i(x, a_j^0, a_{-ij}).$$

The term $R^x_{i \to j}$ can be thought of as the true utility relationship between players $\mathcal{P}_i$ and $\mathcal{P}_j$. With this view, this condition states that the weaker player, $\mathcal{P}_j$, has a greater benefit to his utility when the stronger player, $\mathcal{P}_i$, is not utilizing the same resource.

Theorem 4.1 Consider any single-selection distributed welfare game. If the players' utility functions satisfy Conditions 4.1, 4.2, and 4.3 then an equilibrium exists.

Proof: We begin by renumbering the players in order of strength with $\mathcal{P}_1$ being the strongest player. This is possible because of Condition 4.1.

We will now 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 actions. First, player $\mathcal{P}_1$ selects the action x according to

$$x \in \arg\max_{\tilde{x} \in X} U_1(\tilde{x}, a_{-1}^0) \tag{10}$$

Next, player $\mathcal{P}_2$ selects action y according to

$$y \in \arg\max_{\tilde{y} \in X} U_2(x, \tilde{y}, a_{-12}^0).$$

If $x \neq y$, then by (10) and Condition 4.2 we know that

$$U_1(x, a_2^0, a_{-12}^0) \geq U_1(y, a_2^0, a_{-12}^0) \geq U_1(y, y, a_{-12}^0).$$

Therefore, player $\mathcal{P}_1$ can not improve his utility by switching his strategy, i.e.,

$$U_1(x, y, a^0_{-12}) \geq U_1(z, y, a^0_{-12}),\ \forall z \in X.$$

If $x = y$, then by Condition 4.3, we know that

$$U_2(a^0_1, x, a^0_{-12}) \quad \geq \quad \frac{U_2(x, x, a^0_{-12})}{U_1(x, x, a^0_{-12})}\, U_1(x, a^0_2, a^0_{-12}).$$

Using the above inequality, we can conclude that $U_2(x, x, a^0_{-12}) \geq U_2(x, y, a^0_{-12})$ implies $U_1(x, x, a^0_{-12}) \geq U_1(y, x, a^0_{-12})$. Therefore, player $\mathcal{P}_1$ can not improve his utility by switching his strategy.

If $n = 2$, then $a = (x, y)$ would be an equilibrium. Otherwise this argument could be repeated n times to construct an equilibrium.

□

4.3 Distributed welfare games with anonymous players

We now move to a second important subclass of distributed welfare games – distributed welfare games with anonymous players. In such games, players are anonymous with regard their impact on the global objective and the distribution rule.

In order to formally state this condition, we need some notation. For any resource $x \in X$ and action profile $a \in \mathcal{A}$, let $\sigma_x(a)$ denote the number of players utilizing resource x, i.e., $\sigma_x(a) := |\{\mathcal{P}_i \in \mathcal{P} : x \in a_i\}|$. For a distributed welfare game with anonymous players, the global objective satisfies $\phi_x(a) = \phi_x(a')$ for any resource $x \in X$ and action profiles $a, a' \in \mathcal{A}$ such that $\sigma_x(a) = \sigma_x(a')$.

The natural distribution rule in the case of anonymous players is the equal share utility design (3), and this rule turns out to guarantee the existence of an equilibrium.

Proposition 4.1 If a distributed welfare game has anonymous players then an equilibrium exists under the equal share utility design (3).

To prove this proposition, we show that anonymous distributed welfare games are congestion games [21, 18], which are guaranteed to have an equilibrium.

Congestion games are a specific class of games in which player utility functions have a special structure. In order to define a congestion game, we must specify the action set, $\mathcal{A}_i$, and utility function, $U_i(\cdot)$, of each player. Towards this end, let X denote a finite set of resources. For each resource $x \in X$, there is an associated *congestion function*, $c_x : \{0, 1, 2, ...\} \to \mathbb{R}$, that reflects the cost of using the resource as a function of the number of players using that resource. The action set, $\mathcal{A}_i$, of each player, $\mathcal{P}_i$, is defined as the set of resources available to player $\mathcal{P}_i$, i.e., $\mathcal{A}_i \subset 2^X$. Accordingly, an action, $a_i \in \mathcal{A}_i$, reflects a selection of (multiple) resources, $a_i \subset X$. A player is "using" resource x if $x \in a_i$. For an action profile $a \in \mathcal{A}$, the utility of player $\mathcal{P}_i$ using resources indicated by a_i depends only on the total number of players using the same resources. More precisely, the utility of player $\mathcal{P}_i$ is defined as

$$U_i(a) = \sum_{r \in a_i} c_r(\sigma_r(a)). \tag{11}$$

To see that an anonymous distributed welfare game can be expressed as a congestion game when the equal share utility design is used, we take the following representation:

(a) resources X,

(b) cost functions $c_x(k) = \frac{\phi_x(k)}{k},\ k > 0$, where k is the number of users of resource x, and

(c) utility functions $U_i(a) = \sum_{x \in a_i} c_x(\sigma_x(a))$,

By recalling that a pure Nash equilibrium always exists in congestion games, we obtain Proposition 4.1.

5 The efficiency of equilibria

Given a guarantee of the existence of an equilibrium, the next natural question is *"how efficient are the equilibria?"* The most common measure used in the literature is the *price of anarchy*, e.g., Chapters 17-21 in [20]. The price of anarchy (or more appropriately in this case, the "price of localization") is defined as the worst case ratio between the social welfare at any equilibrium and the optimal social welfare. Formally, if the price of anarchy is $\gamma \leq 1$, then for any equilibrium a^{ne} and optimal allocation a^{opt} satisfy $\phi(a^{ne}) \geq \gamma\, \phi(a^{opt})$.

