

Multi-agent, Multi-target Path Planning in Markov Decision Processes

Farhad Nawaz and Melkior Ornik, *Senior Member, IEEE*

Abstract—Missions for autonomous systems often require agents to visit multiple targets in complex operating conditions. This work considers the problem of visiting a set of targets in minimum time by a team of non-communicating agents in a Markov decision process (MDP). The single-agent problem is at least NP-complete by reducing it to a Hamiltonian path problem. We first discuss an optimal algorithm based on Bellman’s optimality equation that is exponential in the number of target states. Then, we trade-off optimality for time complexity by presenting a suboptimal algorithm that is polynomial at each time step. We prove that the proposed algorithm generates optimal policies for certain classes of MDPs. Extending our procedure to the multi-agent case, we propose a target partitioning algorithm that approximately minimizes the expected time to visit the targets. We prove that our algorithm generates optimal partitions for clustered target scenarios. We present the performance of our algorithms on random MDPs and gridworld environments inspired by ocean dynamics. We show that our algorithms are much faster than the optimal procedure and more optimal than the currently available heuristic.

Index Terms—Agents and Autonomous systems, Markov processes, Stochastic systems, Graph partitioning.

I. INTRODUCTION

A common high-level motion planning problem [1], [2] is for a team of non-communicating agents to visit multiple target states in minimal time. Complex environmental conditions, e.g., underwater regions [3] and planetary bodies [4] may make it difficult for agents to communicate with each other in real time. Motivated by the scenario of non-communicating agents, the high-level planning problem is naturally decoupled to a problem of assigning targets to each agent, and a subsequent single agent path planning problem with multiple target states. Namely, we first seek to find a strategy for a single agent to visit a set of target states in minimum time. Then, we consider a task allocation problem of assigning multiple targets to multiple agents such that the joint travel time to visit the targets by a team of agents is minimal.

The complex and partly unknown environment in many path planning scenarios has motivated the use of stochastic dynam-

ics [5] that model the motion of an agent as a Markov decision process (MDP) [1]. An MDP is a mathematical framework that accounts for the lack of knowledge or complexity of the system dynamics by allowing for stochastic transitions between system states [6]. We interpret the problem of visiting multiple targets by a single agent in an MDP as an extension of a stochastic shortest path (SSP) problem [7]. An SSP problem seeks to find a policy that drives the agent operating on an MDP to a target state with minimal expected cost. Classical dynamic programming methods can be adopted for SSP problems [8]. Since our objective is to visit multiple target states instead of a single target state, a modified approach is required. Hence, we propose a strategy that depends on the history of visited target states.

By pursuing a dynamic programming approach exploited in classical SSP problems [8], we solve our problem for global optimality on a product MDP [9] that includes the history of visited states in the state space. However, we show that our problem is at least NP-complete because it is a generalization of the Hamiltonian path problem [10]. Existing approaches based on approximate dynamic programming are computationally expensive to solve the multi-target problem and are tailored for model-free scenarios [11]. One novel contribution of this work is solving a planning problem that is locally optimal at each time step to trade-off optimality and time complexity, while utilizing the known stochastic model of the environment. We also prove that our algorithm is optimal for some classes of MDPs. A similar approach based on value iteration is used in [12] to explore a set of targets on a gridworld environment with local sensor information. However, our approach is for a general MDP rather than a gridworld without any local sensor data.

Extending to the multi-agent case, we formulate a task allocation problem [13] with the objective of minimizing the time to visit the targets assigned to each agent. The combined target assignment and path planning problem on an MDP is a generalization of the multiple traveling salesman problem (m -TSP) [14] with stochastic transition dynamics. The m -TSP asks to find the shortest possible path for m agents to visit a set of nodes on a complete weighted graph. The objective of task allocation in the m -TSP with stochastic transitions is to assign targets to each agent such that the expected time to visit the targets assigned to each agent is minimized. Existing work [15] related to m -TSP have considered only graphs and minimize the sum of costs between the agents, but we focus on MDPs and minimize the maximum cost between the agents. In our work, we propose a novel partitioning algorithm that

This work was supported by NASA Early Stage Innovations grant no. 80NSSC19K0209, DARPA grant no. HR001120C0065 and Office of Naval Research grant no. N00014-23-1-2505.

F. Nawaz was with the University of Illinois Urbana-Champaign, Urbana, IL 61801, USA. He is now with the University of Pennsylvania, Philadelphia, PA 19104, USA (e-mail: farhadn@seas.upenn.edu).

M. Ornik is with the Department of Aerospace Engineering and the Coordinated Science Laboratory, University of Illinois Urbana-Champaign, Urbana, IL 61801, USA (e-mail: mornik@illinois.edu).

builds upon a heuristic presented in [16] which attempts to solve the task allocation problem in the m -TSP. We compute a complete weighted model graph whose nodes are the states of the original MDP and the weights are the optimal expected time to reach every pair of states. We then adopt the algorithm in [16] for the model graph to generate a partition of the target states that attempts to minimize the time to visit the target states by multiple agents. Each agent then visits the targets assigned to them by following our suboptimal single agent policy to minimize the expected time to visit the target states. Though our partitioning algorithm is suboptimal in general, one contribution of our work is providing conditions of optimality for our heuristic procedure on MDPs where targets are *clustered*. Clusters are defined based on the optimal expected time to reach a state from another state. The presence of clustered target states in an environment is a common scenario in various multi-agent missions [17], [18].

The rest of the paper is organized as follows. Section II provides the necessary mathematical background. In Section III, we formally state the planning problem and the partitioning problem. In Section IV, we discuss the optimal policy for the single agent case. We propose a novel suboptimal single agent planning procedure in Section V and prove that it is optimal for certain classes of MDPs. In Section VI, we propose a partitioning algorithm for the multi-agent case and provide conditions of optimality for our heuristic procedure on clustered target states. Finally, we present the numerical results of our proposed algorithms in Section VII on random graphs, MDPs, and stochastic gridworld environments. In Section VIII, we also compare the performance of our single agent heuristic to the nearest neighbor algorithm for graphs and conclude the paper with possible future directions in Section VIII.

II. PRELIMINARIES

In this paper, $|S|$ denotes the cardinality of set S and 2^S indicates the set of all subsets of S . Notation $\mathbb{E}^\pi[X]$ denotes the expectation of a random variable X when an agent on a Markov decision process (MDP) follows policy π , whereas $\mathbb{E}[X|Y=y]$ denotes the expectation of X given a value y for another random variable Y .

A. Markov Decision Process

An MDP is a mathematical framework to model the motion of agents in a stochastic environment as defined below [6].

Definition 1. A finite Markov decision process is a tuple $\mathcal{M} = (S, A, \mathcal{T})$, where S is a finite set of states, A is a finite set of actions and $\mathcal{T} : S \times A \times S \rightarrow [0, 1]$ is a transition probability function where $\sum_{s' \in S} \mathcal{T}(s, a, s') = 1$ for all $s \in S, a \in A$.

The dynamics of an agent operating in an MDP \mathcal{M} are given as follows. The agent in state $s \in S$ chooses an action $a \in A$ and transitions to a state $s' \in S$ in one time step with probability $\mathcal{T}(s, a, s')$.

A policy π_t for an agent is defined by the probability $\pi_t(a|s)$ that the agent in state s takes action a at time t . A policy π_t applied on an MDP \mathcal{M} generates a random process [19] with

transition probabilities $\pi_t(a|s)\mathcal{T}(s, a, s')$ for all $a \in A, s \in S$, at every time t . A policy is *deterministic* if at every time t and for all states $s \in S$, there exists some $a \in A$ such that $\pi_t(a|s) = 1$. A policy is *stationary* when it is time-independent, i.e., $\pi_t = \pi_{t'}$ for all t, t' . A stationary policy applied on an MDP \mathcal{M} induces a *Markov chain* [19]. Every policy generates many possible *paths* on the MDP. A sample path of an agent is a sequence of states $s_0 s_1 \dots$, where $\pi_t(a|s_t)\mathcal{T}(s_t, a, s_{t+1}) > 0$ for all $t \geq 0$.

B. Value Function

In order to encode the task objective, an MDP is often associated with a reward function $R : S \times A \times S \rightarrow \mathbb{R}$ [6] and the objective is to maximize the long-term expected reward. A value function quantitatively describes how good a state is to satisfy the objective. The *value function* of a *discounted reward infinite horizon* MDP for a stationary deterministic policy π [6] is

$$V^\pi(s) = \mathbb{E}^\pi \left[\sum_{t=0}^{\infty} \gamma^t R(s_t, \pi(s_t), s_{t+1}) \middle| s_0 = s \right], \quad (1)$$

where γ is the discount factor and $\pi(s_t) = a_t$. The objective of maximizing the long-term expected reward is the optimization problem $\max_{\pi} V^\pi(s)$ for all $s \in S$, which returns the *optimal value function* $V^*(s)$. We define an operator T^π [6] for the policy π on some arbitrary value function $V(s)$ by

$$T^\pi(V(s)) = \sum_{s' \in S} \mathcal{T}(s, \pi(s), s') (R(s, \pi(s), s') + \gamma V(s')). \quad (2)$$

A function $V(s)$ is the true value function $V^\pi(s)$ for policy π if and only if $V(s) = T^\pi(V(s))$. Similar to (2), an operator T is given by

$$T(V(s)) = \max_{a \in A} \sum_{s' \in S} \mathcal{T}(s, a, s') (R(s, a, s') + \gamma V(s')). \quad (3)$$

From Bellman's optimality principle [7], $V(s)$ is the optimal value function $V^*(s)$ if and only if $V(s) = T(V(s))$. The operators T^π and T are monotonic, i.e., if $V^1(s) \geq V^2(s)$ for every $s \in S$, then $T^\pi(V^1(s)) \geq T^\pi(V^2(s))$ and $T(V^1(s)) \geq T(V^2(s))$ [6]. The optimal policy π^* that maximizes the long-term expected reward can be obtained from the optimal value function $V^*(s)$ as given below [7]:

$$\pi^*(s) = \arg \max_{a \in A} \sum_{s' \in S} \mathcal{T}(s, a, s') (R(s, a, s') + \gamma V^*(s')). \quad (4)$$

As shown in Theorem 17.8 of [20], the optimal policy is deterministic for all finite MDPs. Hence, we consider only deterministic policies in this work.

In [7], a cost $c(s, a, s') = -R(s, a, s')$ is associated with an MDP. A stochastic shortest path (SSP) problem [7] is a special case of the total cost infinite horizon problem where the agent should reach a goal state with minimum expected cost. In a SSP problem, (i) there is no discounting ($\gamma = 1$), (ii) the target state $s^g \in S$ is *absorbing*, i.e., $\mathcal{T}(s^g, a, s^g) = 1$ for all $a \in A$, and (iii) the target state s^g is *cost-free*, i.e., $c(s^g, a, s) = 0$ for all $a \in A, s \in S$, whereas all other transitions incur a positive cost. A stationary deterministic policy π is *proper* if, when using the policy π , there is a positive probability that the agent will eventually reach the target state for all initial states [7].

C. Cover Time

In simple terms, the cover time for a discrete-time finite-state random process is the time required by an agent to visit all the states, while the hitting time is the time required to visit a particular state. In this subsection, we formally define *hitting time* and *cover time* for a discrete-time finite-state random process $X = \{X_t\}$, where $X_t \in S$ for all $t \geq 0$ [19].

Definition 2. Hitting time of a state $s \in S$ is the first time the agent visits s , starting from $s_0 \in S$. We denote it by the random variable H_{X,s_0}^s and is defined by

$$H_{X,s_0}^s = \min\{t : X_t = s | X_0 = s_0\}. \quad (5)$$

We denote $H_{\mathcal{M},s_0}^{s,\pi}$ as the hitting time of state s starting from s_0 for the MDP \mathcal{M} with policy π , and $\mathbb{E}^\pi[H_{\mathcal{M},s_0}^s] = \mathbb{E}[H_{\mathcal{M},s_0}^{s,\pi}]$ as the expected hitting time.

