A GAME THEORETIC APPROACH TO UAV ROUTING AND INFORMATION COLLECTION

BY

OMKAR THAKOOR

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Computer Science
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2017

Urbana, Illinois

Adviser:

Assistant Research Professor Jugal Garg

ABSTRACT

In recent times, the use of Unmanned aerial vehicles (UAVs) for tasks which involve high endurance or perilous environments, has become increasingly vital. A typical problem is that of information collection, in particular when multiple UAVs are involved, which prompts an important problem of routing these UAVs through the search environment with the goal of maximizing the collected information. Most of the previous line of work assumes a centralized control and full communication among the UAVs, thus posing this as an optimization problem solved via centralized solutions. However, in applications where communication is infeasible, each UAV must individually solve the problem. Assuming a natural scenario of UAVs being compensated for the collected information makes them self-interested agents trying to maximize their payoffs. Consequently, our game-theoretic approach is a natural fit. While our game model is primarily based on the game model used in a previous work [1], it is also significantly generalized, incorporating interesting facets of information fusion and multi-modality-composed information. This game is closely related to the well-studied classes of *congestion-type* and *resource selection* games, but cannot be cast into these classes unless certain critical constraints are relaxed. Our contribution to this literature, is a result on existence of pure Nash equilibria via existence of the Finite Improvement Property, which applies to any *singleton congestion-type games* having a certain class of payoff functions. Finally, to our best knowledge, our results providing theoretically guaranteed tight bounds on the *Price of anarchy* and *Price of stability*, are the first such results in the literature involving a game-theoretic approach to UAV routing.

To my parents, for their love and support.

ACKNOWLEDGMENTS

I would like to express my profound gratitude to Prof. Rakesh Nagi and my adviser Prof. Jugal Garg for providing invaluable ideas, resources, and insights, for meticulously reviewing this report, and for their constant support and encouragement, without which this research would not have been possible.

TABLE OF CONTENTS

LIST OF ABBREVIATIONS

UAV	Unmanned Aerial Vehicle
PoS	Price of Stability
PoA	Price of Anarchy
NE	Nash Equilibrium
wlog	without loss of generality
LHS	Left Hand Side
RHS	Right Hand Side
u.a.r.	uniformly at random

CHAPTER 1

INTRODUCTION

In recent times, the use of Unmanned Aerial Vehicles (UAVs) in modern battlefields has become increasingly vital and beneficial, particularly owing to their utility in environments that are — as characterized by [2] — dull, dirty, and dangerous. That is, they are particularly useful for tasks which involve high endurance and/or perilous environments. One such task is that of collecting information via surveillance of such sensitive region. Moreover, when there are multiple UAVs at disposal, an important problem is that of routing these UAVs in the region in order to avoid collisions and redundant duplication of efforts while maximizing the collected information. This can be posed as an optimization problem which may be effectively solved using centralized algorithms. However, these do not apply to applications where communication is infeasible so that each UAV must individually solve the problem, or scenarios in which the UAVs are self-interested agents trying to maximize their payoffs. Hence, in this paper, we study a problem where UAVs are self-interested competitive agents trying to route themselves in the surveillance region so as to maximize their own collected information — thus, a mere single-objective centralized solution does not apply.

The problem of information collection has been widely studied in the search theory literature. A classical search problem here is to maximize the probability of detecting a hidden target, for instance, as in [3, 4, 5, 6, 7, 8]. On the other hand, [9], associates a potential information gain with each subregion based on an entropy based function and aim to maximize the total gain. In this paper, we resort to a similar model, where the region of surveillance is divided into discrete cells each having an associated information value, which we treat as an abstract entity allowing flexibility for what these values capture; possibly the target-detection probabilities as in the former problem. Our basic model of the search environment and the information collection formulation, is mathematically equivalent to that of [1], in terms of discretization of the search space and time-steps, and payoff definitions, albeit with some other interpretational differences with various parameters. Additionally, our model also incorporates some interesting and useful extensions. The first is that of information fusion, which "can be defined as the combination of multiple sources to obtain improved information (cheaper, greater quality, or greater relevance)" [10] and has been used in robotics and military applications [11]. Secondly, our model allows the information to be multi-modal i.e. present in the form of multiple modalities such as vision, audition, tactition, thermoception etc. These additional aspects make our information collection problem very versatile in terms of applicability. Moreover, to our best knowledge, our results providing theoretically guaranteed tight bounds on the *Price of anarchy* and *Price of stability*, are the first such results in the related literature.

A large body of the previous work on the routing problem such as [12, 13], assume a centralized control and full communication among the UAVs and the central controller. However, the infeasibility of communication can be a very critical constraint in cases such as when surveilling in sensitive areas during covert military operations. Consequently, constructing a centralized solution does not apply. In these situations, each UAV has to construct its own route, since it cannot dynamically obtain information about where the other UAVs are. Furthermore, we consider a natural arrangement that the UAVs are compensated for the information they collect, which makes them self-interested agents trying to maximize the information they collect. As a result, a game-theoretic approach is a perfect fit to tackle this problem. Game-theoretic models have been deployed in numerous other routing problems in transportation and networking applications such as [14] and [15]. The class of games we formulate in this paper, is closely related to the class of congestion games or resource-selection games, and numerous other variants, defined and studied in — most notably [16, 17, 18, 19, 20, 21]. Our class of games has some critical differences with these well-studied classes of games, in terms of cost-sharing protocols, player weights etc., and hence the results established on these classes of games do not directly apply in the general case, making our theoretical results on the existence of pure equilibria and bounds on Price of Stability and Price of Anarchy, interesting and non-trivial.

CHAPTER 2

TWO-PLAYER SINGLE-STEP GAME

2.1 Problem Description and Game Formulation

We define the information collection problem as a game between two players. The surveillance environment is discretized into sub-regions referred to as cells. The problem can be formulated in two settings, which we call *correlated* and *non-correlated.* In the former setting, the information to be collected is in the form of a single modality. Thus, as the UAVs try to maximize the information collected, a cell which is attractive to one UAV is also to other UAVs — more specifically, the intrinsic worth from visiting the cells occurs in an identical order to all the UAVs. This mutuality of the preference among cells leads to the name 'correlated'. This is not necessarily true, when extended to the case of multiple modalities constituting the information in a cell. The UAVs have different sensors for capturing the information of each kind (i.e. modality), and the information of each kind available in a cell could be present in largely varying amounts. Thus, the UAVs may value a cell differently depending on which modalities their sensors are most effective for. Hence, we call this setting non-correlated. Naturally, the correlated setting can be realized as a special case of the non-correlated setting. Hence, we establish results for the non-correlated case and provide stronger implications for the correlated case.

The game is defined as follows. There are two players each corresponding to a UAV and each has a known initial cell. Each player has n pure strategies, corresponding to n different cells the player can move to in the next time-step; n is typically equal to 8 or 9 in the standard problem specifications, but our model is general enough to allow an arbitrary value of n. Also, while we assume the same number of cells accessible to both the players, this constraint can be easily relaxed. Next, We have M different modalities for each of which, each cell has a certain amount of information available. The payoffs of the players — as precisely defined later on — linearly scale with these information values and hence, the values can be taken to be from the range $[0, 1]$ wlog. As the information comprises of M different modalities, we represent it as a vector of dimension M. Each UAV has M different sensors to collect these M kinds of information. The payoff of player i, i.e., the information it can collect, from moving to cell j depends on

1. information available in the cell j,
2. effectiveness of its M sensors denoted by $\vec{\rho_i} = (\rho_{i1}, \ldots, \rho_{im})$, s.t. $0 \geqslant \rho_{im} \geqslant 1,\ \forall m$, and,
3. whether the cell gets shared, i.e., whether the other player also moves to the same cell simultaneously.

This formulation implies that each player has at most two different payoffs possible from each strategy and consequently, we can represent the payoffs with the payoff matrix as follows. Let α_j and α'_j denote the payoffs of player 1 from strategy j when the corresponding cell is unshared and when shared, respectively. Similarly, let β_j and β'_j denote the payoffs of player 2 from strategy j when the corresponding cell is unshared and shared, respectively. The current locations of the players may be different, and thus, there may be some cells which are commonly accessible to both and some which are *private*. For convenience, we order the strategies of players so that if there is a common cell where both players can move, then the index of the corresponding strategy is same for both the players, and smaller than the index of any strategy corresponding to the respective private cells. Depending on their current locations, suppose there are k common cells ($0 \leqslant k \leqslant n$). Then the payoff matrix of the game will be as shown in Table 1.

Observe that the α'_j and β'_j appear only on the diagonal entries ($1 \leqslant j \leqslant k$). This matrix formulation is general enough to allow any definitions of payoff computation, and different formulae for shared payoffs. We formalize these definitions for our game in the following section.

(α'_1, β'_1)	$\cdots$	(α_1, β_k)	(α_1, β_{k+1})	$\cdots$	(α_1, β_n)
(α_2, β_1)	$\cdots$	(α_2, β_k)	(α_2, β_{k+1})	$\cdots$	(α_2, β_n)
$\vdots$	$\cdots$	$\vdots$	$\cdots$	$\vdots$	$\vdots$
(α_k, β_1)	$\cdots$	(α'_k, β'_k)	(α_k, β_{k+1})	$\cdots$	(α_k, β_n)
(α_{k+1}, β_1)	$\cdots$	(α_{k+1}, β_k)	$(\alpha_{k+1}, \beta_{k+1})$	$\cdots$	(α_{k+1}, β_n)
$\vdots$	$\cdots$	$\vdots$	$\cdots$	$\vdots$	$\vdots$
(α_n, β_1)	$\cdots$	(α_n, β_k)	(α_n, β_{k+1})	$\cdots$	(α_n, β_n)

Table 1: Payoff matrix of the 2-player single-step game

2.2 Parameter Specification and Payoff Definition

We denote the information corresponding to a cell as a vector $\vec{v}$, so that each v_k corresponds to the information corresponding to modality k. Next, we have the following parameters:

Sensor effectiveness $\vec{\rho_i}$ of UAV i: If the information available in a cell is $\vec{v}$, then the information obtained by UAV i with sensor effectiveness $\vec{\rho_i}$ is $\vec{\rho_i} \cdot \vec{v}$. This formula defines the payoffs without sharing. When a cell is shared, the total information obtained for type j is computed as follows: the information collected by player 1 is $\rho_{1j} v_j$, and that by player 2 is a fraction ρ_{2j} of what is left, i.e., $(1 - \rho_{1j}) v_j$. Thus, the total information collected is $\rho_{1j} v_j + (1 - \rho_{1j}) v_j \rho_{2j}$. Note that this expression is symmetric for both the players. Then, this total payoff is shared in the ratio of $\rho_{1j} : \rho_{2j}$ by the two players. As a result, UAV i gains $\rho_{ij} v_j (1 - \frac{\rho_{1j}\rho_{2j}}{\rho_{1j}+\rho_{2j}})$. We can write the net information gain for UAV i as $(\vec{r} \odot \vec{\rho_i}) \cdot \vec{v}$ or $\vec{\rho_i} \cdot (\vec{r} \odot \vec{v})$, where $r_j = (1 - \frac{\rho_{1j}\rho_{2j}}{\rho_{1j}+\rho_{2j}})$.

Information fusion parameter $\vec{\gamma_l}$: We incorporate the concept of information fusion by saying that the maximum information gain from a cell may be higher if more than one UAV visit this cell together due to the information fusion. We model this by introducing fusion parameters $\vec{\gamma_l} = (\gamma_{l1}, \ldots, \gamma_{lM})$, where l denotes the number of UAVs visiting a cell simultaneously, thereby involved in the fusion process, and γ_{l_i} denotes the fusion parameter for modality i. We say that in the case of fusion, the information available in a cell can be treated to be $\vec{\gamma_l} \odot \vec{v}$, if l players visit the cell simultaneously. That is, the information of each modality j gets scaled by a constant factor γ_{lj}. Clearly, $\vec{\gamma_1} = (1, \ldots, 1)$, and $\gamma_{lj} \geqslant 1, \forall l \geqslant 2, \forall j$, so that information fusion only increases the information gain. When dealing with a 2-player game, we denote the fusion parameter $\vec{\gamma_2}$ as simply $\vec{\gamma}$ for convenience. Thus, for a cell corresponding to strategy j, having information $\vec{v_j}$, the shared payoffs of the the two players can be written as $\alpha'_j = (\vec{r} \odot \vec{\rho_1} \odot \vec{\gamma_2}) \cdot \vec{v_j}$ and $\beta'_j = (\vec{r} \odot \vec{\rho_1} \odot \vec{\gamma_2}) \cdot \vec{v_j}$. Now, we further assume that the following conditions hold for the fusion process, which are all practical and help impose a reasonable structure to the problem:

- $\forall i,\ \alpha'_i > \alpha_i \Leftrightarrow \beta'_i > \beta_i$
- $\forall i, j,\ (\alpha'_i > \alpha_i \Leftrightarrow \alpha'_j > \alpha_j)$ and $(\beta'_i > \beta_i \Leftrightarrow \beta'_j > \beta_j)$
- $\forall i, j,\ (\alpha_i > \alpha_j \Leftrightarrow \alpha'_i > \alpha'_j)$ and $(\beta_i > \beta_j \Leftrightarrow \beta'_i > \beta'_j)$

Together, these conditions imply that, if the fusion gives better shared payoffs than the unshared ones, it does so for both the players, in all the cells, and the shared payoff values follow the same order as the respective unshared payoff values for the individual players. The same implications hold even for the case when the shared payoffs with fusion, are less than the respective unshared payoffs.

For this game, we establish results for the existence and computation of Nash Equilibria, and bounds on the Price of Stability (PoS) and Price of Anarchy(PoA). These results differ, depending on whether there is moderate fusion ($\alpha_i > \alpha'_i\ \forall i$), or significant fusion ($\alpha_i < \alpha'_i\ \forall i$). First, we establish the results for the former case.

CHAPTER 3

MODERATE INFORMATION FUSION

In this setting, we prove that a pure NE always exists and is computable in linear time. We also prove tight bounds on PoA and PoS. Finally, we show that the number of mixed equilibria has a low upper-bound — linear in the number of strategies, and that the mixed equilibria are Pareto-dominated by the pure equilibria.
Recall that, as described in the previous chapter, we have, $\alpha_i < \alpha'_i \; \forall i$, and $\beta_i < \beta'_i \; \forall i$. That is, sharing a cell is never preferred over visiting it solely, by either player.

3.1 Existence and Computation of Pure Nash Equilibrium

Computing a Nash equilibrium in a two-player game is, in general, PPAD-complete [22], and checking if a pure NE exists, is NP-complete [23]. However, the payoff matrix of the UAV game is special and we show that there always exists a pure Nash equilibrium and it can be computed in $O(n)$ time.

Notation: The cell corresponding to strategy s of player i, is denoted by $cell^i(s)$. The set of cells corresponding to a set of strategies S of player i, is denoted by $cells^i(S)$.

