

Collaborative Multi Agent Physical Search with Probabilistic Knowledge

Noam Hazon Yonatan Aumann Sarit Kraus

Department of Computer Science

Bar-Ilan University

Ramat Gan Israel 52900

{hazonn,aumann,sarit}@cs.biu.ac.il

Abstract

This paper considers the setting wherein a group of agents (e.g., robots) is seeking to obtain a given tangible good, potentially available at different locations in a physical environment. Traveling between locations, as well as acquiring the good at any given location consumes from the resources available to the agents (e.g., battery charge). The availability of the good at any given location, as well as the exact cost of acquiring the good at the location is not fully known in advance, and observed only upon physically arriving at the location. However, a-priori probabilities on the availability and potential cost are provided. Given such as setting, the problem is to find a strategy/plan that maximizes the probability of acquiring the good while minimizing resource consumption. Sample applications include agents in exploration and patrol missions, e.g., rovers on Mars seeking to mine a specific mineral. Although this model captures many real world scenarios, it has not been investigated so far.

We focus on the case where locations are aligned along a path, and study several variants of the problem, analyzing the effects of communication and coordination. For the case that agents can communicate, we present a polynomial algorithm that works for any fixed number of agents. For non-communicating agents, we present a polynomial algorithm that is suitable for any number of agents. Finally, we analyze the difference between homogeneous and heterogeneous agents, both with respect to their allotted resources and with respect to their capabilities.

1 Introduction

In many Multi-Agent Settings (MAS), agents need to explore the environment in order to fulfill their task. For example, consider Rover robots seeking to mine a certain mineral on the face of Mars. While there may be prior knowledge regarding *candidate* mining sites (e.g., based on satellite images), the actual availability at any given location may only be determined upon reaching the location. Furthermore, the

cost associated with the mining, e.g., in terms of battery consumption, may depend on the exact conditions at each site (e.g., soil type, terrain, etc.), and hence fully known only upon reaching the site. In fact, in some cases the cost may be prohibitive - i.e. when the Rover lacks sufficient battery charge. Thus, successful exploration of the environment is crucial for completing the task. However, in physical environments, exploration itself comes at a cost, namely - battery power for travel. Thus, while exploration is essential for mining, the two competes with each other for resources.

In physical environments, exploration itself entails complex tradeoffs, as traveling to one site may increase, or decrease, the distance to other sites. Thus, with multiple agents at hand, geographically subdividing the search space among the different agents may be the way to go. However, if agents have means of communication, then they may not wish to become too distant, as they can call upon each other for assistance. For example, even if a Rover does not have sufficient battery power for mining at a given location, it may be useful for it to travel to the site in order to determine the exact mining cost, and call for other robots that do have the necessary battery power. In this case, the scheduling of the robots' travel times is key, and must be carefully planned.

Finally, agents may be of different types, or with different amounts of resources. For example, Rover robots may be entering the mission with differing initial battery charges. They may also differ in their capabilities, like a team of rovers where some were specifically designed for mining missions, and thus require less battery power for the same mining task.

This paper aims at taking the first steps in understanding the characteristics of such multi-agent physical environment settings, and developing efficient exploration strategies for the like. To the best of our knowledge, it is the first to do so. As a start, we focus on the case where the mining sites are located along a path, as in the case of a perimeter patrol by a team of robots. We note that many multi-agent coverage algorithms convert their complex environment into a simple long path [Spires and Goldsmith, 1998; Gabrieli and Rimón, 2001; Hazon and Kaminka, 2005]. Furthermore, the problem in more general metric spaces can be shown to be NP-complete, even for the planar graphs. We also focus on the case where mining costs are rounded/estimated to one of a constant number of possible options (e.g., one, two or three hours).

We consider two variants of the problem. In the first variant, coined *Max-Probability*, we are provided with a group of agents, each with an initial resource budget (e.g., battery charge), and the goal is to maximize the probability of successfully completing the task (e.g., obtain the mineral). In the second variant, coined *Min-Budget*, we are required to guarantee some pre-determined success probability, and the goal is to minimize the initial resource allotments necessary in order to achieve said success probability.

Of course, Mars rovers are only one example of the general setting of exploration in a physical environment, and the discussion and results of this paper are relevant to any such setting, provided that exploration and fulfilling the task use the same type of resource. Another example would be a setting where agents need to acquire a good, potentially available at one of several shops, but need to pay for transportation from one shop to another.

Results. We separately consider the setting where agents can communicate and the setting where they cannot. For non-communicating agents we show a polynomial algorithm for the *Max-Probability* problem that is suitable for any number of agents. For the *Min-Budget* problem with non-communicating agents, we present a polynomial algorithm for the case that all agents must be allotted identical resources, but show that the problem is NP-hard for the general case (unless the number of agents is fixed). Next we consider agents that can communicate, and can call upon each other for assistance. As noted above, in this case the scheduling of the different agents' moves must also to be carefully planned. We present polynomial algorithms for both the *Max-probability* problem and the *Min-Budget* problem that work for any constant number of agents (but become non-polynomial when the number of agents is not constant). Finally, we extend our results to the case of heterogenous agents with different capabilities.

