k-Agent Sufficiency for Multiagent Stochastic Physical Search Problems

Daniel S. Brown(✉), Steven Loscalzo, and Nathaniel Gemelli

Air Force Information Directorate, Rome, NY 13441, USA
{daniel.brown.81,steven.loscalzo,nathaniel.gemelli}@us.af.mil

Abstract. In many multi-agent applications, such as patrol, shopping, or mining, a group of agents must use limited resources to successfully accomplish a task possibly available at several distinct sites. We investigate problems where agents must expend resources (e.g. battery power) to both travel between sites and to accomplish the task at a site, and where agents only have probabilistic knowledge about the availability and cost of accomplishing the task at any location. Previous research on Multiagent Stochastic Physical Search (mSPS) has only explored the case when sites are located along a path, and has not investigated the minimal number of agents required for an optimal solution. We extend previous work by exploring physical search problems on both paths and in 2-dimensional Euclidean space. Additionally, we allow the number of agents to be part of the optimization. Often, research into multiagent systems ignores the question of how many agents should actually be used to solve a problem. To investigate this question, we introduce the condition of k-agent sufficiency for a multiagent optimization problem, which means that an optimal solution exists that requires only k agents. We show that mSPS along a path with a single starting location is at most 2-agent sufficient, and quite often 1-agent sufficient. Using an optimal branch-and-bound algorithm, we also show that even in Euclidean space, optimal solutions are often only 2- or 3-agent sufficient on average.

Keywords: Stochastic physical search · Planning under uncertainty · Multiagent optimization · k-Agent sufficiency

1 Introduction

We investigate the problem of multiple agents seeking for a single item that may possibly be obtained at one of several locations. We assume that the availability and actual cost to acquire the item at any site is not fully known beforehand, but that a priori probabilistic cost distributions are known. In particular we examine problems where there is a finite resource that must be expended to both travel and obtain the item of interest. We refer to this class of problems as Multiagent Stochastic Physical Search (mSPS). Examples of this type of problem include battery-powered mining or space exploration robots seeking a precious metal deposit or specific mineral sample, hikers seeking a suitable location to

T. Walsh (Ed.): ADT 2015, LNAI 9346, pp. 171–186, 2015.
DOI: 10.1007/978-3-319-23114-3_11

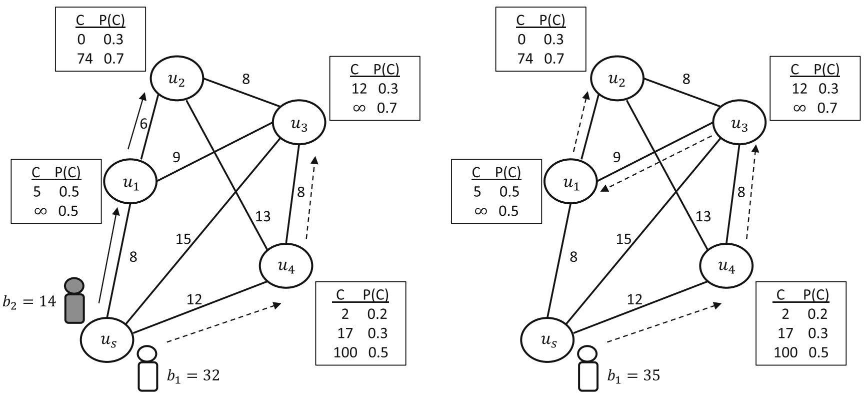

(a) 2 agent solution (b) 1 agent solution

Fig. 1. (a) Example of a two-agent strategy where agent 1 is allocated a starting budget of $b_1 = 32$ and the path $\pi^1 = \langle u_s, u_4, u_3 \rangle$ and agent 2 is allocated a starting budget of 14 and the path $\pi^2 = \langle u_s, u_1, u_2 \rangle$. Probability distributions over cost are shown next to each site. The joint probability of success of the shown multiagent search strategy is $1 - Pr(\text{failure of agent 1}) \cdot Pr(\text{failure of agent 2}) = 1 - Pr(\text{failure at } u_4) \cdot Pr(\text{failure at } u_3) \cdot Pr(\text{failure at } u_1) \cdot Pr(\text{failure at } u_2) = 1 - 0.5 \cdot 0.7 \cdot 0.7 \cdot 0.5 = 0.8775$. (b) Using a single agent, the same probability of success can be achieved with a strategy that requires less budget.

set up a base camp, or tourists using public transportation to explore several different shopping areas for a desired souvenir. What makes the above problems challenging, is that while actual distance costs may be reliably and accurately estimated using satellite imagery, maps, or taxi fares, the actual cost to accomplish the task (or purchase the item of interest) at a specific location may be unknown until an agent actually visits the site. One of the major challenges is that we assume each agent must use a single budget (e.g. battery power, fuel, or currency) to both travel and obtain the item. This adds extra complexity to the problem because it means that taking a different path to a site can change the probability of success—a longer path will consume more budget, reducing the budget available to obtain the item. Figure 1(a) shows an example problem along with a possible two-agent solution that allocates a total budget of 46 and achieves a joint probability of success of 0.8775. However, Fig. 1(b) shows that an equivalent probability of success can be achieved using only a single agent with a lower total budget of 35.