Theorem 1. *There always exists a pure Nash equilibrium in a two player UAV game and it can be computed in $O(n)$ time where n is the number of strategies for each player.*

Proof. We sort the strategies for each player in the non-increasing order of available information in the corresponding cells. Let the order for the first player be $i_1, i_2, \ldots, i_n$, and for the second player be $j_1, j_2, \ldots, j_n$. Further, we assume that in case of equal payoffs, the strategies are ordered in the increasing order of index-values. Since the fusion is moderate, α_{i_1}, and β_{j_1} are the largest possible payoffs for the two players respectively. Let $A_1 = \{i \mid \alpha_i = \alpha_{i_1}\}$ be the set of strategies of player 1 corresponding to cells with the largest possible information gain, and $A_2 = \operatorname{argmax}_{i \notin A_1} \alpha_i$ be the set of strategies of player 1 corresponding to cells with the second largest possible information gain. Similarly, let B_1 and B_2 be the sets of strategies of player 2 corresponding to cells with the largest and the second largest information gain, respectively. We now consider the following cases:

Case 1: $|A_1| > 1$ and $|B_1| > 1$

In this case, it is easy to see that (i, j) is a pure Nash Equilibrium for every $i \in A_1,\ j \in B_1$ such that $cell^1(i) \neq cell^2(j)$. Since A_1 and B_1 have at least 2 strategies in them, there exist many such equilibria, in particular, at least 2 when each has only 2 strategies. Further, $|A_1| > 1 \Rightarrow \exists i \in A_1$ such that player 2 is not moving into $cell^1(i)$, hence, i is in the best response set of player 1. Similarly, the best response of player 2 must lie in B_2. Consequently, there are no other pure equilibria. Also, each of these pure NE, maximizes the total gain (i.e., social welfare).

Case 2: $|A_1| = 1$ and $|B_1| > 1$

a) $cell^1(i_1) \notin cells^2(B_1)$: Same as in Case 1, since B_1 has at least 2 strategies, player 2's best response to any pure strategy of player 1, must lie in B_1. As per the condition here, this best response cannot be i_1, and so, i_1 is always the best response of player 1. Thus, (i_1, j) for every $j \in B_1$, is a pure equilibrium and these are the only equilibria.

b) $cell^1(i_1) \in cells^2(B_1)$: If player 2 does not play i_1, we have the same situation as Case 2a, and thus, (i_1, j) for every $j \in B_1 \backslash \{i_1\}$, is a pure equilibrium. Additionally, there could be other equilibria depending on the following:

   i. $\alpha_{i_2} < \alpha'_{i_1}$: Thus, i_1 is a dominant strategy for player 1, and so it will always play that. The best response of player 2 to this is any $j \in B_1 \backslash \{i_1\}$, and thus all such (i_1, j), are the only equilibria.

ii. $\alpha_{i_2} \geqslant \alpha'_{i_1}$: Thus, if player 2 plays i_1, player 1 can play any strategy $i \in A_2$, and it is easy to see that (i, i_1) is indeed an equilibrium.

Case 3: $|A_1| > 1$ and $|B_1| = 1$
This case is symmetric to Case 2 and a similar analysis follows.

Case 4: $|A_1| = 1$ and $|B_1| = 1$

a) $cell^1(i_1) = cell^2(j_1)$: Thus, the players' most preferred cells coincide. We make further cases as follows, each characterized by two conditions. One is for comparing α_{i_2} with α'_{i_1}, and the other is for comparing β_{i_2} with β'_{i_1}. If both the conditions strict inequalities, only one of the following four cases exists and the only pure Nash equilibria are the ones discussed therein, however, in case of an equality for any of the two respective conditions, two or all of the following four cases and the equilibria described therein, may coexist.

i. $\alpha_{i_2} \leqslant \alpha'_{i_1}$ and $\beta_{j_2} \leqslant \beta'_{j_1}$: In this case i_1 is a dominant strategy for player 1, and j_1 is a dominant strategy for player 2. Hence, (i_1, j_1) is a pure Nash equilibrium.

ii. $\alpha_{i_2} \leqslant \alpha'_{i_1}$ and $\beta_{j_2} \geqslant \beta'_{j_1}$: In this case i_1 is a dominant strategy for player 1, and any strategy $j \in B_2$ is a best response to it, for player 2, and hence, any such (i_1, j) is a pure Nash equilibrium.

iii. $\alpha_{i_2} \geqslant \alpha'_{i_1}$ and $\beta_{j_2} \leqslant \beta'_{j_1}$: Similar to case ii, j_1 is a dominant strategy for player 2, and any strategy $i \in A_2$ is a best response to it, for player 1. Hence, any such (i, j_1) is a pure Nash equilibrium.

iv. $\alpha_{i_2} \geqslant \alpha'_{i_1}$ and $\beta_{i_2} \geqslant \beta'_{i_1}$: This case is an anti-coordination game. (i_1, j) $\forall j \in B_2$, as well as (i, j_1), $\forall i \in A_2$ are all equilibria for this case.

b) $cell^1(i_1) \neq cell^2(j_1)$: In this case, (i_1, j_1) is trivially a pure equilibrium. There may be another pure equilibrium, if the two players each occupy their opponent's most preferred cell and this turns out to be the best response to each other. It is easy to see that this translates to the conditions $cell^2(j_1) \in cells^1(A_2)$, $cell^1(i_1) \in cells^2(B_2)$, and, $\beta_{j_2} \geqslant \beta'_{j_1}, \alpha_{i_2} \geqslant \alpha'_{i_1}$; if these hold, (j_1, i_1) is also a Nash equilibrium.

From the above analysis, it is clear that there always exists a pure Nash equilibrium. To see the computation complexity, first, note that A_1 can be computed in $O(n)$ time by computing the maximum i_1 in $O(n)$ time. Similarly, B_1 can be computed in $O(n)$ time. Further A_2 and B_2 can be similarly computed in $O(n)$ time by computing the second maximum payoffs. Lastly, each of the conditions in the cases above can be computed in $O(n)$ time. Thus, a pure equilibrium can be computed in $O(n)$ time. □

In the next section, we analyze the PoS and PoA for this game.

3.2 Social Welfare

As concluded in the previous section, there always exists a pure Nash equilibrium and it can be computed efficiently. We now investigate how good or bad it can be with respect to the optimal solution which maximizes the social welfare — which is defined as the sum of the individual payoffs of the players. For this, we obtain bounds on pure *Price of Stability* (PoS) and the *Price of Anarchy* (PoA), which are the two well-known metrics used in economics and game theory, to quantify the efficiency of the equilibria. In the following definitions, the *best* and the *worst* NE refer to those which give the least and most social welfare among all equilibria.

Definition 1. *The pure price of stability (PoS) is defined as:*

$$PoS = \frac{\textit{The optimal social welfare}}{\textit{social welfare for the best } \text{pure } \textit{Nash equilibrium}}.$$

First, we begin with the following well-known property.

Lemma 1. *Let* $x_i, y_i, v_i \geqslant 0$ *for* i = *1,2,...,d. Then,*

$$\frac{\sum_{i=1}^{d} x_i v_i}{\sum_{i=1}^{d} y_i v_i} \leqslant \max_{1 \leqslant i \leqslant d} \frac{x_i}{y_i}.$$

We now establish upper bounds on the pure PoS, first for the general parameters, and next, for a *practical* special case of identical sensors, i.e., $\vec{\rho_1} = \vec{\rho_2}$, and no fusion.

Theorem 2. *The pure PoS of the 2-player single-step game with moderate fusion, is at most*

$$1 + \max_{1 \leqslant j \leqslant M} \frac{\max\{\rho_{1j}, \rho_{2j}\}((\rho_{1j} + \rho_{2j}) - (\rho_{1j} + \rho_{2j} - \rho_{1j}\rho_{2j})\gamma_j)}{(\rho_{1j} + \rho_{2j})(\rho_{1j} + \rho_{2j} - \rho_{1j}\rho_{2j})\gamma_j}.$$

Proof. We prove the result by analyzing the same various cases as discussed in the proof of Theorem 1.

In Case 1, all pure equilibria achieve maximum social welfare. Hence the pure PoS here is 1.

In Case 2 — both 2a and 2b, there exist equilibria (i_1, j) for $j \in B_1$, which maximize the social welfare, making the PoS equal to 1 for this case.

In Case 3, it being similar to Case 2, the PoS is again simply 1.

For Case 4a, let $i_1 = j_1 = 1$, and let the information in the cell corresponding to strategy 1 be $\vec{v}$. It implies that $\alpha_1 = \vec{\rho_1} \cdot \vec{v}$ and $\beta_1 = \vec{\rho_2} \cdot \vec{v}$, $\alpha'_1 = (\vec{\rho_1} \odot \vec{r} \odot \vec{\gamma}) \cdot \vec{v}$ and $\beta'_1 = (\vec{\rho_2} \odot \vec{r} \odot \vec{\gamma}) \cdot \vec{v}$. Among the four cases possible here, more than one of them coexisting gives rise to more pure equilibria and a possibly smaller PoS. Hence, to bound the worst value, it suffices to analyze with the assumption of each case existing exclusively, i.e., assuming a relation of strict inequality between α_{i_2} and α'_{i_1}, as well as β_{i_2} with β'_{i_1}.

We analyze Case 4a-i first. We have $\alpha_{i_2} < \alpha'_1$ and $\beta_{j_2} < \beta'_1$, and thus, $(1, 1)$ is the unique equilibrium, with a social welfare of $\alpha'_1 + \beta'_1$. One can see, that the maximum social welfare is achieved for either $\alpha_1 + \beta_{j_2}$ or $\alpha_{i_2} + \beta_1$. Hence, the PoS is at most $\frac{\max\{\alpha_{i_2} + \beta_1, \alpha_1 + \beta_{j_2}\}}{\alpha'_1 + \beta'_1} \leqslant \frac{\max\{\alpha'_1 + \beta_1, \alpha_1 + \beta'_1\}}{\alpha'_1 + \beta'_1}$. We can write,

$$\begin{aligned}
\frac{\alpha'_1 + \beta_1}{\alpha'_1 + \beta'_1} &= \frac{(\vec{\rho_1} \odot \vec{r} \odot \vec{\gamma} + \vec{\rho_2}) \cdot \vec{v}}{((\vec{\rho_1} + \vec{\rho_2}) \odot \vec{r} \odot \vec{\gamma}) \cdot \vec{v}} \\
&\leqslant 1 + \max_{1 \leqslant j \leqslant M} \rho_{2j} \frac{(\rho_{1j} + \rho_{2j}) - (\rho_{1j} + \rho_{2j} - \rho_{1j}\rho_{2j})\gamma_j}{(\rho_{1j} + \rho_{2j})(\rho_{1j} + \rho_{2j} - \rho_{1j}\rho_{2j})\gamma_j}.
\end{aligned}$$

The above inequality, follows from Lemma 1. Similarly, we can write

$$\frac{\alpha_1 + \beta'_1}{\alpha'_1 + \beta'_1} \leqslant 1 + \max_{1 \leqslant j \leqslant M} \rho_{1j} \frac{(\rho_{1j} + \rho_{2j}) - (\rho_{1j} + \rho_{2j} - \rho_{1j}\rho_{2j})\gamma_j}{(\rho_{1j} + \rho_{2j})(\rho_{1j} + \rho_{2j} - \rho_{1j}\rho_{2j})\gamma_j}.$$

Consequently, the PoS for this case is at most

$$1 + \max_{1 \leqslant j \leqslant M} \frac{\max\{\rho_{1j}, \rho_{2j}\}((\rho_{1j} + \rho_{2j}) - (\rho_{1j} + \rho_{2j} - \rho_{1j}\rho_{2j})\gamma_j)}{(\rho_{1j} + \rho_{2j})(\rho_{1j} + \rho_{2j} - \rho_{1j}\rho_{2j})\gamma_j}.$$

For Case 4a-ii, We have $\alpha_{i_2} < \alpha'_1$ and $\beta_{j_2} > \beta'_1$. Any pure equilibrium is of the form $(1, j)$ for some $j \in B_2$, which gives a social welfare of $\alpha_1 + \beta_{j_2}$. Since $\alpha_1 > \alpha'_1 > \{\alpha_i, \alpha'_i\}, \forall i \geqslant 2$, and $\beta_1 > \beta_{j_2} > \{\beta'_k, \beta_j\}, \forall j > 2, k \geqslant 1$, the maximum social welfare is either $\alpha_1 + \beta_{j_2}$ or $\alpha_{i_2} + \beta_1$. Therefore PoS is at most $\max\{1, \frac{\alpha_{i_2}+\beta_1}{\alpha_1+\beta_{j_2}}\} \leqslant \max\{1, \frac{\alpha'_1+\beta_1}{\alpha_1+\beta'_1}\}$. It is easy to see that this is no worse than $\frac{\alpha'_1+\beta_1}{\alpha'_1+\beta'_1}$, and thus, no worse than the bound for Case 4a-i.

Similarly, it can be shown that Case 4a-iii gives a bound no worse than Case 4a-i.

For Case 4a-iv, the two pure Nash equilibria of the game are (i_1, j_2) and (i_2, j_1). It is easy to check that one of them gives the maximum social welfare, hence PoS for this case is 1.

Finally, in Case 4b, among the two possible pure NE, (i_1, j_1) corresponds to the best social welfare, giving a PoS of 1.
Thus, Case 4a-i above gives the worst possible PoS, thereby establishing the desired bound. □

Next, we obtain a stronger result, for a special case.

Corollary 1. *When the two players have identical sensors and the fusion is absent, the pure PoS of the 2-player single-step game is at most* $3/2$.

Proof. Follows from Theorem 2 by setting $\vec{\rho_1} = \vec{\rho_2}$ and $\vec{\gamma} = (1, 1)^T$. □

Next, we obtain similar results for PoA.

Definition 2. *The price of anarchy (PoA) is defined as:*

$$PoA = \frac{\textit{The optimal social welfare}}{\textit{social welfare for the worst Nash equilibrium}}.$$

Unlike the result for PoS, we establish upper bound results for the general PoA, not just pure PoA. First, we do so for the general parameters, and next, for the *practical* special case of the players having identical sensors, i.e. $\vec{\rho_1} = \vec{\rho_2}$, and fusion being absent.

Theorem 3. *The PoA of the 2-player single-step UAV game with moderate fusion, is at most*

$$\max_{1\leqslant j\leqslant M} \frac{\rho_{1j}+\rho_{2j}}{(\rho_{1j}+\rho_{2j}-\rho_{1j}\rho_{2j})\gamma_j}.$$

Proof. Again, we prove the result by analyzing the same various cases as laid out in the proof of Theorem 1.

In Case 1, all the pure equilibria achieve maximum social welfare. Hence the pure PoA is 1 here.

In Case 2a and Case 2b-i as well, all the pure equilibria achieve maximum social welfare, and so the pure PoA here, is 1. For Case 2b-ii, however, i.e. when $\alpha_{i_2} \geqslant \alpha'_{i_1}$ and $i_1 \in B_1$, we know that (i, i_1) is an equilibrium for every $i \in A_2$. This gives a total welfare of $\alpha_{i_2} + \beta_{j_1}$, which is lesser than the maximum welfare of $\alpha_{i_1} + \beta_{j_1}$ achieved by playing (i_1, j) for any $j \in B_1 \backslash \{i_1\}$. Thus, the PoA is $\frac{\alpha_{i_1}+\beta_{j_1}}{\alpha_{i_2}+\beta_{j_1}}$, which is at most $\frac{\alpha_{i_1}+\beta_{j_1}}{\alpha'_{i_1}+\beta_{j_1}}$, under the constraint $\alpha_{i_2} \geqslant \alpha'_{i_1}$. Similarly, the PoA in Case 3 is at most $\frac{\alpha_{i_1}+\beta_{j_1}}{\alpha_{i_1}+\beta_{j_2}}$, which is at most $\frac{\alpha_{i_1}+\beta_{j_1}}{\alpha_{i_1}+\beta'_{j_1}}$.

For Case 4a, let $i_1 = j_1 = 1$, and let the information in the cell corresponding to strategy 1 be $\vec{v}$. It implies that $\alpha_1 = \vec{\rho_1} \cdot \vec{v}$ and $\beta_1 = \vec{\rho_2} \cdot \vec{v}$, $\alpha'_1 = (\vec{\rho_1} \odot \vec{r} \odot \vec{\gamma}) \cdot \vec{v}$ and $\beta'_1 = (\vec{\rho_2} \odot \vec{r} \odot \vec{\gamma}) \cdot \vec{v}$. Among the four cases possible here, more than one of them coexisting gives rise to more pure equilibria and a possibly larger PoA. However, to bound the worst value, it suffices to analyze with the assumption of each case existing exclusively, and considering the worst bound among the four. Thus, we assume a relation of strict inequality between α_{i_2} and α'_{i_1}, as well as β_{i_2} with β'_{i_1}.