Unfortunately, without any assumptions on the social utility function ϕ, the price of anarchy can be arbitrarily bad in distributed welfare games. However, with two, fairly weak, practical conditions on the interaction between ϕ and the utility functions, we can guarantee that the price of anarchy is close to 1.

The first condition is on the structure of the social welfare function.

Condition 5.1 (Submodularity) The social welfare function ϕ is *submodular*. A function $\phi : 2^{\mathcal{A}} \to \mathbb{R}$ is submodular if $\phi(X) + \phi(Y) \geq \phi(X \cap Y) + \phi(X \cup Y)$ for all $X, Y \subseteq 2^{\mathcal{A}}$.

Submodularity corresponds to the notion of a decreasing marginal utility and is a common assumption in many resource allocation problems, e.g., [22, 12]. Further, it is a key property underlying the design of many centralized algorithms for these problems.

The second condition is a a weak form of alignment between the social welfare and player utility functions.

Condition 5.2 (Alignment) A player's utility is at least equal to his marginal contribution to the global objective, i.e., $U_i(a) \geq \phi(a) - \phi(a_i^0, a_{-i}),\ \ \forall \mathcal{P}_i \in \mathcal{P},\ a \in \mathcal{A}$.

The power of these conditions is the following proposition, which provides an efficiency guarantee for all equilibrium assignments regardless of the heterogeneity of the players and resources.

Proposition 5.1 Consider a distributed welfare game that obeys Conditions 5.1 and 5.2. If an equilibrium exists, then the price of anarchy is 1/2.

This proposition follows from noting that Conditions 5.1 and 5.2 are two of the three conditions that define the class of utility games used by Vetta in [22]. Further, the third and final condition in the definition of utility games follows immediately from the definition of distributed welfare games. Thus, we can apply Theorem 3.4 from [22] to obtain Proposition 5.1.

Though Proposition 5.1 is powerful, it is only useful when an equilibrium is known to exist. Thus, its applicability depends on the results we have proven in Section 4. Further, this bound does not show how the efficiency scales with the number of players or the parameters of the social welfare function. Using a more detailed analysis, we can attain such results.

Theorem 5.1 Consider any single-selection distributed welfare game with n anonymous players satisfying Conditions 5.1 and 5.2. If an equilibrium exists, then the price of anarchy is $\frac{n}{2n-1}$.

Notice that Theorem 5.1 shows that the worst-case price of anarchy is increasing as the number of players increases. In fact, as $n \to \infty$ the price of anarchy approaches $1/2$, which matches Proposition 5.1. We can also see that this bound is tight:

Example 5.1 (Tightness of the price of anarchy) Consider a single-sector sensor coverage game with the equal share utility design (3). There are n sensors each having invariant detection probability $p = 1$ and n sectors. One sector, x, has reward 1 and the other sectors have reward $1/n$.

Notice that if all sensors cover x, then the total reward garnered is 1 and the assignment is an equilibrium. However, the optimal assignment has one sensor in each sector and achieves reward $1 + (n-1)/n$. Thus, the price of anarchy is $n/(2n-1)$.

In order to prove Theorem 5.1, we need the following lemma.

Lemma 5.1 Consider a single-selection distributed welfare game with anonymous players that obeys Conditions 5.1 and 5.2. Assume that an equilibrium a^{ne} exists. If $\sigma_x(a^{ne}) > 0$ for any resource $x \in X$, then there exists an optimal allocation a^{opt} such that $\sigma_x(a^{opt}) > 0$.

Proof: Before beginning the proof we define the following function which will also be used in subsequent proofs: for all $k > 0$, let

$$U_x(k) \quad := \quad \frac{1}{k}\,\phi_x(a), \tag{12}$$

where a is any action profile where $\sigma_x(a) = k$. With somewhat an abuse of notation, we will sometimes write $\phi_x(a) = \phi_x(\sigma_x(a))$. Using this notation, $U_i(a) = U_x(\sigma_x(a))$ for any player utilizing resource x.

Let a^{opt} and a^{ne} be an optimal and Nash allocation respectively. Suppose $\sigma_x(a^{opt}) = 0$ and $\sigma_x(a^{ne}) > 0$. Then, there exists a sector y such that $\sigma_y(a^{opt}) > \sigma_y(a^{ne})$. Since a^{ne} is a Nash equilibrium we know that

$$\begin{aligned} U_x(1) &\geq U_x(\sigma_x(a^{ne})), \\ &\geq U_y(\sigma_y(a^{ne})+1), \\ &\geq \phi_y(\sigma_y(a^{ne})+1) - \phi_y(\sigma_y(a^{ne})), \end{aligned}$$

where the first inequality comes from condition (1) and the last inequality comes from condition (2). Since $\sigma_y(a^{opt}) > \sigma_y(a^{ne})$ and a^{opt} is an optimal allocation, we know that

$$\begin{aligned} \phi_y(\sigma_y(a^{ne})+1) - \phi_y(\sigma_y(a^{ne})) &\geq \phi_y(\sigma_y(a^{opt})) - \phi_y(\sigma_y(a^{opt})-1), \\ &\geq \phi_x(1), \\ &= U_x(1). \end{aligned}$$

