

Policy Synthesis for Factored MDPs with Graph Temporal Logic Specifications

Murat Cubuktepe
The University at Texas at Austin
Austin, Texas
mcubuktepe@utexas.edu

Zhe Xu
The University at Texas at Austin
Austin, Texas
zhexu@utexas.edu

Ufuk Topcu
The University at Texas at Austin
Austin, Texas
utopcu@utexas.edu

ABSTRACT

We study the synthesis of policies for multi-agent systems to implement *spatial-temporal* tasks. We formalize the problem as a *factored* Markov decision process subject to so-called *graph temporal logic* specifications. The transition function and the spatial-temporal task of each agent depend on the agent itself and its neighboring agents. The structure in the model and the specifications enable to develop a distributed algorithm that, given a factored Markov decision process and a graph temporal logic formula, decomposes the synthesis problem into a set of smaller synthesis problems, one for each agent. We prove that the algorithm runs in time linear in the total number of agents. The size of the synthesis problem for each agent is exponential only in the number of neighboring agents, which is typically much smaller than the number of agents. We demonstrate the algorithm in case studies on disease control and urban security. The numerical examples show that the algorithm can scale to hundreds of agents.

ACM Reference Format:

Murat Cubuktepe, Zhe Xu, and Ufuk Topcu. 2020. Policy Synthesis for Factored MDPs with Graph Temporal Logic Specifications. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*, Auckland, New Zealand, May 9–13, 2020, IFAAMAS, 9 pages.

1 INTRODUCTION

We consider a system where multiple agents coordinate to achieve a set of *spatial-temporal tasks* defined over an underlying graph modeling the interaction between the agents. For example, in the graph in Figure 1, each node represents an agent, and each agent has its own set of states and actions. We draw edges between *neighboring* agents that exchange their current state information. For each agent, the transition function between its states also depends on the current state of the neighboring agents. The labels at each node of the graph provide information about the agents related to the task. The spatial-temporal task specified for each agent depends both on the agent itself and its neighboring agents. For example, the different nodes in a graph as shown in Figure 1 can model different police officers. The states of each node can represent the intersections that the corresponding police officer is monitoring in a city. An edge between two nodes exists if the two corresponding police officers can share their state information. The task might be “agent 3 or a neighboring agent of 3 should be in the intersection labeled as blue in every two time steps”.

Partially funded by the grants AFRL FA9550-19-1-0169, and ONR N00014-18-1-2829. *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*, B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand. © 2020 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.




Figure 1: An example of an undirected graph. The nodes (1, 2, 3, and 4) of the graph indicate the agents. Each agent has its own set of states and actions. We draw an edge between two agents if they share the current state information. The color of the states gives the label for each agent.

We model the behavior of each agent as a Markov decision process (MDP) [25], which has been widely used to model and solve sequential decision-making problems. We can represent the states of the overall multi-agent system by explicitly enumerating all possible states of the agents, meaning each agent is interacting with all other agents. The resulting state space of the composed MDP will scale exponentially in the number of agents, and the representation will be impractical for policy synthesis. The interaction between the agents is typically *sparse*, and each agent in the system exchanges their current state information with only a few other agents. Examples of such systems appear in biochemical networks [24], smart grids [7], swarm robots [22], and disease control [6].

Because of the sparsity of the interaction between agents, the composed MDP is typically a factored MDP [5, 13], which provides an efficient representation for MDPs with multiple agents. In factored MDPs, the transition probabilities of an agent often depend only on a small number of other agents that are neighbors in the underlying graph. Such a representation alleviates the need for explicitly enumerating all possible states of the agents, and facilitates the synthesis of policies for systems with a large number of agents.

Related work in factored MDPs typically focuses on defining a reward function and computing a policy to maximize the expected reward. A wide range of tasks, such as avoiding certain parts of an environment, cannot be expressed in any reward function [16, 17]. Linear temporal logic [1, 23] is a language that can concisely express such tasks. However, linear temporal logic cannot express spatial-temporal tasks involving multiple agents.

Our first contribution is to represent spatial-temporal tasks in a specification language called graph temporal logic (GTL) [32]. GTL formulas represent tasks such as “the police officer at node 4 or their neighboring officers should visit intersection labeled as blue

in every two hours”. GTL is an extension of linear temporal logic and focuses on the spatial-temporal properties of the labels on a graph. We use GTL formulas to express tasks that concern a set of agents on the graph.

Expressing the spatial-temporal tasks as GTL specifications gives us the following advantages. 1) GTL provides a concise way to constrain the behavior of the agents to satisfy the spatial-temporal tasks, which cannot be done by designing a reward function in general. 2) GTL resembles natural languages, thus requirements specified by practitioners can be translated into GTL specifications. 3) GTL specifications can be converted into a deterministic finite automaton [32], which facilitates the policy synthesis utilizing the techniques in formal methods.

Our second contribution is to develop a centralized algorithm for the factored MDP policy synthesis problem subject to GTL specifications. We first derive the centralized algorithm based on a linear programming problem (LP). The proposed LP is a generalization of the LP formulation for MDPs subject to linear temporal logic specifications [1, 8]. By solving the LP, we synthesize a policy for each agent, which is a function of the states of the agent and its neighboring agents on the graph.

Our third contribution is to develop a distributed algorithm, since the centralized algorithm may not scale to a large number of agents. We develop the distributed algorithm by extending the centralized synthesis algorithm. Specifically, the algorithm decomposes the centralized problem into smaller synthesis problems, one for each agent. The algorithm then solves the smaller synthesis problems in parallel at each iteration. The running time of the algorithm is linear with the number of agents in the graph and exponential with the number of neighbors for each agent, which is typically much smaller due to the sparsity of the graph.

We demonstrate the effectiveness of the algorithm on two examples with a large number of agents. In the first example, we consider disease management in crop fields with GTL specifications [26]. A contaminated crop field can infect its neighbors, and the yield of that crop field decreases. The GTL specifications ensure that some of the crop fields are treated immediately to prevent the spreading of disease. We consider maximizing the expected yield of the crop fields while ensuring that some of the crop fields are treated immediately. We also show that the decentralized algorithm outperforms the centralized algorithm with a large number of agents. Specifically, the running time of the decentralized algorithm scales linearly with the number of agents. In the second example, we consider an urban security problem. The objective of the problem is to assign patrol tasks to police officers such that certain critical locations in the city are sufficiently monitored. We express the task of monitoring the critical locations in GTL specifications. The police officers need to coordinate with each other to satisfy the GTL specifications. The results show that the proposed distributed algorithm scales to hundreds of agents while ensuring that the agents achieve tasks that are specified in GTL formulas.

Related work. Related work on factored MDPs considers maximizing an expected reward (or, minimizing an expected cost). The existing results consider optimizing for a value function using approximate linear programming [9, 12, 13], approximate policy iteration [21, 26], and approximate value iteration [5, 13]. However, optimizing for a value function may not be sufficient to ensure

safety or performance guarantees in tasks that include multiple agents. Reference [16] shows that no reward structure, in general, can capture tasks that are given by temporal logic specifications.