For each of the cases 4a-i, 4a-ii, and 4a-iii, it can be easily seen that all the possible pure equilibria in the respective cases have equal total welfare. Consequently, the PoA is only as bad as the PoS. Thus, the largest possible value is achieved for case 4a-i, which is at most $\frac{\max\{\alpha'_1+\beta_1, \alpha_1+\beta'_1\}}{\alpha'_1+\beta'_1}$.

For Case 4a-iv, the pure Nash equilibria of the game are $(1, j)$, $\forall j \in B_2$ and $(i, 1)$, $\forall i \in B_2$. It is easy to check that one of them gives the maximum social welfare, hence PoA for this case is nothing but the ratio of the larger of the two to the smaller, i.e., $\max\{\frac{\alpha_1+\beta_{j_2}}{\alpha_{i_2}+\beta_1}, \frac{\alpha_{i_2}+\beta_1}{\alpha_1+\beta_{j_2}}\}$. The first of the two terms is at most $\frac{\alpha_1+\beta_1}{\alpha'_1+\beta_1}$. The other term is bounded similarly, and so, the PoA is at most $\max\{\frac{\alpha_1+\beta_1}{\alpha'_1+\beta_1}, \frac{\alpha_1+\beta_1}{\alpha_1+\beta'_1}\}$.

For Case 4b, if (i_1, j_1) is the only equilibrium, the PoA is simply 1, since this equilibrium gives the maximum social welfare of $\alpha_{i_1} + \beta_{j_1}$. Although, (j_1, i_1)

may also be an equilibrium under the constraints $\alpha_{i_1} > \alpha_{j_1} \geqslant \{\alpha'_{i_x}, \alpha_{j_y}\}, \forall x \geqslant 1, y \geqslant 2$, and $\beta_1 > \beta_{i_1} \geqslant \{\beta'_{i_x}, \beta_{j_y}\}, \forall x \geqslant 1, y \geqslant 2$. This other equilibrium has a social welfare of $\alpha_{j_1} + \beta_{i_1} = \alpha_{i_2} + \beta_{j_2}$. Hence, the PoA is $\frac{\alpha_{i_1}+\beta_{j_1}}{\alpha_{i_2}+\beta_{j_2}} \leqslant \frac{\alpha_{i_1}+\beta_{j_1}}{\alpha'_{i_1}+\beta'_{j_1}}$. Now, this can be compared to all the previous cases, and seen to be worse than all of them. Finally, we can expand it as,

$$\begin{aligned}\frac{\alpha_{i_1}+\beta_{j_1}}{\alpha'_{i_1}+\beta'_{j_1}} &= \frac{\vec{\rho_1}\cdot\vec{v_1}+\vec{\rho_2}\cdot\vec{v_2}}{(\vec{\rho_1}\odot r\odot\vec{\gamma})\cdot\vec{v_1}+(\vec{\rho_2}\odot r\odot\vec{\gamma})\cdot\vec{v_2}} \\ &\leqslant \frac{1}{\min\limits_{1\leqslant j\leqslant M} r_j\gamma_j} = \max_{1\leqslant j\leqslant M}\frac{\rho_{1j}+\rho_{2j}}{(\rho_{1j}+\rho_{2j}-\rho_{1j}\rho_{2j})\gamma_j}.\end{aligned}$$

Again, the inequality follows from Lemma 1. This is the bound we get for pure PoA.

Further, we show that the PoA at any mixed equilibrium, is no worse than the bound above for Case 4b. Consider a mixed Nash Equilibrium where player 1 is randomizing from a set of pure strategies S_1, with an expected payoff P. If $i_1 \in S_1$, then P equals the expected payoff from playing i_1. If $i_1 \notin S_1$, P is at least the expected payoff from playing i_1. Now, the expected payoff from playing i_1 is at least the worst payoff from playing i_1, i.e. α'_{i_1}. Thus at any mixed equilibrium, player 1 has a payoff of at least α'_{i_1}. Similarly player 2 has a payoff of at least β'_{j_1}. Since the maximum social welfare is at most $\alpha_{i_1} + \beta_{j_1}$, PoA is at most $\frac{\alpha_{i_1}+\beta_{i_1}}{\alpha'_{i_1}+\beta'_{j_1}}$, giving us the same bound as above, thus proving the required result. □

Next, we obtain a stronger result, for a special case.

Corollary 2. *When the two players have identical sensors and the fusion is absent, the PoA of the 2-player single-step game is at most at most* 2.

Proof. Follows from Theorem 3 by setting $\vec{\rho_1} = \vec{\rho_2}$ and $\vec{\gamma} = (1,1)^T$. □

3.3 Distinct Payoffs and Mixed Equilibria

If multiple strategies can have the same payoffs for a player, the players would be indifferent among those strategies and this can give rise to a large number of mixed Nash equilibria. In this section, we show that if we impose an additional constraint of the information values in all the cells being distinct, we get interesting results regarding the set of mixed equilibria.

First, we can truncate the payoff matrix, by observing that all but one strategy corresponding to cells in the unshared space, are strictly dominated by one of those strategies. Hence, all such strategies can be removed from the matrix. Wlog, for both the players, let strategy $k+1$ to be this one dominating strategy among those corresponding to the unshared space. When any player has a dominant strategy, that strategy would be the best response to any pure or mixed strategy of the opponent. Thus, in such a case, there cannot be any mixed Nash equilibria. In other cases, there may be many; we will first show that this number has a small bound.

Theorem 4. *The number of mixed equilibria is $O(k)$.*

Proof. Let the α-values of the shared cells in the decreasing order be α_{i_1}, α_{i_2}, …, α_{i_k} and similarly, the β-values of the shared cells in the decreasing order be $\beta_{j_1}, \beta_{j_2}, \ldots, \beta_{j_k}$. Consider a mixed Nash equilibrium where, let S_1, S_2 denote the sets of pure strategies mixed by the 2 players respectively. Different pairs of S_1, S_2 correspond to different equilibria. We first establish the following lemma:

Lemma 2. *$\forall x, y$, where, $1 \leqslant x < y \leqslant k$, $i_y \in S_1 \Rightarrow i_x \in S_2$, and $j_y \in S_2 \Rightarrow j_x \in S_1$*

Proof. We have, $\alpha_{i_x} > \alpha_{i_y}$. Thus, if player 2 plays i_x with 0 probability, then i_x is a strictly better response than i_y to any mix of strategies in S_2. Consequently, player 1 will not play i_y in its mixed strategy. Thus, $i_y \in S_1 \Rightarrow i_x \in S_2$. The other implication in the lemma follows similarly. □

Now, let i_l and j_m be the strategies in the shared space having the largest indices in S_1, S_2 respectively. Hence, $S_1 \subseteq \{i_1, \ldots, i_l\} \cup \{k+1\}$ and $S_2 \subseteq \{j_1, \ldots, j_m\} \cup \{k+1\}$. From the lemma above, it follows that $i_l \in S_1 \Rightarrow i_x \in S_2 \; \forall x < l$. Thus, such an equilibrium exists, only if $\{i_1, \ldots, i_{l-1}\} \subseteq \{j_1, \ldots, j_m\}$. Similarly, we must have $\{j_1, \ldots, j_{m-1}\} \subseteq \{i_1, \ldots, i_l\}$. Thus, $l-1 \leqslant |S_2| \leqslant m$ and $m-1 \leqslant |S_1| \leqslant l$. Hence, $|l-m| \leqslant 1$. Accordingly, we have the following cases:

Case 1: $l = m$

S_1 can have at most l ($= m$) strategies, $m-1$ of which get fixed to be $j_1, \ldots, j_{m-1}$ as shown above. S_1 can be formed by including or excluding the remaining m^{th} strategy, as well as, by including or excluding $k+1$ i.e. in at most 4 ways. Similarly S_2 can have at most 4 different values. Thus, we can have at most 16 different pairs $(S_1,\ S_2)$.

Case 2: $l = m+1$
We must have $\{i_1, \ldots, i_{l-1}\} = \{j_1, \ldots, j_m\}$, and by a similar analysis as the previous case, it is easy to see that S_2 can have only two different values based on whether $k+1$ is included or not. S_1, on the other hand, has $m-1$ strategies fixed among $\{i_1, \ldots, i_l\}$, and thereby, it can have 8 different values depending on whether each of the remaining two shared-cell strategies as well as $k+1$ are included or not, and hence, we can have at most 16 different pairs $(S_1,\ S_2)$.

Case 3: $m = l+1$
Similar to the previous case, there can be at most 16 different pairs $(S_1,\ S_2)$.

As l,m can take values from 2 through k, there are roughly $3k$ pairs $(l$,$m)$ such that $|l-m| \leqslant 1$. Since the number of different pairs $(S_1,\ S_2)$ for each possible (l,m) is bounded by a constant as shown, the number of different mixed equilibria is $O(k)$, thus proving the theorem. □

Next, we show that all the mixed equilibria are worse outcomes in terms of Pareto-dominance as compared to the pure equilibria.

Theorem 5. *All the mixed equilibria are Pareto-dominated by the pure equilibria.*

Proof. As above, consider a mixed Nash equilibrium where, S_1, S_2 denote the sets of pure strategies mixed by the two players respectively. In order for (S_1, S_2) to correspond to a mixed NE, S_1 and S_2 must have at least two strategies each. Hence $\exists i_x \in S_1$ s.t. $x \geqslant 2$. We know that a mix of strategies in S_1 is a best-response iff each pure strategy in S_1 is. Hence, the expected payoff from playing this mixed strategy is same as the expected payoff from playing i_x which is obviously bounded by the best payoff possible by playing i_x which is $\alpha_{i_x} \leqslant \alpha_{i_2}$. It is easy to see that player 1's payoff in all the pure equilibria is at least α_{i_2}. Similarly, player 2's payoff from playing the mixed strategy must be worse than β_{i_2} which is the least payoff of player 2 across all the pure equilibria. Thus, playing a mixed strategy gives mutually strictly worse payoffs. Hence, the mixed equilibria are all Pareto-dominated by the pure equilibria. □

This result implies that, with the assumption that both the agents are self-interested and aware of the opponent being self-interested, we can conclude that they would not want to play strategies corresponding to any mixed Nash equilibrium.

CHAPTER 4

SIGNIFICANT INFORMATION FUSION

The setting of significant fusion is one where sharing a cell gives better payoffs than visiting it individually. As mentioned in Section 2, we have assumed that the fusion process has the properties that, if the fusion gives better payoffs, it does so for both the players, in all the cells, and the payoff values follow the same order as the unshared individual payoff values.

4.1 Existence and Computation of Pure Nash Equlibrium

We first obtain that the same results hold for the existence and computation of pure equilibria, as in the moderate fusion case.

Theorem 6. *There always exists a pure Nash equilibrium in a 2-player single-step game with significant fusion and it can be computed in $O(n)$ time where n is the number of strategies for each player.*

Proof. Let i_1 and j_1 be the respective strategies which correspond to cells containing the largest amount of information. Let $P = \{i : \alpha'_i \geqslant \alpha_{i_1}\}$ and $Q = \{j : \beta'_j \geqslant \beta_{j_1}\}$. Note that $i \geqslant k \Rightarrow i \notin P$ and $i \notin Q$ since α'_i and β'_i are not defined for $i \geqslant k$. One can see that, if $P \cap Q$ is non-empty, then $\forall x \in P \cap Q$, we have $\alpha'_x \geqslant \alpha_{i_1}$ and $\beta'_x \geqslant \beta_{j_1}$, and hence, $(x,\ x)$ is a pure equilibrium.

Case 1: $cell^1(i_1) = cell^2(j_1)$: In this case , $P \cap Q$ is necessarily non-empty since $i_1 (= j_1)$ must be in $P \cap Q$. As pointed above, $\forall x \in P \cap Q$, $(x,\ x)$ is a pure equilibrium. It can be easily shown that there exist no other pure equilibria.

Case 2: $cell^1(i_1) \neq cell^2(j_1)$: In this case, $P \cap Q$ may or may not be empty. If $i_1 \in Q$ or $j_1 \in P$, $P \cap Q$ is necessarily non-empty, and again, $\forall x \in P \cap Q$, $(x,\ x)$ is a pure equilibrium, and these are the only pure equilibria. On the other hand, if $i_1 \notin Q$ and $j_1 \notin P$, (i_1, j_1) is a Nash equilibrium. Additionally, $P \cap Q$ may still be non-empty, and if it is, $\forall x \in P \cap Q$, $(x,\ x)$ is a pure equilibrium.

Thus, even in the case of significant fusion, there always exists a pure NE. Further, it is straight-forward to see that computing i_1, j_1, and subsequently, computing P and Q can be done in $O(n)$ time and hence this is no worse than the case with no fusion.

□

Next, we establish bounds on the PoS and the PoA in this setting.

4.2 Social Welfare

As done for the moderate fusion case, we first establish an upper bound on the pure PoS.

Theorem 7. *The pure PoS of the 2-player single-step game with significant fusion, is at most*

$$1 + \max_{1 \leqslant j \leqslant M} \frac{\max\{\rho_{1j}, \rho_{2j}\}((\rho_{1j} + \rho_{2j} - \rho_{1j}\rho_{2j})\gamma_j - (\rho_{1j} + \rho_{2j}))}{(\rho_{1j} + \rho_{2j})^2}.$$

Proof. We prove this by analyzing the various cases formulated in the proof of Theorem 6.

For Case 1, one can see that the strategy (x, x) for any $x \in P \cap Q$ gives mutually better payoffs than any (x, y) for $x \neq y$. Thus, the maximum social welfare is achieved for one of these, all of which are equilibria. Hence, the PoS is simply 1.

In Case 2, if $P \cap Q$ is non-empty, again, the maximum social welfare for is achieved at (x, x) for some $x \in P \cap Q$, which is an equilibrium. Hence, the PoS is simply 1. However, $P \cap Q$ may be empty, and (i_1, j_1) maybe the only equilibrium. In that case, better social welfare maybe achieved at some

(x, x) in $P \cup Q$. Thus, the PoS is at most $\frac{\alpha'_x + \beta'_x}{\alpha_{i_1} + \beta_{j_1}}$. We can then write,

$$\frac{\alpha'_x + \beta'_x}{\alpha_{i_1} + \beta_{j_1}} \leqslant \frac{\max\{\alpha'_{i_1} + \beta_{j_1}, \alpha_{i_1} + \beta'_{j_1}\}}{\alpha_{i_1} + \beta_{j_1}} = 1 + \max_{i \in \{1,2\}} \frac{(\vec{\rho_i} \odot \vec{r} \odot \vec{\gamma} - \vec{\rho_i}) \cdot \vec{v}}{(\vec{\rho_1} + \vec{\rho_2}) \cdot \vec{v}}$$

$$= 1 + \max_{1 \leqslant j \leqslant M} \frac{\max\{\rho_{1j}, \rho_{2j}\}((\rho_{1j} + \rho_{2j} - \rho_{1j}\rho_{2j})\gamma_j - (\rho_{1j} + \rho_{2j}))}{(\rho_{1j} + \rho_{2j})^2}.$$

This gives us the desired bound. □

Next, we similarly obtain an upper bound on PoA.

Theorem 8. *The PoA of the 2-player single-step game with significant fusion, is at most*

$$(1 - \frac{\rho_{1j}\rho_{2j}}{\rho_{1j} + \rho_{2j}})\gamma_j.$$

Proof. We first prove the bound on pure PoA by analyzing the various cases formulated in the proof of Theorem 6.