Definition 3. Cover time of a set of states $\mathcal{V} \subseteq S$ is the time required by the agent to visit all the states $s \in \mathcal{V}$, starting from $s_0 \in S$. We denote it by the random variable $C_{X,s_0}^{\mathcal{V}}$ and is defined as

$$C_{X,s_0}^{\mathcal{V}} = \min\{t : H_{X,s_0}^s \leq t \text{ for all } s \in \mathcal{V}\}. \quad (6)$$

The equation for cover time (6) is equivalent to $C_{X,s_0}^{\mathcal{V}} = \max_{s \in \mathcal{V}} H_{X,s_0}^s$. Similar to hitting time, we denote $C_{\mathcal{M},s_0}^{\mathcal{V},\pi}$ as the cover time of \mathcal{V} starting from s_0 for the MDP \mathcal{M} with policy π , and $\mathbb{E}^\pi[C_{\mathcal{M},s_0}^{\mathcal{V}}] = \mathbb{E}[C_{\mathcal{M},s_0}^{\mathcal{V},\pi}]$ as the expected cover time. Since the MDP \mathcal{M} is known in our work, we use the notation $H_{s_0}^{s,\pi} = H_{\mathcal{M},s_0}^{s,\pi}$ and $C_{s_0}^{\mathcal{V},\pi} = C_{\mathcal{M},s_0}^{\mathcal{V},\pi}$. A Markov chain is *irreducible* if the agent can reach s' starting from s in a finite number of steps with non-zero probability, for every pair of states $s, s' \in S$. If the Markov chain induced by a stationary policy π and MDP \mathcal{M} is irreducible, then the expected cover time is finite [21].

Utilizing the above preliminaries, we formulate the problem of optimal cover time in the subsequent section.

III. PROBLEM FORMULATION

The objective considered in this work is for multiple agents to jointly visit multiple target states in minimal expected time. We make the following assumption on a priori known MDP \mathcal{M} .

Assumption 1. There exists some stationary policy π such that the resulting Markov chain induced by π is irreducible.

If Assumption 1 does not hold, then there exists some $s \in S$ and $\mathcal{V} \subseteq S$ such that the expected cover time $\mathbb{E}^\pi[C_{s_0}^{\mathcal{V}}]$ is infinite for all policies [21].

We consider the following problem statement for a single agent to visit multiple targets.

Problem 1. Let an agent operate in an MDP \mathcal{M} . Let $s_0 \in S$ be its initial state, \mathcal{V} be the set of target states to be covered, and $\mathbb{E}^\pi[C_{s_0}^{\mathcal{V}}]$ be the expected cover time of \mathcal{V} for an agent starting from state s_0 and acting under policy π . Under Assumption 1, find a control policy π^* that solves the optimization problem

$$\pi^* = \arg \min_{\pi} (\mathbb{E}^\pi[C_{s_0}^{\mathcal{V}}]). \quad (7)$$

The expected cover time $\mathbb{E}^\pi[C_{s_0}^{\mathcal{V}}]$ depends both on the initial state $s_0 \in S$ and the target set \mathcal{V} . Hence, it is appropriate to consider the motion of the agent on the product state space $S_p = S \times 2^{\mathcal{V}}$ which is often used in the theory of model checking [9]. The agent's extended state $(s, \bar{\mathcal{V}}) \in S_p$ encodes its current location $s \in S$ and the remaining set of states $\bar{\mathcal{V}} \subseteq \mathcal{V}$ to be visited. Hence, a transition in the product space S_p denotes a change in the state $s \in S$ of the agent and a change in the remaining set of states $\bar{\mathcal{V}} \subseteq \mathcal{V}$ to be visited. The transition probabilities $\mathcal{T}_p : S_p \times A \times S_p \rightarrow [0, 1]$ are defined as

$$\begin{aligned} \mathcal{T}_p((s_1, \bar{\mathcal{V}}_1), a, (s_2, \bar{\mathcal{V}}_2)) &= \begin{cases} \mathcal{T}(s_1, a, s_2) & \text{if } \bar{\mathcal{V}}_2 = \bar{\mathcal{V}}_1 \setminus \{s_2\}, \\ 0 & \text{otherwise,} \end{cases} \\ &\text{for all } (s_1, \bar{\mathcal{V}}_1), (s_2, \bar{\mathcal{V}}_2) \in S_p, a \in A. \end{aligned} \quad (8)$$

The key takeaway from (8) is that the agent can never shrink the remaining set of states to be visited by more than one in a single transition. Then, a stationary deterministic policy π on the product MDP $\mathcal{M}_p = (S_p, A, \mathcal{T}_p)$ is a mapping from the product space to the action space given by $\pi : S_p \rightarrow A$. In our problem, the objective is to visit all states in \mathcal{V} . Thus, the target set is $T_p = \{S \times \{\emptyset\}\}$ and we wish to compute a policy that leads the agent to any state $t_p \in T_p$ in minimal expected time.

Using (6) and [21], given a set $\bar{\mathcal{V}} \subseteq S$ to be covered, the expected cover time $\mathbb{E}^\pi[C_s^{\bar{\mathcal{V}}}]$ when the agent starts from state $s \in S$ with policy π is given by

$$\begin{aligned} \mathbb{E}^\pi[C_s^{\bar{\mathcal{V}}}] &= 1 + \\ &\sum_{s' \in S} \mathcal{T}_p((s, \bar{\mathcal{V}}), \pi(s, \bar{\mathcal{V}}), (s', \bar{\mathcal{V}} \setminus \{s'\})) \mathbb{E}^\pi[C_{s'}^{\bar{\mathcal{V}} \setminus \{s'\}}] \\ &\text{for all } (s, \bar{\mathcal{V}}) \in S_p. \end{aligned} \quad (9)$$

Since we consider only states $(s', \bar{\mathcal{V}} \setminus \{s'\})$ in (9), all the zero probabilities mentioned in the second condition of (8) are not included in (9). Based on (9), we interpret one unit of time as the cost incurred for the immediate one-step state transition and $\mathbb{E}^\pi[C_{s'}^{\bar{\mathcal{V}} \setminus \{s'\}}]$ as the future expected cost for each possible state $s' \in S$ until the agent covers the set $\bar{\mathcal{V}}$. If $\bar{\mathcal{V}} = \emptyset$ or $\bar{\mathcal{V}} = \{s\}$, then the agent has already covered the required set and cover time is zero: $\mathbb{E}^\pi[C_s^{\{s\}}] = \mathbb{E}^\pi[C_s^{\{\emptyset\}}] = 0$ for all $s \in S$. By Assumption 1, there exists a policy $\pi' : S \rightarrow A$ which generates an irreducible Markov chain. We can define a policy $\pi : S_p \rightarrow A$ as $\pi(s, \bar{\mathcal{V}}) = \pi'(s)$ for all $(s, \bar{\mathcal{V}}) \in S_p$ so that the expected cover times $\mathbb{E}^\pi[C_s^{\bar{\mathcal{V}}}]$ are finite for all $(s, \bar{\mathcal{V}}) \in S_p$.

Since we introduced the MDP problem in Section II-B as maximizing the long term expected reward, we multiply (9) by -1 to convert cost into reward. We interpret value function (2) as negative of the expected cover time. The reward function $R_p : S_p \rightarrow \mathbb{R}$ is

$$R_p(s, \bar{\mathcal{V}}) = \begin{cases} 0 & \text{if } \bar{\mathcal{V}} \in \{\emptyset, \{s\}\}, \\ -1 & \text{otherwise.} \end{cases} \quad (10)$$

The reward function described in Section II-B for the product space S_p is $R : S_p \times A \times S_p \rightarrow \mathbb{R}$, but the reward function in (10) is $R_p : S_p \rightarrow \mathbb{R}$. All notions introduced in Section II can be analogously stated for the reward function in (10) as well. From (9) and (10), the recursive value function equation for a given policy π is

$$V^\pi(s, \bar{\mathcal{V}}) = -1 + \sum_{s' \in S} \mathcal{T}_p((s, \bar{\mathcal{V}}), \pi(s, \bar{\mathcal{V}}), (s', \bar{\mathcal{V}} \setminus \{s'\})) V^\pi(s', \bar{\mathcal{V}} \setminus \{s'\})$$

for all $(s, \bar{\mathcal{V}}) \in S_p$, (11)

where $V^\pi(s, \bar{\mathcal{V}}) = -\mathbb{E}^\pi [C_s^{\bar{\mathcal{V}}}]$ is the value function. Therefore, problem (7) produces the same solution as $\pi^* = \arg \max_\pi (V^\pi(s_0, \mathcal{V}))$. Analogously, we can interpret expected cover time (9) as the expected cost and solve the optimization problem (7) to minimize the expected cost.

Once the agent reaches a target state $t_p \in T_p$, from (8), it does not transition anymore, and from (10), the cost is 0. Comparing (2) and (11), $\gamma = 1$ in our work, which matches the SSP problem defined in Section 2.1 of [7]. Therefore, Problem 1 is an extended version of the SSP problem on the product space S_p , with a target set T_p instead of one target state.

Extending the problem from a single agent to a team, we now consider multiple agents operating on the MDP \mathcal{M} , all starting from the same state $s_0 \in S$. In this work, we assume that all the agents follow the same transition dynamics \mathcal{T} . We first define a partition of multiple targets, and then introduce the problem of optimal target assignment to multiple agents before the start of the mission.

Definition 4. *Let there be m agents operating in an MDP \mathcal{M} and \mathcal{V} be the set of target states to be covered. Then, $\mathcal{P} = \{P_1, P_2, \dots, P_m\}$ is a partition of the target states \mathcal{V} for m agents such that $\cup_{i=1}^m P_i = \mathcal{V}$, $P_i \cap P_j = \emptyset$ for all $i, j \in \{1, 2, \dots, m\}$, $i \neq j$.*

For a fixed partition \mathcal{P} , the optimal expected time for m agents to jointly visit the target states \mathcal{V} is

$$\min_{\{\pi_1, \pi_2, \dots, \pi_m\}} \mathbb{E} \left[\max_{P_i} (C_{s_0}^{P_i, \pi_i}) \right], \quad (12)$$

where the policy for agent $i \in \{1, 2, \dots, m\}$ is π_i . Solving the problem in (12) would require us to compute the expectation of the maximum of random variables $C_{s_0}^{P_i, \pi_i}$ and then jointly optimize over the policies $\{\pi_1, \pi_2, \dots, \pi_m\}$. Since it is hard to jointly optimize over the expected value of the maximum of random variables [22], we assume that each agent i uses the policy π_i^* from Problem 1 that minimizes the expected cover time for targets P_i . Hence, we implicitly make the approximation

$$\min_{\{\pi_1, \pi_2, \dots, \pi_m\}} \mathbb{E} \left[\max_{P_i} (C_{s_0}^{P_i, \pi_i}) \right] \approx \max_{P_i} \left(\min_{\pi_i} \mathbb{E} [C_{s_0}^{P_i, \pi_i}] \right) \quad (13)$$

by fixing the policy $\pi_i^* = \arg \min_\pi \mathbb{E}^\pi [C_{s_0}^{P_i}]$ for each agent i so that we can take the maximum out of the expectation as given in (13) and then solve for the partition \mathcal{P} . Since we

interchanged the order of minimizing over the policies and maximizing over the target sets P_i , (13) is not an equality in general and establishing the relationship between the two values remains a topic for future work.

The formal problem statement of allocating targets to multiple agents is given below after making the approximation in (13).

Problem 2. *Let there be m agents operating in an MDP \mathcal{M} . Let $s_0 \in S$ be the initial state of the m agents, and \mathcal{V} be the set of target states to be covered. Let \mathcal{P} be a partition of the target states \mathcal{V} for m agents. Let π_i^* and $\mathbb{E}^{\pi_i^*} [C_{s_0}^{P_i}]$ be the optimal policy and optimal expected cover time for agent $i \in \{1, 2, \dots, m\}$, respectively. Find a partition \mathcal{P}^* that solves the optimization problem*

$$\mathcal{P}^* = \arg \min_{\mathcal{P}} \max_i \left(\mathbb{E}^{\pi_i^*} [C_{s_0}^{P_i}] \right). \quad (14)$$

In subsequent sections, we present the solution methodologies that approximately minimize the expected cover time.