1.1 Related Work

Models of a single agent search process with prior probabilistic knowledge have been studied in the economic literature for years, promoting several reviews [Lippman and McCall, 1976; McMillan and Rothschild, 1994]. They have also been extended to multi-agent environments in [Sarne and Kraus, 2005]. Nevertheless, these economic-based search models assume that the cost associated with observing a given opportunity is stationary (i.e., does not change along the search process). While this assumption facilitates the analysis of search models, it is frequently impractical in the physical world. The use of changing search costs suggests an optimal search strategy structure different from the one used in traditional economic search models: other than deciding when to terminate its search, the agent needs to integrate into its decision making process exploration sequence considerations.

Changing search costs has been previously considered in the MAS domain in the context of Graph Search Problems [Koutsoupias *et al.*, 1996]. Here, the agent is seeking a single item, and a distribution is defined over all probabilities of finding it in each of the graph's nodes [Ausiello *et al.*, 2000]. Nevertheless, upon arriving at a node the success factor is binary: either the item is there or not. Extensions of these

applications to scenarios where the item is mobile are of the same character [Gal, 1980; Koopman, 1980].

The work of [Aumann *et al.*, 2008] is the first to analyze physical search problems with the assumption of prior probabilistic knowledge. Their work provides fundamental results for the single-agent case, showing that a physical search problem is hard on metric spaces and analyzing the case where the locations are aligned along a path like in [Spires and Goldsmith, 1998; Gabriely and Rimon, 2001]. Even under these settings some problems remain hard, unless the number of possible costs is constant. Unfortunately, their extension to the multi-agent case handles only a basic model. In their model all the resources and costs are shared among a group of homogeneous agents, with a simple coordination mechanism. They also assume that communication is available all the time among all the agents. Their assumptions may be realistic for a group of agents that has the same bank account and they charge it simultaneously for the use of movement and purchase. It is not applicable to a variety of physical search tasks where each agent has its own private budget, like in the Rover robots example where each one has its own battery (corresponding to its private budget) that it uses for movement and mining.

1.2 Terminology and Definitions

We are provided with m sites - $S = \{u_1, \dots, u_m\}$, which represent potential mining locations, together with a distance function $dis : S \times S \rightarrow R^+$ - determining the travel costs between any two sites. Since we focus on the case in which the sites are all on a single path we can assume that, WLOG (without loss of generality) all sites are points located on the line, and do away with the distance function dis . Rather, the distance between u_i and u_j is simply $|u_i - u_j|$. Furthermore, WLOG we may assume that the sites are ordered from left-to-right, i.e. $u_1 < u_2 < \dots < u_m$. We are also provided with a cost probability function $p_i(c)$ - stating the probability that the cost of obtaining the good at location i is c . Let D be the set of distinct costs with a non-zero probability, and let $d = |D|$. We assume that d is bounded and that the actual cost at a site is only revealed upon reaching it. In addition, we are provided with k agents and a vector of their initial locations $(u_s^{(1)}, \dots, u_s^{(k)})$, each of which is assumed to be, WLOG, one of the sites u_i (the probability of obtaining the good at this site may be 0). Again, WLOG we may assume that the agents are ordered from left-to-right, i.e. $u_s^{(1)} < u_s^{(2)} < \dots < u_s^{(k)}$. Finally, each agent j has its own initial budget B_j (unlike the shared budget model proposed by [Aumann *et al.*, 2008]).

Given these inputs, the goal is to find a plan that maximizes the probability of obtaining the good (by any of the agents), while minimizing the necessary budget. We assume that the goal is not individualized; the agents seek to obtain only one good and having multiple goods is not beneficial. Furthermore, they do not care which agent will obtain the good. The standard approach in such multi-criterion optimization problems is to optimize one of the objectives while bounding the other. In our case, we get two concrete problem formulations:

1. *Max-Probability*: given initial budgets B_j , for each agent j , maximize the probability of obtaining the good.

2. *Min-Budget*: given a target success probability p_{succ} , minimize the agents' initial budgets necessary to guarantee obtaining the good with a probability of at least p_{succ} .

In the *Min-Budget* problem it is also important to distinguish between two different agents models:

- *Identical budgets*: the initial budgets of all the agents must be the same. The problem is to minimize this initial budget, and we denote the problem as *Min-Budget*^{identical}.
- *Distinct*: the agents' initial budgets may be different. In this case the problem is to minimize the average initial budget, and we denote the problem as *Min-Budget*^{distinct}.

2 Non-Communicating Agents

We first consider the case where agents cannot communicate with each other. In this case agents cannot assist each other. Hence a solution is a *strategy* comprising of a set of ordered lists, one for each agent, determining the sequence of sites this agent must visit.