In many multiagent search problems it is often assumed that the number of agents is fixed and that having multiple searchers is better than a single searcher. However, this may not be the case when searchers start from the same location, when both search and acquisition are costly, and when there is a limit to the total allocatable budget. For example, in many vehicle routing problems, the number of vehicles is part of the problem definition and solutions often assume that all the vehicles will be used [14]. Some vehicle routing problems try to minimize the number of required agents, but these problems do not consider probabilities of success or costs to purchase or acquire an item along a vehicle's route [15]. Work on the multiple traveling salesman (a special case of the mSPS) and its extensions sometimes allows the number of salesman to be variable, but we know of no proven bounds on the maximum number of agents required for different problems [2,11]. In many other multiagent scenarios such as multiagent task allocation, the number of agents is also often assumed to be fixed [5,12]. Previous work on stochastic physical search has either focused on single agent solutions or has assumed that the number of agents is not part of the optimization [1,3,7,9,10].

Our research assumes that there is a given upper bound on the number of agents available to search for the item; however, we do not restrict our solutions to a fixed number of agents. Instead, we algorithmically decide which of the available agents should participate in the search in order to maximize the probability of successfully obtaining the item as well as minimizing the required total budget allocated to the agents. Because our problem is bi-objective we use the standard epsilon constraint method to split the problem into two dual objectives. The *Min-Budget* objective is to minimize the total budget allocated to the agents while guaranteeing a specified minimum probability of success. The *Max-Probability* objective is to maximize the probability of success given an upper bound on the budget that can be allocated. In both of these problems, the agents all start at the same starting location and a solution is an allocation of resources to each agent, along with a search path for each agent.

While stochastic physical search problems capture many real world planning and algorithmic decision problems, very little is known about the solution properties of these problems. Work by Aumann et al. and Hazon et al. has proposed optimal algorithms for the case when sites are located along a path; however, their work never actually computes or analyzes optimal solutions and does not investigate the frequency of optimal solutions that required more than one agent [1,6,7]. Work by Brown et al. and Hudack et al. has examined the Min-Budget and Max-Probability problems on general graphs and 2-dimensional euclidean spaces; however, they only consider the single agent case [3,9]. This paper provides, to the best of our knowledge, the first theoretical and empirical investigation of the solution properties of the multiagent Min-Budget and Max-Probability problems.

We start by examining problems where the item may possibly be obtained from a set of locations on a path. We examine two cases: (1) single-cost and (2) multi-cost. We prove that in both cases, problems are at most 2-agent sufficient, and empirically investigate the frequency of 1-agent sufficient problems. We next investigate solutions to the Min-Budget and Max-Probability problems when locations are in 2-dimensional Euclidean space. We provide a theoretical analysis of when multiple agents are unnecessary and use an exact branch-and-bound algorithm to provide empirical insights into k-agent sufficiency for 2-dimensional problems. We show that in many cases, even when searching in two dimensions, the optimal strategy is to use a very small number of searchers, rarely requiring more than 3 searchers. We conclude by discussing the factors that contribute to a search problem being k-agent sufficient for different values of k along with areas for future research.

2 Problem Definition

A stochastic physical search problem is defined by a graph $G(S^+, E)$ with a set of locations $S^+ = S \cup \{u_s\}$ where $S = \{u_1, \ldots, u_m\}$ is the set of m sites offering an item of interest, u_s is the starting location, and $E \subseteq S^+ \times S^+$ is the set of edges. Each $(i, j) \in E$ has a non-negative cost of travel t_{ij}. For each site $i \in S$ we are given a cost probability mass function $p_i(c)$, which gives the probability that the item will cost c at site u_i. We assume that the actual cost is not revealed until the agent visits the site and that the cost remains fixed thereafter. We further assume that there is a finite number of possible costs in the support of $p_i(c)$, $\forall i \in S$. Finally, we define a set $\mathcal{N}$ of $n^{\max} = |\mathcal{N}|$ agents that are available, but not required, to be used in the search. Each agent n starts at u_s with a starting budget $b^*_n, \forall n \in \mathcal{N}$. We let $b^* = \{b^*_n : n \in N\}$ and let $B^* = \sum_{n \in N} b^*_n$, the total budget allocated to all agents. We assume that once budgets are allocated they are non-transferable, that two or more agents cannot combine their budgets to obtain the item, and that agents cannot share information about sites they have visited with other agents. Following previous work on stochastic physical search problems [7], we assume that success is achieved if any agent is able to purchase the item. We also assume that the item cannot be found at the start site, u_s.

We examine two dual problems (1) Min-Budget: Given a required probability of success p^*_{succ} find the initial budget allocation b^* that satisfies p^*_{succ} and that minimizes B^*. (2) Max-Probability: Given an upper bound on the budget available for allocation of B^* determine the optimal budget allocation b^* so as to maximize the probability of success. A solution to either problem is an allocation of starting budgets $b*$ along with a set of paths Π^* where each individual path $\pi^n \in \Pi^*$ is a sequence of sites in S^+, where π^n_i is the i^{th} site visited along path π^n, and where each path starts at the start site, i.e., $\pi^n_0 = u_s, \forall n \in \mathcal{N}$. We assume that success is achieved if any agent arrives at a site where the actual cost is less than or equal to that agent's remaining budget.

2.1 k-Agent Sufficiency

Before we investigate solutions to the Min-Budget and Max-Probability mSPS problems we define the term k-agent sufficiency as it relates to multiagent optimization problems of the kind investigated in this paper.