IV. OPTIMAL POLICY FOR SINGLE AGENT

The objective of Problem 1 is to find a policy that results in minimal expected cover time for a single agent. In this section, we adopt *policy iteration* [11] to solve for π^* of the SSP problem we discussed in Section III. We describe the optimal policy iteration procedure below, which also serves as the motivation for our heuristic method.

It is shown in [8] that *policy iteration* — when starting with a proper policy — can be adopted for SSP problems. Assumption 1 guarantees the existence of a stationary policy such that the agent can reach s' starting from s in a finite number of steps with non-zero probability, for every pair of states $s, s' \in S$. Hence, we can use Lemma 3 in [23] to start with a proper policy that chooses an action uniformly at random for all $(s, \bar{\mathcal{V}}) \in S_p$ where $\bar{\mathcal{V}} \neq \emptyset$.

The first step of policy iteration is to recursively compute the value function for the policy $\pi^*(s, \bar{\mathcal{V}})$ using (11). Since the policy is proper, as shown in Proposition 1.1(a) of [7], recursively computing the value function using (11) converges to the actual value function for policy π . Then, the new policy is computed using (4) where $\gamma = 1$ for SSP problems. The two steps of value function computation and policy update are repeated until the old policy and new policy are exactly same.

Proposition 1. *Under Assumption 1 and starting from a proper policy, policy iteration on the product MDP \mathcal{M}_p with reward function (10) returns an optimal stationary deterministic policy $\pi^* : S_p \rightarrow A$ for Problem 1.*

Proof. We showed in Section III that Problem 1 is a SSP problem on the product MDP \mathcal{M}_p with multiple target states. In this section, we described the policy iteration procedure to solve for the optimal policy of a SSP problem. As shown in Proposition 3.5 of [7], the new updated policy is strictly better if the old policy is not optimal. Since the number of proper deterministic policies is finite, policy iteration on the MDP \mathcal{M}_p with reward function (10) always returns the optimal policy to reach a target state with $\bar{\mathcal{V}} \in \{\emptyset, \{s\}\}$. \square

Though Proposition 1 returns the optimal policy, finding it comes at a high computational cost as we will show that Problem 1 is NP-complete by reducing it to a Hamiltonian path problem [10]. A graph is a deterministic MDP, where the transition probability function satisfies $\mathcal{T}(s, a, s') \in \{0, 1\}$ for all $s, s' \in S$, $a \in A$. A Hamiltonian path [10] visits all the states of a graph exactly once. If a Hamiltonian path exists for a graph and the set of targets is the state space S , then Problem 1 seeks to find a Hamiltonian path on the graph. However, the problem of determining whether a Hamiltonian path exists for a graph is NP-complete [10] in the number of states. Hence, Problem 1 is at least as hard as minimizing the cover time on a graph. Since the value function is computed for all $(s, \bar{\mathcal{V}}) \in S_p$ in the policy iteration procedure, it consumes $O(|S|^2 (2^{|\mathcal{V}|}))$ number of operations which is exponential time. Since we discussed using the Hamiltonian path problem that Problem 1 is at least NP-complete, the policy iteration procedure indeed achieves exponential time complexity for Problem 1.

In the next section, we propose a method with polynomial time complexity at each time step.

V. SUBOPTIMAL VALUE ITERATION FOR SINGLE AGENT

In this section, we present a heuristic method for Problem 1 by adopting an approximate value function. We prove that our procedure optimally reaches an unvisited target state at every time step and also prove the optimality of our algorithm on a small class of deterministic MDPs.

The intuition behind our procedure is to, at every time t , compute a policy that drives the agent to an unvisited state in \mathcal{V} in minimum expected time. We use the computational procedure in policy iteration as a motivation to compute an approximate value function based on value iteration. Let, at time t , the remaining set of states to be visited be $\bar{\mathcal{V}}_t$. We propose a reward function $\hat{R}_t : S \rightarrow \mathbb{R}$ by

$$\hat{R}_t(s) = \begin{cases} -|\bar{\mathcal{V}}_t| & \text{if } s \notin \bar{\mathcal{V}}_t, \\ -|\bar{\mathcal{V}}_t| + 1 & \text{if } s \in \bar{\mathcal{V}}_t. \end{cases} \quad (15)$$

At every time t , we aim to find a policy that yields maximum long-term expected reward using (15) for the MDP \mathcal{M} . The optimal value function $\hat{V}_t^*(s)$ for such a problem is

$$\hat{V}_t^*(s) = \max_{\hat{\pi}_t} \mathbb{E}^{\hat{\pi}_t} \left[\sum_{k=0}^{\infty} \gamma^k \hat{R}_t(s_k) \mid s_0 = s \right], \quad (16)$$

where $\hat{\pi}_t : S \rightarrow A$ is the deterministic stationary policy generated at time t . Let the optimal policy which maximizes (16) be $\hat{\pi}_t^*$. Then, using Bellman's optimality principle [7], $\hat{V}_t^*(s)$ and $\hat{\pi}_t^*(s)$ can be computed as

$$\hat{V}_t^*(s) = \max_{a \in A} \sum_{s' \in S} \mathcal{T}(s, a, s') \left(\hat{R}_t(s') + \gamma \hat{V}_t^*(s') \right), \quad \text{for all } s \in S, \quad (17)$$

$$\hat{\pi}_t^*(s) = \arg \max_{a \in A} \sum_{s' \in S} \mathcal{T}(s, a, s') \left(\hat{R}_t(s') + \gamma \hat{V}_t^*(s') \right), \quad \text{for all } s \in S. \quad (18)$$

The discount factor is $\gamma \in [0, 1)$. Note that we used a discount factor of $\gamma = 1$ in Section IV to solve for the optimal policy since it was a SSP problem on the product MDP \mathcal{M}_p . We can't use a discount factor of $\gamma = 1$ in this section because the infinite sum in (16) will have a finite value only if $\gamma \in [0, 1)$. Let at time t the state of the agent be $s_t \in S$. If $\gamma = 0$, (18) returns an action that transitions the agent to a state in $\bar{\mathcal{V}}_t$ in exactly one transition with maximal probability. If there is no such one-step transition, i.e., $\mathcal{T}(s_t, a, s') = 0$ for all $s' \in \bar{\mathcal{V}}_t, a \in A$ when $\gamma = 0$, then $\hat{\pi}_t^*(s_t)$ has multiple solutions. We use the convention that if $\hat{\pi}_t^*(s_t)$ is not unique, then the agent chooses one of the actions selected at random since all such actions return the optimal value function $\hat{V}_t^*(s)$ for $\gamma = 0$. Thus, if $\gamma = 0$ and the MDP is a graph, then (18) is exactly the *nearest neighbor* heuristic [24]. However, our heuristic presented in this section aims to find a path of minimal length to an unvisited state. As we increase γ , the emphasis on rewards of future time steps also increases.

From (16) and (15), we note that the approximate value function $\hat{V}_t^*(s)$ changes if and only if $\bar{\mathcal{V}}_t$ changes. Therefore, the benefit of (17) is that the value function depends only on the states $s \in S$ and we recompute $\hat{V}_t^*(s)$ only when the agent visits an unvisited state $s \in \bar{\mathcal{V}}_t$. Unlike the policy iteration procedure described in Section IV, where we need to compute the value function for all states S and all subsets of \mathcal{V} , our new proposed procedure computes the value function only for all states S .

We present our procedure in Algorithm 1. The value function $\hat{V}_t(s)$ is computed for all states in lines 6–12 using value iteration and (17). The policy is computed in line 14 as given in (18). Lines 14–17 are repeated until the agent visits a state in $\bar{\mathcal{V}}_t$. Once the agent visits a state in $\bar{\mathcal{V}}_t$, the set of states to be covered is updated and denoted by $\bar{\mathcal{V}}_{t+1}$. Then, the new value function is computed for $\bar{\mathcal{V}}_{t+1}$. The iterative procedure repeats until the agent has visited all the states in \mathcal{V} .

Now, we seek to prove that Algorithm 1 will terminate in finite time. We use the following claim.

Proposition 2. *Under Assumption 1, there exists some $\bar{\varepsilon} > 0$ such that for all $\varepsilon \in (0, \bar{\varepsilon})$ and any $\gamma \in (0, 1)$, at every time $t \geq 0$, Algorithm 1 generates a policy that drives the agent starting from state $s_t \in S$ to eventually reach some state $s \in \bar{\mathcal{V}}_t$ with probability 1, and reach it with minimal expected time.*

Proof. In Algorithm 1, we use a stopping criterion ε which returns an estimate $\hat{V}_t(s)$ of the optimal value function $\hat{V}_t^*(s)$ that solves $T(\hat{V}_t^*(s)) = \hat{V}_t^*(s)$. The value function computed based on value iteration in lines 7–11 uses the optimal Bellman operator T as described in Section II-B. Since the operator T is monotone [6], the value function computed using value iteration converges to the optimal value function. Thus, as shown on page 27 of [7], there exists a small enough $\bar{\varepsilon} > 0$ such that for all $\varepsilon \in (0, \bar{\varepsilon})$, the policy computed in line 14 is optimal after a finite number of iterations of value iteration present in lines 7–11 of Algorithm 1. The computed policy is optimal with respect to the reward function defined in (15). From (15), at every time t , all states $s \in \bar{\mathcal{V}}_t$ have the same reward $-|\bar{\mathcal{V}}_t| + 1$, which is higher than the reward $-|\bar{\mathcal{V}}_t|$

Algorithm 1: Suboptimal Value Iteration

```

1 Input:  $S, A, \mathcal{T}, \mathcal{V}, \varepsilon, s_0, \gamma$ ;
2 Initialize:  $t = 0, \bar{\mathcal{V}}_0 = \mathcal{V} \setminus \{s_0\}$ ;
3  $\hat{V}_t(s) = -\infty$  for all  $s \in S$ ;
4 while  $\bar{\mathcal{V}}_t \neq \emptyset$  do
5    $\Delta = \infty$ ;
6   while  $\Delta \geq \varepsilon$  do
7     forall  $s \in S$  do
8        $\hat{V}'_t(s) \leftarrow \hat{V}_t(s)$ ;
9        $\hat{V}_t(s) \leftarrow$ 
10         $\max_{a \in A} \sum_{s' \in S} \mathcal{T}(s, a, s') \left( \hat{R}_t(s') + \right.$ 
11         $\left. \gamma \hat{V}_t(s') \right)$ ;
12     end
13      $\Delta \leftarrow \max_s |\hat{V}_t(s) - \hat{V}'_t(s)|$ 
14   end
15   repeat
16      $\hat{\pi}_t^*(s_t) \leftarrow$ 
17      $\arg \max_{a \in A} \sum_{s' \in S} \mathcal{T}(s_t, a, s') \left( \hat{R}_t(s') + \right.$ 
18      $\left. \gamma \hat{V}_t(s') \right)$ ;
19     Implement action  $\hat{\pi}_t^*(s_t)$ ;
20      $t = t + 1$ ;
21     Update set to be covered:  $\bar{\mathcal{V}}_t = \bar{\mathcal{V}}_{t-1} \setminus \{s_t\}$ ;
22   until  $|\bar{\mathcal{V}}_t| = |\bar{\mathcal{V}}_{t-1}| - 1$ ;
23 end

```

assigned for the states $s \in S \setminus \bar{\mathcal{V}}_t$. Hence, the optimal policy computed in line 14 will eventually drive the agent to some state $s \in \bar{\mathcal{V}}_t$ with probability 1, in minimal expected time. \square

Corollary 1. *Under Assumption 1 and any $\gamma \in (0, 1)$, Algorithm 1 will eventually cover the given set of states \mathcal{V} with probability 1.*

Proof. As shown in Proposition 2, the agent reaches a state in $\bar{\mathcal{V}}_t$ in minimal expected time, i.e., $|\bar{\mathcal{V}}_{t'}| = |\bar{\mathcal{V}}_t| - 1$ for some finite time $t' > t$, eventually leading to $|\bar{\mathcal{V}}_t| = 0$. \square

The value function update in line 9 of Algorithm 1 has time complexity $O(|S||A|)$. Lines 7-11 iterates over all possible states and hence it requires $O(|S|^2|A|)$ operations. Then, the time complexity to compute the policy in line 14 is $O(|S||A|)$. Therefore, in a single time step, Algorithm 1 consumes only $O(|S|^2|A| + |S||A|)$ operations, in comparison to $O(|S|^2(2^{|\mathcal{V}|}))$ operations consumed by the optimal policy iteration procedure described in Section IV.