The success probability of a strategy is the probability that at least one of the agents will succeed in its task. Technically, it is easier to calculate the complementary failure probability: the probability that all the agents will not succeed in their tasks. For example, suppose that the sites and agents are located as illustrated in Figure 1, and consider the illustrated strategy. This strategy fails if for both agents and each of the sites they visit the cost of the item is higher than their remaining budget. This will happen with probability $\frac{1}{2} \cdot \frac{1}{4} \cdot \frac{1}{2} \cdot \frac{4}{5} = \frac{1}{20}$. Hence, the success probability of this strategy is $\frac{19}{20}$.




Figure 1: A possible input with a suggested strategy. The numbers on the edges represent traveling costs. A table at each site u_i represents the cost probability function $p_i(c)$. The strategy of each agent is illustrated by arrows.

We start by considering the *Max-Probability* problem. We prove:

Theorem 1 *In the no communication case if the number of possible costs is constant then Max-Probability can be solved in polynomial time for any number of agents.*

The proof is based on the following definitions and lemmata.

Note that multiple strategies may result in the same success probability. In this case we say that the strategies are *equivalent*. In particular there may be more than one optimal strategy.

Definition Let S be a strategy. Agents i and $\bar{i}$ are said to be *separated* by S if each site that is reached by i is not reached by $\bar{i}$.

Lemma 2 *If agents i and $\bar{i}$ are not separated by any optimal strategy. Then in any optimal strategy at least one of these agents must pass the initial location of the other.*

Proof WLOG assume that i is on the right side of $\bar{i}$. Consider an optimal strategy S . Let r be the rightmost site that is reached by i and $\bar{l}$ the leftmost site that is reached by $\bar{i}$. Assume that none of the agents passes the initial location of the other in S . Thus, there is at least one site between their initial locations that is reached by both agents. WLOG assume that $\bar{i}$ reaches at least one site with a higher budget than i 's remaining budget when reaching it, and denote by $\bar{r}^*$ the rightmost such site. Consider the following modified strategy: i goes according to S till the stage it has to reach $\bar{r}^*$. If i did not reach r yet then instead of reaching $\bar{r}^*$ it goes all the way straight to r . Otherwise, it stops just before reaching $\bar{r}^*$. $\bar{i}$ goes according to S till the stage it has to reach $\bar{r}^*$. If $\bar{i}$ did not reach $\bar{l}$ yet then after reaching $\bar{r}^*$ it goes all the way straight to $\bar{l}$. Otherwise, it stops after reaching $\bar{r}^*$. Agents i and $\bar{i}$ are separated by this strategy and it has at least the same success probability as S , in contradiction. ■

Lemma 3 *Suppose that agents i and $\bar{i}$ are not separated by any optimal strategy. Let S be an optimal strategy. Suppose that in S agent i passes the initial location of agent $\bar{i}$ and agent $\bar{i}$ does not stay in its initial location. Then, there is an optimal strategy such that one of the following holds:*

- $\bar{i}$ moves only in one direction which is opposite to the final movement's direction of i . Furthermore, if the final movement's direction of i is right(left) then $\bar{i}$ passes the leftmost(rightmost) site that is reached by i .
- either i or $\bar{i}$ does not move.

Proof WLOG assume that i is on the right side of $\bar{i}$. Let $[l, r]$ be the interval of sites covered by i . Since i passes the initial location of $\bar{i}$, l is located on the left of $u_s^{(\bar{i})}$ and r is located on the right of $u_s^{(\bar{i})}$.

First we show that we may assume that $\bar{i}$ reaches at least one site outside the interval $[l, r]$. If this is not the case, consider two cases. If i 's remaining budget at each site is always as high as $\bar{i}$'s remaining budget then $\bar{i}$ does not have to move and the theorem holds. Otherwise, let $\bar{r}^*$ the rightmost site where $\bar{i}$'s remaining budget is higher than i 's remaining budget. If $\bar{r}^*$ is on the left side of i 's initial location, then as in the proof of Lemma 2, the agents can be separated. If $\bar{r}^*$ is on the right side of i 's initial location and it equals r , there is no need for i to reach r since at each site in $[u_s^{(\bar{i})}, r]$, $\bar{i}$ has at least the same budget as i . Thus, there is an optimal strategy where either i does not move or it moves only to the left, so $\bar{i}$ passes the rightmost site that is reached by i . If $\bar{r}^*$ is on the right side of i but on the left side of r then there is no need for $\bar{i}$ to go beyond $\bar{r}^*$. Since it has more budget than i at this location, $\bar{i}$ can move to l while i moves to r . Thus, again, there is an optimal strategy where either i does not move or it moves only to the right, so $\bar{i}$ passes the leftmost site that is

reached by i . Thus, we may assume that $\bar{i}$ reaches at least one site outside the interval $[l, r]$.