Therefore, combining the two previous set of inequalities we obtain that

$$\phi_y(\sigma_y(a^{opt})) - \phi_y(\sigma_y(a^{opt})-1) = \phi_x(1),$$

meaning that the marginal contribution of a player utilizing resource y for the allocation a^{opt} is equivalent to the marginal contribution if that same player was utilizing x. Consider the following profile a^{opt^*} where

$$\begin{aligned}
\sigma_x(a^{opt^*}) &= 1, \\
\sigma_y(a^{opt^*}) &= \sigma_y(a^{opt^*}) - 1, \\
\sigma_z(a^{opt^*}) &= \sigma_z(a^{opt^*}), \ \ \forall z \neq x, y.
\end{aligned}$$

Then $\phi(a^{opt^*}) = \phi(a^{opt})$; hence a^{opt^*} is also an optimal action profile. This completes the proof.
□

Proof: *(of Theorem 5.1)* Let a^{ne} be any equilibrium and a^{opt} be any optimal allocation satisfying Lemma 5.1. The general idea of this proof is to look at a new action profile $a^{cover} \in (2^X)^n$ which satisfies

$$\sigma_x(a^{cover}) = \max(\sigma_x(a^{opt}), \sigma_x(a^{ne}))$$

for all $x \in X$. From Lemma 5.1, we know that we can cover both a^{opt} and a^{ne} using

$$n + \sum_{x \in X : \sigma_x(a^{ne}) > 0} (\sigma_x(a^{ne}) - 1)$$

players. We accomplish this by placing the n players according to a^{ne} and then placing the remaining players according to a^{opt}, noting that at least one player utilizing each resource $x \in X : \sigma_x(a^{ne}) > 0$ is already accounted for. Given this definition and the submodularity of the social welfare function, we have

$$\phi(a^{ne}) \leq \phi(a^{opt}) \leq \phi(a^{cover}).$$

Further, we have that

$$\begin{aligned}
&\phi(a^{cover}) \\
&\quad \leq \phi(a^{ne}) + \sum_{x \in X : \sigma_x(a^{ne}) > 0} (\sigma_x(a^{ne}) - 1) \frac{\phi_x(a^{ne})}{\sigma_x(a^{ne})}, \\
&\quad = \phi(a^{ne}) + \sum_{x \in X : \sigma_x(a^{ne}) > 0} \left(\frac{\sigma_x(a^{ne}) - 1}{\sigma_x(a^{ne})} \right) \phi_x(a^{ne}), \\
&\quad \leq \phi(a^{ne}) + \left(\frac{n-1}{n} \right) \sum_{x \in X : \sigma_x(a^{ne}) > 0} \phi_x(a^{ne}), \\
&\quad = \phi(a^{ne}) \left(2 - \frac{1}{n} \right).
\end{aligned}$$

$$\frac{\phi(a^{ne})}{\phi(a^{opt})} \geq \frac{\phi(a^{ne})}{\phi(a^{cover})} = \frac{n}{2n-1}$$

□

6 Implications for sensor coverage

We will now return to the sensor coverage game to see the implications of our general results.

6.1 Utility designs that guarantee equilibrium existence

The sensor coverage game with single-sector sensors is an example of a single-selection distributed welfare game. In Section 4, we derived three sufficient conditions that guarantee the existence of an equilibrium in these games (Theorem 4.1). So, the task now is to design distribution rules for the sensor coverage game that satisfy these three sufficient conditions. It turns out that Conditions 4.1 – 4.3 are satisfied by two natural distribution rules. Refer to the appendix for the verification of the conditions.

Corollary 6.1 The marginal share utility design (9) and the proportional share utility design (4) both satisfy Conditions 4.1 – 4.3. Thus, they both guarantee the existence of an equilibrium in the single-sector sensor coverage game.

The fact that these two, intuitive, local utility designs both guarantee the existence of an equilibrium is more surprising than might appear at first glance. For instance, Example 4.3 illustrates that the marginal contribution sharing rule does not always guarantee the existence of an equilibrium. To highlight the sensitivity of these rules, consider the following example:

Example 6.1 (Sensitivity of local utility design) *Consider the following variation of the proportional share utility design:*

$$g_i(x, a) = \frac{p_i^\beta}{\sum_{\mathcal{P}_j : a_j = x} p_j^\beta}.$$

Note that when, $\beta = 1$, this is equivalent to the proportional share utility design. When $\beta \to 0$, this rule becomes the equal share utility design, and when $\beta \to \infty$, this rule gives the entire reward to the sensor with the highest detection probability. So, we have already seen that an equilibrium exists when $\beta = 1$ and doesn't exist when $\beta = 0$ or $\beta = \infty$. We will now see that when $\beta \neq 1$, an equilibrium may not exist.

Consider 2 sectors x and y and 2 sensors 1 and 2. The rewards are $R_x = R_y = 1$ and the detection probabilities are $p_1 > p_2$. Now, we are going to set a scenario where if $\beta > 1$ sensor 1 chases sensor 2 and if $\beta < 1$ sensor 2 chases sensor 1.