The work in [27] considers the problem of coordinating multiple homogeneous agents subject to constraints on the number of agents achieving a task given in temporal logic specifications. However, they do not allow agents to be heterogeneous and do not differentiate between different agents on different tasks. Recent work in [15, 18, 19] proposes a spatial-temporal logic for swarm robots. The proposed logic is less expressive than GTL, and their solution approach involves solving a mixed-integer linear program, which is significantly more challenging—in theory and in practice—than the optimization problems that arise in the proposed approach. Reference [29] proposes a framework for potential-based collision avoidance of multi-agent systems with temporal logic specifications. However, the proposed controller is centralized, and the controller does not scale to a large number of agents.

2 PRELIMINARIES

A *probability distribution* over a finite set X is a function $\mu: X \rightarrow [0, 1] \subseteq \mathbb{R}$ with $\sum_{x \in X} \mu(x) = \mu(X) = 1$. The set X of all distributions is $\text{Distr}(X)$. For two sets A, B we define $A \subseteq B$ if B contains all elements of A .

Definition 2.1 (Undirected graph). Let $G = (V, E)$ be an undirected graph, where $V = \{1, \dots, M\}$ is a finite set of agents and E is a finite set of edges. We use $e = \{1, 2\}$ to denote that the edge $e \in E$ connects 1 and $2 \in V$. For agents i, j , we call j a neighboring agent of i , if there is an edge e that connects i and j . Let $N(i) \subseteq V$ be the set of agent i and the neighboring agents of the agent i . We denote $N(i, j) = N(i) \cap N(j)$ and $N(i \setminus j) = N(i) \setminus N(j)$. The number of agents in V is given by $M = |V|$.

Factored MDP

Definition 2.2 (Factored MDP). A factored Markov decision process (MDP) $\mathcal{M} = (S, s^I, \text{Act}, \mathcal{P}, \pi, \mathcal{L}, R)$ on the undirected graph $G = (V, E)$ is defined by a finite set S of states, which is given as a Cartesian product of the states for each agent i in the graph, i.e., $S = S_1 \times S_2 \times \dots \times S_M$, an initial state $s^I \in S$, a finite set $\text{Act} = \text{Act}_1 \times \text{Act}_2 \times \dots \times \text{Act}_M$ of actions, a transition function $\mathcal{P} = \mathcal{P}_1 \times \mathcal{P}_2 \times \dots \times \mathcal{P}_M$, where for each agent i , $\mathcal{P}_i: S_{N(i)} \times \text{Act}_{N(i)} \rightarrow \text{Distr}(S_{N(i)})$ gives the transition function for agent i , where $S_{N(i)}$ and $\text{Act}_{N(i)}$ denote the Cartesian product of the sets of states and actions of the agents in $N(i)$ respectively, a finite set π of atomic propositions, a labeling function $\mathcal{L} = \mathcal{L}_1 \times \mathcal{L}_2 \times \dots \times \mathcal{L}_M$, where $\mathcal{L}_i: S_{N(i)} \rightarrow 2^\pi$ that labels each state $s \in S_{N(i)}$ with a subset of atomic propositions $\mathcal{L}_i(s) \subseteq \pi$ and a reward function $R = R_1 \times R_2 \times \dots \times R_M$, where $R_i: S_{N(i)} \times \text{Act}_{N(i)} \rightarrow \mathbb{R}_{\geq 0}$ assigns a reward to state-action pairs for agent i . We use $s(t)$ to denote the states of all agents in V at time index t .

We give two examples to illustrate the concepts in factored MDPs. First example shows the relationship between the agents in the graph. Second example illustrates how the transition probabilities between states of an agent depend on the neighboring agents.

Example 2.3. We show an example of a factored MDP in Figure 2. Consider a graph $G = (V, E)$, where $V = \{1, 2, 3, 4\}$ and




Figure 2: An example of a factored MDP. For each agent i in V , S_i depicts the state space of the agent i , and s_i^j depicts the state j of the agent i . The arrows between the states of the agent i shows the transitions between states of the agent v . Let the current states be s_1^1, s_2^2 , and s_3^3 for the agents 1, 2, and 3 respectively. Then, the transition probability of agent 1 at state s_1^1 to s_1^2 and s_1^3 is a function of s_1^1, s_2^2 , and s_3^3 .

$E = \{\{1, 2\}, \{1, 3\}, \{2, 3\}, \{2, 4\}, \{3, 4\}\}$. The black nodes in Figure 2 represent the agents in the graph G . We denote the edges between the agents in the graph with green lines in Figure 2. For $i \in V$, we denote the state space and action space of the agent i as S_i and Act_i , and we denote the state j in S_i as s_i^j . We treat s_i^j and s_k^l to be different states if $i \neq k$ or $j \neq l$. For this example, $N(1) = \{1, 2, 3\}, N(4) = \{2, 3, 4\}$, and $N(1, 4) = \{2, 3\}$.

Example 2.4. We show an example of the transition probabilities between states of the agent 1 in a factored MDP in Figure 3. For this example, we assume that $S_{N(1)} = \{s_1, s_2, s_3\}$, for $s_1 \in S_1, s_2 \in S_2, s_3 \in S_3$ and $Act_{N(1)} = Act_1$, meaning the transition probability of the agent 1 is a function of the states of agent 1 and the neighboring agents, and the action of agent 1. The transition probabilities between the states of S_1 are given with red lines in Figure 3. For example, in Figure 3, the transition probabilities between states in S_1 for a given action α_1^1 is a function of $s_1 \in S_1, s_2 \in S_2$ and $s_3 \in S_3$.

Definition 2.5 (Policy). A memoryless and randomized policy for a factored MDP \mathcal{M} is a function $\sigma = \sigma_1 \times \dots \times \sigma_M$, where for each agent i , $\sigma_i: S_{N(i)} \rightarrow \text{Distr}(Act_{N(i)})$. The set of all policies is $Pol_{\mathcal{M}}$.

Applying a policy $\sigma \in Pol_{\mathcal{M}}$ to a factored MDP \mathcal{M} yields an induced factored Markov chain \mathcal{M}_σ .

Definition 2.6 (Factored induced MC). For a factored MDP $\mathcal{M} = (S, s^I, Act, \mathcal{P}, \pi, \mathcal{L}, R)$ and a policy $\sigma \in Pol_{\mathcal{M}}$, the factored MC induced by \mathcal{M} and σ is $\mathcal{M}_\sigma = (S, s^I, Act, \mathcal{P}_\sigma, \pi, \mathcal{L}, R)$, where

$$\mathcal{P}_\sigma(s'|s) = \sum_{\alpha \in Act(s)} \sigma(s, \alpha) \cdot \mathcal{P}(s'|s, \alpha) \quad \forall s, s' \in S.$$

Definition 2.7 (Trajectory). A finite or an infinite sequence $\varrho_\sigma = s(0)s(1)s(2) \dots$ of states generated in \mathcal{M} under a policy $\sigma \in Pol_{\mathcal{M}}$ is called a trajectory. The state $s(0)$ denotes the initial state s^I .




Figure 3: An example of transition probabilities between states for the agent 1 on the underlying graph of a factored MDP if the state in the neighboring agents are $s = 2$ for the agent 2 and $s = 4$ for the agent 3 if an action $\alpha = 1$ is taken for the agent 1.