WLOG assume that i 's final movement's direction is left and suppose that $\bar{i}$ reaches at least one site outside the interval $[l, r]$ to the left of l . If $\bar{i}$'s budget at l is higher than i 's remaining budget there, then it is also higher at $u_s^{(\bar{i})}$, and again the agents can be separated. If $\bar{i}$'s budget at l is not higher than i 's remaining budget, then $\bar{i}$ does not have to move since i can reach the same sites to the left of l .

Now suppose that $\bar{i}$ moves to the right (which is the opposite direction of i 's final movement) and passes $u_s^{(i)}$, but it also changes its direction. The only reason for $\bar{i}$ to change directions is to reach a site on the left side of its initial location, with a higher budget than i has at this site, or to reach a site that i does not reach at all. In both cases $\bar{i}$ must reach each site in $[l, u_s^{(i)}]$ with at least the same budget as i has at the same location, so either S is not optimal, or we can modify S by letting only $\bar{i}$ to move while i does not move at all. ■

Using these lemmata we observe that for any two agents, there are only a constant number of possible cases where the agents are not separated by the optimal strategies. Figure 2 illustrates the three core cases (the others are symmetrical). Here, agents 1 and 3 are non-separated agents. Note that every agent between them, like agent 2, does not have to move at all in the optimal strategy.




Figure 2: The only three cases where a pair of agents may not be separated.

Therefore we can use a dynamic programming approach to find an optimal strategy whereby all the agents are separated, but we also check the non-separated strategies individually.

Recall that in our problem the objective is to maximize the success probability, given the initial budgets. Technically, it is easy to work with the failure probability instead of the success probability.

Definition $\text{fail}[u_i, j]$ is the minimal failure probability if the

only reachable sites are in the interval $[u_1, u_i]$, and only agents $1, \dots, j$ are allowed to move. $\text{act}[u_i, j]$ is the optimal strategy achieving $\text{fail}[u_i, j]$, under the same conditions¹.

Note that where $u_i < u_s^{(j)}$, $\text{fail}[u_i, j]$ is not defined. Given $\text{act}[u_i, j]$, $\text{fail}[u_i, j]$ can be easily computed in $O(m)$ steps. For technical reasons we add another agent, 0, with a budget of zero and set its initial location to the leftmost site, i.e. $u_s^{(0)} = u_1$. $\text{fail}[u_i, 0] = 1$ for all i , and this agent doesn't affect the failure probability of any policy.

We are now ready to prove Theorem 1

Proof of Theorem 1 We use dynamic programming to calculate $\text{fail}[u_m, k]$ and $\text{act}[u_m, k]$. For $\text{fail}[u_i, 1]$ and $\text{act}[u_i, 1]$, which is the single agent case, we employ the polynomial algorithm of [Aumann *et al.*, 2008].

Given any agent $\bar{j}$ we first consider the case where $u_i = u_s^{(\bar{j})}$. In this case in the optimal strategy $\bar{j}$ moves only to the left, or not at all. Let $u_l^{(\bar{j})}$ be the leftmost site visited by $\bar{j}$ with the optimal strategy for the given interval, and agent l be the one such that $u_s^{(l)} \leq u_l^{(\bar{j})}$ (l may equal 0). Each agent t such that $l < t < \bar{j}$ does not move in the optimal strategy. Otherwise, agents t and $\bar{j}$ are not separated and according to Lemma 3 agent t must pass the rightmost site $u_s^{(\bar{j})}$, which is not possible. The same argument shows that each agent t such that $t \leq l$ does not reach $u_l^{(\bar{j})}$. Therefore $\text{act}[u_i, \bar{j}]$ is composed of $\text{act}[u_{l-1}^{(\bar{j})}, l]$, which are already known, together with the movement of agent $\bar{j}$ to $u_l^{(\bar{j})}$. Thus, computing $u_l^{(\bar{j})}$ takes $O(m)$ steps.

Next, consider the case where $u_i > u_s^{(\bar{j})}$. In this case, in the optimal strategy $\bar{j}$ may move in both directions, or not move at all. Let $u_l^{(\bar{j})}$ be the leftmost site visited by $\bar{j}$ with the optimal strategy for this interval, and agent l is the one such that $u_s^{(l)} \leq u_l^{(\bar{j})}$. First note that each agent t , $t \leq l$, and $\bar{j}$ are separated by the optimal policy, or $\bar{j}$ does not move. Otherwise, according to Lemma 2 t must pass the initial location of $\bar{j}$ but according to Lemma 3 $\bar{j}$ must reach a site outside the interval $[u_s^{(l)}, u_s^{(\bar{j})}]$ which does not occur. Since $\bar{j}$ passes the initial locations of every agent t , $l < t < \bar{j}$, if one of them moves it goes only in the opposite direction of the final movement direction of $\bar{j}$ according to Lemma 3, and as illustrated in Figure 2. Since they all must move in the same direction, according to the same Lemma at most one of them moves in the optimal policy. Therefore, to compute $\text{act}[u_i, \bar{j}]$ we check only the following options, and choose the best one:

1. $\bar{j}$ does not move, and $\text{act}[u_i, \bar{j}] = \text{act}[u_i, \bar{j} - 1]$.
2. Each agent t , $t \leq l$, does not move. Thus, $\text{act}[u_i, \bar{j}]$ is composed of $\text{act}[u_{l-1}^{(\bar{j})}, l]$, with the optimal movement of agent $\bar{j}$ in the interval $[u_l^{(\bar{j})}, u_i]$.