The inherent suboptimality of Algorithm 1 is a consequence of the value function $\hat{V}_t(s)$ which does not always capture the complete information about the minimum expected cover time. However, as we will show in Section VII using numerical results that $\hat{V}_t(s)$ often produces a reasonable approximation of the minimal expected cover time. Additionally, we can prove that the policy from (18) is optimal for path graphs, cycle graphs, and complete graphs. These undirected graphs are illustrated in Fig. 1; we define them as follows. If the states (vertices) are denoted by $S = \{1, 2, \dots, n\}$ and the




Fig. 1: An illustration of the graphs where Theorem 1 applies, with possible initial states in red.

terminal vertices are 1 and n , the *path graph* has edges $E_p = \{\{i, i + 1\} \mid i = 1, 2, \dots, n - 1\}$, and the *cycle graph* has edges $E_p \cup \{\{1, n\}\}$, whereas the *complete graph* has an edge between every pair of vertices.

Theorem 1. *Algorithm 1 always produces optimal paths with any $\gamma \in (0, 1)$ for Problem 1 to cover any given set of targets for the following deterministic MDPs: (i) path graphs with initial state as either of the terminal vertices, (ii) cycle graphs, and (iii) complete graphs.*

Proof. The result follows directly from Proposition 2 by noting that Algorithm 1 updates the target set each time the agent visits a target. For complete graphs, since the agent can reach any state from any other state in exactly one step, Algorithm 1 trivially visits the target states optimally, without visiting any state that is not a target. In path graphs with terminal initial conditions, there is only one direction the agent can start to move and hence that path is optimal. In cycle graphs, the agent has two options for the initial direction of motion. From Proposition 2, the agent chooses the direction which takes it to the nearest unvisited target. Once that target is visited, the rest of the problem is like a path graph. \square

In the next section, we consider the task of jointly visiting multiple targets by a team of agents. We discuss the optimal solution to the problem of partitioning targets to multiple agents as defined in Problem 2 and propose a suboptimal procedure.

VI. PARTITIONING TARGETS TO MULTIPLE AGENTS

Multi-agent planning in our framework has two stages. The first stage partitions the targets to agents as defined in Problem 2, and the second stage consists of optimal planning for each agent described in Problem 1. In Section IV, we showed that optimal planning for single agent is NP-complete. However, even if Problem 1 is solved to optimality for all agents operating on a deterministic MDP, Problem 2 reduces to a graph partitioning problem, which is provably NP-hard [25].

Previous work [16] which solves a partitioning problem similar to Problem 2 deals with a multi-agent traveling salesman problem on a complete graph. Though our work also aims to solve a partitioning problem, it deals with a Markov decision process and not necessarily a complete graph. Hence, we base our approach to Problem 2 on ideas from [16], but substantially adapting their method for a stochastic environment. We also

present some theoretical conditions under which our algorithm generates optimal partitions on an environment with clustered targets.

A. Average length of a Hamiltonian path

Since the traveling salesman problem is defined on a weighted complete graph, we compute a weighted complete directed graph $\mathcal{G} = (S, E)$ for the MDP \mathcal{M} such that the weights are $w(s_1, s_2) = \min_{\pi} \mathbb{E}^{\pi} [C_{s_1}^{\{s_2\}}] = \min_{\pi} \mathbb{E}^{\pi} [H_{s_1}^{s_2}]$ for all $s_1, s_2 \in S$. The weights $w(s_1, s_2)$ are the optimal expected hitting times for every pair of states $s_1, s_2 \in S$ and all the weights are finite because of Assumption 1. The weights can be computed using the policy iteration procedure described in Section IV with $\mathcal{V} = \{s_2\}$ and $s_0 = s_1$ for all $s_1, s_2 \in S$ which requires $O(|S|^3|A|)$ number of operations. Hence, the time to compute graph \mathcal{G} is polynomial in the number of states and actions.

As discussed in Section IV, since computing the optimal cover time is NP-complete, we adopt the idea used in [16] to approximate the optimal cover time. Let \mathcal{G}_i be the sub-graph of \mathcal{G} induced by P_i , where P_i are the target states assigned to agent i as defined in Problem 2. We denote the average length of a Hamiltonian path which covers all the states in \mathcal{G}_i starting from $s_0 \in S$ by $L_a(\mathcal{G}_i)$. We use the average length of a Hamiltonian path as a heuristic for the optimal cover time $\mathbb{E}^{\pi_i^*} [C_{s_0}^{P_i}]$ for all $i \in \{1, 2, \dots, m\}$. Since we consider $s_0 \in S$ as the initial state throughout this work, we don't explicitly mention s_0 in $L_a(\mathcal{G}_i)$. Let $|P_i| = n_i$ for all $i \in \{1, 2, \dots, m\}$. Since there are $n_i(n_i - 1)$ edges on the sub-graph induced by the states P_i , and there are $n_i - 1$ edges on a path to cover all the states in P_i ,

$$\begin{aligned} L_a(\mathcal{G}_i) &= (n_i - 1) \left(\frac{\sum_{s_1 \in P_i} \sum_{s_2 \in P_i} w(s_1, s_2)}{n_i(n_i - 1)} \right) + \\ &\quad \frac{\sum_{s \in P_i} w(s_0, s)}{n_i}, \\ \Rightarrow L_a(\mathcal{G}_i) &= \frac{\sum_{s_1 \in P_i} \sum_{s_2 \in P_i} w(s_1, s_2) + \sum_{s \in P_i} w(s_0, s)}{n_i}. \end{aligned} \quad (19)$$

Using (19), we can compute $L_a(\mathcal{G}_i)$ for each $i \in \{1, 2, \dots, m\}$ in $O(n_i^2)$. Since $\sum_{i=1}^m n_i = |\mathcal{V}|$ and $\sum_{i=1}^m n_i^2 < |\mathcal{V}|^2$, we can compute $L_a(\mathcal{G}_i)$ for all $i \in \{1, 2, \dots, m\}$ in $O(|\mathcal{V}|^2)$. Hence, instead of attempting to solve (14) in Problem 2, we try to find a partition \mathcal{P}' that solves the below optimization problem using $L_a(\mathcal{G}_i)$ from (19):

$$\mathcal{P}' = \arg \min_{\mathcal{P}} \max_i (L_a(\mathcal{G}_i)). \quad (20)$$

The problem in (20) is the same problem in (14), but replaces the optimal cover time $\mathbb{E}^{\pi_i^*} [C_{s_0}^{P_i}]$ with the average length of the Hamiltonian path $L_a(\mathcal{G}_i)$ on the sub-graph \mathcal{G}_i .

The optimization problem in (20) is NP-hard because it reduces to the popular number partition problem (NPP) [26], which is NP-hard as shown in [16]. The NPP seeks to divide a set \mathcal{S} of positive integers with a fixed even sum K into two subsets $\{\mathcal{S}_1, \mathcal{S}_2\}$ such that the sum of numbers in each subset is exactly $\frac{K}{2}$. Similarly, given a complete graph \mathcal{G}

and $\mathcal{V} \subseteq S \setminus \{s_0\}$, (20) asks us to find a partition $\{\mathcal{G}_1, \mathcal{G}_2\}$ where $\mathcal{G}_1 \cup \mathcal{G}_2 = \mathcal{V}$ such that $\max\{L_a(\mathcal{G}_1), L_a(\mathcal{G}_2)\} = \frac{K}{2}$ for $m = 2$ agents. Thus, Problem (20) is NP-hard and we adopt the heuristic presented in [16] for the graph \mathcal{G} .

B. Partitioning by transfers and swaps

The heuristic partitioning procedure searches for a series of transfers and swaps of states between pairs of sub-graphs $(\mathcal{G}_i, \mathcal{G}_k)$ to decrease $\max\{L_a(\mathcal{G}_i), L_a(\mathcal{G}_k)\}$. We denote the sum of edge weights of a sub-graph \mathcal{G}_i by

$$W(\mathcal{G}_i) = \sum_{s_1 \in P_i} \sum_{s_2 \in P_i} w(s_1, s_2) + \sum_{s \in P_i} w(s_0, s), \quad (21)$$

and the contribution of a state $s \in S$ to $W(\mathcal{G}_i)$ by

$$\Delta W(\mathcal{G}_i, s) = \sum_{s_1 \in P_i} w(s_1, s) + \sum_{s_1 \in P_i} w(s, s_1) + w(s_0, s). \quad (22)$$

From (21) and (19), we can write the average length of Hamiltonian path on a sub-graph \mathcal{G}_i by

$$L_a(\mathcal{G}_i) = \frac{W(\mathcal{G}_i)}{n_i}, \quad (23)$$

and the maximum average length of Hamiltonian path for a partition \mathcal{P} by

$$M_a(\mathcal{P}) = \max_i L_a(\mathcal{G}_i). \quad (24)$$

Let us analyze the transfer of a state $s_i \in P_i$ from \mathcal{G}_i to \mathcal{G}_k . After the transfer, for the sub-graphs \mathcal{G}'_i and \mathcal{G}'_k ,

$$\begin{aligned} L_a(\mathcal{G}'_i) &= \frac{W(\mathcal{G}_i) - \Delta W(\mathcal{G}_i, s_i)}{n_i - 1}, \\ L_a(\mathcal{G}'_k) &= \frac{W(\mathcal{G}_k) + \Delta W(\mathcal{G}_k, s_i)}{n_k + 1}. \end{aligned} \quad (25)$$

Transfers from \mathcal{G}_i to \mathcal{G}_k are more useful when $L_a(\mathcal{G}_i) \gg L_a(\mathcal{G}_k)$. A transfer might still be beneficial if $\max\{L_a(\mathcal{G}_i), L_a(\mathcal{G}_k)\}$ is reduced. However, when $L_a(\mathcal{G}_i)$ and $L_a(\mathcal{G}_k)$ are nearly equal, then a transfer might not reduce $\max\{L_a(\mathcal{G}_i), L_a(\mathcal{G}_k)\}$. In this case, swapping two states between the sub-graphs might reduce $\max\{L_a(\mathcal{G}_i), L_a(\mathcal{G}_k)\}$. Let $s_i \in P_i$ be moved from \mathcal{G}_i to \mathcal{G}_k , and $s_k \in P_k$ be moved from \mathcal{G}_k to \mathcal{G}_i simultaneously. Then, the sum of the edge weights of the new sub-graphs \mathcal{G}'_i and \mathcal{G}'_k are

$$\begin{aligned} W(\mathcal{G}'_i) &= W(\mathcal{G}_i) - \Delta W(\mathcal{G}_i, s_i) + \Delta W(\mathcal{G}_i, s_k) - \\ &\quad w(s_i, s_k) - w(s_k, s_i), \\ W(\mathcal{G}'_k) &= W(\mathcal{G}_k) - \Delta W(\mathcal{G}_k, s_k) + \Delta W(\mathcal{G}_k, s_i) - \\ &\quad w(s_i, s_k) - w(s_k, s_i). \end{aligned} \quad (26)$$

From (26), $L_a(\mathcal{G}'_i)$ and $L_a(\mathcal{G}'_k)$ can be computed using (23). Since there are n_i potential transfers from \mathcal{G}_i , and $n_i n_k$ potential swaps between \mathcal{G}_i and \mathcal{G}_k , the best transfer and swap can be computed in $O(|\mathcal{V}|^2)$.