Given an induced factored MC \mathcal{M}_σ , starting from the initial state s^I , the state visited at step t is given by a random variable $X(t)$. The probability of reaching state s' from state s in one step, denoted $\mathbb{P}(X(t+1) = s'|X(t) = s)$ is equal to $\mathcal{P}_\sigma(s'|s)$. We can extend one-step reachability over a set of trajectories ϱ_σ , i.e., $\mathbb{P}(X(n) = s(n), \dots, X(0) = s(0)) = \mathbb{P}(X(n) = s(n)|X(n-1) = s(n-1)) \cdot \mathbb{P}(X(0) = s(0), \dots, X(n-1) = s(n-1))$. We denote the set of all trajectories in \mathcal{M} under the policy σ by $Tr_\sigma(\mathcal{M})$.

Graph Temporal Logic

We review the theoretical framework of graph temporal logic (GTL), which is introduced in [32]. For a more detailed discussion about the syntax of GTL, we refer the reader to [32].

We denote Y as the set of states for the edges, which is given as a Cartesian product of the states for each edge e_i , i.e., $Y = Y_1 \times Y_2 \times \dots \times Y_{|E|}$, where Y_i is the state space of the edge e_i . Let $\mathbb{T} = \{1, 2, \dots\}$ be a discrete set of time indices. We use $y(t)$ to denote the states of all edges in E at time index t .

Definition 2.8 (Graph-temporal trajectory). [32] A graph-temporal trajectory on a graph G is a tuple $g = (s, y)$, where $s: \mathbb{T} \rightarrow S$ assigns a label for each node $i \in V$ at each time index $t \in \mathbb{T}$, and $y: \mathbb{T} \rightarrow Y$ assigns a label for each edge $e_i \in E$ at each time index $t \in \mathbb{T}$.

Definition 2.9 (Node and edge propositions). [32] An atomic node proposition is a Boolean valued map on S , which is a predicate. An edge proposition is a Boolean valued map on Y .

Let π be an atomic node proposition, and $O(\pi)$ be the subset of S for which the atomic proposition π is true. Similarly, let ρ be an edge proposition, and $O(\rho)$ be subset of Y for which ρ is true.

A graph-temporal trajectory $g = (s, y)$ satisfies an atomic node proposition π at a node v and at a time index t , which we denote as $(g, v, t) \models \pi$, if and only if $s_v(t) \in O(\pi)$. Similarly, $g = (s, y)$ satisfies an edge proposition ρ at an edge e and at a time index t , which we denote as $(g, e, t) \models \rho$, if and only if $y_e(t) \in O(\rho)$.

Example 2.10. Referring back to the example in Figure 1, we show an example with 4 police officers. The node and edge labels are from a graph-temporal trajectory g for a time index 1 and 2. The node labels represent the intersection that the corresponding police officer is monitoring, and the edge labels represent the distance between two neighboring officers. The atomic node proposition $\pi = (x = \text{blue} \vee \text{orange})$ is satisfied by g at nodes 1, 2, and 4 at time index 1, and at nodes 2, and 3 at time index 2. The edge proposition




Figure 4: An example of a graph-temporal trajectory of time length 2 on an undirected graph, with the red items indicating node labels, and the blue numbers indicating edge labels. The first item in the labels is for time index 1, and the second item in the labels is for time index 2.

$\rho = (y \leq 2)$ is satisfied by g at $e_1, e_3,$ and e_4 at time index 1, and at $e_1, e_2,$ and e_5 at time index 2.

Definition 2.11 (Neighboring operation). [32] Given a graph-temporal trajectory $g = (s, y)$ on a graph G , an edge proposition ρ , and a subset $V' \subseteq V$ of nodes, we define the neighboring operation $\circ_{\rho} : 2^V \times \mathbb{T} \rightarrow 2^V \times \mathbb{T}$ as

$$\circ_{\rho}(V', t) = (\{v | \exists v' \in V', \exists e \in E, e = \{v', v\}, (g, e, t) \models \rho\}, t).$$

$\circ_{\rho}(V', t)$ are the set of nodes that can be reached from V' through an edge e if the edge proposition ρ is true by g at time index t .

Example 2.12. For the graph-temporal trajectory g on the graph G at time index 1 in Figure 4,

$$\begin{aligned} \circ_{y \leq 2}(\{1\}, 1) &= (\{2\}, 1), \\ \circ_{y \leq 2} \circ_{y \leq 2}(\{1\}, 1) &= \circ_{y \leq 2}(\{2\}, 1) = (\{1, 4\}, 1). \end{aligned}$$

A graph-temporal trajectory g at a node v at a time index t satisfies the formula $\exists^N(\circ_{\rho_n} \cdots \circ_{\rho_1})\varphi$ if there exist at least N nodes in $(\circ_{\rho_n} \cdots \circ_{\rho_1})(v, t)$ where the formula φ is satisfied by a graph-temporal trajectory g at time index t . We also define that a graph-temporal trajectory g satisfies a GTL formula φ at a node v , denoted as $(g, v) \models \varphi$, if g satisfies φ at node v at time index 0.

For the GTL formulas that we consider, specifically *syntactically co-safe* (resp. *safe*) formulas, we can build a DFA $\mathcal{A}^{\varphi, v}$ (resp. $\mathcal{A}^{\neg\varphi, v}$) over \mathcal{AP} that accepts precisely the graph-temporal trajectories that satisfy (resp. violate) the GTL formula φ at node v [32]. We focus on syntactically co-safe and syntactically safe formulas for the remainder of the paper, which can be satisfied or violated with a graph-temporal trajectory of finite time length.

For simplicity, we focus on the node labels only in this paper. In this case, the graph-temporal trajectory g reduces to a trajectory ϱ (as defined in Definition 2.7). We use $(\varrho, v, t) \models \varphi$ to denote that a trajectory ϱ satisfies the GTL formula φ at node v at time index t .

3 PROBLEM FORMULATION

To solve the factored MDP policy synthesis problem, we synthesize a policy for each agent that satisfies the GTL specification. We construct a factored MDP \mathcal{M} that captures all trajectories of \mathcal{M} satisfying a GTL formula φ by taking the product of \mathcal{M} and the DFA \mathcal{A}^{φ} corresponding to φ , which represents the GTL specifications as reachability specifications.

Definition 3.1 (Product factored MDP). Let $\mathcal{M} = (S, s^I, Act, \mathcal{P}, \pi, \mathcal{L}, R)$ be a factored MDP and $\mathcal{A} = (Q, q_I, \Sigma, \delta, Acc)$ be a DFA. The product factored MDP is a tuple $\mathcal{M}^P = (S^P, s^{I,P}, Act, \mathcal{P}^P, \mathcal{L}^P, R^P, Acc^P)$ with a finite set $S^P = S \times Q$ of states, an initial state $s^{I,P} = (s^I, q) \in S^P$ that satisfies $q = \delta(q_I, \mathcal{L}(s^{I,P}))$, a finite set Act of actions, a labeling function $\mathcal{L}^P((s, q)) = \{q\}$, a transition function $\mathcal{P}_i^P((s, q), \alpha, (s', q')) = \mathcal{P}_i(s, \alpha, s')$ if $q' = \delta(q, \mathcal{L}(s'_i))$, and $\mathcal{P}_i^P((s, q), \alpha, (s', q')) = 0$ otherwise, a reward function R^P that satisfies $R^P(s, q, \alpha) = R(s, \alpha)$ for $s \in S$ and $\alpha \in Act$, the acceptance condition $Acc^P = \{A_1^P, \dots, A_k^P\}$, where $A_i^P = S^P \times A_i$ for all $A_i \in Acc$. We assume all accepting states are absorbing.