Definition 1. *A mSPS problem is k-agent sufficient if an optimal solution exists such that $|B^+| = k$ where $B^+ = \{b^*_n \in b^* \mid b^*_n > 0, \forall n \in \mathcal{N}\}$.*

The following result is true for all mSPS problems.

Proposition 1. *If an SPS problem has zero travel costs between all sites, then it is 1-agent sufficient.*

Proof. Assume you have two agents i and j with starting budgets $b_i, b_j > 0$. Since the agents can travel between sites without incurring costs, an equivalent probability of success can be achieved with a single agent, given less starting budget $b' = \max(b_i, b_j) \leq b_i + b_j$. □

Thus, for the remainder of the paper we assume that all travel costs are non-zero.

3 mSPS Along a Path

We first investigate the Min-Budget and Max-Probability problems where the set of locations in S^+ are restricted to be along a path. We note that the following discussion on paths is not purely academic, as many multi-agent coverage algorithms convert their complex environment into a path and many perimeter monitoring and border control tasks could also be represented by sites along a path [4,7,8,13].

To simplify our analysis, we follow the methodology used by Hazon et al. [7], and assume WLOG (without loss of generality) that all locations are along a line such that the travel cost between any two sites u_i and u_j is $t_{ij} = |u_i - u_j|$. We also assume WLOG that the sites are ordered from left to right such that $u_1 \leq u_2 \leq \cdots \leq u_m$. We first examine the case when there is only one possible cost to obtain the item. Despite the simplicity of this problem, we show that the results for k-sufficiency are non-trivial. We then examine the case where there are multiple possible item costs.

Before examining the single and multi-price cases, we note the following.

Proposition 2. *When u_s is the leftmost (rightmost) location, then the problem is 1-agent sufficient and the optimal strategy only moves to the right (left).*

Proof. Any other strategy to cover the same locations would use at least as much budget and achieve no greater probability of success. □

Thus, the most interesting cases, in terms of k-agent sufficiency, are those where the start site u_s is towards the middle of the path.

3.1 Single Price

We first assume that all sites either offer the item for a cost of c_0 or do not offer the item at all (this can simply be modeled as a cost of ∞). All we are given are the a priori probabilities p_i that the item is available for cost c_0 at site i.

We first note the following useful lemma and definition proposed by Aumann et al. [1].

Lemma 1. *Consider a price c_0 and suppose that an agent's optimal strategy starting at point u_s covers the interval $[u_\ell, u_r]$ while the remaining budget is at least c_0. Then WLOG we may assume that the agent's optimal strategy is either $(u_s \rightarrowtail u_r \rightarrowtail u_\ell)$ or $(u_s \rightarrowtail u_\ell \rightarrowtail u_r)$.*

Definition 2. *Agents i and j are said to be* separated *by a strategy if each site in S that is reached by i is not reached by j.*

We now prove the following lemma and theorem which give us our first k-agent sufficiency condition.

Lemma 2. *With a single price, there exists an optimal multiagent strategy where every agent is separated.*

Proof. Assume by contradiction that two agents i and j are not separated in every optimal strategy. Consider the intervals covered by these two agents, $[l_i, r_i]$ and $[l_j, r_j]$, respectively. Let $[L, R] = [l_i, r_i] \cup [l_j, r_j]$ be the full combined coverage area of the two agents. In the case that $[l_i, r_i] \subset [l_j, r_j]$ we can safely remove agent i from the strategy, resulting in a strategy with the same probability of success but lower budget. Otherwise, WLOG we can assume based on Lemma 1 that only i reaches L and only j reaches R. However, now the separated strategy of $u_s \rightarrowtail L$ for agent i and $u_s \rightarrowtail R$ for j guarantees at least the same probability of success with no more budget. This contradicts our assumption. □

Theorem 1. *For a path with a single possible price, there is always an optimal strategy with fewer than 3 agents. If using two agents is optimal, then only one agent moves left and only one agent moves right in the optimal strategy.*

Proof. This follows as a direct result of Lemmas 1 and 2. □

The work of Aumann et al. [1] provides an $O(m)$ algorithm for the single agent single price problem. Based on the result of Theorem 1 we can easily adapt the algorithm given by Aumann et al. to obtain an $O(m)$ algorithm for the multiagent single item cost Min-Budget and Max-Probability cases by simply checking each possible single agent coverage region to see if dividing the region between two agents results in lower budget or higher probability of success. Each of these checks can be done in constant time.

3.2 Single Price k-Agent Sufficiency

We now examine when the single price problem is 1-agent sufficient. When there are multiple agents, each one has to carry at least c_0 of budget to enable purchasing when the item is available. Thus the question of 1-agent sufficiency is directly related to the ratio of travel distances and c_0. We note that for increasing values of c_0, there exists a point at which c_0 is so high that it dominates the total travel cost.