We present our heuristic partitioning algorithm in Algorithm 2. In practice, we use a heuristic procedure [27] to generate the initial partition \mathcal{P} which will be described in Section VII when presenting the numerical results. A substantial difference of our procedure from the algorithm

presented in [16] is the computation of model graph $\mathcal{G}(S, E)$. Also, [16] does either a swap or a transfer at every iteration, but Algorithm 2 could possibly do both swaps and transfers at each iteration. Since there are only finite number of swaps and transfers possible, and we recheck a partition \mathcal{P} only if $M_a(\mathcal{P})$ has strictly reduced in the previous iteration, Algorithm 2 terminates in finite time, generating a local optimum for the problem in (20).

We will prove that Algorithm 2 will generate optimal partitions \mathcal{P}^* in Problem 2 for MDPs with *clustered* target states as defined below. The clusters are defined based on the hitting time between states. We denote w_c to be the *worst case hitting time* for pairs of target states within the same cluster, and w_l to be the *best case hitting time* for pairs of target states in different clusters.

Definition 5. Given an MDP $\mathcal{M}(S, A, \mathcal{T})$, we define a set of target states \mathcal{V} to be clustered if \mathcal{V} can be partitioned to m clusters as $\mathcal{R}(w_c, w_l, \mathcal{V}) = \{R_1, R_2, \dots, R_m\}$ where each cluster $R_i = \{s_i^1, s_i^2, \dots, s_i^n\}$ has n targets, $\cup_{i=1}^m R_i = \mathcal{V} \subseteq S$, and $R_i \cap R_j = \emptyset$ for all $i, j \in \{1, 2, \dots, m\}, i \neq j$. Additionally,

$$H_{s_1}^{s_2, \pi_{s_1, s_2}^*} \in [w'_c, w_c] \text{ for all } s_1, s_2 \in R_i, s_1 \neq s_2, \quad (27)$$

$$i \in \{1, 2, \dots, m\},$$

$$H_{s_1}^{s_2, \pi_{s_1, s_2}^*} \in [w_l, w'_l] \text{ for all } s_1 \in R_i, s_2 \in R_j, \quad (28)$$

$$i, j \in \{1, 2, \dots, m\}, i \neq j,$$

where policy $\pi_{s_1, s_2}^* : S \times S \rightarrow A$ is the optimal policy that solves Problem 1 with $\mathcal{V} = \{s_2\}$ and $s_0 = s_1$, and $w_c \leq w_l$. We also define w_1 and w_2 so that

$$H_{s_0}^{s, \pi_{s_0, s}^*} \in [w_1, w_2] \text{ for all } s \in \mathcal{V}. \quad (29)$$

On graph $\mathcal{G}(S, E)$, since the weights are defined as $w(s_1, s_2) = \min_{\pi} \mathbb{E} [H_{s_1}^{s_2, \pi}] = \mathbb{E} [H_{s_1}^{s_2, \pi_{s_1, s_2}^*}]$ for all $s_1, s_2 \in S$, the bounds in (27), (28) and (29) also apply to the weights $w(s_1, s_2)$. From (27) and (28), the bound for weights of states within the same cluster is

$$w(s_1, s_2) \in [w'_c, w_c] \text{ for all } s_1, s_2 \in R_i, s_1 \neq s_2, \quad (30)$$

$$i \in \{1, 2, \dots, m\},$$

and the bound for weights of states in different clusters is

$$w(s_1, s_2) \in [w_l, w'_l], \text{ for all } s_1 \in R_i, s_2 \in R_j, \quad (31)$$

$$i, j \in \{1, 2, \dots, m\}, i \neq j.$$


In Fig. 2, we depict a deterministic MDP (graph) with $m = 3$ clusters and $n = 4$ targets in each cluster which satisfies the bounds in Definition 5 where $w_c = 2, w_l = 26, w_1 = 13, w_2 = 15$. Such large difference between w_c and w_l will be more clear from the below theorem. Having defined an MDP with clustered target states, we expect the partition \mathcal{R} to be optimal for Problem 2 when w_l is sufficiently larger than w_c . Indeed, we prove this claim in the following theorem.

Algorithm 2: Partitioning by transfers and swaps

```

1 Input:  $S, A, P, \mathcal{V}, m, s_0$ ;
2 Compute  $\mathcal{G}(S, E)$  using policy iteration;
3 Initialize: A partition  $\mathcal{P}$ ; all subgraphs are unchecked
   for swap and for transfer;
4 repeat
5   forall pairs of unchecked subgraphs  $(\mathcal{G}_i, \mathcal{G}_k)$  do
6      $S_{min} = \max(L_a(\mathcal{G}_i), L_a(\mathcal{G}_k));$ 
7     Compute  $L_a(\mathcal{G}'_i)$  and  $L_a(\mathcal{G}'_k)$  using (26);
8     if  $\max(L_a(\mathcal{G}'_i), L_a(\mathcal{G}'_k)) < S_{min}$  then
9       | Swap the states  $(s_i, s_k)$  between  $(\mathcal{G}_i, \mathcal{G}_k)$ ;
10    else
11      |  $(i, k) \leftarrow$  checked for swap;
12    end
13    Compute  $L_a(\mathcal{G}'_i)$  and  $L_a(\mathcal{G}'_k)$  using (25);
14    if  $\max(L_a(\mathcal{G}'_i), L_a(\mathcal{G}'_k)) < S_{min}$  then
15      | Transfer the state  $s_i$  from  $\mathcal{G}_i$  to  $\mathcal{G}_k$ ;
16    else
17      |  $(i, k) \leftarrow$  checked for transfer;
18    end
19    Update the partition  $\mathcal{P}$ ;
20  end
21 until  $M_a(\mathcal{P})$  has not reduced;
22 Output: Partition  $\mathcal{P}$ 

```



all $i, j \in \{1, 2, \dots, m\}$. The set of targets R'_i and R'_j each contain a state that belongs to cluster j and i , respectively. All partitions other than \mathcal{R} can be represented by assigning more pairs of states to the wrong clusters. We denote $\pi_{s_0, \mathcal{V}}^* : S \times 2^{\mathcal{V}} \rightarrow A$ to be the optimal policy that solves Problem 1 with target set $\mathcal{V} \subseteq S$ and initial state $s_0 \in S$. From (27) and (28), the bounds on cover times are

$$C_{s_0}^{R_j, \pi_{s_0, R_j}^*} \leq w_2 + (n-1)w_c \text{ for all } j \in \{1, 2, \dots, m\}, \quad (33)$$

$$C_{s_0}^{R'_i, \pi_{s_0, R'_i}^*} \geq w_1 + (n-2)w'_c + w_l > w_1 + w_l, \quad (34)$$

$$C_{s_0}^{R'_j, \pi_{s_0, R'_j}^*} \geq w_1 + (n-2)w'_c + w_l > w_1 + w_l.$$

Since hitting time and cover time are random variables, the above bounds also hold in expectation. For \mathcal{R} to be a partition with expected cover time less than that of \mathcal{R}' , the following should hold. From (33) and (34), if

$$\begin{aligned} w_1 + w_l &\geq w_2 + (n-1)w_c, \\ \Rightarrow w_l &\geq (n-1)w_c + (w_2 - w_1), \end{aligned} \quad (35)$$

then the expected cover time of partition \mathcal{R} is less than that of \mathcal{R}' . There may exist two pairs of states in two wrong clusters such that the partition is $\mathcal{R}'' = \{R_1, R_2, \dots, R'_i, R''_j, \dots, R_m\}$ where $R''_i = \{s_i^1, s_i^2, \dots, s_i^k, s_i^l, \dots, s_i^n\}$ and $R''_j = \{s_j^1, s_j^2, \dots, s_j^k, s_j^l, \dots, s_j^n\}$. Then,

$$C_{s_0}^{R''_i, \pi_{s_0, R''_i}^*} \geq w_1 + (n-3)w'_c + 2w_l > w_1 + w_l, \quad (36)$$

$$C_{s_0}^{R''_j, \pi_{s_0, R''_j}^*} \geq w_1 + (n-3)w'_c + 2w_l > w_1 + w_l.$$

Similarly, from (36), the condition in (35) should hold for the expected cover time of \mathcal{R} to be less than the expected cover time of any partition with atleast a pair of states in the wrong cluster. Since all partitions other than \mathcal{R} can be represented by assigning more pairs of states to the wrong clusters, \mathcal{R} is the optimal partition for Problem 2 if (35) is satisfied. \square

Having proved that the clustered partition \mathcal{R} is the optimal partition, we prove that Algorithm 2 generates \mathcal{R} under certain conditions.

Theorem 3. *Under Assumption 1, given an initial partition \mathcal{P} of the target states \mathcal{V} where $|P_i| = n$ for all $i, j \in \{1, 2, \dots, m\}$, if there exists a clustered partition \mathcal{R} as in Definition 5 and*

$$w_l > 3nw_c + w'_c + \frac{w_2 - w_1}{2}, \quad (37)$$

then Algorithm 2 will generate the partition \mathcal{R} from an initial partition \mathcal{P} .

Proof. We are operating on the complete graph $\mathcal{G} = (S, E)$ computed in line 2 of Algorithm 2. Let the set of targets associated with agent i be

$$P_i = \cup_{j=1}^m Q_i^j \text{ where } Q_i^j \subseteq R_j, \cup_{i=1}^m Q_i^j = R_j, \sum_{j=1}^m |Q_i^j| = n, \quad (38)$$

$$Q_i^j \cap Q_k^l = \emptyset, \text{ for all } i, j, k, l \in \{1, 2, \dots, m\}, i \neq j \neq k \neq l.$$

The set Q_i^j is the set of target states which belong to cluster j and are associated with agent i . Consider a swap operation between two arbitrary sets $P_i, P_k \in \mathcal{P}$ and two target states $s_i^j \in Q_i^j, s_k^l \in Q_k^l$. Let the sets after the swap operation be P'_i, P'_k and the change in the sum of weights be given by $\Delta W(\mathcal{G}_i) = W(\mathcal{G}'_i) - W(\mathcal{G}_i)$. Therefore, from (26),

$$\Delta W(\mathcal{G}_i) = -\Delta W(\mathcal{G}_i, s_i^j) + \Delta W(\mathcal{G}_i, s_k^l) - w(s_i^j, s_k^l) - w(s_k^l, s_i^j), \quad (39)$$

$$\Delta W(\mathcal{G}_k) = -\Delta W(\mathcal{G}_k, s_k^l) + \Delta W(\mathcal{G}_k, s_i^j) - w(s_i^j, s_k^l) - w(s_k^l, s_i^j).$$

The contribution of a state s_i^j to a sub-graph \mathcal{G}_i can be written from (22) as

$$\begin{aligned} \Delta W(\mathcal{G}_i, s_i^j) &= \sum_{s_1 \in Q_i^j} w(s_1, s_i^j) + \sum_{s_1 \in Q_i^j} w(s_i^j, s_1) + \\ &\sum_{s_1 \in P_i \setminus Q_i^j} w(s_1, s_i^j) + \sum_{s_1 \in P_i \setminus Q_i^j} w(s_i^j, s_1) + w(s_0, s_i^j). \end{aligned} \quad (40)$$

For ease of writing, let us denote $|Q_i^j| = n_i^j$ for all $i, j \in \{1, 2, \dots, m\}$. Thus, from (29), (40) and (31),

$$\Delta W(\mathcal{G}_i, s_i^j) \geq 2(w'_c(n_i^j - 1) + w_l(n - n_i^j)) + w_1. \quad (41)$$

Similarly, from (29), (30) and (40),

$$\Delta W(\mathcal{G}_i, s_k^l) \leq 2(w_c(n_i^l) + w'_l(n - n_i^l)) + w_2. \quad (42)$$

Thus, from (29), (30), (31) and (40), we can write the bounds for sub-graph \mathcal{G}_k as

$$\begin{aligned} \Delta W(\mathcal{G}_k, s_k^l) &\geq 2(w'_c(n_k^l - 1) + w_l(n - n_k^l)) + w_1, \\ \Delta W(\mathcal{G}_k, s_i^j) &\leq 2(w_c(n_k^j) + w'_l(n - n_k^j)) + w_2. \end{aligned} \quad (43)$$