A memoryless and randomized policy for a product factored MDP \mathcal{M}^P is a function $\sigma^P = \sigma_1^P \times \sigma_2^P \times \dots \times \sigma_M^P$, where $\sigma_i^P : S_{N(i)}^P \rightarrow \text{Distr}(Act_{N(i)})$. A memoryless policy σ^P is a finite-memory policy σ' in the underlying factored MDP \mathcal{M}^P . Given a state $(s, q) \in S^P$, we consider q to be a memory state and define $\sigma'(s) = \sigma^P(s, q)$.

Let $\sigma^P \in \text{Pol}_{\mathcal{M}^P}$ be a policy for \mathcal{M}^P and let $\sigma' \in \text{Pol}_{\mathcal{M}}$ be the policy on \mathcal{M} constructed from σ^P through the procedure explained above. The trajectories of the factored MDP \mathcal{M} under the policy σ' satisfy the GTL specification φ with a probability of at least λ at node v and at a time index t , if and only if the trajectories of the product factored MDP \mathcal{M}^P under the policy σ^P reach some accepting states in \mathcal{M}^P with a probability of at least λ [1].

Definition 3.2 (Occupancy measure). The occupancy measure o_i^{σ} of a policy σ_i for a set of neighboring agents $N(i)$ of a factored MDP \mathcal{M} is defined as

$$\begin{aligned} o_i^{\sigma}(\hat{s}_{N(i)}, \hat{\alpha}_{N(i)}) &= \mathbb{E} \left[\sum_{t=0}^{\infty} \mathbb{P}(\hat{s}^{N(i)}(t) = \hat{s}_{N(i)}, \right. \\ &\quad \left. \hat{\alpha}_{N(i)}(t) = \hat{\alpha}_{N(i)} | s_{N(i)}(t) = s_{N(i)}^I) \right], \end{aligned} \quad (1)$$

where $\hat{s}_{N(i)}(t) = \{s_{j_1}(t), \dots, s_{j_{|N(i)|}}(t)\} \in S_{N(i)}$ and $\hat{\alpha}_{N(i)}(t) = \{\alpha_{j_1}(t), \dots, \alpha_{j_{|N(i)|}}(t)\} \in Act_{N(i)}$ denote the state and action of the agent i and the neighboring agents in \mathcal{M} at time index t . The equality $\hat{s}_{N(i)}(t) = \hat{s}_{N(i)}$ means all elements of $\hat{s}_{N(i)}(t)$ and $\hat{s}_{N(i)}$ are the same. The occupancy measure $o_i^{\sigma}(\hat{s}_{N(i)}, \hat{\alpha}_{N(i)})$ is the expected number of times to take the action α_{j_k} at the state s_{j_k} for all $k \in N(i)$ under the policy σ_i .

Problem Statement

In this section, we formally state the policy synthesis problem of a factored MDP subject to GTL specifications. Our objective is to synthesize a policy that induces a stochastic process with a maximum expected reward whose trajectories satisfy the given GTL specification with at least a desired probability. To this end, we synthesize a policy $\sigma^P \in \text{Pol}_{\mathcal{M}^P}$ that reaches and stays in the accepting states in \mathcal{M}^P with probability of at least λ .

PROBLEM 1. Given a factored product MDP \mathcal{M}^P , λ, k , and a set of agents \hat{V} , compute a policy $\sigma^P \in \text{Pol}_{\mathcal{M}^P}$ that solves the problem

$$\text{maximize}_{\sigma^P \in \text{Pol}_{\mathcal{M}^P}} \mathbb{E} \left[\sum_{t=0}^{\infty} R(s(t), \alpha(t)) \right] \quad (2)$$

$$\text{subject to } \mathbb{P}_{\mathcal{M}^P}^{\sigma^P}((\varrho^{\sigma^P}, v, k) \models \varphi) \geq \lambda, \quad (3)$$

where $\mathbb{P}_{\mathcal{M}^P}^{\sigma^P}((\varrho^{\sigma^P}, v, k) \models \varphi)$ denotes the probability of satisfying the GTL specification φ with the trajectory ϱ^{σ^P} for some agents $v \in \hat{V}$.

\hat{V} is the set of agents that have a GTL specification that is required to be satisfied at the time index k in the factored product MDP \mathcal{M}^P under the policy σ^P . Without loss of generality, we assume $k = 1$.

4 POLICY SYNTHESIS

We now describe the proposed approach to synthesize a policy for each agent to solve Problem 1. We first give a centralized formulation based on a linear programming problem. We then develop a distributed approach based on the centralized formulation.

4.1 Centralized Approach

In this section, we propose a linear programming problem (LP) for solving Problem 1. Our solution is based on the multi-objective dual LP formulation to compute a policy that maximizes the expected reward while satisfying a graph temporal logic specification φ in a factored MDP with $|V| = 1$, which is an MDP [8, 25]. Let \bar{S} be the set of all states in the MDP \mathcal{M} that are not in Acc . Then, we define the variables of the dual LP formulation as follows.

- $o(s, \alpha) \in [0, \infty)$ for each state $s \in \bar{S}$ and action $\alpha \in Act$ defines the occupancy measure of a state-action pair for the policy σ , i.e., the expected number of times of taking action α in state s .
- $o(s) \in [0, 1]$ for each state $s \in Acc$ defines the occupancy measure of an accepting state $s \in Acc$, which is equal to the probability of reaching the accepting state s .

Note that the variables $o(s)$ are defined in the interval $[0, 1]$ instead of $[0, \infty)$, as they represent the probability of reaching an accepting state $s \in Acc$. The variable $o(s, \alpha)$ represents the expected number of taking action α in a non-accepting state s , thus it can exceed 1. We refer to [8, 28] for further explanation of the variables.

The dual LP is given by

$$\text{maximize } \sum_{s \in \bar{S}} o(s)R(s) \quad (4)$$

subject to

$$\forall s \in \bar{S}, \quad \sum_{\alpha \in Act} o(s, \alpha) = \sum_{s' \in \bar{S}} \sum_{\alpha \in Act} \mathcal{P}(s|s', \alpha) o(s', \alpha) + \mu(s), \quad (5)$$

$$\forall s \in Acc, \quad o(s) = \sum_{s' \in \bar{S}} \sum_{\alpha \in Act} \mathcal{P}(s|s', \alpha) o(s', \alpha) + \mu(s), \quad (6)$$

$$\sum_{s \in Acc} o(s) \geq \lambda, \quad (7)$$

where $\mu(s) = 1$ if $s = s^I$ and $\mu(s) = 0$ if $s \neq s^I$. The constraints (5) and (6) ensure that the expected number of times transitioning to a state $s \in \bar{S}$ is equal to the expected number of times to take action α that transitions to a different state $s' \in S$. These constraints are also referred to as *flow constraints* [25]. The constraint (7) ensures that the specification φ is satisfied with a probability of at least λ .

For any optimal solution o to the LP in (4)–(7),

$$\sigma(s, \alpha) = \frac{o(s, \alpha)}{\sum_{\alpha' \in Act} o(s, \alpha')} \quad (8)$$

is an optimal policy, and o is the occupancy measure of σ [8, 25].