Note that the utilities to each sensor when both are in the same sector are

$$U_i(x, x) = U_i(y, y) = \frac{p_i^\beta}{p_1^\beta + p_2^\beta}(1 - (1 - p_1)(1 - p_2))$$

If the sensors are in different sectors, the utilities are

$$U_i(x, y) = U_i(x, y) = p_i$$

Now, consider player 1. The difference in utility between being together and being apart is equal to:

$$\begin{aligned} U_1(x, x) - U_1(y, x) &= \frac{p_1^\beta}{p_1^\beta + p_2^\beta}(1 - (1 - p_1)(1 - p_2)) - p_1 \\ &= \frac{p_2 p_1^\beta}{p_1^\beta + p_2^\beta}\left(1 - p_1 - \left(\frac{p_1}{p_2}\right)^{1-\beta}\right) \end{aligned}$$

Similarly, for player 2, we have the difference in utility between being together and being apart is equal to:

$$U_2(x,x) - U_2(y,x) = \frac{p_1 p_2^\beta}{p_1^\beta + p_2^\beta}\left(1 - p_2 - \left(\frac{p_2}{p_1}\right)^{1-\beta}\right)$$

Notice that if $\beta = 1$ both players prefer to be apart regardless of p_1, p_2. However, if $\beta \neq 1$, the sensors exhibit chasing behavior. In particular, set $p_1 = 0.5$ and $p_2 = \epsilon > 0$. Then, by letting $\epsilon \to 0$, we can observe chasing behavior. In particular, as $\epsilon \to 0$, when $\beta < 1$ we see that sensor 1 prefers to be away from sensor 2 and sensor 2 prefers to be with sensor 1, and when $\beta > 1$ the sensors have the opposite preferences.

Given the sensitivity of the local utility designs in the single-sector case, it is clear that designing a distribution rule that will guarantee existence of an equilibrium in the general case of heterogeneous sensors is a difficult open question. However, by applying Proposition 4.1 for anonymous multi-selection distributed welfare games we attain the following result. In this setting, anonymous sensors means that the sensors all have the same invariant detection probability, i.e., *homogenous* as opposed to *heterogenous* sensors.

Corollary 6.2 In the multi-sector homogenous sensor coverage game, the equal share utility design (3) guarantees the existence of an equilibrium.

6.2 Bounding the efficiency of equilibria

Recall that in our discussion of the efficiency of distributed welfare games, we placed two extra conditions on distributed welfare games that allowed us to attain bounds on the price of anarchy. It is immediate to observe that ϕ is submodular in the sensor coverage game (Condition 5.1). We refer the reader to the appendix for a verification of Condition 5.2.

Corollary 6.3

(i) In a single-sector or multi-sector sensor coverage game, under the proportional share utility design (4) and the marginal share utility designs (9), the price of anarchy is 1/2.[5]

(ii) In a single-sector homogeneous sensor coverage game with n sensors, under equal share utility design (3), the price of anarchy is $\frac{n}{2n-1}$.

Corollary 6.3 does not address how the sensor detection probabilities affects the price of anarchy. The following theorem seeks to rectify this.

Theorem 6.1 Consider a single-sector homogeneous sensor coverage game with n sensors each having invariant detection probability p. Under the equal share utility design (3) the price of anarchy is bounded by

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

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

[5]Note that we have not guaranteed the existence of an equilibrium in the multi-sector heterogeneous sensor coverage game under either the proportional share of the marginal share utility designs, but that both rules guarantee existence of an equilibrium in the multi-sector homogenous sensor coverage game.

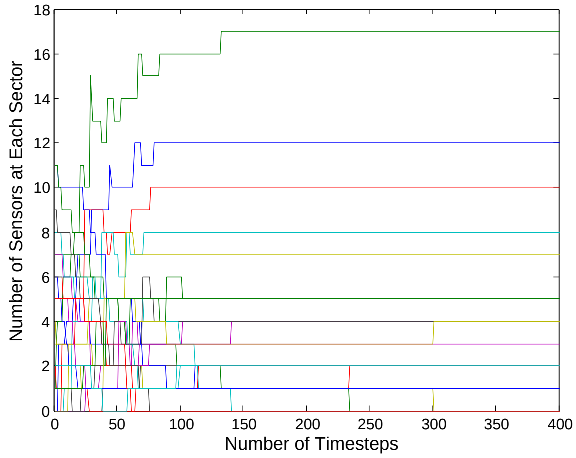

Figure 1: Evolution of Number of Sensors in Each Sector over Mission

Note that this bound is an increasing function of the detection probability of the sensors, which is intuitive since Example 5.1 relies on having $p(x, a) = 1$ while if $p(x, a) = 0$ the price of anarchy is trivially 1.

6.3 Simulation experiments

To this point, we have explored equilibrium behavior in the sensor coverage game. The question that remains is how can autonomous sensors reach an equilibrium in a distributed fashion? While not focusing on this question in detail, we will illustrate the potential of the theory of learning in games as a local control mechanism for coordinating group behavior.

We consider a sensor coverage game with 100 single-sector homogenous sensors with invariant detection probability $p = 0.25$. The mission space is $X = \{x_1, ..., x_{25}\}$. The reward 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., $\mathcal{A}_i = X$ and uses the equal share utility design (3).

There is a large body of literature analyzing distributed learning algorithms in congestion games, or equivalently potential games [24, 17, 16, 15, 14]. We will apply *fading memory 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 [16] for the details of the learning rule. We use the following discount factor and inertia: $\lambda = 0.5$ and $\epsilon = 0.02$.

Figure 1 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 2 illustrates the evolution of the social welfare in addition to the efficiency gap between the equilibrium and the optimal. From Theorem 6.1, we know that the price of anarchy must be greater than 0.541. Our simulation illustrates that Theorem 6.1 provides a very conservative estimate of the efficiency of the attained equilibrium since the price of anarchy we observe is 0.936.