From (39), (41), (42) and (43),

$$\begin{aligned} \frac{\Delta W(\mathcal{G}_i)}{2} &\leq (w'_l - w_l)n - n_i^l(w'_l - w_c) + n_i^j(w_l - w'_c) + \\ &w'_c - w_l + \frac{w_2 - w_1}{2}, \\ \frac{\Delta W(\mathcal{G}_k)}{2} &\leq (w'_l - w_l)n - n_k^j(w'_l - w_c) + n_k^l(w_l - w'_c) + \\ &w'_c - w_l + \frac{w_2 - w_1}{2}. \end{aligned} \quad (44)$$

Using Fig. 2 as an illustration, if we fix w_l to be the smallest weight between states of different clusters, and w_c to be the largest weight between states of same clusters, then the bound for w'_l is

$$w'_l \leq w_l + 2w_c. \quad (45)$$

From (44) and (45),

$$\begin{aligned} \frac{\Delta W(\mathcal{G}_i)}{2} &\leq 2nw_c - n_i^l(w'_l - w_c) + n_i^j(w_l - w'_c) + \\ &w'_c - w_l + \frac{w_2 - w_1}{2}. \end{aligned} \quad (46)$$

Since $-w'_l \leq -w_l$, and $w_l - w'_c \leq w_l$, (46) can be written as

$$\frac{\Delta W(\mathcal{G}_i)}{2} \leq 2nw_c - n_i^l(w_l - w_c) + n_i^j w_l + w'_c - w_l + \frac{w_2 - w_1}{2}. \quad (47)$$

Adding and subtracting $n_i^j w_c$ to (47),

$$\frac{\Delta W(\mathcal{G}_i)}{2} \leq 2nw_c - n_i^l(w_l - w_c) + n_i^j(w_l - w_c) + n_i^j w_c + w'_c - w_l + \frac{w_2 - w_1}{2}. \quad (48)$$

Since $n_i^j \leq n$, (48) can be written as

$$\frac{\Delta W(\mathcal{G}_i)}{2} \leq 3nw_c - w_l + (n_i^j - n_i^l)(w_l - w_c) + w'_c + \frac{w_2 - w_1}{2}. \quad (49)$$

Similarly, using the same argument as (46), (47), and (48),

$$\frac{\Delta W(\mathcal{G}_k)}{2} \leq 3nw_c - w_l + (n_k^l - n_k^j)(w_l - w_c) + w'_c + \frac{w_2 - w_1}{2}. \quad (50)$$

From (49) and (50),

$$n_i^j \leq n_i^l \text{ and } n_k^l \leq n_k^j \text{ and } w_l > 3nw_c + w'_c + \frac{w_2 - w_1}{2} \Rightarrow \Delta W(\mathcal{G}_i) < 0 \text{ and } \Delta W(\mathcal{G}_k) < 0. \quad (51)$$

From (38), since $\sum_{j=1}^m n_i^j = \sum_{i=1}^m n_i^j = n$ for all $i, j \in \{1, 2, \dots, m\}$, there always exists some (i, j, k, l) such that $n_i^j \leq n_i^l$ and $n_k^l \leq n_k^j$. If $w_l > 3nw_c + w'_c + \frac{w_2 - w_1}{2}$, then (51) is satisfied, and from Algorithm 2 and (26), the states s_i^j and s_k^l are swapped between \mathcal{G}_i and \mathcal{G}_k . After the swap, n_i^j and n_k^l are decremented by one, whereas n_i^l and n_k^j are incremented by one. Therefore, from (49) and (50), $\Delta W(\mathcal{G}'_i) < \Delta W(\mathcal{G}_i)$ and $\Delta W(\mathcal{G}'_k) < \Delta W(\mathcal{G}_k)$ where \mathcal{G}'_i and \mathcal{G}'_k are the sub-graphs after the swap. Hence, the states are swapped until $n_i^j = n_k^l = 0$ or $n_i^l = n_k^j = n$. If $n_i^l = n$, then all the states of cluster l are associated with agent i . If $n_i^j = 0$, then all the states $s_i^j \in Q_i^j$ which belong to cluster j and were initially associated with agent i are now with agent k . The states are swapped for all $i, j, k, l \in \{1, 2, \dots, m\}$ until the boundary cases of $n_i^j = n_k^l = 0$ or $n_i^l = n_k^j = n$ are reached. Since $\sum_{j=1}^m n_i^j = \sum_{i=1}^m n_i^j = n$ for all $i, j \in \{1, 2, \dots, m\}$, all the states of each cluster j are associated with only one agent i , thus resulting in the partition \mathcal{R} . An agent cannot have two clusters because the initial partition has n targets associated with each agent which is the same as the number of targets in each cluster. Hence, only swap operations happen between agents because a transfer operation increases $W(\mathcal{G}_i)$ for some agent i . \square

Therefore, Algorithm 2 generates optimal partitions $\mathcal{P}^* = \mathcal{R}$ for Problem 2 if the bounds on the hitting time of the states in the MDP independently satisfy the conditions in Theorem 2 and Theorem 3. One such MDP was illustrated in Fig. 2.

We assumed in Theorem 3 that the number of targets assigned to each agent in the initial partition are all equal. If we do not have any prior knowledge about the MDP dynamics and location of targets, one naive way to partition would indeed be to assign equal number of targets to each agent. We expect a partition with equal sized target assignments to be optimal

for Problem 2 if the optimal hitting times and cover times are independent and uniformly distributed. However, this question formally remains open for future work.

We also assume in our work that the clusters defined in Definition 5 have equal number of targets. We believe the ideas presented in our proof of Theorem 3 can be extended to clusters with different number of targets. We could bound the change in sum of weights $\Delta W(\mathcal{G}_i), \Delta W(\mathcal{G}_k)$ for every pair of sub-graphs $(\mathcal{G}_i, \mathcal{G}_k)$ in terms of the number of targets n_j in cluster j where $i, j, k \in \{1, 2, \dots, m\}$. We could then derive conditions for a swap operation between the sub-graphs similar to (51). We reserve this more detailed analysis for future work.

Though Theorem 2 and Theorem 3 apply only to MDPs with clustered targets, we show in the subsequent section with numerical experiments that Algorithm 2 generates optimal partitions \mathcal{P}^* which solves Problem 2 for some MDPs that not necessarily have clustered targets. We also validate the algorithms for single agent described in Section IV and Section V using numerical experiments.

VII. NUMERICAL RESULTS

In this section, we validate our algorithms on different environments for both single-agent and multi-agent scenarios. Since previous work was on heuristics for single agent planning on graphs [24], we compare the performance and results of our single agent planning algorithm on random graphs. Our work is primarily motivated by stochastic dynamics and hence we present the numerical results of our algorithms on random MDPs. We also validate our algorithm on more realistic gridworld environment motivated by ocean dynamics.

A. Random graphs

In this section, we implement our heuristic procedure on graphs for single agent path planning and partitioning of targets to multiple agents.

1) *Single agent*: In this section, we present our results for the single agent case. The paths produced by Algorithm 1 to cover all states of some deterministic MDPs are given in Fig. 3. We compare the performance of Algorithm 1 and the nearest neighbor heuristic [24]. We use policy iteration to compute the optimal cover time, but we could use any algorithm that solves the SSP problem to optimality on the product MDP [8]. We compute the average value of the cover time by performing 1000 runs of the nearest neighbor heuristic. However, we do not need to run policy iteration and Algorithm 1 1000 times to get the cover time; instead, just running it once produces the cover time of the deterministic policy. We use $\varepsilon = 10^{-20}$ and $\gamma = 0.01$ for Algorithm 1. We choose such an ε close to zero so that the computation of value function terminates after some reasonable number of iterations, while at the same time the computed approximate value function is close to the actual value function. If $\varepsilon = 0$, the iterative procedure might not stop in reasonable time because of floating point errors, although the value function has converged. We choose $\gamma = 0.01$ in Algorithm 1 so that the algorithm incorporates some look-ahead on how good the states are in future time. If $\gamma = 0$, there is no look-ahead as described in Section V.




Fig. 3: Paths produced by Algorithm 1 to cover all states of some graphs and illustration of Algorithm 2 with $m = 2$ agents. The initial state is in red, and the states assigned to each agent are shown in green and blue, respectively.

Hence, Algorithm 1 would perform exactly as the nearest neighbor heuristic. As we increase γ , the rate of convergence to the optimal value function at each time step decreases considerably without much gain in the overall optimality of the cover time for the MDPs presented in this section.

We choose the graphs in Fig. 3 to be significantly different than those in Fig. 1 to show that Algorithm 1 is optimal for a larger class of graphs like those in Figs. 3b and 3c. The runtime of Algorithm 1 for the graphs in Fig. 3 is 100-1000 times faster than the optimal algorithm, with just a 5% suboptimality on average. Though the nearest neighbour heuristic is 10 times faster than our heuristic, the average suboptimality is approximately 72%.

To illustrate our heuristic on more general graphs, we simulate 1000 runs of the nearest neighbor heuristic and Algorithm 1, with $\varepsilon = 10^{-20}$, $\gamma = 0.01$, for random connected graphs. We present the cover time values and the average runtime for 1000 system runs of Algorithm 1 and nearest neighbor heuristic in Table I. The policy iteration procedure consumes an immense amount of runtime in MATLAB to compute the optimal cover time for larger than 11 target states, which we denote as ‘timeout’ in Table I. However, Algorithm 1 and the nearest neighbor heuristic run much quicker than the optimal algorithm. The nearest neighbor heuristic naturally has better runtime than Algorithm 1, however, it yields significantly worse cover times. On the other hand, the cover times of Algorithm 1 are always optimal or almost-optimal, whereas the average runtime is still on the order of 100-1000 times faster than the optimal algorithm. The average variance of the cover time for nearest neighbour heuristic is 54.84 for the graphs in Table I. We performed all the numerical experiments using MATLAB R2020a on a computer with an Intel Core i7 2.6 GHz processor and 16GB RAM.

2) Multiple agents: In this section, we present our results of the partitioning algorithm for the multi-agent case. We use the greedy vertex m -center algorithm [27] to generate the initial partition \mathcal{P} in line 3 of Algorithm 2 for m agents.

In Fig. 3, we present the partitions generated by Algorithm 2 for the same graphs discussed in Section VII-A.1 with $m = 2$ agents. In Fig. 4, we illustrate Algorithm 2 for random graphs

TABLE I: Cover time comparison on random graphs for a single agent.

Graph		Optimal algorithm (policy iteration)		Algorithm 1		Nearest neighbor heuristic	
$ S $	$ \mathcal{V} $	Optimal cover time	Runtime (sec)	Cover time	Runtime (sec)	Average cover time	Average runtime (sec)
91	10	10	547.948	11	3.389	12.878	0.005
50	10	10	192.798	10	0.243	12.489	0.003
62	8	9	27.946	9	0.946	13.403	0.003
97	10	11	726.563	12	0.616	15.977	0.004
200	10	11	4112.927	11	6.515	13.835	0.014
150	8	11	263.947	11	2.537	13.389	0.011
130	9	9	357.842	10	1.771	11.196	0.007
170	10	10	2897.701	10	3.709	13.483	0.013
250	9	10	2086.298	11	11.993	12.797	0.019
200	11	11	41473.81	11	6.94	13.484	0.013
1000	9	12	9167.786	13	15.894	26.087	0.09
180	12	N/A	timeout	12	7.427	15.988	0.017
200	15	N/A	timeout	16	9.271	21.232	0.02
500	50	N/A	timeout	55	35.153	101.084	0.051
1000	100	N/A	timeout	113	78.472	212.96	0.101
500	80	N/A	timeout	89	70.285	185.17	0.06

with $m = 3$ agents and $\mathcal{V} = S$. In Table II, we compare the cover time $\max_i C_{s_0}^{P_i, \pi_i^*}$ between the optimal partition \mathcal{P}^* for Problem 2 and the partition \mathcal{P} generated by Algorithm 2 for the graph scenarios in Fig. 3 and Fig. 4. We use a naive method to obtain the optimal partition \mathcal{P}^* by brute-force search of the optimal cover time of all possible subsets of \mathcal{V} . The optimal cover time can be computed using policy iteration or any other algorithm that solves the SSP problem to optimality on a product MDP [8]. From Table II, the partition generated by Algorithm 2 is optimal for graphs in Figs. 3b, 3c, 4b and 4d. The average runtime of brute-force search is 61.292 seconds, whereas Algorithm 2 consumes 0.01 seconds on average for the graphs in Table II.