Theorem 2. *Suppose that the optimal strategy covers the interval $[u_\ell, u_r]$ while the remaining budget is at least c_0. If $c_0 \geq \max(|u_s - u_\ell|, |u_r - u_s|)$ then the problem is 1-agent sufficient.*

Proof. Assume by contradiction that the optimal solution requires two agents, i and j. Let b_i and b_j be the starting budgets of i and j, where $t_i = b_i - c_0$ and $t_j = b_j - c_0$ are the portions of the budgets allocated for travel. By assumption and by Theorem 1 we have

$$c_0 \geq \max(|u_s - u_\ell|, |u_r - u_s|) = \max(t_i, t_j). \tag{1}$$

For a single searcher to cover both search paths it must have as a minimum travel budget

$$t' = 2\min(t_i, t_j) + \max(t_i, t_j) \tag{2}$$

to enable an out and back trip on the shorter leg, followed by an out trip on the longer leg. Thus the budget b' for a single agent is given by

$$\begin{aligned} b' &= t' + c_0 && (3)\\ &= 2\min(t_i, t_j) + \max(t_i, t_j) + c_0 && (4)\\ &\leq t_i + t_j + \max(t_i, t_j) + c_0 && (5)\\ &\leq t_i + t_j + 2c_0 && (6)\\ &= b_i + b_j && (7) \end{aligned}$$

Which contradicts our assumption that two agents were required. □

Figure 2 shows an example of how the number of agents allocated changes for both Min-Budget and Max-Probability as c_0 is increased. To obtain these results we uniform randomly generated 25 sites along a 100 unit long interval and let u_s be the median site. Probabilities of success p_i are randomly chosen between 0 and 0.5 for each site. For the Min-Budget problem, we examined several different values for the required probability of success. We know that the best success probability is achieved by visiting all the sites, giving

$$p_{succ}^{max} = 1 - \prod_{i \in S} 1 - p_i. \tag{8}$$

To vary the solutions we set the required probability of success equal to $\rho \cdot p_{succ}^{max}$ for different values of ρ. The results for Min-Budget are shown in Fig. 2(a). The x-axis shows the cost of the item and the y-axis shows the percentage of 1000 random instances that were 1-agent sufficient. When $\rho = 1$ all the sites must be visited. We see that as expected, when $\rho = 1$, most problems are not 1-agent sufficient, until the cost gets close to $\max(|u_s - u_1|, |u_m - u_s|) \approx 50$, when 1-agent sufficiency is guaranteed by Theorem 2. However, as soon as all of the sites are not required (i.e. $\rho < 1$), the probability of an instance being 1-agent sufficient dramatically increases.

Figure 2(b) shows the results for Max-Probability. Because there is only a single purchase cost we can think of this cost as a fixed start-up cost that is incurred for each agent used. Thus, for low values of c_0 and low starting budget, it is more beneficial to divide and conquer and send one agent left and one agent right. On the other hand, for large values of c_0 the start-up cost to use a second

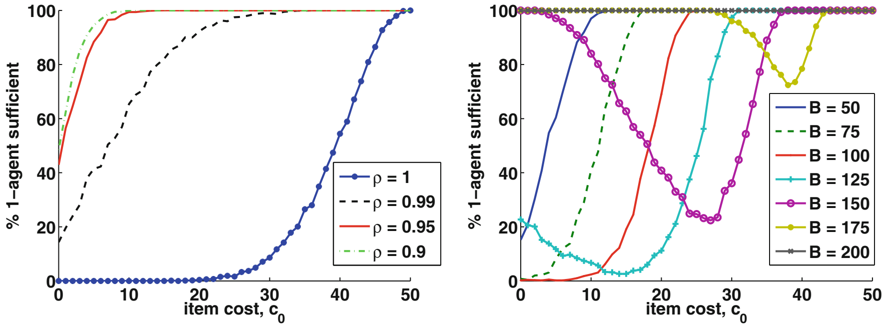

(a) Min-Budget (b) Max-Probability

Fig. 2. Percentage of 1000 random 25-site, single-price problems that are 1-agent sufficient. Sites are randomly placed in the interval $[0, 100]$. (a) Min-Budget results for different required probabilities of success $p^*_{succ} = \rho \cdot (1 - \prod_{i \in S} 1 - p_i)$ where p_i is the probability the item is available for the single cost. (b) Max-Probability results for different total budget allotments B.

agent is so high that most optimal solutions only require a single agent. Note that given a starting budget of 100 with $c_0 = 0$ the optimal solution is to always use two agents with each agent coverage one half of the solution space. However, as c_0 increases the two agent solution coverage region decreases because agents cannot reach the farther endpoints and still have enough to purchase the item. Eventually, c_0 is so high that giving two agents c_0 plus budget to travel requires more budget than using a single agent. Given a starting budget of 150 and $c_0 = 0$ one agent has enough to traverse the entire interval. In this case we see that as c_0 increases, eventually two agents can cover a larger region (resulting in higher probability of success) than one agent on its own. However, past a certain point the start-up cost of c_0 starts to dominate the travel costs and single agent solutions become more common.

3.3 Multiple Prices on a Path

We now consider the case where there can be a large number of different realizable costs at the sites. Note that Lemma 1 is not always true for multiple purchase prices. Figure 3 shows a simple example where multiple direction changes are optimal. However, we do have the following result:

Lemma 3. *In the multi-price case, if agent i and agent j both only travel left (right), then one of them is unnecessary.*

Proof. WLOG assume agent i travels past the leftmost (rightmost) point visited by j. Let b_i and b_j be the starting budgets of i and j and let $t_i = |l_i - u_s|$ and $t_j = |l_j - u_s|$ be the distances traveled by each agent, respectively. Consider the

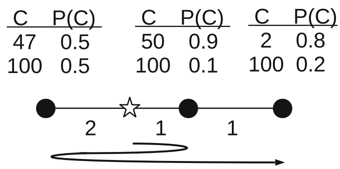

Fig. 3. When there are multiple realizable costs, optimal paths can often include several changes of direction. Shown is the optimal strategy for a single agent solving Max-Probability with starting budget of 51, or equivalently solving Min-Budget with a required probability of success equal to 0.99.

new strategy where only one agent i' travels with budget

$$b_{i'} = t_i + (b_i - t_i + b_j - t_j) \tag{9}$$
$$= b_i + b_j - t_j \tag{10}$$
$$\leq b_i + b_j. \tag{11}$$

This single-agent strategy guarantees no decrease in probability of success with no increase in total budget. □

We also utilize the following lemma, adapted from Aumann et al. [1].