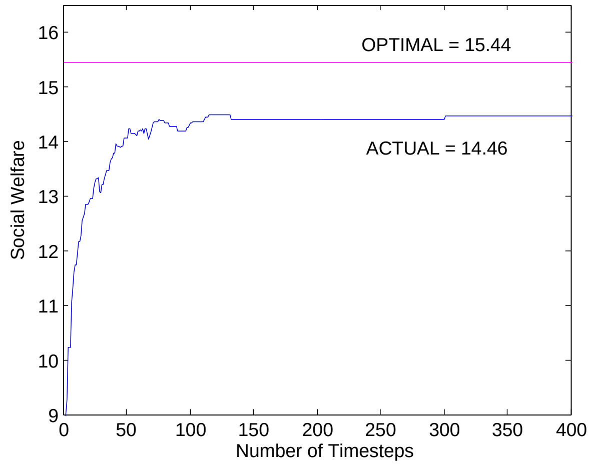

Figure 2: Evolution of Welfare Function over Mission

7 Conclusions and open questions

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 social 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. It should be noted that the structure of the social welfare function ϕ for sensor coverage parallels those for many other problems, e.g., weapon targeting, fault detection, and surveillance.

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] A. Ageev and M. Sviridenko. Pipage rounding: a new method for constructing algorithms with proven performance guarantee. *Journal of Combinatorial Optimization*, 2004.

[2] R. Ahuja, T. Magnanti, and J. Orlin. *Network flows: theory, algorithms, and applications*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1993.

[3] R. K. Ahuja, A. Kumar, K. Jha, and J. B. Orlin. Exact and heuristic methods for the weapon-target assignment problem. http://ssrn.com/abstract=489802, 2003.

[4] 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*, 129:584–596, September 2007.

[5] S.S. Dhillon, K. Chakrabarty, and S.S. Iyengar. Sensor placement for grid coverage under imprecise detections. In *Proc. of Conf. on Information Fusion*, pages 1581–1587, 2004.

[6] 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)*, pages 667–676, 2006.

[7] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, Cambridge, MA, 1991.

[8] D. Ghose, M. Krichman, J.L. Speyer, and J.S. Shamma. Modeling and analysis of air campaign resource allocation: A spatio-temporal decomposition approach. *IEEE Transactions on Systems, Man, and Cybernetics*, **32**(3):403–418, 2002.

[9] T. Ibaraki and N. Katoh. *Resource allocation problems: algorithmic approaches*. MIT Press, Cambridge, MA, USA, 1988.

[10] S.S. Iyengar and R.R. Brooks. *Distributed sensor networks*. Chapman & Hall, Boca Raton, FL, USA, 2005.

[11] R. Johari and J. N. Tsitsiklis. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research*, 29(3):407–435, August 2004.

[12] A. Krause and C. Guestrin. Near-optimal obervation selection using submodular functions. In *Proc. of Conf. on Artifical Intelligence*, 2007.

[13] W. Li. and C. G. Cassandras. Sensor networks and cooperative control. *European Journal of Control*, 2005. to appear.

[14] 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.

[15] J. R. Marden, G. Arslan, and J. S. Shamma. 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.

[16] 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.

[17] D. Monderer and L. Shapley. Fictitious play property for games with identical interests. *Games and Economic Theory*, 68:258–265, 1996.

[18] D. Monderer and L. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.

[19] R.A. Murphey. Target-based weapon target assignment problems. In P.M. Pardalos and L.S. Pitsoulis, editors, *Nonlinear Assignment Problems: Algorithms and Applications*, pages 39–53. Kluwer Academic Publishers, 1999.

[20] N. Nissan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic game theory*. Cambridge University Press, New York, NY, USA, 2007.

[21] R. W. Rosenthal. The network equilibrium problem in integers. *Networks*, 3(1):53–59, 1973.

[22] 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.*, pages 416–425, 2002.

[23] D. Wolpert and K. Tumor. An overview of collective intelligence. In J. M. Bradshaw, editor, *Handbook of Agent Technology*. AAAI Press/MIT Press, 1999.

[24] H. P. Young. *Individual Strategy and Social Structure*. Princeton University Press, Princeton, NJ, 1998.

[25] H. P. Young. *Strategic Learning and its Limits*. Oxford University Press, 2005.

[26] Y. Zou and K. Chakrabarty. Uncertainty-award and coverage-oriented deployment for sensor networks. *J. of Parallel and Dist. Comp.*, 64(7):788–798, 2004.

A Proofs for the sensor coverage example

A.1 Proof of Corollary 6.1

Corollary 6.1 will follow from the following lemmas, which verify Conditions 4.1–4.3 for both the proportional share utility design and the marginal cost ratio utility design in the case of the single-sector sensor coverage game.

Lemma A.1 *The proportional share utility design satisfies Conditions 4.1–4.3 in the single-sector sensor coverage game.*

Proof: We can immediately see that consistency of strengths condition is satisfied (Condition 4.1):

$$\frac{U_1(x,x,a_{-12})}{U_2(x,x,a_{-12})} = \frac{p_1}{p_2}$$

Next, we will prove that the relative strengths condition holds (Condition 4.3). We again start by defining $S(x) := \{\mathcal{P}_i \in S : a_i = x, i > 2\}$.

$$\begin{aligned} \sigma_x &= \sum_{S_j \in S(x)} p_j \\ \pi_x &= \prod_{S_j \in S(x)} (1 - p_j) \end{aligned}$$