For Case 1, (x, x) for every $x \in P \cap Q$ is an equilibrium. As shown earlier, the maximum social welfare is achieved for one of these. Hence, the minimum of these leads to the worst PoA, that is,

$$\frac{\max_{x \in P \cap Q} \alpha'_x + \beta'_x}{\min_{x \in P \cap Q} \alpha'_x + \beta'_x} \leqslant \frac{\alpha'_{i_1} + \beta'_{j_1}}{\alpha_{i_1} + \beta_{j_1}} \leqslant \max_{1 \leqslant j \leqslant M} (1 - \frac{\rho_{1j}\rho_{2j}}{\rho_{1j} + \rho_{2j}})\gamma_j.$$

In Case 2, if $P \cap Q$ is non-empty, again, the maximum social welfare is achieved at (x, x) for some $x \in P \cap Q$. Further, it is easy to see that any such equilibrium gives a social welfare of $\alpha'_x + \beta'_x$, which is better than $\alpha_{i_1} + \beta_{j_1}$ due to the fact that $x \in P \cap Q$. This makes (i_1, j_1) the worst equilibrium. Hence, the PoA is $\frac{\alpha'_x + \beta'_x}{\alpha_{i_1} + \beta_{j_1}} \leqslant \frac{\alpha'_{i_1} + \beta'_{j_1}}{\alpha_{i_1} + \beta_{j_1}}$, thus giving the same bound as above. On the other hand, if $P \cap Q$ is empty, then (i_1, j_1) is the only equilibrium. Consequently, the PoA is only as bad as the PoS, which is at most $\frac{\max\{\alpha'_{i_1} + \beta_{j_1}, \alpha_{i_1} + \beta'_{j_1}\}}{\alpha_{i_1} + \beta_{j_1}}$ as analyzed in proving Theorem 7. Again it can be easily seen that this is no worse than $\frac{\alpha'_{i_1} + \beta'_{j_1}}{\alpha_{i_1} + \beta_{j_1}}$, which is the bound for the previous cases.

Finally, we show that the same bound applies to the mixed equilibria. The payoff of player 1 must be at least the payoff it would get by switching to pure strategy i_1. The minimum value of the latter is simply α_{i_1}. Similarly, the payoff of player 2 from any mixed strategy, must be at least β_{j_1}. Thus, the worst social welfare at any mixed equilibrium is $\alpha_{i_1} + \beta_{j_1}$, whereas the maximum possible is nothing but $\alpha'_{i_1} + \beta'_{j_1}$, thus giving us the same bound of $\frac{\alpha'_{i_1}+\beta'_{j_1}}{\alpha_{i_1}+\beta_{j_1}}$ as the pure PoA. □

CHAPTER 5

MULTI-PLAYER GAME: CORRELATED SETTING

In this section, we consider a more general scenario consisting of p players (UAVs), albeit with the constraint of their payoffs being correlated. Formally, we have a set of players $P = \{1, \ldots, p\}$. Player i has a sensor effectiveness ρ_i. The region of surveillance consists of a set of cells $C = \{c_1, \ldots, c_k\}$. If a set of players S choose a certain cell $c \in C$ having information $v(c)$, the aggregate payoff is given by $\gamma_{|S|}\left(1 - \prod_{i \in S}(1 - \rho_i)\right) v(c)$. Further, each player $i \in S$, gets an individual share of the payoff proportional to its ρ_i, i.e. $\frac{\rho_i}{\sum_{i \in S} \rho_i} \gamma_{|S|}\left(1 - \prod_{i \in S}(1 - \rho_i)\right) v$. This can be written in a general form of $\rho_i M_c(S)$, where M_c for each cell c, is a function $M_c : 2^P \to \mathbb{R}$. It is easy to see, that for our game, $M_c(S) = \gamma_{|S|} \frac{\left(1 - \prod_{i \in S}(1 - \rho_i)\right)}{\sum_{i \in S} \rho_i} v(c)$. For our next result, however, we do not need this explicit definition of M_c and we will show that our result holds true for a certain class of functions in general.

5.1 Existence of Pure Nash Equilibrium

First, we state the Finite Improvement Property (FIP) [24]: Any sequence of strategy-tuples in which each strategy-tuple differs from the preceding one in only one coordinate (such a sequence is called a path), and the unique deviator in each step strictly increases the payoff it receives (an improvement path), is finite. Obviously, any maximal improvement path, is terminated by an equilibrium.

Next, we state a natural monotonicity definition for functions defined on sets:

Definition 3. *Let A be a set. A function $f : 2^A \to \mathbb{R}$ is monotonically non-decreasing, if $\forall A_1, A_2 \subseteq A, A_1 \subseteq A_2 \Rightarrow f(A_1) \leqslant f(A_2)$. Similarly, f is monotonically non-increasing, if $A_1 \subseteq A_2 \Rightarrow f(A_1) \geqslant f(A_2)$.*

Next, we prove the existence of a pure equilibrium by showing that the game admits the finite improvement property.

Theorem 9. *If the function $\{M_c | c \in C\}$ are all monotonically non-increasing, or all monotonically non-decreasing functions, then the game admits the Finite Improvement Property.*

Proof. To prove this, we extend the argument used in [17] — which is used there for proving the FIP for symmetric congestion games. Suppose, for the sake of proving contradiction, that there is an infinite improvement path. Since there are only finitely many joint strategies, this essentially means that there is an improvement cycle, say, of size l, given by $\sigma_1, \sigma_2, \ldots, \sigma_l, \sigma_1$, where, σ_j is the joint-strategy in the j^{th} step. Further, let $S_j(c)$ denote the set of players going to cell c in the j^{th} step. Let $C^{\#} \subseteq C$ denote the set $\{c \mid \exists i, j,\ S_i(c) \neq S_j(c)\}$, i.e., those cells which are not occupied by the same set of players throughout the whole improvement cycle. First, we prove for the case where $\{M_c | c \in C\}$ are all monotonically non-increasing.
Wlog, suppose the improvement cycle and the cells, are enumerated such that $\min_{1 \leqslant j \leqslant l, c \in C^{\#}} M_c(S_j(c)) = M_{c_1}(S_l(c_1))$. Now, since $c_1 \in C^{\#}$, $\exists j$ such that $S_j(c_1) \neq S_l(c_1)$. Consider largest such j, i.e. $S_{j+1}(c_1) = S_l(c_1)$. Since, each M_c is monotonically non-increasing, and since the minimum value of M_{c_1} is attained for $S_l(c_1)$, thus, also for $S_{j+1}(c_1)$, it follows that $S_j(c_1) \subset S_{j+1}(c_1)$, Thus, the unique deviator between σ_j and σ_{j+1}, wlog say player 1, must be changing his strategy to c_1 from some other cell c_i, say. Thus, $c_i \in C^{\#}$, and further, for this deviation to be an improvement for Player 1, it must be that $\rho_1 M_{c_i}(S_j(c_i)) < \rho_1 M_{c_1}(S_{j+1}(c_1))$, and hence, $M_{c_i}(S_j(c_i)) < M_{c_1}(S_{j+1}(c_1)) = M_{c_1}(S_l(c_1))$. This contradicts the assumption that, $\min_{1 \leqslant j \leqslant l, c \in C^{\#}} M_c(S_j(c)) = M_{c_1}(S_l(c_1))$, and consequently, there cannot exist an improvement cycle, thus proving the finite improvement property.
The case where $\{M_c | c \in C\}$ are all monotonically non-decreasing, can be shown similarly. Wlog, suppose the improvement cycle and the cells, are enumerated such that $\max_{1 \leqslant j \leqslant l, c \in C^{\#}} M_c(S_j(c)) = M_{c_1}(S_1(c_1))$. Now, since $c_1 \in C^{\#}$, $\exists j$ such that $S_j(c_1) \neq S_1(c_1)$. Consider the smallest such j, i.e. $S_{j-1}(c_1) = S_1(c_1)$. Since, every M_c is monotonically non-decreasing, and since the minimum value of M_{c_1} is attained for $S_1(c_1)$, thus, also for $S_{j-1}(c_1)$, it follows that $S_j(c_1) \subset S_{j-1}(c_1)$, Thus, the unique deviator between σ_{j-1} and σ_j,

wlog, player 1, must be changing his strategy to c_1 from some other cell, say c_i. Thus, $c_i \in C^{\#}$, and further, for this deviation to be an improvement for player 1, it must be that $\rho_1 M_{c_i}(S_j(c_i)) > \rho_1 M_{c_1}(S_{j-1}(c_1))$, and hence, $M_{c_i}(S_j(c_i)) > M_{c_1}(S_{j-1}(c_1)) = M_{c_1}(S_1(c_1))$. This contradicts the assumption that, $\max\limits_{1 \leqslant j \leqslant l, c \in C^{\#}} M_c(S_j(c)) = M_{c_1}(S_1(c_1))$, and consequently, there cannot exist an improvement cycle, thus proving the Finite improvement property. □

In our UAV game, it is intuitive to see that if there is no fusion, each M_c is a non-increasing function since sharing a cell with an additional player should decrease payoff — a rigorous proof is given by Lemma 11 in the appendix. In general, whenever the fusion parameters γ_l are small enough so that M_c are all non-increasing functions, which is the case of Moderate Fusion. On the other hand, when γ_l are large enough so that every M_c is a non-decreasing function, it is the case of Significant Fusion. The theorem above, implies that, in the cases of moderate fusion and significant fusion, there always exists a pure Nash equilibrium.

Next, we give an algorithm to efficiently compute a pure equilibrium, in the case of Significant Fusion.

5.2 Computing a Pure Nash Equilibrium for Significant Fusion

The case of significant fusion is when, more the number of UAVs in the same cell, larger is the individual payoff for each UAV there. That is, $\forall S \subset P, p \in P \backslash S, i \in S, c \in C$, we have $\rho_i M_c(S) \leqslant \rho_i M_c(S \cup \{p\})$. We use a function $A : C \to P$ to denote the set of players which can access a given cell in one step.

The algorithm consists of a number of iterations. Starting with the set of all players, and the set of all cells, in each iteration, some players are assigned a particular cell as their strategy to play and the set of remaining players and the set of remaining cells is carried forward to the next iteration. In each iteration i, for each cell c in C^i, we compute $M_c(A(c) \cap P^i)$, where P^i is the set of remaining players in that iteration, and C^i is the set of remaining cells. Then, we choose the cell c_i^{max} for which the value thus computed is maximum, and assign this cell as the strategy for all the players in $A(c_i^{max}) \cap P^i$. Subsequently, we update the set of remaining players $P^{i+1} = P^i \backslash (A(c_i^{max}) \cap P^i)$ and $C^{i+1} = C^i \backslash \{c_i^{max}\}$ and move on to the next iteration.

Algorithm 1 Algorithm to compute a pure NE

$P^1 \leftarrow P, \quad C^1 \leftarrow C, \quad i \leftarrow 1$
while $P^i \neq \phi$ **do** ▷ Terminate if no players remaining
  $c_i^{max} \leftarrow Null, \quad maxScore \leftarrow -\infty$
  for $c \in C^i$ **do**
    $Q^i(c) = A(c) \cap P^i$
    $score \leftarrow M_c(Q^i(c))$
    if $score > maxScore$ **then**
      $maxScore \leftarrow score, \quad c_i^{max} \leftarrow c$
    end if
  end for ▷ c_i^{max} is computed.
  for $p \in Q^i(c_i^{max})$ **do**
    `Assign cell` c_i^{max} `to` p
  end for
  $P^{i+1} = P^i \backslash Q^i(c_i^{max})$
  $C^{i+1} = C^i \backslash \{c_i^{max}\}$
  $i \leftarrow i + 1$
end while

Proof of Correctness :

Wlog, let the cells be enumerated such that $\forall i,\ c_i^{max} = c_i$. We first note that $i < j \Leftrightarrow P^j \subset P^i \Rightarrow Q^j(c_j) \subseteq Q^i(c_j)$. Now, Suppose player 1 is assigned the cell c_i by the algorithm. Hence, $1 \notin A(c_j), \forall j < i$, since otherwise, it would have been assigned a cell before the i^{th} iteration. Thus, the only cells player 1 could possibly move to, are $\{c_j\}_{j>i}$. The payoff of player 1 before deviation, by playing c_i, is $\rho_1 M_{c_i}(Q^i(c_i))$, which is at least $\rho_1 M_{c_j}(Q^i(c_j))$, since $c_i^{max} = c_i$. Further, $\forall j > i,\ 1 \in A(c_j) \Rightarrow 1 \in Q^i(c_j)$. This further implies that since $Q^j(c_j) \subseteq Q^i(c_j)$, it must also be that, $(Q^j(c_j) \cup \{1\}) \subseteq Q^i(c_j)$, whenever 1 can access c_j. Hence, by the monotonically non-decreasing behavior of each M_c, Player 1's payoff after deviation to c_j, i.e., $\rho_1 M_{c_j}(Q^j(c_j) \cup 1)$, is at most $\rho_1 M_{c_j}(Q^i(c_j))$, which is at most its payoff before deviation. Hence, it has no incentive to switch, and the same applies to all the players. Hence, the algorithm does produce a Nash equilibrium. It is easy to check that the algorithm takes time $O(|C| \cdot |P|)$.

Next, we establish bounds on the Price of Stability and Price of Anarchy, when all players are identical.

5.3 PoS and PoA for Homogeneous Fleet

Having a homogeneous fleet refers to all the UAVs having identical sensor effectiveness ρ. In such a case, the payoff of a player simply depends on the number of players it shares a cell with, and not the actual subset of players. We denote the individual payoff of a player when n players share a cell c, by $v^n(c) = v(c)\frac{1-(1-\rho)^n}{n}$. Our result is as follows:

Theorem 10. *The Price of Stability and the Price of Anarchy, in singleton congestion games with payoffs defined as above, is at most* $2 - 1/p$.

Proof. Let σ^{eq} be an equilibrium and let σ^{wf} be a joint strategy which gives the maximum social welfare. Suppose, starting with σ^{eq}, σ^{wf} is achieved by a series of deviations, where each deviation refers to a player switching from a cell c_i to a cell c_j. Since the players are identical, only the cells involved in a deviation matter, and not the player who deviates. Thus, these deviations can be represented by a directed graph over the cells, where each cell is a vertex and a deviation from one cell to another is represented as a directed edge. Note that since only the number of players in a cell matters in computing any payoffs, any path in the graph of length more than 1, between nodes u and v say, can simply be replaced by a single edge (u, v), since both of them equivalently result in the number of players at cell u decreasing by 1, the number of players at cell v increasing by 1, and any other cells on the path being unaffected. Thus, any graph G can be reduced to G^*, one which doesn't have any paths of length more than 1, and similarly no cycles. Thus, the reduced graph will only have sources, sinks and isolated vertices. Figure 1 illustrates this with an example. Vertices such as c_3 with a larger in-degree than out-degree in G become sinks in G^*. Similarly, vertices like c_1 having a larger out-degree in G become sources in G^*, and the remaining ones like c_7 where the two degrees are equal in G, become isolated in G^*. Now, consider the group of players who are in a cell c at equilibrium. We consider three

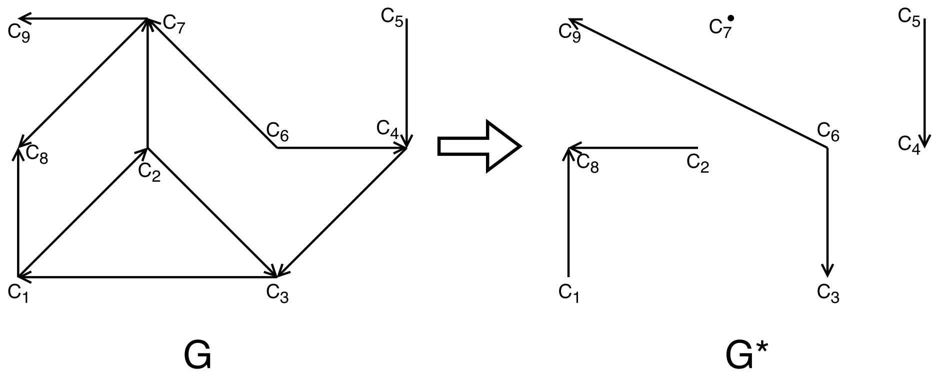

Figure 1: Reduction of a deviation graph

cases for c:

1. c is an isolated vertex in G^*: It is easy to see that the payoff of every player here remains the same.