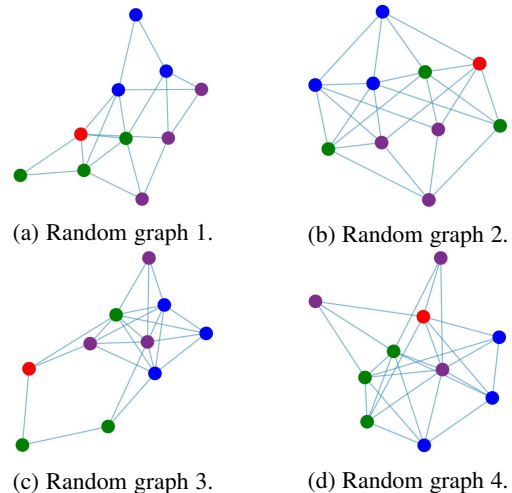


Fig. 4: Illustration of Algorithm 2 on random graph scenarios with $|\mathcal{V}| = 10$ for $m = 3$ agents. The initial state is in red, and the states assigned to each agent are shown in green, blue and violet, respectively.

TABLE II: Cover time comparison for multiple agents for the graphs in Fig. 3 and Fig. 4.

Graph	Optimal partition \mathcal{P}^*		Partition \mathcal{P} by Algorithm 2	
	Cover time $\max_i C_{s_0}^{P_i, \pi_i^*}$	Runtime (sec)	Cover time $\max_i C_{s_0}^{P_i, \pi_i^*}$	Runtime (sec)
Fig. 3a	3	31.21	4	0.003
Fig. 3b	3	32.13	3	0.003
Fig. 3c	5	40.213	5	0.005
Fig. 4a	3	81.213	4	0.019
Fig. 4b	3	80.15	3	0.013
Fig. 4c	3	79.91	4	0.009
Fig. 4d	3	84.217	3	0.021

B. Random MDPs

In this section, we implement our algorithms on MDPs for single agent path planning and target assignment to multiple agents.

1) *Single agent:* In this section, we validate Algorithm 1 for the single-agent case with $\varepsilon = 10^{-20}$ and $\gamma = 0.01$, on general MDPs. We test the algorithms on MDPs where the number of states and targets are selected at random, and for each state $s \in S$ and action $a \in A$, the transition probabilities $\mathcal{T}(s, a, s')$ form a discrete uniform distribution with all $s' \in S$ as the support. We select the initial state and the set to be covered at random. We present the optimal expected cover time computed using policy iteration, and the average cover time and average runtime for 1000 system runs of Algorithm 1 in Table III. The average variance of the cover time obtained by Algorithm 1 for the MDPs in Table III is 10.48. The average suboptimality of Algorithm 1 for the MDPs in Table III is 19.658%. This suboptimality is offset by a considerable reduction in runtime. The average runtime of the MDPs in Table III for the optimal algorithm using policy iteration is 13404.225 sec, and Algorithm 1 is 22.262 sec.

2) *Multiple agents:* In this section, to illustrate the heuristic partitioning procedure on more general MDPs, we implement Algorithm 2 on the random MDPs presented in Table III. In Table IV, we present the optimal expected cover time computed using the naive method described in Section VII-

TABLE III: Cover time comparison on random MDPs for a single agent.

MDP		Optimal algorithm (policy iteration)		Algorithm 1		
$ S $	$ V $	Expected cover time	Runtime (sec)	Average cover time	Variance of cover time	Average runtime (sec)
50	10	36.732	1573.755	38.98	5.949	0.141
70	8	45.394	395.643	49.253	7.801	0.19
80	10	37.603	4166.8411	42.91	5.642	0.193
100	9	51.517	1850.285	60.878	11.133	0.484
100	10	52.7734	6400.20	57.847	8.397	0.514
200	10	68.355	59338.536	86.148	14.679	5.266
200	9	56.348	3862.91	74.056	6.173	2.661
150	8	59.714	741.04	72.89	10.234	0.17
170	10	73.287	43615.027	93.675	18.311	1.035
120	10	57.024	14654.176	73.824	12.493	0.772
1000	9	90.067	10081.21	111.59	12.64	5.823
500	10	N/A	timeout	110.659	14.756	13.96
500	50	N/A	timeout	859.12	19.615	47.965
1000	12	N/A	timeout	643.974	18.531	31.039

A.2, and the expected cover time of the partitions generated by Algorithm 2 for random MDPs with $m = 3$ agents. We implement Algorithm 1 for each agent using the partition generated by Algorithm 2 as well as the optimal partition. We present the average cover time and average runtime for 1000 system runs of Algorithm 1 in Table IV. The cover time for one run of Algorithm 1 for multiple agents is computed as $\max_i(C_{s_0}^{P_i})$ where P_i are the targets assigned to agent i and $C_{s_0}^{P_i}$ is the cover time of agent i .

We present the results only for 3 agents because of the saturation of cover time with a larger number of agents for the number of targets in Table IV. In Fig. 5, we plot the ratio of the average cover time for the multi-agent case to the single-agent case for three MDP scenarios, when the number of agents are increased. The average cover time is computed for 1000 system runs of Algorithm 1 on the partition \mathcal{P} generated by Algorithm 2. The average runtime of Algorithm 2 marginally increases with an increase in the number of agents m , but it also eventually saturates. We keep the full analysis of the effect of number of agents on the cover time and performance of Algorithm 2 for future work.

From Table IV, the expected cover time of the optimal single-agent algorithm on the partition \mathcal{P} generated by Algorithm 2 is either optimal or almost-optimal when compared to the optimal partition \mathcal{P}^* . The average cover time is similar when Algorithm 1 is implemented on the partition generated by Algorithm 2 and the optimal partition. Hence, the partition generated by Algorithm 2 is the same or similar to the optimal partition. Since Algorithm 1 is implemented independently for each agent in the multi-agent case, there is a slight increase of suboptimality in the cover time of Algorithm 1 when compared to the single agent case. Though the average cover time when implementing Algorithm 1 is more than one half of the optimal cover time for $m = 3$ agents, the runtime of our heuristic partitioning and path planning procedure is 10^5 times faster than the optimal partitioning and optimal path planning method.

C. Ocean dynamics

In this section, we implement our heuristic procedures for single agent path planning, and assignment of multiple-targets to multiple agents on stochastic gridworlds modeled by realistic ocean currents [1].




Fig. 5: Effect of saturation of average cover time with increasing number of agents when Algorithm 1 is implemented on the partition \mathcal{P} by Algorithm 2.

1) *Single agent*: In this section, we simulate Algorithm 1 with transition dynamics motivated by an autonomous underwater vehicle operating in an oceanic environment [1], modeled as a 20×20 grid. The vehicle's transitions are to move one step north, west, south, or east of its current state, except at the edges of the environment. The set to be covered constitutes regions associated with formation and evolution of algal blooms [1]. The uncertainty and high variability of ocean currents in space and time demand the use of stochastic transition dynamics. We use the MDP environment in [28] to generate realistic transition probabilities from a Gaussian distribution of ocean models [1]. We illustrate Problem 1 on the resulting gridworld using Algorithm 1 with $\varepsilon = 10^{-20}, \gamma = 0.4$ in Fig. 6. We increase the value of γ compared to previous experiments because there is a considerable decrease in suboptimality for the MDPs considered in this section, even though there is a slight decrease in rate of convergence to the optimal value function at each time step. The cover time for the path presented in Fig. 6 is 88. The two “blips” in the final stretch of the path before completing the mission are a consequence of stochastic dynamics. The average cover time for 1000 runs of Algorithm 1 for the problem in Fig. 6 is 84.87 with a variance of 21.267, and the average runtime is 9.0119 sec. The optimal expected cover time for the scenario in Fig. 6 is 71.8534 which is computed using policy iteration with a runtime of 5760 sec.

2) *Multiple agents*: In this section, we implement Algorithm 2 to assign targets to agents motivated by a team of autonomous underwater vehicles operating in an ocean [1], with the same environment and agents used in Section VII-C.1. We illustrate Problem 2 on the resulting gridworlds using Algorithm 2 in Fig. 7 and Fig. 8.

In Fig. 7, we depict the assignment of targets to agents on a 10×10 stochastic gridworld using Algorithm 2. The scenario in Fig. 7a has clustered target states, and the initial state of the agents is approximately equidistant from all the clusters. Therefore, Algorithm 2 has naturally generated the clustered partition which is also verified to be the optimal partition. The optimal partition is obtained by doing a brute-force search of all possible optimal expected cover times computed using policy iteration. In Fig. 7b, the initial state and the set of target






Fig. 6: A representative path produced by Algorithm 1 to cover 10 states on a 20×20 gridworld with stochastic dynamics. The red dot is the initial state, green squares are the targets, and the path of the agent is denoted by brown lines.

states are selected at random. Algorithm 2 has again generated the optimal partition for the scenario in Fig. 7b where the optimality is verified by brute-force search.

In Fig. 8, we illustrate the assignment of targets to agents on 20×20 stochastic gridworlds using Algorithm 2 and also implement Algorithm 1 with $\varepsilon = 10^{-20}, \gamma = 0.7$ to synthesize each agent's policy to visit the targets assigned to them. In Fig. 8a, we depict the same scenario presented in Fig. 6 but with $m = 3$ agents. Algorithm 2 has generated the optimal partition for Fig. 8a. The optimal partition can be verified by a brute force method which requires 1586 sec whereas Algorithm 2 requires 3.742 sec. The cover time for the paths presented in Fig. 8a is 27. In Table IV, we present the expected cover time of the optimal partition and the partition generated by Algorithm 2 for the scenarios in Fig. 7 and Fig. 8. We also implement Algorithm 1 as described in Section VII-B.2 for multiple agents and compare the respective runtime. From Table IV, Algorithm 2 is 10^4 times faster than the naive procedure, but generating the optimal partition for the gridworld scenarios of Fig. 7 and Fig. 8a. We note that the average cover time of Algorithm 1 on the heuristic partition generated by Algorithm 2 is better than the average cover time of Algorithm 1 on the optimal




(a) $m = 4$ agents with $|\mathcal{V}| = 9$ targets.




(b) $m = 3$ agents with $|\mathcal{V}| = 10$ targets.

Fig. 7: Illustration of Algorithm 2 on 10×10 gridworld environments with stochastic transition dynamics. The black dot is the initial state, and the target assignment for each agent are denoted by blue, violet, red, and brown squares.



(a) $m = 3$ agents with $|\mathcal{V}| = 10$ targets.



(b) $m = 4$ agents with $|\mathcal{V}| = 40$ targets.

Fig. 8: A representative path generated by Algorithm 1 after the partitioning by Algorithm 2 on 20×20 gridworld environments with stochastic dynamics. The black dot is the initial state, the path and targets for each agent are denoted by blue, violet, red and brown lines and squares, respectively.

partition for the scenario in Fig. 8a. This could be because the suboptimality of Algorithm 1 on the optimal partition might be worse than the suboptimality of Algorithm 1 on the partition generated by Algorithm 2 for the planning mission in Fig. 8a. Algorithm 1 has an average suboptimality of 24 % for the gridworld environments in Table IV when compared to a 50 % suboptimality for the random MDPs in Table IV. Most of the real-world environments have sparse dynamics [1], [5] like the gridworld environment, and our algorithms have a lower suboptimality for such scenarios when compared to random MDPs. In particular, for the scenario in Fig. 8a, the average cover time using Algorithm 1 for $m = 3$ agents is only 5% greater than the optimal expected cover time. This suboptimality is offset by a considerable reduction in runtime as given in Table IV.