Now,

$$\begin{aligned} U_1(y,x,a_{-12})\frac{p_2}{p_1} &= \frac{p_2}{p_1 + \sigma_y}\left(1 - (1-p_1)\pi_y\right) \\ &= U_2(x,y,a_{-12})\left(\frac{p_2+\sigma_y}{p_1+\sigma_y}\right)\left(\frac{1-(1-p_1)\pi_y}{1-(1-p_2)\pi_y}\right) \end{aligned}$$

So, it is sufficient to show that

$$\frac{1-(1-p_1)\pi_y}{1-(1-p_2)\pi_y} \leq \frac{p_1+\sigma_y}{p_2+\sigma_y}$$

which is equivalent to showing that

$$\frac{(p_1-p_2)\pi_y}{1-(1-p_2)\pi_y} \leq \frac{p_1-p_2}{p_2+\sigma_y}$$

After some further algebra, we attain the following equivalent condition

$$(1+\sigma_y)\pi_y \leq 1 \tag{13}$$

Working with the LHS, we see that

$$\begin{aligned}
(1+\sigma_y)\pi_y &= (1+\sum_{S_j \in S(x)} p_j) \prod_{S_j \in S(x)} (1-p_j) \\
&\leq \prod_{S_j \in S(x)} (1+p_j) \prod_{S_j \in S(x)} (1-p_j) \\
&= \prod_{S_j \in S(x)} (1-p_j^2) \\
&\leq 1
\end{aligned}$$

which completes the proof that the relative strength condition is satisfied.

Finally, we will show that the no exploitation condition is satisfied (Condition 4.2). We will adjust our definition of $S(x)$ as follows: $S(x) := \{\mathcal{P}_i \in S : a_i = x, i \neq 2\}$. Similarly, we adjust the definitions of σ_x and π_x. Now, calculating directly, we see

$$\begin{aligned}
U_1(x,x,a_{-12}) &= \frac{p_1}{p_2+\sigma_x}(1-(1-p_2)\pi_x) \\
&= U_1(x,0,a_{-12}) \left(\frac{\sigma_x}{p_2+\sigma_x}\right) \left(\frac{1-(1-p_2)\pi_x}{1-\pi_x}\right)
\end{aligned}$$

So, it is sufficient to show that

$$\frac{1-(1-p_2)\pi}{1-\pi} \leq \frac{p_2+\sigma_x}{\sigma_x}$$

which is equivalent to showing that

$$\frac{p_2\pi}{1-\pi} \leq \frac{p_2}{\sigma_x}$$

After some algebra, we then attain the following equivalent condition

$$(1+\sigma_y)\pi_y \leq 1$$

which we have already argued is satisfied. So, we have verified the exploitation condition.
□

Lemma A.2 *The marginal cost ratio utility design satisfies Conditions 4.1–4.3 in the single-sector sensor coverage game.*

Proof: First, notice that the marginal contribution of a sensor is

$$\phi(x, a_{-i}) - \phi(a_i^0, a_{-i}) = p_i \prod_{j \neq i : a_j = x} (1 - p_j)$$

After some algebra, the marginal contribution ratio rule can be written as

$$g_i(x, a_i) = \frac{\frac{p_i}{1-p_i}}{\sum_{j : a_j = x} \frac{p_j}{1-p_j}}$$

So, we have that

$$\frac{U_1(x, x, a_{-12})}{U_2(x, x, a_{-12})} = \frac{p_1(1-p_2)}{p_2(1-p_1)}$$

This verifies the consistency of strengths condition (Condition 4.1).

Next, we will verify the relative strengths condition (Condition 4.3). First, define $S(y) := \{\mathcal{P}_i \in S : a_i = y, i > 2\}$ and

$$\pi(y) = \prod_{j \in S(y)} (1 - p_j)$$

$$\sigma(y) = \sum_{j \in S(y)} \frac{p_j}{1 - p_j}$$

Now,

$$\begin{aligned} & U_1(y, x, a_{-12}) \frac{p_2(1-p_1)}{p_1(1-p_2)} \\ = \; & \left(\frac{\frac{p_2}{1-p_2}}{\frac{p_1}{1-p_1} + \sigma(y)} \right) (1 - (1 - p_1)\pi(y)) \\ = \; & U_2(x, y, a_{-12}) \left(\frac{\frac{p_2}{1-p_2} + \sigma(y)}{\frac{p_1}{1-p_1} + \sigma(y)} \right) \left(\frac{1 - (1-p_1)\pi(y)}{1 - (1-p_2)\pi(y)} \right) \end{aligned}$$

So, it is sufficient to show that

$$\frac{1 - (1-p_1)\pi(y)}{1 - (1-p_2)\pi(y)} \leq \frac{\frac{p_1}{1-p_1} + \sigma(y)}{\frac{p_2}{1-p_2} + \sigma(y)}$$

which is equivalent to

$$\frac{(p_1 - p_2)\pi(y)}{1 - (1-p_2)\pi(y)} \leq \frac{p_1 - p_2}{(1-p_1)p_2 + (1-p_1)(1-p_2)\sigma(y)}$$

Some algebra now yields the following equivalent condition

$$\pi(y)(1 - p_1 p_2 + (1-p_1)(1-p_2)\sigma(y)) \leq 1$$

So, it is sufficient to have

$$\pi(y)(1+\sigma(y)) \leq 1 \tag{14}$$

Let us now work with the LHS. For simplicity, we will denote $S(y) = \{1, \dots, m\}$ and let $\pi(y)_{-1\dots i} = (1-p_m)\dots(1-p_{i+1})$. Further, note that

$$\begin{aligned}
\pi(y) &= \prod_{j \in S(y)} (1-p_j) \\
&= \pi(y)_{-1} - p_1 \pi(y)_{-1} \\
&= \pi(y)_{-12} - p_2 \pi(y)_{-12} - p_1 \pi(y)_{-1} \\
&= 1 - p_m - p_{m-1}(1-p_m) - \dots - p_2 \pi(y)_{-12} - p_1 \pi(y)_{-1}
\end{aligned}$$

Now, we can rewrite the LHS as

$$\begin{aligned}
\pi(y)(1+\sigma(y)) &= \pi(y) + \sum_{j \in S(y)} p_j \pi(y)_{-j} \\
&\leq 1
\end{aligned}$$

which completes the verification of the relative strengths condition.