2. c is a sink in G^*: Thus, there are at least as many players in c at σ^{wf}, as there were at the equilibrium σ^{eq}. Hence, the payoff of these players is bounded above by their payoff at the equilibrium.

3. c is a source in G^*: Consider a player i who is in cell c at σ^{eq}, and is in c' at σ^{wf}. By the nature of G^*, c' must be a sink in G^*. Hence, there are at least as many players in c' at σ^{wf}, as there were at σ^{eq}. Thus, if player i were to be the unique deviator at equilibrium from c to c', it would have gotten at least as much a payoff as it gets by playing c' at σ^{wf}. Further, since σ^{eq} is an equilibrium, the payoff of player i at σ^{eq} is at least as much as it would get by deviating to any other cell, in particular c', and in turn, greater than its payoff at σ^{wf}. Now, suppose there were x players in cell c at σ^{eq}, of which y are not in c at σ^{wf}, while $x - y$ players continue to be in c at σ^{wf}. Then, the payoffs of the y deviating players, at σ^{wf} is at most as much as their payoff at σ^{eq}, i.e. $v(c)\frac{1-(1-\rho)^x}{x}$. The payoff of any of the remaining $x - y$ players, improves from $v(c)\frac{1-(1-\rho)^x}{x}$ to $v(c)\frac{1-(1-\rho)^{x-y}}{x-y}$. Hence, if all the players leave, i.e., $y = x$, the total welfare of this group of x players cannot increase. On the other hand, if $y < x$, the total social welfare for this group of players, can increase by a factor of at most $\frac{\left(1-(1-\rho)^{x-y}\right)+y(1-(1-\rho)^x)/x}{1-(1-\rho)^x}$. It can be easily shown that this expression is a monotonically increasing function of ρ for $\rho \in (0, 1]$ and has the maximum value of $1 + y/x$ when $\rho = 1$. Further, it is easy to see that the maximum value of $1 + y/x$, for $y < x \leqslant p$ is $2 - 1/p$, which is achieved for $x = p, y = p - 1$. Thus, the maximum possible gain in the total welfare of this group is when all but one players leave the cell.

Thus, for a group of players which are in a particular cell at equilibrium, the sum of their payoffs either remains the same, decreases, or increases by a factor of at most $2 - 1/p$ as analyzed for the three cases above. Hence, the total social welfare of all the players, which is the sum of welfares of all such groups, can increase by a factor of at most as much as any of the individual groups, which is nothing but $2-1/p$, giving us the required bound. (Note that the analysis above holds for any equilibrium, and thus, the worst equilibrium in particular, giving the bound on PoA.)
Further, this bound can be shown to be tight for the Price of Stability as well, with a simple example. Let there be p players with $\rho = 1$ for each player. Let there be p cells with every cell being a valid strategy for every player. Let the information available in various cells be as follows : $v(c_1) = p, v(c) = 1 - \epsilon \ \forall c \neq c_1; \epsilon > 0$. It is easy to see that c_1 is a dominant strategy for every player, giving a unique equilibrium $(c_1, \ldots, c_1)$. The total social welfare here is p. It is easy to see that the social welfare is maximum for the joint strategy $(c_1, c_2, \ldots, c_p)$ which equals $p + (p-1)(1-\epsilon)$. Hence the Price of stability is $1 + (1 - 1/p)(1 - \epsilon)$ which approaches $2 - 1/p$ as $\epsilon \to 0$. $\square$

CHAPTER 6

MULTI-STEP GAME

We model the problem of UAV surveillance as a game between the UAVs. We formulate two different games having a vital distinction, in the following setting. $\mathcal{P}$ is a finite set of p players, each corresponding to a UAV. The geographical region of surveillance is represented as a directed graph, where $\mathcal{C}$, the set of vertices - more commonly referred to as cells - represent various small sub-regions, and the directed edges of the graph capture the connectivity between these cells. Moving along any edge and surveilling the subsequent cell altogether corresponds to one time-step. The number of time-steps for which the game lasts is denoted by l. The game is for each player to move in this network for l time-steps, while capturing the information from the cells visited along the route, with the goal of maximizing this information captured. Thus, the set of strategies for player i, denoted by S_i, is nothing but a set of walks of length l starting from player i's initial cell. The set of 'joint strategy profiles', or simply 'outcomes', is denoted by $S = \bigtimes_{i \in \mathcal{P}} S_i$. Each cell has an associated information value denoted by a function $v : C \to \mathbb{R}^+$. Each player $i \in \mathcal{P}$ has a sensor effectiveness denoted by $\rho_i \in [0, 1]$, which determines how much information the player can collect from what is available in the cell it visits. Finally, the payoff of a player depends on the outcome, and is denoted by $\pi_i : S \to \mathbb{R}^+$ for every player i. The net payoff of a player is the sum of the payoffs it gets by visiting the cells on its walk. With a slight abuse of notation, we denote player i's payoff from cell $c \in C$, when the outcome is $s \in S$, by $\pi_i(s, c)$, so that $\pi_i(s) = \sum_{c \in \mathcal{C}} \pi_i(s, c)$. We naturally define $\pi_i(s, c)$ to be zero if i does not visit c at all when playing s_i. However, when it does visit the cell (possibly more than once), the value $\pi_i(s, c)$ can be defined in two different ways depending on the logistics of the real-world scenario, giving rise to two different games as follows:

With temporal aspect: In this case, any player gets an instant payoff after visiting a cell (in a manner described below), and these payoffs get accumulated constituting its net payoff. Consider a cell c initially having a value $v(c)$. The first player to visit c, say player 1, gets a payoff $\rho_1 v(c)$, and we say that the *value left* in c is $(1-\rho_1)v(c)$. Similarly, each player i, on visiting a cell c, gets a payoff that is ρ_i fraction of the value left in c at the time of its visit, leaving behind $(1-\rho_i)$ fraction of that value. Thus, if a sequence of k players say $(x_1, x_2, \ldots, x_k)$ visit c *one after the other*, then the i^{th} visitor x_i gets a payoff of $\rho_{x_i}\left(\prod_{j<i}(1-\rho_{x_j})\right)v(c)$ corresponding to *that* visit; if the same player is also the j^{th} visitor for some $j \neq i$, it will get a payoff for each such visit defined similarly. The combined payoff of all these players, from visiting c is

$$\left(1-\prod_{j\leqslant k}(1-\rho_{x_j})\right)v(c). \tag{1}$$

Note that this combined payoff is independent of the order of the players. Thereby, if $\{x_1, x_2, \ldots, x_k\}$ is a set of players visiting c *simultaneously*, that is, in the same time-step, then we say that they altogether capture the same amount of information as they would if they visit it *one after the other* in some order, and further, we define the payoff of player x_i as the share of this combined payoff proportional to ρ_{x_i}. Thus, this is equal to

$$\frac{\rho_{x_i}}{\sum_{j\leqslant k}\rho_{x_j}}\left(1-\prod_{j\leqslant k}(1-\rho_{x_j})\right)v(c). \tag{2}$$

Thus, the payoff of a player from a visit to a cell depends on which players visit the cell before it and which players visit simultaneously.

Without temporal aspect: In this case, the payoff from visiting a cell is determined at the end of the game, regardless of the order in which the players visit the cell. Since the order is immaterial, we can represent the visitors of a cell c as a multiset, say $\mathcal{P}'$, having support in $\mathcal{P}$ and an associated multiplicity function denoted by $m_{\mathcal{P}'}(\cdot)$. In case of no ambiguity, we drop the subscript and denote the multiplicity function as simply $m(\cdot)$. The payoff of a visitor from a single visit is precisely as in (2), and thus with possibly multiple visits, the payoff of player $i \in \mathcal{P}'$ is given by

$$\frac{\rho_i \; m(i)}{\sum\limits_{j \in \mathcal{P}'} \rho_j \; m(j)} \left(1 - \prod_{j \in \mathcal{P}'} (1 - \rho_j)^{m(j)} \right) v(c). \tag{3}$$

Next, we establish results for both these games on existence of pure equilibria, bound on PoA etc. For the latter, we establish the smoothness of these games, which is defined as follows [25]:

Definition 4. *(λ, μ)–**smoothness**: A payoff-maximization game — one where each player has a payoff function $\pi_i(s)$ that it strives to maximize — is called (λ, μ)–smooth if*

$$\forall s, s^* \in S, \; \sum_{i \in \mathcal{P}} \pi_i(s_i^*, s_{-i}) \geqslant \lambda \sum_{i \in \mathcal{P}} \pi_i(s^*) - \mu \sum_{i \in \mathcal{P}} \pi_i(s). \tag{4}$$

We first analyze the game with temporal aspect.

6.1 Multi-step Game with Temporal Aspect

In this section, we analyze the game with temporal aspect. As discussed before, the payoff of a player from visiting a cell is not merely dependent on which players visit the cell, but also on the order in which they visit the cell. The combined payoff, however, when a sequence of players visit a cell c (some of them possibly simultaneously), does not depend on their order, and can be easily computed as in (2). Let A, B be multisets with support in $\mathcal{P}$, with m_A, m_B the respective multiplicity functions, and for a cell c, let $\pi_B^A(c)$ denote the combined payoff which the visitors in B would obtain by visiting cell c (as many times as the respective multiplicities in B), when preceded by all (and only) the visitors as represented by A. Note that the multiset-representation is sufficient for this to be well-defined, since the order of visitors in A among themselves, and similarly that of visitors in B among themselves, does not matter when computing the said *combined* payoff. Indeed, the exact expression can be easily obtained to be

$$\pi_B^A(c) = \left(1 - \prod_{j \in B}(1-\rho_j)^{m_B(j)}\right) \prod_{j \in A}(1-\rho_j)^{m_A(j)} v(c). \tag{5}$$

Here, the entity $\prod_{j \in A}(1-\rho_j)^{m_A(j)}$ denotes the fraction of $v(c)$ left in c after visitors in A have visited, and the fraction of it collected by B is computed similarly.

With this notation, the following observations are immediate:

Lemma 3. *Let A, B, D be multisets with support in $\mathcal{P}$. Then, $\pi_B^A(c) + \pi_D^{A \uplus B}(c) = \pi_{B \uplus D}^A(c)$.*

Proof. Note that $B \uplus D$ is the multiset sum of B and D and thus represents the combined visitors in B as well as D. Thus, the result follows from the definition of $\pi_B^A(c)$, as both the sides equal the combined payoff of visitors in B and D when they are preceded by visitors in A. □

Lemma 4. *Let A, B, B' be multisets with support in $\mathcal{P}$ s.t. $B' \subseteq B$. Then, $\pi_{B'}^A(c) \leqslant \pi_B^A(c)$.*

Proof. Applying Lemma 3 on $A, B', B \backslash B'$ respectively, we get, $\pi_B^A(c) - \pi_{B'}^A(c) = \pi_{B \backslash B'}^{A \uplus B'}(c) \geqslant 0$. Rearranging gives the required result. □

Lemma 5. *Let A, A', B be multisets with support in $\mathcal{P}$ s.t. $A' \subseteq A$. Then, $\pi_B^A(c) \leqslant \pi_B^{A'}(c)$.*

Proof. Using (5), it is easy to see that $\pi_B^A(c)/\pi_B^{A'}(c) = \prod_{j \in A \backslash A'}(1-\rho_j)^{m_A(j) - m_{A'}(j)} \leqslant 1$. Rearranging gives the required result. □

Next, suppose c is a cell, and A is a multiset with support in $\mathcal{P}$. For each player j, let $S'_j \subseteq S_j$ denote the subset of strategies of player j in which j visits c exactly $m_A(j)$ times. Then, $S' = \times_j S'_j$ is the set of outcomes for which the multiset A precisely captures which players visit cell c and how often. In case of such an outcome, we refer to A as the *visitor set* for c. Also, for any multiset A and a player $i \in A$, let $A|_i$ denote the multiset $m_A(i) \otimes \{i\}$, i.e. $A|_i$ only contains i — with the same multiplicity as A — and let $A|_{-i}$ denote the multiset $A \backslash A_i$. We now show another important result.

Lemma 6. *Let A be any multiset with support in $\mathcal{P}$, and let $i \in A$. Let c be any cell and let $S' \subseteq S$ be the set of outcomes for which A is the visitor set for c. Then,*

$$\forall s \in S', \;\; \pi_i(s, c) \geqslant \pi_{A|_i}^{A|_{-i}}(c).$$

In other words, keeping the visitor set of a cell fixed, the payoff of a player from all its visits to the cell is no worse than the payoff it would get when all its visits are preceded by all the other visitors in the visitor set. While the result appears very intuitive, a rigorous proof can be found in the appendix.

Using these, we now show that this game is (1,1)–smooth. Let s and s^* be any two outcomes. For every player i, let q^i denote the outcome (s_i^*, s_{-i}). For any cell c, let multisets A_c, and A_c^* denote the visitor sets for cell c when the outcomes are s, s^* respectively. Note that when the outcome is q^i, the visitor set of cell c can be written as $A_c^*|_i \uplus A_c|_{-i}$. With this notation, we can write,

$$\begin{aligned} \sum_{i \in \mathcal{P}} \pi_i(s) &= \sum_{i \in \mathcal{P}} \sum_{c \in \mathcal{C}} \pi_i(s, c) \\ &= \sum_{c \in \mathcal{C}} \sum_{i \in \mathcal{P}} \pi_i(s, c) \\ &= \sum_{c \in \mathcal{C}} \pi_{A_c}^{\varnothing}(c). \end{aligned} \tag{6}$$

similarly, we have,

$$\sum_{i \in \mathcal{P}} \pi_i(s^*) = \sum_{c \in \mathcal{C}} \pi_{A_c^*}^{\varnothing}(c). \tag{7}$$

Finally, for the outcomes q^i, we can write,

$$\begin{aligned} \sum_{i \in \mathcal{P}} \pi_i(q^i) &= \sum_{i \in \mathcal{P}} \sum_{c \in \mathcal{C}} \pi_i(q^i, c) \\ &= \sum_{c \in \mathcal{C}} \sum_{i \in \mathcal{P}} \pi_i(q^i, c). \end{aligned} \tag{8}$$

Now, if player i does not visit cell c when playing strategy s_i^*, equivalently, if it is not contained in the visitor set A_c^*, its payoff from visiting c is simply zero. Thus, (8) becomes,

$$\sum_{i\in\mathcal{P}} \pi_i(q^i) = \sum_{c\in\mathcal{C}} \sum_{i\in A_c^*} \pi_i(q^i, c)$$
$$\geqslant \sum_{c\in\mathcal{C}} \sum_{i\in A_c^*} \pi_{A_c^*|_i}^{A_c|_{-i}}(c). \tag{9}$$

Here, (9) follows from Lemma 6. Finally, adding (6) and (9) and subtracting (7) gives,

$$\sum_{i\in\mathcal{P}} \pi_i(q^i) + \sum_{i\in\mathcal{P}} \pi_i(s) - \sum_{i\in\mathcal{P}} \pi_i(s^*) \geqslant \sum_{c\in\mathcal{C}} \left(\pi_{A_c}^{\varnothing}(c) + \sum_{i\in A_c^*} \pi_{A_c^*|_i}^{A_c|_{-i}}(c) - \pi_{A_c^*}^{\varnothing}(c) \right). \tag{10}$$

Now, we show that each term of the summation on the R.H.S of (10) is always non-negative, via the following lemma:

Lemma 7. *For a cell $c \in \mathcal{C}$, let A and A^* be the visitor sets of c for outcomes $s, s^* \in S$ respectively. Then,*

$$\pi_A^{\varnothing}(c) + \sum_{i\in A^*} \pi_{A^*|_i}^{A|_{-i}}(c) \geqslant \pi_{A^*}^{\varnothing}(c). \tag{11}$$

Proof. Included in the Appendix. □

Thus, it follows from (11) and (10) that,

$$\sum_{i\in\mathcal{P}} \pi_i(q^i) + \sum_{i\in\mathcal{P}} \pi_i(s) - \sum_{i\in\mathcal{P}} \pi_i(s^*) \geqslant 0$$
$$\sum_{i\in\mathcal{P}} \pi_i(q^i) + \sum_{i\in\mathcal{P}} \pi_i(s) \geqslant \sum_{i\in\mathcal{P}} \pi_i(s^*).$$

Thus, by Definition 4, we have the desired result:

Theorem 11. *The multi-step game with temporal aspect is $(1,1)$–smooth.*

As shown in [25], a (λ, μ)–smooth payoff-maximization game has a price of anarchy at most $\frac{1+\mu}{\lambda}$, and this bound applies to the PoA with respect to all equilibrium concepts(mixed, correlated and not just pure. Thus,

Corollary 3. *The multi-step game with temporal aspect has a price of anarchy at most 2.*

Thus, the game has a constant PoA bound independent of the number of players, number of time-steps. We now show with an example, that this bound is tight.