4.2 LP-based Policy Synthesis of Factored MDPs

We now describe our centralized approach for the policy synthesis problem for factored MDPs subject to GTL specifications. Let $\bar{S}_{N(i)}^P$ be the set of all states in the product factored MDP \mathcal{M}^P that are not in A_i^P for each agent i . Then, we define the variables of the LP for policy synthesis as follows.

- $o_i(\hat{s}_{N(i)}, \hat{\alpha}_{N(i)}) \in [0, \infty)$ for each set of neighboring states $\hat{s}_{N(i)} \in \bar{S}_{N(i)}^P$ and actions $\hat{\alpha}_{N(i)} \in Act_{N(i)}$ defines the occupancy measure of a state-action pair for σ_i .
- $o_i(\hat{s}_{N(i)}) \in [0, 1]$ for each state $\hat{s}_{N(i)} \in A_i^P$ defines the probability of reaching an accepting state $s \in A_i^P$.

The objective of the LP is given by

$$\text{maximize } \sum_{i=1}^M \sum_{\hat{s}_{N(i)} \in \bar{S}_{N(i)}^P} \sum_{\hat{\alpha}_{N(i)} \in Act_{N(i)}} o_i(\hat{s}_{N(i)}, \hat{\alpha}_{N(i)}) R_i(\hat{s}_{N(i)}, \hat{\alpha}_{N(i)}). \quad (9)$$

For each agent $i \in V$, and state $\hat{s}_{N(i)} \in \bar{S}_{N(i)}^P$, the constraints

$$\sum_{\hat{\alpha}_{N(i)} \in Act_{N(i)}} o_i(\hat{s}_{N(i)}, \hat{\alpha}_{N(i)}) - \mu(\hat{s}_{N(i)}) = \sum_{\hat{s}'_{N(i)} \in \bar{S}_{N(i)}^P} \sum_{\hat{\alpha}_{N(i)} \in Act_{N(i)}} \mathcal{P}_i^P(\hat{s}_{N(i)} | \hat{s}'_{N(i)}, \hat{\alpha}_{N(i)}) o_i(\hat{s}'_{N(i)}, \hat{\alpha}_{N(i)}) \quad (10)$$

denote the flow constraints, similar to the constraints (5).

For each agent $i \in V$, and state $\hat{s}_{N(i)} \in A_i^P$, the constraints

$$o_i(\hat{s}_{N(i)}) - \mu(\hat{s}_{N(i)}) = \sum_{\hat{s}'_{N(i)} \in \bar{S}_{N(i)}^P} \sum_{\hat{\alpha}_{N(i)} \in Act_{N(i)}} \mathcal{P}_i^P(\hat{s}_{N(i)} | \hat{s}'_{N(i)}, \hat{\alpha}_{N(i)}) o_i(\hat{s}'_{N(i)}, \hat{\alpha}_{N(i)}) \quad (11)$$

denote the flow constraints for the accepting states, analogous to the constraints (6).

For agents $i, j \in V$ such that $N(i) \cap N(j) \neq \emptyset$, we ensure that the occupancy measure is consistent in the states and actions of agents $k \in N(i, j)$. Thus, the agents take account of its neighboring agents' occupation measures during the policy computation. For each set of states $\hat{s}_{N(i,j)} \in S_{N(i,j)}$ and actions $\hat{\alpha}_{N(i,j)} \in Act_{N(i,j)}$, the constraints

$$\sum_{\hat{s}_{N(i)} \supseteq \hat{s}_{N(i,j)}} \sum_{\hat{\alpha}_{N(i)} \supseteq \hat{\alpha}_{N(i,j)}} o_i(\hat{s}_{N(i)}, \hat{\alpha}_{N(i)}) = \sum_{\hat{s}_{N(j)} \supseteq \hat{s}_{N(i,j)}} \sum_{\hat{\alpha}_{N(j)} \supseteq \hat{\alpha}_{N(i,j)}} o_j(\hat{s}_{N(j)}, \hat{\alpha}_{N(j)}) \quad (12)$$

ensure that the time spent in the set of states $\hat{s}_{N(i,j)}$ and taking the set of actions $\hat{\alpha}_{N(i,j)}$ is equal for the policies of agents i and j .

Finally, the constraints

$$\sum_{\hat{s}_{N(i)} \in A_i^P} o_i(\hat{s}_{N(i)}) \geq \lambda_i \quad (13)$$

encode the specification constraints for each agent $i \in V$, similar to the constraints (7). We illustrate the constraints (12) by an example.

Example 4.1. Consider the factored MDP in Figure 2 and the agents 1 and 2. $N(1, 2) = \{1, 2, 3\}$, $N(1 \setminus 2) = \{3\}$ and $N(2 \setminus 1) = \{4\}$.

Algorithm 1: Distributed Method with Primal-Splitting ADMM

Initialize: o_i^0 and κ_i^0 ($i = 1, 2, \dots, M$);

for $k = 0, 1, \dots, I$ **do**

for $i = 1, 2, \dots, M$ **do**

$z_i^{k+1} \leftarrow -\frac{1}{M} \left(\sum_{i=1}^M A_i o_i^k - \kappa_i^k / \beta \right) + \left(A_i o_i^k - \kappa_i^k / \beta \right)$.

$o_i^{k+1} \leftarrow \operatorname{argmin}_{o_i} f_i(o_i) + \frac{\beta}{2} \left\| A_i o_i - z_i^{k+1} - \kappa_i^k / \beta \right\|_2^2$.

$\kappa_i^{k+1} \leftarrow \kappa_i^k - \beta (A_i o_i^{k+1} - z_i^{k+1})$.

$\operatorname{res}_p \leftarrow \sum_{i=1}^M \|A_i o_i^k - z_i^k\|_2^2$.

$\operatorname{res}_d \leftarrow \sum_{i=1}^M \beta \|v_i^k - v_i^{k-1}\|_2^2$.

if $\operatorname{res}_p \leq \gamma$ **and** $\operatorname{res}_d \leq \gamma$ **then**

return o_i for $i = 1, 2, \dots, M$.

return o_i for $i = 1, 2, \dots, M$.

Therefore, to ensure that the occupancy measure is consistent for agents 1 and 2, we add the constraints

$$\sum_{s_3 \in S_{N(1,2)}} \sum_{\alpha_3 \in \operatorname{Act}_{N(1,2)}} o_1(\hat{s}_{N(1,2)}, s_3, \hat{\alpha}_{N(1,2)}, \alpha_3) = \sum_{s_4 \in S_{N(2,1)}} \sum_{\alpha_4 \in \operatorname{Act}_{N(2,1)}} o_2(\hat{s}_{N(1,2)}, s_4, \hat{\alpha}_{N(1,2)}, \alpha_4). \quad (14)$$

for $\hat{s}_{N(1,2)} \in S_{N(1,2)}$ and $\hat{\alpha}_{N(1,2)} \in \operatorname{Act}_{N(1,2)}$.

The LP in with the objective (9) and the constraints (10)–(13), computes a policy for each agent i that satisfies the GTL specification and maximizes the expected reward. However, the LP in (9)–(13) can be time consuming to solve if the number of agents M is large. In the next section, we propose a distributed approach that runs in time linear in M .