Lemma 4. *Consider two agents i and j that start at u_s. Then, there is an optimal strategy such that one of the following holds:*

- *j moves only in one direction which is opposite to i's final movement direction. Furthermore, if i's final movement direction is right (left) then j passes the leftmost (rightmost) site that is reached by i.*
- *either i or j is unnecessary.*

Using Lemmas 3 and 4, we can now prove the following theorem, which is the multi-cost analogue of Theorem 1.

Theorem 3. *For the same-start mSPS problem on a path, there is always an optimal strategy with fewer than 3 agents.*

Proof. Assume by contradiction that 3 agents are necessary in every optimal strategy. Denote these agents i, j, and k. Consider agents i and j. WLOG by Lemma 4 assume that i only moves left in the optimal solution, i passes the leftmost site reached by j, and j's final movement direction is right.

Now consider the results of Lemma 4 applied to agents j and k. There are two cases: (1) Agent k only moves left and passes the leftmost site reached by j. Then either $[l_i, r_i] \subset [l_k, r_k]$ or $[l_k, r_k] \subset [l_i, r_i]$. In either case, by Lemma 3 one of the agents is unnecessary. (2) Agent j only moves right and passes the rightmost site reached by agent k, and k's final movement direction is left. If k passes the leftmost point covered by i, then i is unnecessary by Lemma 3. Otherwise, i passes the leftmost site visited by k. Consider two cases: (a) at l_k, i

has remaining budget less than or equal to k. In this case i is not needed and k can travel to l_i. (b) i has more budget than k at l_k. In this case i also has more budget available than k at all sites left of u_s so k only has to travel right. Thus, by Lemma 3 either j or k are unnecessary. □

3.4 Multi-price *k*-Agent Sufficiency

We now investigate when the multi-price mSPS problem along a path is 1-agent sufficient. Once again we have the obvious cases that if travel costs are all zero, then the problem is 1-agent sufficient and all the sites are located to one side of the start site. We also have the following multi-price analogue of Theorem 2.

Theorem 4. *Consider a strategy that has two agents with paths π^i and π^j and budgets b_i and b_j, respectively. WLOG let $b_i = t_i + c_i^{pur}$ and $b_j = t_j + c_j^{pur}$ where t_i is the budget needed to travel along π_i and c_i^{pur} is the remainder that is allocated to purchase. If $\max(t_i, t_j) \le \min(c_i^{pur}, c_j^{pur})$, then the problem is 1-agent sufficient.*

The proof is almost identical to the proof of Theorem 2.

Figure 4 shows an example of how the number of agents allocated changes for Min-Budget and Max-Probability over different cost profiles for multiple costs along a 100 unit path with 10 sites. Unlike the single-price case, the multi-price problem appears to be NP-Hard. Aumann et al. examine the case where all agents have access to a shared budget and show that even this case is NP-Hard [1]. In this paper, we assume that a distinct, non-sharable, initial budget must be allocated to each agent for both the Max-Probability and Min-Budget problems, adding another dimension of complexity to the problem. However, we were able to use a simplified version of the branch-and-bound algorithm described in Sect. 4 to obtain exact results for smaller sized problems.

Similar to the previous empirical results, we see that as the item costs increase, 1-agent sufficient solutions become more common for the Min-Budget case, but as ρ is increased (i.e. more sites are required to be visited) solutions tend towards two agents unless item costs dominate travel costs. For Max-Probability we see that very small or very large starting budgets lead to solutions with fewer agents, but there is always a dip between the low and high budgets where it becomes more beneficial to use two agents. The scaling and location of this dip is determined by the distribution over item costs, with higher item costs penalizing multiagent solutions.

4 mSPS in 2-D Euclidean Space

The results above have all assumed that the sites are located along a simple path. We now assume that sites are located in 2-dimensional Euclidean space, where the cost to travel between sites is the euclidean distance between sites, i.e. $t_{ij} = \|u_i - u_j\|_2$. When solving both the Min-Budget and Max-Probability problems in Euclidean space, we have the following result:

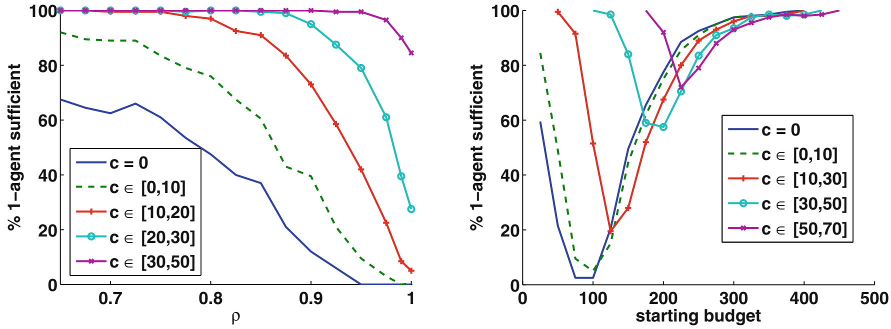

(a) Min-Budget (b) Max-Probability with 10 sites

Fig. 4. Percentage of 200 random 10-site, single-price problems that are 1-agent sufficient where the item costs are either all zero or uniform randomly chosen from some interval. (a) Min-Budget results over different required probabilities of success $p^*_{succ} = \rho \cdot (1 - \prod_{i=1}^{m} 1 - p_i)$ where p_i is the probability the item is available for the single cost at site i. (b) Max-Probability results over total starting budgets.