Finally, we will verify that the no exploitation condition is satisfied (Condition 4.2). To begin, we will adjust our definition of $S(x)$ as follows: $S(x) := \{\mathcal{P}_i \in S : a_i = x, i \neq 2\}$. Similarly, adjust the definitions of $\pi(x)$ and $\sigma(x)$.

Now, calculating directly, we obtain:

$$\begin{aligned}
U_1(x, x, a_{-12}) &= \frac{\frac{p_1}{1-p_1}}{\frac{p_2}{1-p_2} + \sigma(x)}(1 - (1-p_2)\pi(x)) \\
&= U_1(x, 0, a_{-12}) \left(\frac{\sigma_x}{\frac{p_2}{1-p_2} + \sigma_x} \right) \left(\frac{1-(1-p_2)\pi(x)}{1-\pi(x)} \right)
\end{aligned}$$

So, it is sufficient to show that

$$\frac{1-(1-p_2)\pi(x)}{1-\pi(x)} \leq \frac{\frac{p_2}{1-p_2} + \sigma_x}{\sigma_x}$$

equivalently, we need to show that

$$\frac{p_2 \pi(x)}{1-\pi(x)} \leq \frac{p_2}{(1-p_2)\sigma_x}$$

Some algebra then yields the following equivalent condition

$$\pi(x)(1 + (1-p_2)\sigma(x)) \leq 1$$

Working with the LHS we have

$$\pi(x)(1 + (1-p_2)\sigma(x)) \leq \pi(x)(1+\sigma(x)) \leq 1$$

where the last inequality we have shown already above.
□

A.2 Proof of Corollary 6.3

Corollary 6.3 follows from the following lemmas, which verify Condition 5.2 for both the proportional share utility design and the marginal cost ratio utility design in the case of the single-sector sensor coverage game. These proofs are very similar to those in the previous section. Before we proceed, it is useful to notice that it is enough to verify Condition 5.2 in the single-sector sensor coverage game – the multi-sector result follows immediately.

Lemma A.3 *The proportional share utility design satisfies Condition 5.2 in the single-sector sensor coverage game.*

Proof: Define $S(x) := \{\mathcal{P}_j \in S : a_j = x, j \neq i\}$.

$$\sigma_x = \sum_{S_j \in S(x)} p_j$$
$$\pi_x = \prod_{S_j \in S(x)} (1 - p_j)$$

Then, we have that

$$\phi(x, a_{-1}) - \phi(0, a_{-i}) = p_i \pi_x$$

So, the condition becomes

$$\frac{p_i}{p_i + \sigma_x}(1 - (1 - p_i)\pi_x) \geq p_i \pi_x$$

Equivalently, we have

$$(1 - (1 - p_i)\pi_x) \geq \pi_x(p_i + \sigma_x)$$

Some algebra, then yields the following equivalent condition

$$(1 + \sigma_x)\pi_x \leq 1$$

which we have already proven is satisfied, see (13).
□

Lemma A.4 *The marginal cost ratio utility design satisfies Condition 5.2 in the single-sector sensor coverage game.*

Proof: Define $S(x) := \{\mathcal{P}_j \in S : a_j = x, j \neq i\}$.

$$\pi(x) = \prod_{j \in S(x)} (1 - p_j) \tag{15}$$
$$\sigma(x) = \sum_{j \in S(x)} \frac{p_j}{1 - p_j} \tag{16}$$

Then, we have that

$$\phi(x, a_{-1}) - \phi(0, a_{-i}) = p_i \pi(x) \tag{17}$$

$$\tag{18}$$

So, the condition becomes

$$\frac{\frac{p_i}{1-p_i}}{\frac{p_i}{1-p_i} + \sigma(x)}(1 - (1-p_i)\pi(x)) \geq p_i \pi(x) \tag{19}$$

Some algebra yields the following equivalent condition

$$\pi(x)(1 + (1-p_i)\sigma(x)) \leq 1 \tag{20}$$

which we have already argued is satisfied, , see (14).
□

A.3 Proof of Theorem 6.1

Theorem 6.1 follows from the following sequence of lemmas.

Lemma A.5 *Consider a single-sector homogeneous sensor coverage game with n sensors each having fixed detection probability p. The price of anarchy is bounded by*

$$\left(\max_{a+b \leq n, a \geq 0, b \geq 1} \left\{ \frac{a}{a+b} + \frac{1-(1-p)^b}{1-(1-p)^{a+b}} \right\} \right)^{-1}$$

where the maximum is taken over integer a, b.

Proof: We will again describe the optimal placement in terms of the Nash placement. For each sector x covered by the Nash, either every sensor is also present at that sector in the optimal placement, or some number m_x "move" in the optimal. The sensors that move, can provide an additional reward that is bounded by their contribution in the Nash, but will drop the reward gathered from sector x. Note that $0 \leq m_x \leq \sigma_x(a^{ne}) - 1$.