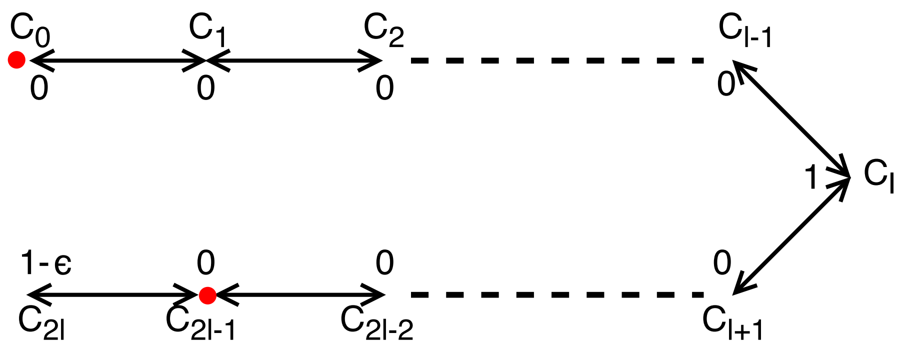

Figure 2: A 2-player game with temporal aspect. The only equilibrium has a social welfare of 1, the maximum possible being 2-ϵ.

Consider the game as shown in Figure 2. Let the number of time-steps be $l(\geqslant 2)$. The graph is a simple path as shown with cells c_i and c_{i+1} being neighbors of each other for each i. Initially, the information in all the cells is 0, except for c_l and c_{2l} - these 2 cells have an information of 1 and $1-\epsilon$ respectively, where ϵ is a small positive constant. There are two players with sensor effectiveness $\rho_1 = \rho_2 = 1$, initially in cells c_0 and c_{2l-1} respectively. Now, player 2 cannot grab information from both c_l and c_{2l} within l time-steps. Further, it follows that the path $c_{2l-1} \to c_{2l-2} \to \ldots \to c_l \to c_{l-1}$ is a dominant strategy for player 2, in which case it gets a payoff of 1 and in response, player 1 gets 0 from any strategy. Clearly, this is a pure equilibrium, leading to a social welfare of 1. On the other hand, if player 2 captures information from (only) c_{2l} allowing player 1 to capture from c_l which it can reach in the l^{th} time-step, then the social welfare can reach its maximum value of $2-\epsilon$. Since there is only 1 equilibrium, the PoS, as well as the PoA for this game is $2-\epsilon$, i.e., it approaches 2 as ϵ approaches 0. Thus, we have,

Theorem 12. *The multi-step game with temporal aspect has a PoS and PoA of 2.*

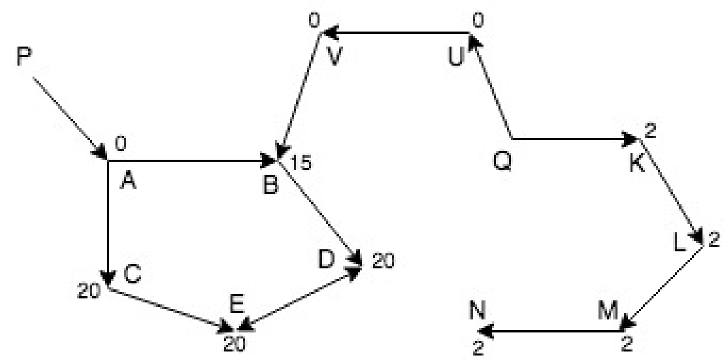

Figure 3: Counter-example where pure NE does not exist (Game with temporal aspect)

6.1.1 Non-existence of Pure NE

Unlike the single-step game, the general multi-step game may not always have a pure Nash Equilibrium, as demonstrated by the following example. Consider the game as shown in Figure 3. The connectivity between cells is given by the directed edges and the information initially available in each cell is shown. The number of time-steps is 4. Players 1 and 2, with sensor effectiveness $\rho_1 = \rho_2 = 1$ are initially in cells P and Q respectively. Thus, Player 1 has two strategies - paths $P \to A \to C \to E \to D$ and $P \to A \to B \to D \to E$. Let these be called 'Left' and 'Right' respectively. Player 2 similarly has 2 strategies, say 'Up' and 'Down' corresponding to paths $Q \to U \to V \to B \to D$ and $Q \to K \to L \to M \to N$. Then, the payoff matrix is given by,

$$\begin{array}{c} \\ Left \\ Right \end{array} \begin{array}{c} \begin{array}{cc} Up & Down \end{array} \\ \left[\begin{array}{cc} (50,25) & (60,8) \\ (55,0) & (55,8) \end{array} \right] \end{array}$$

It is easy to check, that there is no pure equilibrium in this case.

6.2 Multi-step Game without Temporal Aspect

In this section, we analyze the game without temporal aspect. When there is no temporal aspect, the payoff of a player from a visit to a cell is merely dependent on which players visit the cell over the complete course of the game and how many times, regardless of the order in which they visit the cell. Thereby, the combined payoff of players from their visits to a cell c (some of them possibly simultaneously), also does not depend on the order of visits, and can be easily computed using (3). Owing to the different setting in this game than the one with temporal aspect, we opt for a slightly different notation. Let A, B be multisets with support in $\mathcal{P}$, with m_A, m_B the respective multiplicity functions, such that $B \subseteq A$. Then, for a cell c, let $\theta^A_B(c)$ denote the combined payoff which the visitors in B would obtain by visiting cell c (as many times as the respective multiplicities in B), when the complete set of visitors for c is given by A. Naturally, this is only meaningful when $B \subseteq A$. With this notation, the following observations are immediate, and principally analogous to Lemma 3, 4, 5 respectively.

Lemma 8. *Let A, B, D be multisets with support in $\mathcal{P}$ s.t. $B \uplus D \subseteq A$. Then, $\theta^A_B(c) + \theta^A_D(c) = \theta^A_{B \uplus D}(c)$.*

Proof. Follows from definition, as both sides equal the combined payoff of visitors in B and D when the complete set of visitors is given by A. □

Lemma 9. *Let A, B, B' be multisets with support in $\mathcal{P}$ s.t. $B' \subseteq B \subseteq A$. Then, $\theta^A_{B'}(c) \leqslant \theta^A_B(c)$.*

Proof. Applying Lemma 8 on $A, B', B \backslash B'$ respectively, we get, $\theta^A_B(c) - \theta^A_{B'}(c) = \theta^A_{B \backslash B'}(c) \geqslant 0$. Rearranging gives the required result. □

Lemma 10. *Let A be a multiset with support in $\mathcal{P}$, and i be any player. Then, $\theta^A_{A|_i}(c) \geqslant \pi^{A|_{-i}}_{A|_i}(c)$.*

Proof. Note that in a game with temporal aspect, in the case when the visitor set of a cell c is A, one possible outcome s corresponds to all the visitors in A visiting in the same time-step, and thus, $\theta^A_{A|_i}(c)$ is a possible payoff of player i from cell c when its visitor set is fixed to A. Consequently, the result follows from Lemma 6. □

To show smoothness for this game, we proceed similarly as in the game with temporal aspect. Let s and s^* be any two outcomes. For every player i, let q^i denote the outcome $({s^*}_i, s_{-i})$. For any cell c, let multisets A_c, and A^*_c denote the visitor sets for cell c when the outcomes are s, and s^* respectively. Analogous to (6), (7), and (9), we can write,

$$\sum_{i\in\mathcal{P}} \pi_i(s) = \sum_{c\in\mathcal{C}} \theta^{A_c}_{A_c}(c), \tag{12}$$

$$\sum_{i\in\mathcal{P}} \pi_i(s^*) = \sum_{c\in\mathcal{C}} \theta^{A^*_c}_{A^*_c}(c), \tag{13}$$

$$\sum_{i\in\mathcal{P}} \pi_i(q^i) = \sum_{c\in\mathcal{C}} \sum_{i\in A^*_c} \theta^{A_c|_{-i} \uplus A^*_c|_i}_{A^*_c|_i}(c). \tag{14}$$

Next, adding (12) and (14) and subtracting (13) gives,

$$\sum_{i\in\mathcal{P}} \pi_i(q^i) + \sum_{i\in\mathcal{P}} \pi_i(s) - \sum_{i\in\mathcal{P}} \pi_i(s^*) \geqslant \sum_{c\in\mathcal{C}} \left(\theta^{A_c}_{A_c}(c) + \sum_{i\in A^*_c} \theta^{A_c|_{-i} \uplus A^*_c|_i}_{A^*_c|_i}(c) - \theta^{A^*_c}_{A^*_c}(c) \right). \tag{15}$$

Finally, note that $\theta^{A_c}_{A_c}(c) = \pi^{\varnothing}_{A_c}(c)$ and $\theta^{A^*_c}_{A^*_c}(c) = \pi^{\varnothing}_{A^*_c}(c)$ by definition. Further, applying Lemma 10 on $A_c|_{-i} \uplus A^*_c|_i$, we get, $\theta^{A_c|_{-i} \uplus A^*_c|_i}_{A^*_c|_i}(c) \geqslant \pi^{A_c|_{-i}}_{A^*_c|_i}(c)$. With this, (15) becomes,

$$\sum_{i\in\mathcal{P}} \pi_i(q^i) + \sum_{i\in\mathcal{P}} \pi_i(s) - \sum_{i\in\mathcal{P}} \pi_i(s^*) \geqslant \sum_{c\in\mathcal{C}} \left(\pi^{\varnothing}_{A_c}(c) + \sum_{i\in A^*_c} \pi^{A_c|_{-i}}_{A^*_c|_i}(c) - \pi^{\varnothing}_{A^*_c}(c) \right). \tag{16}$$

As seen in the previous section, the RHS of (16) is non-negative by Lemma 7, and in turn, so is the LHS, thus proving the desired result:

Theorem 13. *The multi-step game without temporal aspect is (1, 1)-smooth.*

Again, using the result in [25] as mentioned in the previous section, we have,

Corollary 4. *The multi-step game without temporal aspect has a price of anarchy at most 2.*

Next, we demonstrate the tightness of this bound via an example.

Consider the game as shown in Figure 4. The set of players is $\{1, \ldots, p\}$, with $\rho_1 = \rho_2 = \ldots = \rho_p = 1$. The number of time-steps is l. The cells in the network and the connectivity among them is as shown. The information initially available in cell c is 1, while in cells $c^2_l, \ldots c^p_l$, it is $v = \frac{l}{(p-1)l+1} - \epsilon$

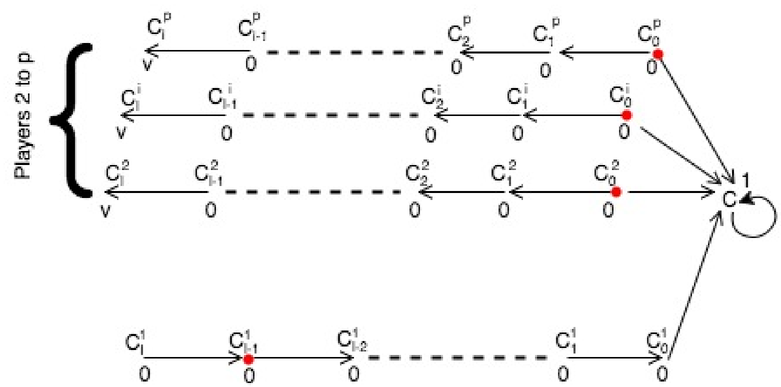

Figure 4: A game without temporal aspect where PoA can get arbitrarily close to 2 with a high number of players and number of time-steps

where ϵ is a small positive constant. The information available is 0 in all other cells. Player 1 is initially in cell c^1_{l-1}, while every other player $i > 1$ is initially in cell c^i_0. Now, the only strategy for player 1 is the path $c^1_{l-1} \rightarrow c^1_{l-2} \rightarrow \ldots \rightarrow c^1_0 \rightarrow c$. Now, the outcome where every other player $i > 1$, plays the strategy $c^i_0 \rightarrow c \rightarrow c \rightarrow \ldots \rightarrow c \rightarrow c$ gives every player $i > 1$, a payoff which evaluates to $\frac{l}{(p-1)l+1}$. Thus, no player i wants to deviate to the other possible strategy $c^i_0 \rightarrow c^i_1 \rightarrow \ldots \rightarrow c^i_l$ as it gives a smaller payoff of $\frac{l}{(p-1)l+1} - \epsilon$. Thus, the aforesaid outcome is a pure Nash Equilibrium, which has a social welfare of 1. However, it can be seen that the social welfare increases as more and more players switch to the respective alternative strategy, and in the extreme case of every player $i > 1$ switching to the respective strategy $c^i_0 \rightarrow c^i_1 \rightarrow \ldots \rightarrow c^i_l$, the social welfare reaches the maximum value of $2 - \frac{1}{(p-1)l+1} - (p-1)\epsilon$, giving the same value of PoA. Thus, as $\epsilon \rightarrow 0$, it approaches $2 - \frac{1}{(p-1)l+1}$, which in turn, can become arbitrarily close to 2 if the parameters p *or* l become arbitrarily large, showing that the bound of 2 is tight.

6.2.1 Non-existence of Pure NE

Unlike the single-step game, the general multi-step game without temporal aspect may not always have a pure Nash Equilibrium, as demonstrated by the following example.