¹There may be more than one strategy with the same failure probability, $\text{act}[u_i, j]$ is one of them

The previous two options assumes that $\bar{j}$ and every other agent are separated. Otherwise:

3. One agent t , $t \leq l$, moves. Let $u_l^{(\bar{t})}$ be the leftmost site visited by either agent t or $\bar{j}$, with the optimal strategy, and agent l is the one such that $u_s^{(l)} \leq u_l^{(\bar{t})}$. $\text{act}[u_{\bar{j}}, \bar{j}]$ is composed of $\text{act}[u_{l-1}^{(\bar{t})}, l]$, with the optimal movement of the two agents $\bar{j}$ and t in the interval $[u_l^{(\bar{t})}, u_i]$.

There are at most m possible options for $u_l^{(\bar{j})}$. In each option we check for at most k agents m possible options for $u_l^{(\bar{t})}$. Therefore for each agent j and site u_i $\text{act}[u_i, j]$ can be found in $O(m^2k)$ steps, and $\text{act}[u_m, k]$ can be found in $O(m^3k^2)$ time steps using $O(mk)$ space. ■

For the *Min-Budget* problem we obtain:

Theorem 4 *In the no communication setting, if the number of costs is constant, then Min-Budget^{identical} can be solved in polynomial time for any number of agents.*

Proof By Theorem 1, given a budget $\bar{B}$, we can calculate the maximum achievable success probability. Thus we can run a binary search over the possible values of $\bar{B}$ to find the minimal one that still guarantees a success probability p_{succ} . The maximum required budget is $2 \cdot |u_1 - u_m|$, which is part of the input. Thus the binary search will require a polynomial number of steps. ■

Theorem 5 *If the number of agents is a parameter, Min-Budget^{distinct} with no communication is NP-Hard even for a single possible cost.*

Proof Aumann et al [Aumann et al., 2008] consider the shared budget case and prove that when the number of agents is not constant the *Min-Budget* problem is NP-hard. In the *Min-Budget^{distinct}* problem the objective is to minimize the average budget, which is the same as minimizing the total budget. Thus, the hardness of the problem follows from that of the *Min-Budget* problem in [Aumann et al., 2008]. ■

3 Communicating Agents

Once communication is added agents can call upon each other for assistance and the relative scheduling between the agents moves must also be considered. In this case a *solution* is an ordered list of *moves*, where each move is a pair stating an agent and its next destination.

The success probability of a solution is now calculated according to the moves order. For example, suppose that the sites and agents are located as illustrated in Figure 3.

Consider the following solution: agent 2 first goes to u_4 and then agent 1 goes to u_2 . Agent 2 is the only one which can succeed at u_4 , with a probability of 0.8. With probability 0.2 it will not succeed and agent 1 has a probability of 0.2 to succeed at u_2 . Hence, the success probability is $0.8 + 0.2 \cdot 0.2 = 0.84$. If we switch the moves order we get a probability of 0.9 to succeed at u_2 with the first move, since agent 2 will be called for assistance if the cost required is less than 100. If




Figure 3: A possible input with suggested moves. The numbers on the edges represent traveling costs. A table at each site u_i represents the cost probability function $p_i(c)$. The moves are illustrated by arrows.

not, agent 2 will move to u_5 as before. Hence, this solution success probability is $0.9 + 0.1 \cdot 0.8 = 0.98$.

When the number of agents is not fixed, *Max-Probability*, *Min-Budget^{identical}* and *Min-Budget^{distinct}* are not known to be solvable in polynomial time. However, in many physical environments where several agents cooperate in exploration and search, the number of agents is relatively small. In this case we can show that all the three problems can be solved in polynomial time. We show:

Theorem 6 *In the setting of communicating agents, if the number of agents and the number of different costs is fixed then Max-Probability, Min-Budget^{identical} and Min-Budget^{distinct} can be solved in polynomial time.*

For brevity, we focus on the *Max-Probability* problem. The same algorithm and similar analysis work also for the other two problems.

First note that in *Max-Probability*, we need to maximize the probability of obtaining the good given the initial budgets B_i , but there is no requirement to minimize the actual resources consumed. Thus, at any site, if agents can obtain the good for a cost no greater than its remaining budget, the search is over. Furthermore, if the cost is beyond the agent's available budget, but there is another agent with enough budget to both travel from its current location and to obtain the good, then this agent is called upon and the search is also over. Otherwise, the good will not be obtained at this site under any circumstances. Thus, the basic strategy structure, which determines which agent go where, remains the same. Unless the search has to be terminated, the decision of one agent where to go next is not affected by the knowledge gained by others.