This gives that

$$\begin{aligned} \phi(a^{opt}) &\leq \sum_{x \in X : \sigma_x(a^{ne}) > 0} \max_{m_x \in [0, \sigma_x(a^{ne})-1]} \left\{ m_x U_x(a^{ne}) + (1-(1-p)^{\sigma_x - m_x}) R(x) \right\} \\ &= \sum_{x \in X : \sigma_x(a^{ne}) > 0} \max_{m_x \in [0, \sigma_x(a^{ne})-1]} \left\{ \frac{m_x}{\sigma_x(a^{ne})} + \frac{1-(1-p)^{\sigma_x(a^{ne}) - m_x}}{1-(1-p)^{\sigma_x(a^{ne})}} \right\} (1-(1-p)^{\sigma_x(a^{ne})}) R(x) \end{aligned}$$

Letting $a_x = m_x$ and $b_x = \sigma_x(a^{ne}) - m_x$, we have

$$\begin{aligned}
\phi(a^{opt}) &\leq \sum_{x\in X:\sigma_x(a^{ne})>0} \max_{a_x+b_x=\sigma_x(a^{ne}),a_x\geq 0,b_x\geq 1}\left\{\frac{a_x}{a_x+b_x}+\frac{1-(1-p)^{b_x}}{1-(1-p)^{a_x+b_x}}\right\}(1-(1-p)^{\sigma_x(a^{ne})})R(x) \\
&\leq \sum_{x\in X:\sigma_x(a^{ne})>0} \max_{a+b\leq n,a\geq 0,b\geq 1}\left\{\frac{a}{a+b}+\frac{1-(1-p)^{b}}{1-(1-p)^{a+b}}\right\}(1-(1-p)^{\sigma_x(a^{ne})})R(x) \\
&= \max_{a+b\leq n,a\geq 0,b\geq 1}\left\{\frac{a}{a+b}+\frac{1-(1-p)^{b}}{1-(1-p)^{a+b}}\right\}\sum_{x\in X:\sigma_x(a^{ne})>0}(1-(1-p)^{\sigma_x(a^{ne})})R(x) \\
&= \max_{a+b\leq n,a\geq 0,b\geq 1}\left\{\frac{a}{a+b}+\frac{1-(1-p)^{b}}{1-(1-p)^{a+b}}\right\}\phi(a^{ne})
\end{aligned}$$

which completes the proof.

□

To obtain a more explicit form of the price of anarchy, we will first relax the constraints and then we will characterize the maximal a, b.

Lemma A.6

$$\max_{a+b\leq n,a\geq 0,b\geq 1}\left\{\frac{a}{a+b}+\frac{1-(1-p)^{b}}{1-(1-p)^{a+b}}\right\} \leq \max_{a+b=n,a\geq 0,b\geq 1}\left\{\frac{a}{a+b}+\frac{1-(1-p)^{b}}{1-(1-p)^{a+b}}\right\}$$

where the LHS is taken over integer a, b and the RHS is taken over real-valued a, b.

Proof: We will start by relaxing the integer optimization to include real-valued a, b.

Next, suppose, that $a_m + b_m = m < n$ are the maximizers under the constraint that $a + b = m$. We will show that $a_n = na_m/m$, $b_n = nb_m/m$ lead to a larger value than a_m, b_m. Combining this with the observation that $a_n + b_n = n$ then completes the proof.

$$\frac{a_n}{a_n+b_n}+\frac{1-(1-p)^{b_n}}{1-(1-p)^{a_n+b_n}} = \frac{a_m}{a_m+b_m}+\frac{1-(1-p)^{n/m}(1-p)^{b_m}}{1-(1-p)^{n/m}(1-p)^{a_m+b_m}}$$

Now, it is enough to show that

$$\frac{1-(1-p)^{n/m}(1-p)^{b_m}}{1-(1-p)^{n/m}(1-p)^{a_m+b_m}} \geq \frac{1-(1-p)^{b_m}}{1-(1-p)^{a_m+b_m}}$$

A bit of algebra shows that this holds as long as $(1-p)^{n/m} \leq 1$, which is always true in our setting since $p \in [0, 1]$.

□

Now, we know that $a + b = n$ and $b = n - a$. So, we need only calculate a.

Lemma A.7

$$a^* = \text{argmax}_{0\leq a\leq n-1}\left\{\frac{a}{n}+\frac{1-(1-p)^{n-a}}{1-(1-p)^{n}}\right\} = \begin{cases} n-1, & p=1; \\ n-\frac{\log\left(\frac{n\log(1/(1-p))}{1-(1-p)^n}\right)}{\log(1/(1-p))}, & p<1. \end{cases}$$

Proof: For the case of $p = 1$, the result is immediate. In the case when $p \neq 1$, we will determine the maximizer by simply differentiating. Differentiating with respect to a gives:

$$\frac{1}{n} - \frac{(1-p)^{n-a}\log(1/(1-p))}{1-(1-p)^n}$$

Setting the derivative equal to zero, then gives

$$(1-p)^{n-a} = \frac{n\log(1/(1-p))}{1-(1-p)^n}$$

Solving for a, we obtain

$$a = n - \frac{\log\left(\frac{n\log(1/(1-p))}{1-(1-p)^n}\right)}{\log(1/(1-p))}$$

which completes the proof.

□