Consider the game as shown in Figure 5. The connectivity between cells is given by the directed edges and the information initially available in each cell is shown. The number of time-steps is 3. Players 1 and 2, with sensor

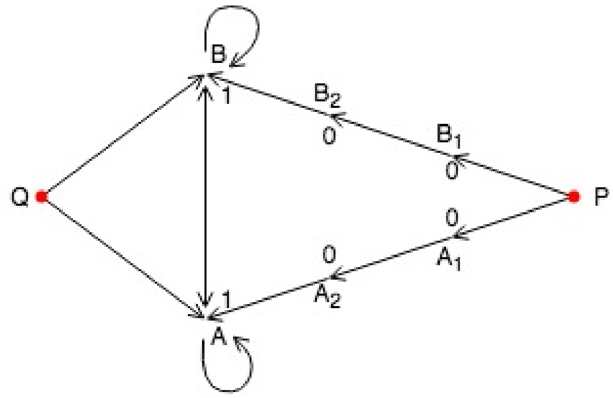

Figure 5: Counter-example where pure NE does not exist (Game without temporal aspect)

effectiveness $\rho_1 = \rho_2 = 0.8$ are initially in cells P and Q respectively. Thus, Player 1 has two strategies: paths $P \to A_1 \to A_2 \to A$ and $P \to B_1 \to B_2 \to B$. Let these be called 's_A' and 's_B' respectively. Player 2 has 8 strategies, however, since the sequence of the visits does not matter, there are 4 distinct ones. Let these be called 'A_3B_0' ,'A_2B_1' ,'A_1B_2' ,'A_0B_3', where 'A_iB_j' denotes a strategy which visits A i times and B j times. Then, the payoff matrix for this game is given by,

$$\begin{array}{c} \\ s_A \\ s_B \end{array} \begin{array}{cccc} A_3B_0 & A_2B_1 & A_1B_2 & A_0B_3 \\ \left[\begin{array}{c} (0.2496, 0.7488) \\ (0.8000, 0.9920) \end{array} \right. & \begin{array}{c} (0.3307, 1.4613) \\ (0.4800, 1.4400) \end{array} & \begin{array}{c} (0.4800, 1.4400) \\ (0.3307, 1.4613) \end{array} & \left. \begin{array}{c} (0.8000, 0.9920) \\ (0.2496, 0.7488) \end{array} \right] \end{array}$$

It is easy to check, that there is no pure equilibrium in this case.

CHAPTER 7

EXPERIMENTATION

In a competitive game environment like in this problem, a typical solution is to play the strategy corresponding to a Nash equilibrium. As seen from the previous sections, however, a pure equilibrium may not exist in the general multi-step game in either of the game settings. Since computing mixed equilibria is an intractable problem, we need to devise easily implementable strategies with reasonable performance guarantees. In this section, we propose such heuristics/algorithms which the UAVs can deploy as solutions to the routing problem. We simulate plausible problem scenarios with randomly generated game parameters and statistically compare these solutions on the grounds of social welfare optimality.

7.1 Setup

The randomly generated game instances are in the following setting. The set of cells is a 10×10 grid. To allow arbitrary connectivity constraints, we have a cell connectivity parameter δ which works as follows: For each cell, each of the cells within a Chebychev distance of 1 (i.e. row-wise, column-wise, or diagonal-wise adjacent cells and the cell itself) is independently chosen to be its out-neighbor with probability δ; if no cells get chosen after having gone through all the adjacent cells, we repeat the process until there is at least one out-neighbor for the cell, so that the graph does not have sinks where the UAVs can get stuck. We set the number of UAVs $p = 5$, the number of time-steps $l = 10$, and cell-connectivity parameter $\delta = 0.8$.

In every game instance, the sensor effectiveness parameters of the UAVs are initialized to values chosen uniformly at random (u.a.r.) from $[0, 1]$. The edges of the graph are chosen randomly as per the cell connectivity parameter δ as described above. The initial positions of the UAVs are chosen uniformly at random from the grid. The cell information values are chosen differently for different problem scenarios as follows:

> 1 *peak*: In this scenario, the distribution of the information values across cells as a function of their location is a generalized normal distribution. That is, we first pick a peak cell c from the grid u.a.r. and then pick information values in all cells as a function of distance d from the peak. This is represented as $v = f(d) = \alpha 2^{-\beta d^{\gamma}}$. Thereby, the peak cell gets a value of α and the value decreases exponentially with distance from the peak, as specified. The parameters α, β, and γ are chosen u.a.r. from ranges $[0.8, 1]$, $[0, 0.1]$, $[0.75, 1.25]$, the last of which (γ) acts as the shape parameter, and the ranges for the parameters were fixed in order to keep the information values in a good range.
>
> 2 or more *peaks*: In this kind of a scenario, we fix a small k such as 2,3 etc. Subsequently, we pick k peak cells from the grid u.a.r. and the information value function across all cells is composed of generalized normal distributions centered around the k peaks respectively. For any cell, the closest peak c_i ($i \leqslant k$) is determined, and the value in this cell is given by $v = f_i(d) = \alpha_i 2^{-\beta_i d_i^{\gamma}}$. Thus, peak c_i has a value α_i and the values diminish in the other cells as we move away from the respective peaks. The parameters α_i's, β_i's, and γ_i's are again chosen u.a.r. from the same ranges as scenario 1.
>
> No *peaks*: In this scenario, all the information values in the cells are simply chosen u.a.r. from $[0, 1]$.

As per the problem scenario, we generate $n = 1000$ game instances with all the parameters randomly chosen as described above. We let the UAVs apply certain heuristics to route in the search space and collect information accordingly. Finally, we compute the social welfare achieved for each game instance as the output of the heuristics, and the two sequences of 1000 outputs for the two strategies are compared is described in a following section.
Next, we describe the various heuristics we implement and compare.

7.2 Heuristics

We implement and compare the following heuristics:

1. Greedy myopic: At each time-step this heuristic is to simply go to the neighbor which currently has the maximum available information — disregarding any other UAVs in the environment.

2. Multi-horizon greedy (parameters h, ϵ): This heuristic considers walks of length h as strategies, while still disregarding any other UAVs in the environment. Since there are other UAVs however, the payoff that a UAV would expect from a strategy if it were the sole player, would be an over-estimate. To rectify this, the payoff from a cell x steps away is discounted by a ratio of ϵ^x when getting estimates for a strategy — due to the fact that the actual payoff received from a cell in expectation would be a smaller and smaller fraction of the optimistic estimate as the cell gets more and more time-steps away, allowing more UAVs to get the payoff first. Finally, the strategy with the highest such (discounted) estimate is chosen. For our experimentation, we set $h = 5$, and $\epsilon = 0.8$.

3. One-step NE: As a pure NE is guaranteed to exist in single-step games, this heuristic is to compute the pure NE that is obtained via Best-response dynamics starting from the initial outcome of everyone choosing the greedy myopic strategy. Each player computes this pure NE and plays the strategy corresponding to this NE.

4. Multi-step NE (parameter h): In this heuristic, a UAV tries to compute a pure NE via Best-response dynamics, for individual strategies being walks of length h. Since this is not guaranteed to exist, if a pure NE is not found within $2pl$ rounds of Best-response dynamics, the player reduces the horizon to $h - 1$ and repeats until a pure NE is found, for horizon say h', and plays the strategy corresponding to this NE. Note that $h' \geqslant 1$ since a pure NE always exists in the single-step game.

We compare the latter two against the former two respectively, so as to conclude that the heuristics involving the Nash equilibria, leveraging our result of its guaranteed existence in the single-step game, perform as good or better than the greedy behavior based strategies serving as benchmarks.

We now discuss how we compare the heuristics and the results obtained.

7.3 Evaluation and Results

As aforementioned, we first compute a sequence of 1000 outputs (social welfares) for any two heuristics h_1 and h_2 to be compared. Next, to compare h_1 versus h_2, we compute a metric which we call as the *rf-curve*. To do so, we compute 2-D points (r, f) based on the two sequences of outputs by varying r between 0.5 and 1.5, as follows: For each r, we compute the fraction of games f in which the ratio of h_1's output to that of h_2 is at least r. It is easy to check that f as a function of r, is monotone. Also, we note that the rf-curve definition is asymmetric — the rf-curve for h_2 versus h_1, can be very different. The two curves are also related though as can be easily seen — if (r, f) is a point on one, $(1/r, 1 - f)$ is a point on the other, and vice versa. For each pair of heuristics being compared, we simultaneously plot both the rf-curves for the ease of visual clarity. The two heuristics can be said to be roughly equally effective if the two rf-curves are close to being coincidental. On the other hand, if one of the rf-curves quite consistently dominates the other, the respective heuristic can be concluded to be more effective. Moreover, the extent of the improvement is proportional to the separation between the two rf-curves throughout the range of r.

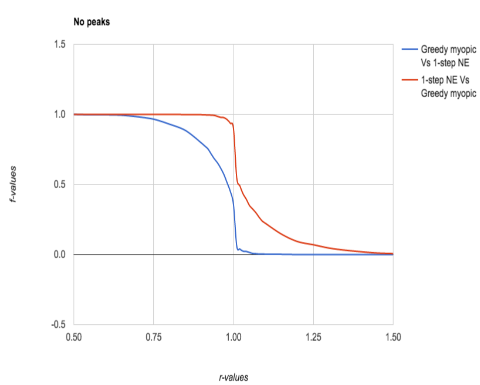

(a) No peaks

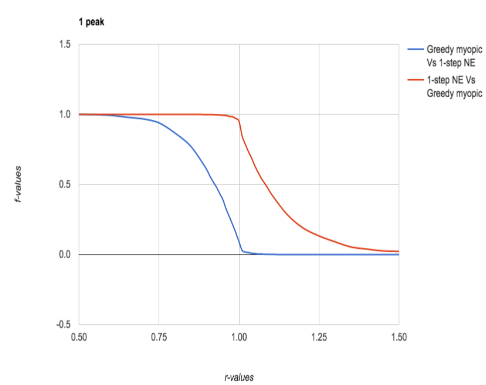

(b) One peak

Figure 6: Comparing the heuristics 'Greedy myopic' and '1-step NE' for scenarios 'No peaks' and '1 peak'

Figure 6 shows the results of comparison between the heuristics 'Greedy myopic' and '1-step NE' for scenarios 'No peaks' and '1 peak'. The plots for scenarios '2 peaks' and '3 peaks' closely resembled the one for '1 peak', and thus, are not included. It is easy to see that '1-step NE' is more effective than

‘Greedy myopic’ in both the scenarios; more so in the ‘1 peak’ scenario as reflected in the visibly higher degree of separation between the two rf-curves. In the ‘No peaks’ scenario, in 97.7% of the games, its output is no worse than 97% that of ‘Greedy myopic’, and in 89.9% it is as good or better, while it improves by 10% or more in 22.3% of the games. On the other hand, in the ‘1 peak’ scenario, in 95.2% of the games, its output is at least as good as that of ‘Greedy myopic’, while it improves by 10% or more in 43.6% of the games, and by 20% or more in 19.0% of the games.

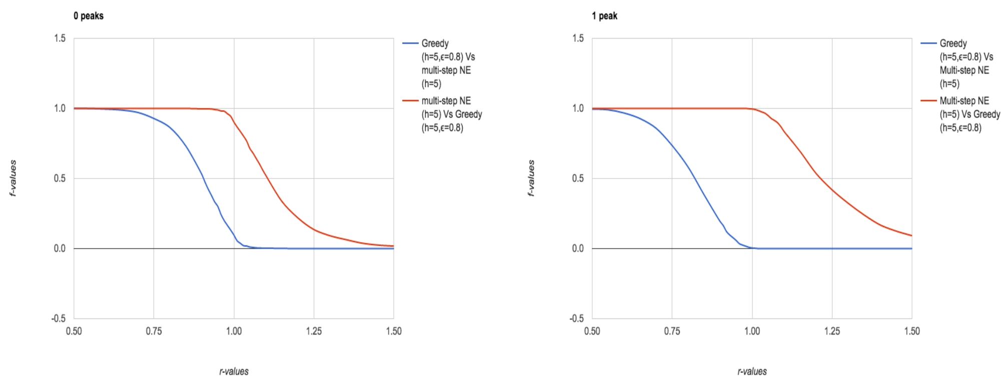

(a) No peaks (b) One peak

Figure 7: Comparing the heuristics ‘Multi-horizon greedy ($h = 5$, $\epsilon = 0.8$)’ and ‘Multi-step NE ($h = 5$)’ for scenarios ‘No peaks’ and ‘1 peak’

Figure 7 shows the results of comparison between the heuristics ‘Multi-horizon greedy ($h = 5$, $\epsilon = 0.8$)’ and ‘multi-step NE($h = 5$)’ for scenarios ‘No peaks’ and ‘1 peak’. Again, the plots for scenarios ‘2 peaks’ and ‘3 peaks’ closely resembled the one for ‘1 peak, and thus, are not included. It is easy to see that ‘multi-step NE’ is more effective than ‘Multi-horizon greedy’ in both the scenarios; more so in the ‘1 peak’ scenario as reflected in the visibly higher degree of separation between the two rf-curves. In the ‘No peaks’ scenario, in 98.0% of the games, its output is no worse than 97% that of the greedy benchmark, while it improves by 10% or more in 52.2% of the games, and by 20% or more in 21.9% of the games. On the other hand, in the ‘1 peak’ scenario, in 99.6% of the games, its output is at least as good as that of the greedy benchmark, while it improves by 10% or more in 83.5% of the games, by 20% or more in 53.6% of the games, and by 40% or more in 16.9% of the games.

CHAPTER 8

SUMMARY

Our main results are as follows. We first study the basic setting of two players and single time-step, albeit with the general formulation involving multiple modalities and information fusion. We prove that a pure NE always exists and is computable in linear time, in the case of *moderate* fusion as well as *significant* fusion. We also prove tight bounds on PoA and PoS in both these fusion cases. In case of moderate fusion, we also show that the number of mixed equilibria has an upper-bound that is linear in the number of strategies, and that the mixed equilibria are Pareto-dominated by the pure equilibria.

By restricting the game to single modality, we study the multi-player single-step case, and prove that a pure NE always exists in this game for both the fusion cases, and more generally, in all the singleton congestion games having a particular class of payoff functions. For the significant fusion case, we provide an efficient algorithm to compute a pure NE which runs in time linear in the number of players times the number of strategies. Finally, assuming no fusion and symmetric players, we prove a tight bound of $2 - 1/p$ on the PoA where p is the number of players.

Finally, by further considering no fusion, we study the multi-player multi-step game, in two differently defined classes of games. For each of them, we provide concrete counter-examples to show that pure NE may not exist. Further, for both the classes, we prove the (λ, μ)–smoothness for $\lambda = \mu = 1$ which leads to an upper bound of 2 on PoA and PoS, which we show to be tight with concrete examples.

Finally we provide simple heuristics for the UAV routing problem, which leverage on the result of existence of pure NE in the single-step game. By simulating a large number of games via randomly generating game parameters, in various plausible problem scenarios, we provide empirical bounds for how good a social welfare can be achieved with these heuristics as compared to the benchmark heuristics based on greedy behavior.

REFERENCES

[1] P. Sujit and D. Ghose, "Multiple agent search of an unknown environment using game theoretical models," in *American Control Conference, 2004. Proceedings of the 2004*, vol. 6. IEEE, 2004, pp. 5564–5569.

[2] R. Braybrook, "Three "D" missions–dull, dirty and dangerous." *Armada International*, vol. 28, no. 1, pp. 10–12, 2004.

[3] M. Kress and J. O. Royset, "Aerial search optimization model (ASOM) for UAVs in special operations," *Military Operations Research*, vol. 13, no. 1, pp. 23–33, 2008.

[4] K. P. Tognetti, "An optimal strategy for a whereabouts search," *Operations Research*, vol. 16, no. 1, pp. 209–211, 1968.

[5] J. B. Kadane, "Discrete search and the Neyman-Pearson lemma," *Journal of Mathematical Analysis and Applications*, vol. 22, no. 1, pp. 156–171, 1968.

[6] J. B. Kadane, "Optimal whereabouts search," *Operations Research*, vol. 19, no. 4, pp. 894–904, 1971.

[7] L. D. Stone, *Theory of Optimal Search*. Academic Press, Inc., 1975.

[8] R. Ahlswede and I. Wegener, *Search Problems*. Hoboken, New Jersey, USA: John Wiley & Sons, 1987.

[9] H. Ortiz-Pena, M. J. Hirsch, M. Karwan, R. Nagi, and M. Sudit, "Representation of potential information gain to measure the price of anarchy on isr activities," *Proc. of the SPIE Next Generation Analyst*, April 2012.