Proposition 3. *The optimal solution for both the Min-Budget and Max-Probability Euclidean mSPS problems consists of distinct non-overlapping paths.*

Proof. We prove this by contradiction. Assume in the optimal solution that there is a site $u_k \in S$ that is visited by at least two agents a_1 and a_2. Because agents cannot share budget to purchase an item, the probability of success obtained at that site will remain unchanged if the agent with lower budget does not travel to u_k. WLOG assume agent a_1 is thus chosen not to visit the site and instead goes straight to its next site on its path. By the triangle inequality this path length is less than or equal to the original path length resulting in a strategy with no more budget and at least equivalent probability of success, resulting in the desired contradiction. □

Thus, for the Euclidean problem, we can safely exclude overlapping paths from our search space. This separation principle allows us to find optimal solutions to the mSPS problem using an extension of the branch-and-bound formulation proposed by Brown et al. for solving the single-agent SPS problem on general graphs [3]. Any graph more complex than a simple path has been shown to be NP-complete [7]; however, we were able to find optimal solutions for small problems up to about 20 sites in 2-d Euclidean space. This allows us to examine the characteristics of optimal solutions to the mSPS problem in a more general and applicable setting.

We note that for each site $i \in S$ we can use the possible costs at i to form a partition over all possible budgets that may be brought to site i. For example, the cost profile shown in Fig. 5 partitions the budget space into three intervals with the corresponding probabilities of failure if an agent arrives at that site with any budget in that interval. There are three possibilities when the agent arrives at the site: (1) the agent's budget is in the interval $[0, 3)$ and it cannot

C	P(C)
3	0.7
10	0.3

(a)

Budget interval	p_{fail}
$[0, 3)$	1.0
$[3, 10)$	0.3
$[10, \infty)$	0

(b)

Fig. 5. A cost profile (a) partitions the interval $[0, \infty)$ into several possible budget intervals (b), each with an associated probability of failure.

obtain the item, (2) the agent's budget is in the interval $[3, 10)$ and it has enough left to obtain the item at the lower cost but not the higher cost, and so will fail to obtain the item with probability 0.3, and (3) the agent has sufficient budget to obtain the item for any of the possible costs.

Both the Min-Budget and Max-Probability branch-and-bound algorithms need to determine the optimal budget allocation over $n^{\max}$ agents. While the number of all possible budget allocations is infinite, we can ignore most of these intervals and only focus on each budget interval $[c_\ell, c_u)$ induced by the possible costs at each site. Thus rather than branching on individual budget values we branch on possible budget intervals for each agent. This can still result in an exponential number of branches, but does allow exact solutions. We refer the reader to [3] for the full details of the successor function and bounding criteria for the single agent case. To extend the work by Brown et al. to the multiagent case, we simply added a budget interval for each available agent. The successors for each state are found as follows: iterate over all unvisited sites and all available agents; add the site to the agent's path; and update the agent's budget interval, the total budget required for all agents, and the joint probability of failure.

5 *k*-Agent Sufficiency in 2-Dimensions

We introduce the following definition that allows us to characterize a certain class of 1-agent sufficient mSPS problems for the 2-dimensional Euclidean case

Definition 3. *A mSPS problem is purchase-dominated if*

$$\min\{c : P_i(c) > 0, i \in S\} > \max_{i \in S^+, j \in S^+} t_{ij}, \tag{12}$$

i.e., the minimum purchase cost at any site is greater than the maximum travel cost between any two sites.

Theorem 5. *If a mSPS problem in Euclidean space is purchase-dominated, then it is* 1*-agent sufficient.*

Proof. WLOG, assume that we have an optimal solution that requires two agents i and j with paths π^i, π^j of corresponding travel costs t_i and t_j. Additionally, each agent may need some additional budget to use for purchasing, c_i^{pur} and

c_j^{pur}. We show that we can achieve the same probability of success using a single agent. To do this assume that a single agent first visits all sites in π^i and then visits all sites (except for the start site) in π^j with a corresponding total path cost equal to $t_i + t_j - t_{u_s,\pi_1^j} + t_{\pi_\omega^i,\pi_1^j}$, where ω is the index to the last element of the path. Additionally, this agent may need some extra budget to allow for purchasing. This agent will need c_i^{pur} to get the probability of failure p_{fail}^i on path π^i. The agent also needs $\min(0, c_j^{pur} - c_i^{pur})$ to get the probability of failure p_{fail}^j on path π^j. Thus the single agent case requires

$$
\begin{aligned}
b' &= t_i + t_j - t_{u_s,\pi_1^j} + t_{\pi_\omega^i,\pi_1^j} + c_i^{pur} + \min(0, c_j^{pur} - c_i^{pur}) && (13)\\
&\le t_i + t_j + t_{\pi_\omega^i,\pi_1^j} + \max(c_i^{pur}, c_j^{pur}) && (14)\\
&< t_i + t_j + \min(c_i^{pur}, c_j^{pur}) + \max(c_i^{pur}, c_j^{pur}) && (15)\\
&= t_i + t_j + c_i^{pur} + c_j^{pur} && (16)\\
&= B^* && (17)
\end{aligned}
$$