Let $c_1 > c_2 > \dots > c_d$ be the set of costs. For each agent j and for each c_i there is an interval $I_i^{(j)} = [u_\ell, u_r]$ of sites covered while the agent's remaining budget is at least c_i . Furthermore, for each j and for all i , $I_i^{(j)} \subseteq I_{i+1}^{(j)}$. Thus, consider for each agent the *incremental* area covered when its remaining budget c_i but less than c_{i-1} , $\Delta_i^{(j)} = I_i^{(j)} - I_{i-1}^{(j)}$ (with $\Delta_1^{(j)} = I_1^{(j)}$). Each $\Delta_i^{(j)}$ is a union of an interval at left of $u_s^{(j)}$ and an interval at the right of $u_s^{(j)}$ (both possibly empty). Since there is communication, an agent may continue to reach sites even if it does not have any chance of obtaining

the good there, in order to reveal the cost for the use of other agents. Thus, the optimal strategy may define also an interval $I_{d+1}^{(j)} = [u_\ell, u_r]$ of sites covered while the remaining budget of j is greater than 0. For brevity, we denote $\bar{d}$ instead of $d+1$. The next lemma, which is the multi agent Max-Probability analogue of Lemma 2 in [Aumann *et al.*, 2008] states that there are only two possible optimal strategies to cover each $\Delta_i^{(j)}$:

Lemma 7 Consider the optimal solution and the incremental areas for each agent j , $\Delta_i^{(j)}$ ($i = 1, \dots, \bar{d}$) defined by this solution. For $i \in 1, \dots, \bar{d}$, let $u_{\ell_i}^{(j)}$ be the leftmost site in $\Delta_i^{(j)}$ and $u_{r_i}^{(j)}$ the rightmost site. Suppose that in the optimal strategy the covering of $\Delta_i^{(j)}$ starts at location $u_{s_i}^{(j)}$. Then, WLOG we may assume that the optimal strategy for each j is either $(u_{s_i}^{(j)} \rightarrow u_{r_i}^{(j)} \rightarrow u_{\ell_i}^{(j)})$ or $(u_{s_i}^{(j)} \rightarrow u_{\ell_i}^{(j)} \rightarrow u_{r_i}^{(j)})$. Furthermore, the starting point for covering $\Delta_{i+1}^{(j)}$ is the ending point of covering $\Delta_i^{(j)}$.

Proof Any strategy other than the ones specified in the lemma would reach all the sites covered by the optimal solution with at most the same available budget. ■

Corollary 8 For k constant, one needs only to consider a polynomial number of options for the set of moves of the agents.

Proof By the previous lemma, the moves of each agent are fully determined by the leftmost and rightmost sites of each $\Delta_i^{(j)}$, together with the choice for the ending points of covering each area. For each j there are $m^{2\bar{d}}/(2\bar{d})!$ possible choices for the external sites of the $\Delta_i^{(j)}$'s, and there are a total of $2\bar{d}$ options to consider for the covering of each. Thus, the total number of options is polynomial (in m). ■

It thus remains to consider the scheduling between the moves, i.e. their *order*. Theoretically, with n moves there are $n!$ different possible orderings. We show, however, that for any given set of moves, we need only to consider a polynomial number of possible orderings.

Consider a given set of moves M , determining the sets $\Delta_i^{(j)}$. Note that for each agent, M fully determines the order of the moves of this agent. A subset M' of M is said to be a *prefix* of M , if for each agent the moves in M' are a prefix of the moves of this agent in M . A subset M' is a *suffix* of M if $M - M'$ is a prefix. We now inductively define the notion of a *cascading order*:

1. The trivial order on moves of a single agent is cascading.
2. Let M be a set of moves, and let c_{i_0} be the highest cost that any agent can pay. An order S on M is *cascading* if M and S can be decomposed $M = M_{pre} \cup M_{mid} \cup M_{post}$ and $S = S_{pre} \circ S_{mid} \circ S_{post}$, such that:
 - M_{pre} is a prefix of M consisting only of moves of agents with budget less than c_{i_0} and S_{pre} is a cascading order on M_{pre} .

- There exists an agent j' with budget at least c_{i_0} such that M_{mid} consists of all the moves of j' in $\Delta_{i_0}^{(j')}$ and S_{mid} is the (one possible) order on these moves.
- M_{post} are the remaining moves in M and S_{post} is a cascading order on them.

We prove (by induction) that cascading orders are optimal.

Lemma 9 For any set of moves M there exists a cascading order with optimal success probability.