4.3 Distributed Approach

In this section, we discuss how we can solve the LP in (9)–(13) in a distributed manner. We utilize alternating direction method of multipliers (ADMM) [2, 11] to solve a large-scale factored MDP synthesis problem by decomposing them into a set of smaller problems. We pose the multi-block problem into an equivalent two-block problem, and apply the *primal-splitting ADMM* algorithm [30] to the equivalent problem to ensure convergence as ADMM does not necessarily converge to an optimal solution with $M > 2$ agents [4].

4.3.1 Primal-Splitting ADMM. The LP in (9)–(13) with M agents can be written as following optimization problem

$$\text{minimize} \quad \sum_{i=1}^M f_i(o_i) \quad (15)$$

$$\text{subject to} \quad \sum_{i=1}^M A_i o_i = 0, \quad (16)$$

where $f_i(o_i)$ is the negative of the objective in (9) and encodes the constraints in (10)–(11) and (13) for each agent i . The constraints in (16) depict the constraints in (12) in a compact form. The matrices A_i encodes the coefficients in (12). Specifically, for agent i , the objective $f_i(o_i)$ in (15) can be expressed as following

$$\text{minimize} \quad f_i(o_i) = (g_1(o_i) + \mathcal{I}_0(g_2(o_i)) + \mathcal{I}_0(g_3(o_i)) + \mathcal{I}_+(g_4(o_i)))$$

where $g_1(o_i)$ is the negative of the objective in (9), and $g_2(o_i) =$

$$\sum_{\hat{\alpha}_{N(i)} \in \operatorname{Act}_{N(i)}} o_i(\hat{s}_{N(i)}, \hat{\alpha}_{N(i)}) - \mu(\hat{s}_{N(i)}) - \sum_{\hat{s}'_{N(i)} \in \mathcal{S}_{N(i)}^p} \sum_{\hat{\alpha}_{N(i)} \in \operatorname{Act}_{N(i)}} \mathcal{P}_i^p(\hat{s}_{N(i)} | \hat{s}'_{N(i)}, \hat{\alpha}_{N(i)}) o_i(s'_i, \hat{\alpha}_{N(i)}),$$

i.e., the evaluation of the constraints (10) for agent i . Similarly, $g_3(o_i)$ and $g_4(o_i)$ of the evaluation of the constraints (11) and (13). \mathcal{I}_0 is the indicator function of $\{0\}$, i.e., $\mathcal{I}_0(\zeta) = 0$ if $\zeta = 0$, and $\mathcal{I}_0(\zeta) = \infty$ otherwise, and \mathcal{I}_+ is the indicator function of nonpositive reals, i.e., $\mathcal{I}_+(\zeta) = 0$ if $\zeta \geq 0$, and $\mathcal{I}_+(\zeta) = \infty$ otherwise [3, p. 218].

We introduce a set of auxiliary variables $z_i, i = 1, \dots, M$ and write the optimization problem in (15)–(16) as following

$$\text{minimize} \quad \sum_{i=1}^M f_i(o_i) \quad (17)$$

$$\text{subject to} \quad A_i o_i = z_i, \quad i = 1, \dots, M, \quad (18)$$

$$\sum_{i=1}^M z_i = 0. \quad (19)$$

The optimization problems in (15)–(16) and in (17)–(19) are equivalent in the sense that they share the same optimal solution set for the variables $o_i, i = 1, \dots, M$, and achieve the same optimal objective value.

We solve the optimization problem in a distributed manner by primal-splitting ADMM [30], which is given in Algorithm 1, where $\beta > 0$ is the algorithm parameter, κ_i is the dual variable for the constraints in (18), and I is the number of iterations. The proposed method achieves an $O(1/k)$ convergence rate after k iterations [30].

The primal residual for the primal feasibility is given by res_p and the dual residual for the dual feasibility is given by res_d in Algorithm 1. The primal residual can be seen as the feasibility residual of the policy with respect to the specification, and the dual residual is the optimality residual of the policy with respect to the objective [2]. We stop the algorithm until it runs for I iterations, or if the residuals are below a threshold γ .

4.3.2 Complexity Analysis. Computationally, the most expensive step of Algorithm 1 is to solve a LP for the o_i update for $i = 1, \dots, M$. The number of variables and constraints in each LP for the o_i update is exponential in $N(i)$, and each LP for the o_i can be solved in time polynomial in the number of variables and constraints via interior-point methods [20]. Therefore, the computation time for each o_i update is exponential in $N(i)$, and the computation time of the Algorithm 1 is linear in M .

On the other hand, the number of variables and constraints in the optimization problem (15)–(16) is linear in M and exponential in $N(i)$. If we solve the optimization problem in (15)–(16) by an interior point method algorithm, then the overall complexity of the algorithm will be *polynomial*, or cubic in M in worst case. In our examples, we show that the running time for solving the optimization problem in (15)–(16) directly does not scale linearly in M .

5 NUMERICAL EXAMPLES

We demonstrate the proposed approach on two domains: (1) disease management in crop fields, (2) urban security. The simulations were

Table 1: The average yield with 100 fields, different values of λ , p and ζ , and 50% of the fields with a GTL specification.

$(p, \zeta) \backslash \lambda$	0.9	0.8	0.7	0.6
(0.1, 0.1)	3.61	4.15	4.57	5.05
(0.2, 0.2)	4.92	5.81	6.71	7.13
(0.5, 0.5)	7.05	8.12	8.38	8.43
(0.8, 0.8)	8.60	8.95	8.99	9.00

performed on a computer with an Intel Core i5-7200u 2.50 GHz processor and 8 GB of RAM with Gurobi 9.0 [14] as the LP solver.

5.1 Disease Management in Crop Fields

We consider the policy synthesis of a factored MDP for disease management in crop fields, which was introduced in [26]. If a crop field is contaminated, it can infect its neighbors and the yield of that field decreases. However, if a field is left fallow, it has a probability ξ of recovering from contamination. The decisions of each year for each field i include two actions ($Act_i = \{1, 2\}$) for each field: cultivate normally ($\alpha_i = 1$) or leave fallow ($\alpha_i = 2$).

The topology of the fields is represented by an undirected graph, where each node in the graph represents one crop field. An edge is drawn between two fields if the two fields share a common border and can infect each other. The number of neighbors for each field is four. Each crop field is in one of three states: $s_i = 1$ corresponds to the case where the field is uninfected. $s_i = 2$ and $s_i = 3$ correspond to different degrees of infection, with $s_i = 3$ corresponding to a higher degree of infection. The probability that a field moves from state $s_i = 1$ to state $s_i = 2$ or $s_i = 2$ to state $s_i = 3$ with $\alpha_i = 1$ is $\mathcal{P}_i = \mathcal{P}_i(\epsilon, p, n_i) = \epsilon + (1 - \epsilon)(1 - (1 - p)^{n_i})$, where ϵ and p are fixed parameters and n_i is the number of the neighbors of i that are infected. If the field i is in $s_i = 3$, then it remains in $s_i = 3$ with a probability 1 with $\alpha_i = 1$. The maximal yield $r = 10$ is achieved by an uninfected, cultivated field, and the yield decreases linearly with the level of infection, from maximal reward r to minimal reward $1 + r/10$. A field left fallow produces a reward of 1.