[10] E. F. Nakamura, A. A. F. Loureiro, and A. C. Frery, "Information fusion for wireless sensor networks: Methods, models, and classifications," *ACM Computing Surveys*, vol. 39, no. 3, 2007.

[11] R. R. Brooks and S. S. Iyengar, *Multi-sensor Fusion: Fundamentals and Applications with Software*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1998.

[12] C. Schumacher, P. R. Chandler, M. Pachter, and L. S. Pachter, "Optimization of air vehicles operations using mixed-integer linear programming," *Journal of the Operational Research Society*, vol. 58, no. 4, pp. 516–527, 2006.

[13] C. C. Murray and M. H. Karwan, "An extensible modeling framework for dynamic reassignment and rerouting in cooperative airborne operations," *Naval Research Logistics*, vol. 57, no. 7, pp. 634–652, 2010.

[14] A. Hayrapetyan, E. Tardos, and T. Wexler, "A network pricing game for selfish traffic," in *Proceedings of the Twenty-fourth Annual ACM Symposium on Principles of Distributed Computing*, ser. PODC '05. New York, NY, USA: ACM, 2005, pp. 284–291.

[15] T. Roughgarden and E. Tardos, "How bad is selfish routing?" *J. ACM*, vol. 49, no. 2, pp. 236–259, Mar. 2002.

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

[17] I. Milchtaich, "Congestion games with player-specific payoff functions," *Games and Economic Behavior*, vol. 13, no. 1, pp. 111 – 124, 1996.

[18] H. Ackermann, H. Röglin, and B. Vöcking, "Pure nash equilibria in player-specific and weighted congestion games," in *International Workshop on Internet and Network Economics.* Springer, 2006, pp. 50–61.

[19] K. Bhawalkar, M. Gairing, and T. Roughgarden, "Weighted congestion games: the price of anarchy, universal worst-case examples, and tightness," *ACM Transactions on Economics and Computation*, vol. 2, no. 4, p. 14, 2014.

[20] T. Harks, M. Klimm, and R. H. Möhring, "Characterizing the existence of potential functions in weighted congestion games," in *International Symposium on Algorithmic Game Theory.* Springer, 2009, pp. 97–108.

[21] V. Gkatzelis, K. Kollias, and T. Roughgarden, "Optimal cost-sharing in general resource selection games," *Operations Research*, vol. 64, no. 6, pp. 1230–1238, 2016.

[22] X. Chen, X. Deng, and S.-H. Teng, "Settling the complexity of computing two-player nash equilibria," *Journal of the ACM (JACM)*, vol. 56, no. 3, p. 14, 2009.

[23] G. Gottlob, G. Greco, and F. Scarcello, "Pure nash equilibria: hard and easy games," *Journal of Artificial Intelligence Research*, vol. 24, pp. 357–406, 2005.

[24] D. Monderer and L. Shapley, "Potential games," *Games and Economic Behavior*, vol. 14, no. 1, pp. 124–143, 1996.

[25] T. Roughgarden, "Intrinsic robustness of the price of anarchy," *Journal of the ACM (JACM)*, vol. 62, no. 5, p. 32, 2015.

APPENDIX A

Proof of Lemma 6

We first establish an intermediate result required for the proof.

Lemma 11. *Let* $0 < y \leqslant 1$*, and let* X *be a set of* $n(\geqslant 0)$ *numbers such that* $\forall x \in X, 0 < x \leqslant 1$*. Then,*

$$\prod_{x \in X}(1-x) \leqslant \frac{1-(1-y)\prod_{x \in X}(1-x)}{y+\sum_{x \in X} x} \leqslant \frac{1-\prod_{x \in X}(1-x)}{\sum_{x \in X} x}.$$

Proof. Let $S = \sum_{x \in X} x$, and $P = \prod_{x \in X}(1-x)$. We want to prove,

$$P \leqslant \frac{1-(1-y)P}{y+S} \leqslant \frac{1-P}{S}.$$

Now, consider the first inequality.

$$\begin{aligned} & P \leqslant \frac{1-(1-y)P}{y+S} \\ \Leftrightarrow \quad & (y+S)P \leqslant 1-(1-y)P \\ \Leftrightarrow \quad & (1+S)P \leqslant 1. \end{aligned}$$

Next, the second inequality is,

$$\begin{aligned} & \frac{1-(1-y)P}{y+S} \leqslant \frac{1-P}{S} \\ \Leftrightarrow \quad & (1-P+P\rho)S \leqslant (y+S)(1-P) \\ \Leftrightarrow \quad & yPS \leqslant y(1-P) \\ \Leftrightarrow \quad & (1+S)P \leqslant 1. \end{aligned} \tag{17}$$

Thus, proving (17) proves both the inequalities. Now, by the AM-GM inequality, we have,

$$\begin{aligned}\frac{1}{n}\left(\sum_{x\in X}(1-x)\right) &\geqslant \left(\prod_{x\in X}(1-x)\right)^{1/n}\\ (n-S)/n &\geqslant P^{1/n}\\ (1-S/n)^n &\geqslant P. \qquad (18)\end{aligned}$$

Also, using the binomial theorem,

$$(1+S/n)^n = \sum_{i=0}^{n}\left(\binom{n}{i}(S/n)^i\right) \geqslant 1+S. \qquad (19)$$

Hence, combining (18) and (19), we get,

$$(1+S)P \leqslant (1+S/n)^n(1-S/n)^n = (1-S^2/n^2)^n \leqslant 1.$$

Thus, this proves (17) as required, and thereby, this Lemma. □

With this, we now prove Lemma 6.

Lemma 6. *Let A be any multiset with support in $\mathcal{P}$, and let $i \in A$. Let c be any cell and let $S' \subseteq S$ be the set of outcomes for which A is the visitor set for c. Then,*

$$\forall s \in S',\ \pi_i(s,c) \geqslant \pi_{A|_i}^{A|_{-i}}(c).$$

Proof. Let s be an outcome with the visitor set for c being A. Suppose player i visits c m times when the visitor set is A. The outcome s can be naturally associated with two well-defined sequences $X_1, \ldots, X_m$ and $Y_1, \ldots, Y_m$ as follows. For each j, X_j denotes the set of players visiting c in the same time-step as the j^{th} visit of player i, and Y_j denotes the multiset of all the visitors visiting strictly before. Naturally, each X_j must contain player i, and each Y_j must contain i with a multiplicity of $j-1$. Also, we must have $\varnothing \subseteq Y_1 \subset Y_2 \ldots \subset Y_m \subseteq A\backslash\{i\}$ by definition. Let v be the information initially available in c. The combined payoff of the visitors in X_j, as per our notation, is $\pi_{X_j}^{Y_j}(c)$. Player i gets a share of it proportional to ρ_i, and this summed over all the visits gives its total payoff from visiting c:

$$\pi_i(s,c) = \sum_{j=1}^{m} \frac{\rho_i}{\sum\limits_{k \in X_j} \rho_k} \pi_{X_j}^{Y_j}(c). \tag{20}$$

For each j, let v_j denote the information available in c just before the visitors in X_j visit, which evaluates to $v_j = v \prod\limits_{k \in Y_j} (1-\rho_k)^{m_{Y_j}(k)}$. Then, in terms of v_j, we can write

$$\pi_{X_j}^{Y_j}(c) = v_j \left(1 - \prod_{k \in X_j} (1-\rho_k) \right).$$

Hence, (20) can be written as

$$\pi_i(s,c) = \sum_{j=1}^{m} \rho_i v_j \frac{\left(1 - \prod\limits_{k \in X_j} (1-\rho_k) \right)}{\sum\limits_{k \in X_j} \rho_k}. \tag{21}$$

Now, let s' be another outcome with similarly defined sequences $X'_1, \ldots, X'_m$ and $Y'_1, \ldots, Y'_m$ such that the only difference from s is that for each j, the players visiting c in the same time-step as the j^{th} visit of player i in the outcome s, now visit strictly before it, in s'. Formally, for each j, we have $Y'_j = Y_j \uplus (X_j \backslash \{i\})$ and $X'_j = \{i\}$. Thus, the visitor set remains A for s'. Now, by definition, v_j is the information available in c after the visitors in Y_j have visited it. Hence the information left after players in $X_j \backslash \{i\})$ subsequently visit it, is $v_j \prod\limits_{k \in X_j \backslash \{i\}} (1-\rho_k)$. Hence, as player i's j^{th} visit to c follows, it gets a payoff that is ρ_i fraction of the value available, and this summed over all the visits gives the total payoff of player i from visiting c, for the outcome s':

$$\pi_i(s',c) = \sum_{j=1}^{m} \left(\rho_i v_j \prod_{k \in X_j \backslash \{i\}} (1-\rho_k) \right). \tag{22}$$

Now, applying the first inequality from Lemma 11 on $y = \rho_i, X = \{\rho_k | k \in X_j \backslash \{i\}\}$, we get,

$$\pi_i(s',c) \leqslant \sum_{j=1}^{m} \rho_i v_j \frac{\left(1 - (1-\rho_i) \prod\limits_{k \in X_j \backslash \{i\}} (1-\rho_k) \right)}{\rho_i + \sum\limits_{k \in X_j \backslash \{i\}} \rho_k}$$

$$= \sum_{j=1}^{m} \rho_i v_j \frac{\left(1 - \prod_{k \in X_j} (1 - \rho_k)\right)}{\sum_{k \in X_j} \rho_k}$$

$$= \pi_i(s, c). \tag{23}$$

Further, let s'' be another outcome with similarly defined sequences $X''_1, \ldots, X''_m$ and $Y''_1, \ldots, Y''_m$ and the visitor set for c being still A, such that all the visits of player i are strictly after all the visits of all the players. Formally, for each j, we have $Y''_j = A|_{-i} \uplus ((j-1) \otimes \{i\})$ and $X''_j = \{i\}$. Further, $\pi_i(s'', c) = \pi_{A|_i}^{A|_{-i}}(c)$. Now, since s' and s'' are such that all the visits of player i are unaccompanied, we can write its payoff from the j^{th} visit as simply $\pi_{\{i\}}^{Y'_j}(c)$ and $\pi_{\{i\}}^{Y''_j}(c)$ respectively. Now, the multiplicity of i is the same (= j-1) in Y'_j and Y''_j, whereas all other players reside in Y''_j with the maximum multiplicity possible for the visitor set A. Thus, $Y'_j \subseteq Y''_j$. Hence, using Lemma 5, we get,

$$\begin{aligned}
\forall j: \quad \pi_{\{i\}}^{Y'_j}(c) &\geqslant \pi_{\{i\}}^{Y''_j}(c) \\
\sum_{j=1}^{m} \pi_{\{i\}}^{Y'_j}(c) &\geqslant \sum_{j=1}^{m} \pi_{\{i\}}^{Y''_j}(c) \\
\pi_i(s', c) &\geqslant \sum_{j=1}^{m} \pi_{\{i\}}^{A|_{-i} \uplus ((j-1) \otimes \{i\})}(c) && \text{(By definition of } Y''_j\text{)} \\
&\geqslant \sum_{j=1}^{m} \left(\pi_{A|_{-i} \uplus (j \otimes \{i\})}^{\varnothing}(c) - \pi_{A|_{-i} \uplus ((j-1) \otimes \{i\})}^{\varnothing}(c) \right) \\
& && \text{(Using Lemma 3)} \\
&\geqslant \pi_{A|_{-i} \uplus (m \otimes \{i\})}^{\varnothing}(c) - \pi_{A|_{-i}}^{\varnothing}(c) \\
&\geqslant \pi_{(m \otimes \{i\})}^{A|_{-i}}(c) && \text{(Using Lemma 3)} \\
&\geqslant \pi_{A|_i}^{A|_{-i}}(c). && (24)
\end{aligned}$$

Thus, it follows from (23) and (24), that, $\pi_i(s, c) \geqslant \pi_{A|_i}^{A|_{-i}}(c)$. Since this holds for any outcome s for which the visitor set for cell c is A, the lemma is proved. □

Proof of Lemma 7

Lemma 7. *For a cell $c \in \mathcal{C}$, let A and A^* be the visitor sets of c for outcomes $s, s^* \in S$ respectively. Then,*

$$\pi_A^{\varnothing}(c) + \sum_{i \in A^*} \pi_{A^*|_i}^{A|_{-i}}(c) \geqslant \pi_{A^*}^{\varnothing}(c). \tag{11}$$

Proof. Proof We will show that,

$$\pi_A^{\varnothing}(c) + \sum_{i \in A^* \setminus A} \pi_{A^*|_i}^{A|_{-i}}(c) \geqslant \pi_{A \cup A^*}^{\varnothing}(c). \tag{25}$$

(Note that for multisets, $A^* \setminus A$ is defined to contain those elements which have a greater multiplicity in A^* than in A; and their multiplicity in $A^* \setminus A$ is precisely the difference of multiplicities in A^* and A.)
It is easy to see that the LHS of (11) is no less than that of (25) since the latter possibly excludes some terms in the summation, and each term is non-negative by definition. It is also easy to see that the RHS of (11) is no greater than that of (25) by Lemma 4. Thus, it suffices to prove (25) to prove this lemma.

Now we prove (25) by induction on the number of players in $A^* \setminus A$, denoted by, say, a.
The base case $a = 0$ is when $A^* \subseteq A$, i.e., $A^* \setminus A = \varnothing$. This holds trivially, as both the sides of the inequality to be proven, become equal to $\pi_A^{\varnothing}(c)$.
Assume, as inductive hypothesis, that (25) holds whenever $a < a_0$, for some $a_0 \in \mathbb{Z}^+$.
Now, consider the case when $a = a_0$. Arbitrarily fix some $x \in A^* \setminus A$. Then, $(A^* \setminus A)|_{-x}$, or equivalently, $A^*|_{-x} \setminus A$ has $a_0 - 1$ distinct elements. We can now write,

$$\pi_A^{\varnothing}(c) + \sum_{i \in A^* \setminus A} \pi_{A^*|_i}^{A|_{-i}}(c)$$
$$= \left(\pi_A^{\varnothing}(c) + \sum_{\substack{i \in A^* \setminus A \\ i \neq x}} \pi_{A^*|_i}^{A|_{-i}}(c) \right) + \pi_{A^*|_x}^{A|_{-x}}(c)$$

$$
\begin{aligned}
&= \left(\pi_A^{\varnothing}(c) + \sum_{i \in A^*|_{-x} \setminus A} \pi_{(A^*|_{-x})|_i}^{A|_{-i}}(c)\right) + \pi_{A^*|_x}^{A|_{-x}}(c) \\
&\geqslant \pi_{A \cup (A^*|_{-x})}^{\varnothing}(c) + \pi_{A^*|_x}^{A|_{-x}}(c) && \text{(using the Ind. Hyp.)} \\
&\geqslant \pi_{(A \cup A^*)|_{-x}}^{\varnothing}(c) + \pi_{A^*|_x}^{A|_{-x}}(c) && \text{(since } (A \cup A^*)|_{-x} \subseteq A \cup (A^*|{-x})\text{)} \\
& && \text{(and using Lemma 4)} \\
&\geqslant \pi_{(A \cup A^*)|_{-x}}^{\varnothing}(c) + \pi_{A^*|_x}^{(A \cup A^*)|_{-x}}(c) && \text{(since } (A \cup A^*)|_{-x} \supseteq A|_{-x}\text{)} \\
& && \text{(and using Lemma 5)} \\
&= \pi_{(A \cup A^*)|_{-x}}^{\varnothing}(c) + \pi_{(A \cup A^*)|_x}^{(A \cup A^*)|_{-x}}(c) && \text{(since } x \in A^* \setminus A \Rightarrow (A \cup A^*)|_x = A^*|_x\text{)} \\
&= \pi_{A \cup A^*}^{\varnothing}(c).
\end{aligned}
$$

Hence, this completes the inductive step and the proof by induction for (25), as required. □