Proof The proof is by induction on the number of agents and the number of moves in M . If there is only one agent moving in M then the order is cascading. Otherwise, consider any other order S on M and let A_{i_0} be the set of agents with budget at least c_{i_0} . Let j' be the first agent in A_{i_0} to cover its $\Delta_{i_0}^{(j')}$ and let t_0 be the time it completes covering it. M_{pre} includes all the moves taken by agents not in A_{i_0} prior to t_0 ;

M_{mid} includes all the moves of j' in $\Delta_{i_0}^{(j')}$; and M_{post} the rest of the moves in M . We show that we do not decrease the success probability by first making all moves of M_{pre} then all those of M_{mid} , and finally those of M_{post} . By the inductive hypothesis S_{pre} , S_{mid} and S_{post} are optimal for M_{pre} , M_{mid} and M_{post} , respectively and the result follows.

Before t_0 all agents in A_{i_0} have a higher budget than any agent not in A_{i_0} . Thus, before t_0 agents of A_{i_0} will never call upon those not in A_{i_0} . Thus, it cannot decrease the success probability if we let the agents not in A_{i_0} take their moves first. Thus, we can allow to first perform all moves of M_{pre} .

Also, before t_0 no agents of A_{i_0} needs to call upon each other for assistance (since they are all in the same resource bracket). Thus, we may allow them to take their moves independently without decreasing the success probability. In particular, we can allow j' to complete its covering of $\Delta_{i_0}^{(j')}$ before any other member of A_{i_0} moves. Thus, we get that first having the moves of M_{pre} and then of M_{mid} does not decrease the success probability. The moves of M_{post} are the remaining moves. ■

Finally we show that the number of cascading orders is polynomial:

Lemma 10 For fixed k and d and any set of moves M there are a polynomial number of cascading orders on M .

Proof Set $f(n, k, d, \ell)$ be the number of cascading orders with k agents, n moves, d costs and ℓ agents in A_{i_0} . We prove by induction that f is a polynomial in n . Since $\ell \leq k$, the result follows. Clearly, for any ℓ , $f(n, k, 0, \ell) = \ell!$ (all of which are useless). Then, by the definition of cascading orders $f(n, k, d, \ell) \leq \ell n^{k-\ell} f(n, k-\ell, d-1, k-\ell) f(n, k, d, \ell-1)$ (the $n^{k-\ell}$ being for the choice of M_{pre}). By the inductive hypothesis $f(n, k-\ell, d-1, k-\ell)$ and $f(n, k, d, \ell-1)$ are polynomials in n . Thus, so is $f(n, k, d, \ell)$. ■

Together with Corollary 8 we get that the total number of options to consider is polynomial, proving the *Max-Probability* part of Theorem 6. The proof for the other two problems is similar.

4 Heterogenous Agents

The analysis so far assumes that all agents are of the same type, with identical capabilities. Specifically, the cost of obtaining the good at any given site is assumed to be the same for all agents. However, agents may be of different *types* and hence with different capabilities. For example, some agents may be equipped with a drilling arm, which allows them to consume less battery power while mining. In this section we consider such situations of *heterogenous agents*, and show that the results can be extended to such settings. Due to lack of space we do not provide the full proofs here, but only the core proof directions.

While agents may have different capabilities, in many cases it is reasonable to assume that if one agent is more capable than the other at one site, it is also more capable at all other sites (or at least no less capable). Hence the following definition:

Definition We say that agents are *inconsistent* if there exist budgets B, B' , agents j, j' , and sites i, i' , such that at location i with budget B

$$\Pr[j \text{ can obtain the good}] < \Pr[j' \text{ can obtain the good}]$$

but at location i' with budget B'

$$\Pr[j \text{ can obtain the good}] > \Pr[j' \text{ can obtain the good}]$$

Theorem 11 *In the no communication setting, if the number of different costs for each agent is constant, then Max-Probability and Min-Budget^{identical} can be solved in polynomial time with any number of heterogenous agents, provided that the agents are consistent.*

The algorithm is essentially the same dynamic programming algorithm described in Section 2. The consistency assumption is necessary for lemmata 2 and 3 to remain true.

In the case where agents can communicate, we can do away with the consistency assumption. Clearly, however, we do need to assume that upon reaching a site, agents can assess the cost for obtaining the good for all other agents. Otherwise, communication would be meaningless as agents would not know which other agent to call. We obtain:

Theorem 12 *In the setting of communicating agents, with a constant number of agents, and a constant number of different costs for each agent, Max-Probability, Min-Budget^{identical} and Min-Budget^{distinct} can be solved in polynomial time even with inconsistent heterogenous agents.*

The algorithm and proof remain essentially the same as those for the homogenous agents case.