Forster and Gilligan [10] consider an optimal long-term disease control strategy of a crop field and determines that the optimal long-term control strategy requires treating most of the contaminated fields immediately to eradicate the disease. Based on this control strategy, we consider the specification $\varphi = \neg \diamond \square_{\leq 2} d \wedge \neg \diamond \square_{\leq 3} \exists^2 \bigcirc d$ for the critical fields, where d means the field is infected. A field satisfies the specification φ if the field is never infected in two time steps in a row, and there should not be two neighboring fields that are infected three time steps in a row. This specification ensures that if a field or its neighboring field gets infected, it will be treated in a very short time to eradicate the disease and to prevent the disease from spreading, which is motivated by the control strategy of [10]. Note that, in φ , the temporal parameter is 2 for the critical fields, and is 3 for the neighboring fields, as we intend to make sure that the critical fields are less infected than the non-critical fields, and also prevent spreading the disease to the critical fields.

We consider the situation where 50% of the fields are considered as critical fields, and we require that all the critical fields satisfy the specification φ with a probability of at least λ . Table 1 shows the average yield for different values of λ , p and ζ , with $\epsilon = 0.1$ and the




Figure 5: Norm of primal residual and dual residual versus the number of iterations for the crop fields problem with different values of the parameter β .

parameter $\beta = 1$ in Algorithm 1 after 500 iterations. The average yield of the fields increases with a lower probability λ of satisfying the specification φ . However, the fields may be contaminated with a higher probability with a lower probability λ of satisfying the specification φ . For example, with a larger value of λ , e.g., $\lambda = 0.9$, the critical fields (with a GTL specification) are infected on average with a probability of 0.03, whereas the non-critical fields (without a GTL specification) are infected with a probability of 0.07, showing that the critical fields are healthy with a higher probability. On the other hand, the probability of fields being healthy is lower with a lower threshold of λ , and the expected yield is higher. This shows the trade-off between the average yield from the fields and the probability of the critical fields being uninfected.

In Figure 5, we show the convergence rate of the approach with different values of the parameter β . The results in Figure 5 show that there is a clear trade-off between the magnitude of the parameter β and the residuals during the iterations. With a higher value of β , e.g., $\beta = 100$, the primal residual converges to a smaller value, e.g. to 10^{-4} after 500 iterations. However, the dual residual only converges to a value of 10^{-2} , and it is higher compared to the dual residuals with lower values of β . We see that with the parameter selection of $\beta = 1$, both the primal residual and the dual residual achieve relatively high accuracy, showing that the feasibility and optimality can be achieved with an accuracy of about 10^{-3} after 500 iterations. On the other hand, with a lower value of β , e.g., $\beta = 0.01$, we obtain a lower dual residual, which means we obtain better performance in the objective value. However, the primal residual achieves lower accuracy with $\beta = 0.01$.

We also show the scalability of the approach by varying the number of fields. We show the running time of the Algorithm 1 and the centralized method, which is solving the LP in (9)–(13) directly, for different number of fields in Figure 6. The results in Figure 6 shows that the running time of the approach scales linearly with the number of crop fields, and the centralized approach does not scale linearly with the number of crop fields. The Algorithm 1 is slower than the centralized method with fewer crop fields due to solving many smaller LP problems iteratively. However, with an increasing number of crop fields, the running time of the centralized method scales super-linearly and runs slower than Algorithm 1.




Figure 6: The running time of distributed method and centralized method with a different number of crop fields after 500 iterations with the distributed method. The shaded region shows the standard deviation of the running times.

5.2 Urban Security

We consider an urban security problem, where a criminal plans his next move randomly based on the information on the nearby locations that are protected by police patrols [31]. There are M police officers assigned to monitor the locations. Each police officer coordinates with a sub-group of police officers in monitoring.

Figure 7 shows 35 intersections in San Francisco, CA with 7 rows and 5 columns [31]. We use a factored MDP to describe the network of the set of intersections. We show the number of crimes that occurred in October 2018 within 500 feet of each intersection in Figure 7¹. A police officer obtains a reward that is equal to the number of crimes in an intersection if the police officer monitors that intersection. In this way, we incentivize the police officers to monitor intersections with higher crime rates.

On the other hand, there exists a set of critical intersections (see the intersections with “*” next to the number of crimes in Figure 7). To ensure that the critical intersections are monitored often by a police officer, we consider a GTL specification for the critical intersections. Specifically, for each critical intersection s_{crit} , we assign a police officer i with the specification $\varphi = \square \diamond_{\leq 3} (s_{crit} \vee \exists^1 \bigcirc s_{crit})$. The specification φ means that each critical intersection s_{crit} should be visited at least once in every three time steps by either the police officer i or a neighboring police officer of i . This specification means that the critical intersections are monitored sufficiently often by a police officer, and it also aims to prevent the police officers from monitoring the intersections with the highest crime rates at all times.

We consider an example with 15 police officers, where each police officer is responsible for 9 intersections on a 3×3 grid. For example, a police officer is responsible for intersections between $5 \leq x_2 \leq 7$ and $3 \leq x_1 \leq 5$, and another police officer is responsible for intersections between $5 \leq x_2 \leq 7$ and $2 \leq x_1 \leq 4$. The objective of the police officers is to maximize the expected reward by monitoring states with the highest crime rates while satisfying φ .

For $\lambda = 0.9$, Figure 7 shows the assignment for the police officers with $\beta = 1$ after 500 iterations. The computation time of the approach was 164.3 seconds. Different colors at each intersection show the average number of police officers monitoring that intersection for a given hour. The results show the GTL specification φ is satisfied by the obtained assignment. We also observe that the police officers monitor the intersections with higher crime rates

¹The crime data can be found in <https://www.crimemapping.com/map/ca/sanfrancisco>.




Figure 7: The 35 intersections in the northeast of San Francisco. The map is obtained from Google maps. The numbering of the states starts from the bottom left corner and goes from left to right in every row. The number beside each location indicates the number of crimes in that area. We denote the critical intersections by labeling time with a “*” next to the number of crimes.

(critical or non-critical) to maximize the expected reward while satisfying the GTL specification φ . Without the GTL specification, the police officers would monitor the intersections with the highest crime rate to obtain maximal expected reward, and ignore the critical intersections. By enforcing the GTL specification, the results in Figure 7 shows that the critical intersections are visited sufficiently often (satisfying φ), while the intersections with higher crime rates are visited as often as possible to maximize the expected reward.

6 CONCLUSIONS AND FUTURE WORK

We proposed a method for the distributed synthesis of policies for multi-agent systems to implement spatial-temporal tasks. We expressed the spatial-temporal tasks in a specification language called graph temporal logic. For such systems, we decomposed the synthesis problem into a set of smaller synthesis problems, one for each agent. We showed that the algorithm runs in time linear in the number of agents and scales to hundreds of agents.

For future work, we will extend the framework to allow the edges of the graph to be time-varying. We will consider scenarios where the edges of the graph are functions of the agent’s local states, and the agents may share the state information with its neighboring agents in certain parts of the state spaces. For example, the agents may share their state information when they are close to each other and need to coordinate with each other to achieve a task.