Resulting in single agent strategy with no more budget and at least equivalent probability of success as the strategy with two agents. □

5.1 Results for Clustered Sites

When sites are located in 2-dimensional Euclidean space, there is the potential for many widely separated clusters of sites, which may result in solutions that require more than 2 agents. To examine the effect of clustered sites on the optimal number of agents we ran two experiments, one for Min-Budget and one for Max-Probability. For both experiments we generated data sets consisting of 5 well-separated cluster centers identified in a 100-by-100 region and generated 15 site locations according to varying cluster tendency (*ct*), or the probability that a site will be near a cluster center. For $ct = 0.0$ all sites were uniformly randomly generated in the region, and for increasing *ct* it becomes more likely that sites are located in close proximity to the cluster centers until at $ct = 1.0$, there are no uniform-randomly generated sites. The start site is always placed in the center of the region. For the Min-Budget experiment we used $\rho = 0.95$ and generated random item costs in the intervals $[0]$, $[10, 30]$ and $[30, 50]$. For the Max-Probability experiment we used item costs of 0 and explored total budgets of 100, 600, and 1000. We ran 100 replicates for each setting.

The results are shown in Fig. 6. We see that for Min-Budget, the number of agents used grows as the item costs decrease. Additionally, we see that as *ct* increases, there is a slight increase in the average number of agents used. For Max-Probability there are two trends that are largely insensitive to the clustering tendency: the 100 budget case, and the 1000 budget case. With budget 100, the optimal solutions tend to include two agents to visit sites that are widely separated, but the limited budget tends not to be split any further. With 1000 budget, one agent can typically visit any useful sites. With budget 600 there is a different trend: the number of agents is greatest at $ct = 0.25$ which corresponds

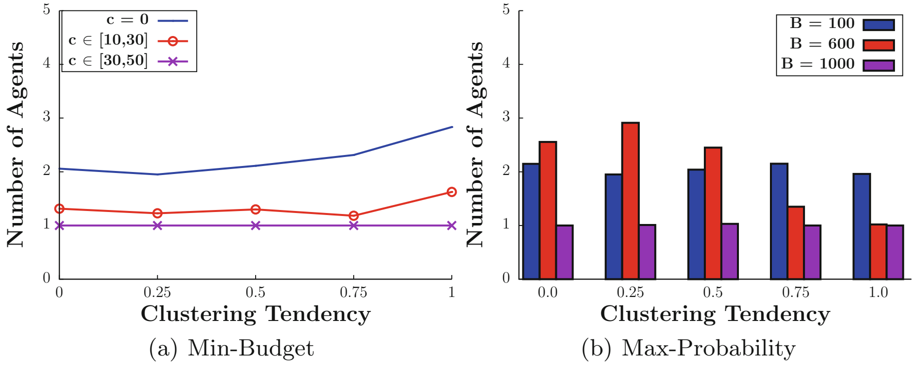

Fig. 6. Average number of agents used in optimal plans for Min-Budget and Max-Probability for 15 sites clustered into 5 clusters in a 100 unit by 100 unit region, with variable clustering tendencies.

to forcing many well-separated sites, so a divide and conquer strategy is used. At $ct = 1.0$, the tight clusters of points make it easier for a single agent to travel within a cluster with very low travel cost.

Our empirical results for these and other settings revealed a tendency towards very few agents in the optimal solution. To obtain a better intuition for this phenomenon, we analyzed increasing numbers of equidistant sites located on a circle centered on the start site.

Theorem 6. *Given an mSPS problem in 2-dimensional Euclidean space with all sites S equidistant from the start site, the problem never requires more than 5 agents.*

Proof. We assume that all sites must be visited in the optimal solution, if all sites are not required, then this can only result in fewer required agents. Additionally, we assume that item costs are all zero, since having positive item costs can never increase the number of agents required in the zero-cost case.

Consider $|S|$ sites equally spaced around a circle of radius r. The case of $|S| = 1$ trivially only requires one agent. Consider two sites as shown in Fig. 7(a). In this case $B^* = 2r$ for two agents, but $B^* = 3r$ for one agent so two agents are necessary. The three site case is shown in Fig. 7(b). In this case using three agents is optimal since the removal of an agent from the solution causes another agent to travel a distance of $\sqrt{2}r + r > 2r$. The cases for 4 and 5 sites are similar—sending an agent along the radius of the circle is cheaper than sending an agent along an edge of the inscribed regular polygon. Figure 7(c) and (d) show the case for $|S| = 6$. This is the break even point where traveling along an edge of the hexagon requires the same budget as traveling along the radius, thus the problem is 1-agent sufficient. For $|S| > 6$ the edges of the inscribed regular $|S|$-gon will be strictly less than r so these cases are all 1-agent sufficient. Even if we relax the assumption that sites are equally spaced, we still only need at most 5 agents. □

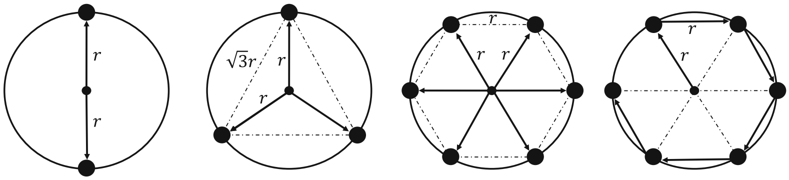

(a) (b) (c) (d)

Fig. 7. Geometric argument that if all sites are equidistant from the start site, then an optimal solution will never require more than 5 agents.