5 Extending our Results - Discussion

In this paper we focused on the case where the sites are located along a path (either closed or a non-closed). There are many settings where this assumption holds. For example, the assumption faithfully captures the setting of perimeter patrol applications (see [Williams and Burdick, 2006;

Elmaliach *et al.*, 2008]). Also, as pointed out in the introduction, many coverage algorithms convert their complex environment into a simple path. However, many physical environments may only be represented by a planar graph. [Aumann *et al.*, 2008] showed that physical search problems are NPC even on trees and even with a single agent, but finding heuristic is of practical interest nonetheless. It seems that the first steps in building such heuristic will be to utilize our results. For example, one should try to avoid repeated coverage as much as possible and to restrict the number of cases where such coverage is necessary, as we showed in theorem 1. Another idea is to convert the complex graph structure into a path, where each site on the path represents a region of strongly-connected nodes on the original graph. Many graphs which represent real physical environments consist of some regions with strongly-connected nodes, but few edges connect these regions (for example, cities, which has many roads inside, but are connected with few highways). A heuristic algorithm for these graphs may use our algorithm to construct a strategy for the sites along the path, and use an additional heuristic for visiting the sites inside a region.

We also considered only the case where mining costs are rounded/estimated to one of a constant number of possible options. We believe that this assumption is appropriate since the given input for our problems includes prior probabilistic knowledge. Usually, this data comes from some sort of estimation so it is reasonable to assume that the number of options is fixed. Nevertheless, if the number of costs will not be a constant it can be rounded to a fixed number of costs, which yields a PTAS (polynomial-time approximation scheme) for our problems.

We also assumed that the agents are after only one good. As soon as we allow more than one good that must be obtained, our results do not hold anymore, and it seems that the problems turn to be in NPC.

6 Conclusions and future work

This paper considers multi-agent physical search with prior probabilistic knowledge. Each agent is equipped with its own budget, which is used both for exploration and for fulfilling the task. We showed that for non-communicating agents there exists a polynomial algorithm that is suitable for any number of agents. This result emphasizes the difference between our model and the shared budget model proposed of [Aumann *et al.*, 2008] where all problem variants are NP-Complete. For agents that do communicate, we presented a polynomial algorithm that works for any *fixed* number of agents. We also extended the analysis to heterogenous agents.

There are still many interesting open problems. First and foremost, the complexity of the problem with a non-constant number of communicating agents is open. Also, considering agents with differing travel capabilities would be interesting. Finally, metric spaces beyond the line remain an open challenge.

Acknowledgments

This work was supported in part by the Israeli Science Foundation under Grant 1685 and in part by the National Science

Foundation under Grant 0705587. Sarit Kraus is also affiliated with UMIACS.

References

- [Aumann *et al.*, 2008] Y. Aumann, N. Hazon, S. Kraus, and D. Sarne. Physical search problems applying economic search models. In *AAAI-08*, Chicago, 2008.
- [Ausiello *et al.*, 2000] G. Ausiello, S. Leonardi, and A. Marchetti-Spaccamela. On salesmen, repairmen, spiders, and other traveling agents. In *The Italian Conference on Algorithms and Complexity*, pages 1–16, 2000.
- [Elmaliach *et al.*, 2008] Y. Elmaliach, A. Shiloni, and G. A. Kaminka. A realistic model of frequency-based multi-robot fence patrolling. In *AAMAS 2008*, volume 1, pages 63–70, 2008.
- [Gabriely and Rimon, 2001] Y. Gabriely and E. Rimon. Spanning-tree based coverage of continuous areas by a mobile robot. *Annals of Math and AI*, 31:77–98, 2001.
- [Gal, 1980] S. Gal. *Search Games*. Academic Press, New York, 1980.
- [Hazon and Kaminka, 2005] N. Hazon and G. A. Kaminka. Redundancy, efficiency, and robustness in multi-robot coverage. In *ICRA 2005*, 2005.
- [Koopman, 1980] B. O. Koopman. *Search and Screening: General Principles with Historical Applications*. Pergamon Press, New York, 1980.
- [Koutsoupias *et al.*, 1996] E. Koutsoupias, C. H. Papadimitriou, and M. Yannakakis. Searching a fixed graph. In *ICALP '96*, pages 280–289, London, UK, 1996.
- [Lippman and McCall, 1976] S. Lippman and J. McCall. The economics of job search: A survey. *Economic Inquiry*, 14:155–189, 1976.
- [McMillan and Rothschild, 1994] J. McMillan and M. Rothschild. Search. In R. Aumann and S. Amsterdam, editors, *Handbook of Game Theory with Economic Applications*, chapter 27, pages 905–927. Elsevier, 1994.
- [Sarne and Kraus, 2005] D. Sarne and S. Kraus. Cooperative exploration in the electronic marketplace. In *AAAI'05*, pages 158–163, 2005.
- [Spires and Goldsmith, 1998] S. V. Spires and S. Y. Goldsmith. Exhaustive geographic search with mobile robots along space-filling curves. In *1st Int. Workshop on Collective Robotics*, pages 1–12. Springer-Verlag, 1998.
- [Williams and Burdick, 2006] K. Williams and J. Burdick. Multi-robot boundary coverage with plan revision. In *ICRA 2006*, pages 1716–1723, May 2006.