REFERENCES

- [1] Christel Baier and Joost-Pieter Katoen. 2008. *Principles of Model Checking*. The MIT Press.
- [2] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, Jonathan Eckstein, et al. 2011. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Foundations and Trends® in Machine Learning* 3, 1 (2011), 1–122.
- [3] Stephen Boyd and Lieven Vandenbergh. 2004. *Convex Optimization*. Cambridge University Press, New York, NY, USA.
- [4] Caihua Chen, Bingsheng He, Yinyu Ye, and Xiaoming Yuan. 2016. The Direct Extension of ADMM for Multi-Block Convex Minimization Problems is not Necessarily Convergent. *Mathematical Programming* 155, 1-2 (2016), 57–79.
- [5] Qiang Cheng, Qiang Liu, Feng Chen, and Alexander T Ihler. 2013. Variational Planning for Graph-based MDPs. In *Advances in Neural Information Processing Systems*. 2976–2984.
- [6] M Choisy, N Peyrard, and R Sabbadin. 2007. A Probabilistic Decision Framework to Optimize the Dynamics of a Network Evolution: Application to the Control of Childhood Diseases. In *Proceedings of European Conference on Complex Systems (ECCS)*.
- [7] Florian Dörfler, Mihailo R Jovanović, Michael Chertkov, and Francesco Bullo. 2014. Sparsity-promoting Optimal Wide-area Control of Power Networks. *IEEE Transactions on Power Systems* 29, 5 (2014), 2281–2291.
- [8] Vojtěch Forejt, Marta Kwiatkowska, Gethin Norman, David Parker, and Hongyang Qu. 2011. Quantitative Multi-Objective Verification for Probabilistic Systems. In *Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 112–127.
- [9] Nicklas Forsell and Régis Sabbadin. 2006. Approximate Linear-programming Algorithms for Graph-based Markov Decision Processes. In *Proceedings of the 2006 conference on ECAI 2006: 17th European Conference on Artificial Intelligence August 29–September 1, 2006, Riva del Garda, Italy*. IOS Press, 590–594.
- [10] Graeme A Forster and Christopher A Gilligan. 2007. Optimizing the Control of Disease Infestations at the Landscape Scale. *Proceedings of the National Academy of Sciences* 104, 12 (2007), 4984–4989.
- [11] Daniel Gabay and Bertrand Mercier. 1975. *A Dual Algorithm for the Solution of Non-linear Variational Problems via Finite Element Approximation*. Institut de Recherche d’informatique et d’automatique.
- [12] Carlos Guestrin, Daphne Koller, and Ronald Parr. 2002. Multiagent Planning with Factored MDPs. In *Advances in Neural Information Processing Systems*. 1523–1530.
- [13] Carlos Guestrin, Daphne Koller, Ronald Parr, and Shobha Venkataraman. 2003. Efficient Solution Algorithms for Factored MDPs. *Journal of Artificial Intelligence Research* 19 (2003), 399–468.
- [14] LLC Gurobi Optimization. 2019. Gurobi Optimizer Reference Manual. (2019). <http://www.gurobi.com>
- [15] Iman Haghighi, Sadra Sadraddini, and Calin Belta. 2016. Robotic Swarm Control from Spatio-Temporal Specifications. In *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, 5708–5713.
- [16] Ernst Moritz Hahn, Mateo Perez, Sven Schewe, Fabio Somenzi, Ashutosh Trivedi, and Dominik Wojtczak. 2019. Omega-Regular Objectives in Model-Free Reinforcement Learning. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 395–412.
- [17] Michael L Littman, Ufuk Topcu, Jie Fu, Charles Isbell, Min Wen, and James MacGlashan. 2017. Environment-Independent Task Specifications via GLLT. *arXiv preprint arXiv:1704.04341* (2017).
- [18] Zhiyu Liu, Bo Wu, Jin Dai, and Hai Lin. 2017. Distributed Communication-Aware Motion Planning for Multi-Agent Systems from STL and SpaTeL Specifications. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE, 4452–4457.
- [19] Zhiyu Liu, Bo Wu, Jin Dai, and Hai Lin. 2018. Distributed Communication-Aware Motion Planning for Networked Mobile Robots Under Formal Specifications. *arXiv preprint arXiv:1812.01204* (2018).
- [20] Yurii Nesterov and Arkadii Nemirovskii. 1994. *Interior-Point Polynomial Algorithms in Convex Programming*. Vol. 13. Siam.
- [21] Nathalie Peyrard and Régis Sabbadin. 2006. Mean Field Approximation of the Policy Iteration Algorithm for Graph-based Markov Decision Processes. In *Proceedings of the 2006 conference on ECAI 2006: 17th European Conference on Artificial Intelligence August 29–September 1, 2006, Riva del Garda, Italy*. IOS Press, 595–599.
- [22] Daniel Pickem, Paul Glotfelter, Li Wang, Mark Mote, Aaron Ames, Eric Feron, and Magnus Egerstedt. 2017. The Robotarium: A Remotely Accessible Swarm Robotics Research Testbed. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 1699–1706.
- [23] Amir Pnueli. 1977. The Temporal Logic of Programs. In *Proc. 18th Annu. Symp. Found. Computer Sci.* Washington, D.C., USA, 46–57. <https://doi.org/10.1109/SFCS.1977.32>
- [24] Thomas P Prescott and Antonis Papachristodoulou. 2014. Layered Decomposition for the Model Order Reduction of Timescale Separated Biochemical Reaction Networks. *Journal of theoretical biology* 356 (2014), 113–122.
- [25] Martin L Puterman. 2014. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons.
- [26] Régis Sabbadin, Nathalie Peyrard, and Nicklas Forsell. 2012. A Framework and a Mean-field Algorithm for the Local Control of Spatial Processes. *International Journal of Approximate Reasoning* 53, 1 (2012), 66–86.
- [27] Yunus Emre Sahin, Petter Nilsson, and Necmiye Ozay. 2017. Provably-Correct Coordination of Large Collections of Agents with Counting Temporal Logic Constraints. In *Proceedings of the 8th International Conference on Cyber-Physical Systems*. ACM, 249–258.
- [28] Yagiz Savas, Melkior Ornik, Murat Cubuktepe, Mustafa O Karabag, and Ufuk Topcu. 2019. Entropy Maximization for Markov Decision Processes Under Temporal Logic Constraints. *IEEE Trans. Automat. Control* (2019).
- [29] Christos K Verginis and Dimos V Dimarogonas. 2018. Motion and Cooperative Transportation Planning for Multi-Agent Systems Under Temporal Logic Formulas. *arXiv preprint arXiv:1803.01579* (2018).
- [30] Xiangfeng Wang, Mingyi Hong, Shiqian Ma, and Zhi-Quan Luo. 2013. Solving Multiple-Block Separable Convex Minimization Problems Using Two-Block Alternating Direction Method of Multipliers. *arXiv preprint arXiv:1308.5294* (2013).
- [31] Bo Wu, Murat Cubuktepe, Suda Bharadwaj, and Ufuk Topcu. 2019. Reward-Based Deception with Cognitive Bias. *arXiv preprint arXiv:1904.11454* (2019).
- [32] Z. Xu, A. J Nettekoven, A. A. Julius, and U. Topcu. 2019. Graph Temporal Logic Inference for Classification and Identification. In *Proc. IEEE Conf. Decision and Control*.