In Fig. 8b, we depict a scenario with $|\mathcal{V}| = 40$ clustered target states and $m = 4$ agents. Since some clusters are farther from the initial state than other clusters, assigning just one agent to the farthest cluster and one other agent to the nearest cluster might not be optimal. Accordingly, the target assignment obtained by Algorithm 2 is not the clustered partition. Two agents (brown and violet) are assigned some targets in the cluster closer to the initial state, along with some targets in clusters that are further away. The cover time results given below validate our argument that the clustered partition might not be optimal for the scenario in Fig. 8b. The cover time for the paths presented in Fig. 8b is 48. The average cover time for 100 runs of Algorithm 1 for the partition in Fig. 8b is 53.52 and the variance is 18.832. The average cover time for 100 runs of Algorithm 1 for the partition with natural clustering is 61.97 and the variance is 17.821. The average runtime for 100 runs of the complete mission to jointly visit the targets by multiple agents is 16.82 sec. On the other hand, it is infeasible to solve Problem 2 by a

brute force procedure. The size of the MDP and target set in Fig. 8b requires 1.0995×10^{12} expected cover time values to be computed at every step of the optimal policy iteration procedure to solve Problem 1. Then, all the 1.0995×10^{12} expected cover time values and the corresponding subsets of the target set should be searched for the optimal partition that solves Problem 2.

VIII. CONCLUSION AND FUTURE WORK

In this paper, we formulated a path planning problem for a team of agents to jointly visit a set of target states in minimum expected time on a Markov decision process. We showed that the decision version of the planning problem on a graph for the single-agent case is the Hamiltonian path problem which is NP-complete. We also showed that the path planning problem with multiple targets is a SSP problem on a product MDP and thus policy iteration can be adopted as the optimal algorithm that is exponential in the number of target states. We utilized policy iteration as a motivation to propose a suboptimal algorithm based on an approximate value function which consumes polynomial number of operations at each time step. We proved that our heuristic algorithm eventually visits the target states with probability 1, and generates optimal solutions for certain classes of deterministic MDPs, i.e., graphs.

For the multi-agent case, we formulated a problem of assigning multiple targets to multiple agents such that the largest optimal expected cover time among the multiple agents is minimal. The average length of a Hamiltonian path on a model graph was used as a heuristic to approximate the optimal cover time. We adopted a heuristic partitioning procedure that attempts to solve the m -TSP on the model graph, which also solves our multi-agent problem on the original MDP. We proved that the heuristic procedure generates optimal

TABLE IV: Cover time comparison on the random MDPs presented in Table III with $m = 3$ agents, and the gridworld environments in Fig. 7 and Fig. 8.

MDP		Optimal single-agent algorithm on optimal partition \mathcal{P}^*		Optimal single-agent algorithm on partition \mathcal{P} by Algorithm 2		Algorithm 1 on optimal partition \mathcal{P}^*		Algorithm 1 on partition \mathcal{P} by Algorithm 2	
$ S $	$ \mathcal{V} $	Expected cover $\max_i \mathbb{E} \pi_i^* \left[C_{s_0}^{P_i} \right]$	Runtime (sec)	Expected cover $\max_i \mathbb{E} \pi_i^* \left[C_{s_0}^{P_i} \right]$	Runtime (sec)	Average cover time	Average runtime (sec)	Average cover time	Average runtime (sec)
50	10	17.77	1760.475	19.293	1573.819	26.749	186.796	27.754	0.069
70	8	23.074	407.727	23.074	395.649	34.622	12.176	35.663	0.083
80	10	17.6	4357.343	17.63	4166.866	26.818	190.694	26.766	0.196
100	9	24.937	1927.184	25.094	1850.321	38.63	77.076	38.906	0.163
100	10	25.453	6672.765	25.453	6400.207	38.575	272.768	39.382	0.183
200	10	32.572	59529.463	34.491	59338.545	48.773	191.459	50.725	0.478
200	9	28.8825	3889.93	29.523	3862.914	41.333	27.497	41.316	0.424
150	8	28.882	745.099	29.523	741.047	46.612	4.273	47.933	0.201
170	10	34.5429	43801.439	35.206	43615.032	54.621	186.73	54.488	0.284
120	10	27.076	14830.499	27.907	14654.199	42.301	176.576	43.091	0.231
1000	9	35.175	10693.451	36.798	10511.28	49.25	51.982	51.02	2.364
500	10	N/A	timeout	N/A	timeout	N/A	timeout	62.953	7.67
500	50	N/A	timeout	N/A	timeout	N/A	timeout	351.067	30.698
1000	12	N/A	timeout	N/A	timeout	N/A	timeout	292.49	20.12
Fig. 7a		19.685	3437	19.685	2027.843	25.601	1410.286	25.68	1.129
Fig. 7b		23.562	4160.2	23.562	2653.843	31.325	1507.505	32.057	1.148
Fig. 8a		29.947	7346	29.947	5762.169	32.007	1589.528	31.63	5.86
Fig. 8b		N/A	timeout	N/A	timeout	N/A	timeout	53.52	16.82

partitions for clustered target states. We validated our heuristic procedures on more general MDPs, as well as gridworld environments motivated by realistic considerations of ocean currents.

In subsequent work we aim to provide theoretical guarantees on the suboptimality of our heuristic algorithms. An interesting direction is to find the relationship between jointly optimizing for the policies of multiple agents and implementing the optimal policy for each agent using the optimal partition. One immediate extension of our work would be to derive conditions of optimality with our heuristic algorithm for clusters with different number of targets. It would be interesting to explore how the initial partition could affect the partition generated by our heuristic procedure. The optimality of equal sized initial partition when the hitting times are assumed to be uniformly distributed is another interesting theoretical question. We think an analysis in the decrease in cover time with increasing number of agents is an interesting direction. We also see a need for developing a unified framework to integrate our partitioning algorithm and the path planning heuristic for complex real-life problems. Such problems could have different features in the optimization objective such as resilience, dynamic travel time and different initial state for each agent.

ACKNOWLEDGMENT

We thank Pranay Thangeda for helping us generate stochastic gridworlds using models of ocean currents. We also thank Anakin Dey for helping us review the writing of the manuscript.

REFERENCES

- [1] W. H. Al-Sabban, L. F. Gonzalez, and R. N. Smith, "Extending persistent monitoring by combining ocean models and Markov decision processes," in *OCEANS 2012 MTS/IEEE Conference*, pp. 1–10, 2012.
- [2] X. Zhang, J. Chen, B. Xin, and Z. Peng, "A memetic algorithm for path planning of curvature-constrained UAVs performing surveillance of multiple ground targets," *Chinese Journal of Aeronautics*, vol. 27, no. 3, pp. 622–633, 2014.
- [3] B. Garau, M. Bonet, A. Alvarez, S. Ruiz, and A. Pascual, "Path planning for autonomous underwater vehicles in realistic oceanic current fields: Application to gliders in the western Mediterranean sea," *Journal of Maritime Research*, vol. 6, no. 2, pp. 5–22, 2009.
- [4] B. Hockman and M. Pavone, *Stochastic Motion Planning for Hopping Rovers on Small Solar System Bodies*. Springer, 2020.
- [5] S. Lim, H. Balakrishnan, D. Gifford, S. Madden, and D. Rus, "Stochastic motion planning and applications to traffic," *International Journal of Robotics Research*, vol. 30, no. 6, pp. 699–712, 2011.
- [6] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 2014.
- [7] D. P. Bertsekas, *Dynamic Programming and Optimal Control, Vol. II*. Athena Scientific, 3rd ed., 2007.
- [8] D. P. Bertsekas and J. N. Tsitsiklis, "An analysis of stochastic shortest path problems," *Mathematics of Operations Research*, vol. 16, no. 3, pp. 580–595, 1991.
- [9] C. Baier and J.-P. Katoen, *Principles of model checking*. MIT Press, 2008.
- [10] D. P. Mehendale, "Hamiltonian graphs and the traveling salesman problem," *arXiv preprint arXiv:0704.3672v10*, 2016.
- [11] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. MIT Press, 2018.
- [12] A. Larach, C. Daoui, and M. Baslam, "A Markov decision model for area coverage in autonomous demining robot," *International Journal of Informatics and Communication Technology*, vol. 6, pp. 105–116, 2017.
- [13] B. P. Gerkey and M. J. Matarić, "A formal analysis and taxonomy of task allocation in multi-robot systems," *International Journal of Robotics Research*, vol. 23, no. 9, pp. 939–954, 2004.
- [14] M. Yousefikhoshbakht, F. Didehvar, and F. Rahmati, "Modification of the ant colony optimization for solving the multiple traveling salesman problem," *Romanian Journal of Information Science and Technology*, vol. 16, no. 1, pp. 65–80, 2013.
- [15] Z. Ren, S. Rathinam, and H. Choset, "Conflict-based steiner search for multi-agent combinatorial path finding," *Proceedings of Robotics: Science and Systems, New York City, NY, USA*, 2022.
- [16] I. Vandermeulen, R. Groß, and A. Kolling, "Balanced task allocation by partitioning the multiple traveling salesperson problem," in *2019 International Conference on Autonomous Agents and Multiagent Systems*, pp. 1479–1487, 2019.
- [17] Y. Yan and Y. Mostofi, "Efficient clustering and path planning strategies for robotic data collection using space-filling curves," *IEEE Transactions on Control of Network Systems*, vol. 4, no. 4, pp. 838–849, 2017.
- [18] S. Aggarwal and N. Kumar, "Path planning techniques for unmanned aerial vehicles: A review, solutions, and challenges," *Computer Communications*, vol. 149, pp. 270–299, 2020.
- [19] P. Winkler and D. Zuckerman, "Multiple cover time," *Random Structures & Algorithms*, vol. 9, no. 4, pp. 403–411, 1996.
- [20] M. Mohri, A. Rostamizadeh, and A. Talwalkar, *Foundations of Machine Learning*. MIT Press, 2018.
- [21] D. A. Levin and Y. Peres, *Markov Chains and Mixing Times*. American Mathematical Society, 2017.
- [22] X. Wu and X. Zhou, "Stochastic scheduling to minimize expected maximum lateness," *European Journal of Operational Research*, vol. 190, no. 1, pp. 103–115, 2008.
- [23] M. Guillot and G. Stauffer, "The stochastic shortest path problem: a polyhedral combinatorics perspective," *European Journal of Operational Research*, vol. 285, no. 1, pp. 148–158, 2020.
- [24] S. Dhakal and R. Chiong, "A hybrid nearest neighbour and progressive improvement approach for travelling salesman problem," in *2008 International Symposium on Information Technology*, pp. 1–4, 2008.
- [25] K. Andreev and H. Racke, "Balanced graph partitioning," *Theory of Computing Systems*, vol. 39, no. 6, pp. 929–939, 2006.
- [26] I. P. Gent and T. Walsh, "Analysis of heuristics for number partitioning," *Computational Intelligence*, vol. 14, no. 3, pp. 430–451, 1998.
- [27] H. Ding, H. Yu, and Z. Wang, "Greedy strategy works for k -center clustering with outliers and coresets construction," *arXiv preprint arXiv:1901.08219*, 2019.
- [28] F. Blahoudek, T. Brázdil, P. Novotný, M. Ornik, P. Thangeda, and U. Topcu, "Qualitative controller synthesis for consumption Markov decision processes," in *32nd International Conference on Computer Aided Verification - Part II*, pp. 421–447, 2020.



Farhad Nawaz is a Ph.D. student in Electrical and Systems Engineering at the University of Pennsylvania. He received his M.S. degree in Aerospace Engineering from the University of Illinois Urbana-Champaign in May 2021. His research interests lie in the areas of combining machine learning and control theory for problems in uncertain and complex dynamical systems. He envisions developing intelligent control frameworks for autonomous systems.



Melkior Ornik is an assistant professor in the Department of Aerospace Engineering and the Coordinated Science Laboratory at the University of Illinois Urbana-Champaign. He received his Ph.D. degree from the University of Toronto in 2017. His research focuses on developing theory and algorithms for learning and planning of autonomous systems operating in uncertain, complex and changing environments, as well as in scenarios where only limited knowledge of the system is available.