6 Conclusions and Future Work

Often, research into multiagent systems ignores the question of how many agents should actually be used to solve a problem. To investigate this question, we introduced the condition of k-agent sufficiency, as it relates to Multiagent Stochastic Physical Search. We provided the first theoretical and empirical analysis of k-agent sufficiency for mSPS when sites are along a path and in 2-d Euclidean space. We showed that mSPS along a path with a single starting location is always 2-agent sufficient, and quite often 1-agent sufficient. We also showed that even in 2-d Euclidean space with a single starting location, optimal solutions usually require at most 3 agents on average. Our results show strong evidence that optimal solutions to the mSPS problem in 2-d Euclidean space never requires more than 5 agents even if sites located in widely-separated clusters. Using a geometric argument we show why this is true when sites are equidistant from the start site. We conjecture that in general, optimal solutions to mSPS problems in 2-d Euclidean space require at most 5 agents. We hope that these results will inspire other researchers in multiagent planning and optimization to consider cases where multiple agents are not always necessary or even desirable, rather than simply showing that an algorithm or solution method scales to x agents.

We note that there are many assumptions not considered in this paper which may cause the number of agents in an optimal solution to increase. Some of these assumptions include no communication during search, starting all agents from the same initial location, having a limit on the maximum budget per agent, having needs for redundancy or collaboration, and problems where the objective is to minimize time. Future work should examine these extensions to see if there are still k-agent sufficiency conditions. We also hope to leverage the k-agent sufficiency results shown in this paper to develop efficient heuristics and approximation algorithms for solving difficult mSPS problem instances.

Acknowledgments. We would like to thank the Air Office of Scientific Research and the Air Force Research Laboratory's Machine Intelligence for Mission Focused Autonomy Program for funding this research. The views expressed in this article are those of the authors and do not reflect the official policy or position of the USAF, DoD, or the US Government.

References

1. Aumann, Y., Hazon, N., Kraus, S., Sarne, D.: Physical search problems applying economic search models. In: Proceedings of the 23rd AAAI Conference on Artificial Intelligence. AAAI Press (2008)
2. Bektas, T.: The multiple traveling salesman problem: an overview of formulations and solution procedures. Omega **34**(3), 209–219 (2006)
3. Brown, D.S., Hudack, J., Banerjee, B.: Algorithms for stochastic physical search on general graphs. In: Workshops at the Twenty-Ninth AAAI Conference on Artificial Intelligence. AAAI Press (2015)
4. Gabriely, Y., Rimon, E.: Spanning-tree based coverage of continuous areas by a mobile robot. Ann. Math. Artif. Intell. **31**(1–4), 77–98 (2001)
5. Gerkey, B.P., Matarić, M.J.: A formal analysis and taxonomy of task allocation in multi-robot systems. Int. J. Robot. Res. **23**(9), 939–954 (2004)
6. Hazon, N., Aumann, Y., Kraus, S.: Collaborative multi agent physical search with probabilistic knowledge. In: Twenty-First International Joint Conference on Artificial Intelligence (2009)
7. Hazon, N., Aumann, Y., Kraus, S., Sarne, D.: Physical search problems with probabilistic knowledge. Artif. Intell. **196**, 26–52 (2013)
8. Hazon, N., Kaminka, G.A.: Redundancy, efficiency and robustness in multi-robot coverage. In: Proceedings of the 2005 IEEE International Conference on Robotics and Automation, ICRA 2005, pp. 735–741. IEEE (2005)
9. Hudack, J., Gemelli, N., Brown, D., Loscalzo, S., Oh, J.C.: Multiobjective optimization for the stochastic physical search problem. In: Ali, M., Kwon, Y.S., Lee, C.-H., Kim, J., Kim, Y. (eds.) IEA/AIE 2015. LNCS, vol. 9101, pp. 212–221. Springer, Heidelberg (2015)
10. Kang, S., Ouyang, Y.: The traveling purchaser problem with stochastic prices: exact and approximate algorithms. Eur. J. Oper. Res. **209**(3), 265–272 (2011)
11. Sariel-Talay, S., Balch, T.R., Erdogan, N.: Multiple traveling robot problem: a solution based on dynamic task selection and robust execution. IEEE/ASME Trans. Mechatron. **14**(2), 198–206 (2009)
12. Shehory, O., Kraus, S.: Methods for task allocation via agent coalition formation. Artif. Intell. **101**(1), 165–200 (1998)
13. Spires, S.V., Goldsmith, S.Y.: Exhaustive geographic search with mobile robots along space-filling curves. In: Drogoul, A., Fukuda, T., Tambe, M. (eds.) CRW 1998. LNCS, vol. 1456, pp. 1–12. Springer, Heidelberg (1998)
14. Toth, P., Vigo, D.: The Vehicle Routing Problem. Society for Industrial and Applied Mathematics, Philadelphia (2001)
15. Vokřínek, J., Komenda, A., Pěchouček, M.: Agents towards vehicle routing problems. In: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems, pp. 773–780. International Foundation for Autonomous Agents and Multiagent Systems (2010)