

Adaptive Bi-Level Multi-Robot Task Allocation and Learning under Uncertainty with Temporal Logic Constraints

Xiaoshan Lin

University of Minnesota, Twin Cities
Minneapolis, United States
lin00668@umn.edu

Roberto Tron

Boston University
Boston, United States
tron@bu.edu

ABSTRACT

This work addresses the problem of multi-robot coordination under unknown robot transition models, ensuring that tasks specified by Time Window Temporal Logic are satisfied with user-defined probability thresholds. We present a bi-level framework that integrates (i) high-level task allocation, where tasks are assigned based on the robots' estimated task completion probabilities and expected rewards, and (ii) low-level distributed policy learning and execution, where robots independently optimize auxiliary rewards while fulfilling their assigned tasks. To handle uncertainty in robot dynamics, our approach leverages data collected during task execution to iteratively refine the expected task completion probabilities and rewards, enabling adaptive task allocation without explicit robot transition models. We theoretically validate the proposed algorithm by showing that it ensures tasks are completed at the desired probability thresholds with high confidence. Finally, we demonstrate the efficacy of our framework through comprehensive simulations.

KEYWORDS

Multi-Robot Systems; Task Allocation; Temporal Logic; Reinforcement Learning

ACM Reference Format:

Xiaoshan Lin and Roberto Tron. 2025. Adaptive Bi-Level Multi-Robot Task Allocation and Learning under Uncertainty with Temporal Logic Constraints. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025)*, Detroit, Michigan, USA, May 19 – 23, 2025, IFAAMAS, 9 pages.

1 INTRODUCTION

Efficient and high-quality task allocation plays a critical role in cooperative multi-robot applications, such as on-demand ridesharing and delivery [11], assembly lines [13], and warehouse logistics [43]. Recently, multi-robot task allocation with temporal constraints has attracted growing attention due to its significance in time-sensitive applications. For instance, [33] formulates a distributed constraint optimization problem to allocate tasks with time windows and ordering requirements to heterogeneous robots. [5] considers uncertainty in completing time-windowed tasks and combines high-level task allocation with low-level scheduling to minimize task incompleteness. [27] addresses tasks with precedence constraints, proposing an algorithm that ensures the probability of task failure

remains below a user-specified threshold under uncertainties in robot traits. Research in this area typically considers tasks that can be executed independently.

Multi-robot systems often involve more complicated scenarios where tasks are logically and temporally correlated. For instance, in warehouse logistics, robots need to coordinate to retrieve items from multiple storage locations based on orders, inventory, and replacement options, then assemble and deliver them within a set time frame, with each step depending on the successful completion of the previous one. Temporal logic (TL) provides a rigorous framework for specifying such temporal order and dependencies between tasks.

Different methods have been proposed to coordinate robots to satisfy a global temporal logic formula, including graph search on product automata [29], sampling-based methods [14, 22], and optimization-based approaches [2, 15, 42]. Other works assign local temporal logic tasks to individual robots, employing methods such as receding horizon planning [35], path-finding on product graph [8], and integer programming-based planning [32]. Nevertheless, these studies do not address uncertainty, which is crucial for time-sensitive tasks. For instance, environmental factors like traffic jams can disrupt task execution within the required time window.

Some prior studies [1, 3, 10, 12] employ reinforcement learning (RL) to maximize the probability of satisfying temporal logic constraints under unknown transition dynamics, primarily focusing on single-robot systems. Multi-agent reinforcement learning (MARL) has been explored for joint policy learning with temporal logic objectives [9, 39, 44]. Additionally, semi-decentralized methods [31], deep RL [24], and model-based RL [21] have been explored to improve learning efficiency. However, these MARL approaches either do not address system model uncertainty or assume full system knowledge, limiting their applicability in uncertain environments.

Overall, there are research gaps in existing multi-robot task allocation, planning, and learning approaches, as they either assume the tasks are independent, neglect the uncertainty in system models, or unrealistically assume full knowledge of uncertainties. Additionally, prior studies often aim to maximize robustness degree [24, 39], minimize incomplete tasks [5], or optimize some cost functions [14, 22]. While these approaches are effective, they often lack quantifiable guarantees on task satisfaction.

Motivated by these gaps, this paper addresses multi-robot task allocation to ensure a specified probability of satisfying Time Window Temporal Logic tasks under unknown transition dynamics. Additionally, a secondary objective is to maximize auxiliary reward functions for each robot based on user preferences. Fig. 1 illustrates an example where robots must transport materials within specific time windows with high probability while also performing auxiliary



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

tasks such as identifying and monitoring traffic-prone intersections or returning to stations for future assignments.

Our proposed formulation, which incorporates primary tasks and auxiliary objectives, has broader applicability in various scenarios, such as (i) **search-and-rescue**, where robots deliver supplies within time-critical windows while minimizing energy use or maximizing coverage to identify other areas needing help, and (ii) **environmental monitoring**, where drones monitor areas during specific intervals, such as wildlife peaks, while collecting additional data or conserving power for longer missions.

In summary, this work makes the following contributions:

- We propose a framework for multi-robot task allocation under unknown robot transition probabilities, integrating high-level task allocation with low-level policy learning to ensure probabilistic satisfaction of time-window temporal logic tasks.
- We incorporate auxiliary reward functions into high-level task allocation as a secondary objective, allowing the optimization of auxiliary rewards to impose user preferences without compromising the primary temporal logic tasks.
- We use data-driven methods to estimate expected rewards and task satisfaction probabilities of robots while following their low-level policies, which enable adaptive high-level task allocation using these refined reward functions and satisfaction probabilities without requiring explicit robot transition models.

The most relevant works in the literature are [1, 5, 20, 27, 30, 31]. [5] employs a similar bi-level framework but assumes independent tasks with single time windows and known task completion probabilities, whereas we consider complex spatio-temporal dependencies and unknown task completion probabilities. In [27], the Sequential Probability Ratio Test (SPRT) is used to verify whether the probability of task failure remains below a user-defined threshold. This approach iteratively recomputes task assignments until passing the SPRT, whereas we formulate an optimization problem whose solution guarantees probabilistic task satisfaction. Both [1] and [20] consider similar problem formulations, aiming to satisfy a temporal logic task with a desired probability. While these works focus on a single agent with a single task, we address multi-robot systems with multiple tasks. [31] considers a similar communication model in which each agent interacts with a centralized coordinator. However, their approach involves the coordinator in lower-level policy learning, whereas the coordinator is solely used for high-level task allocation in our framework. [30] uses RL to estimate uncertain task execution durations and incorporates the learned durations into an auction-based allocation policy to minimize execution time. Similarly, our approach uses RL to estimate the reward functions and improve task allocation. Our work specifically focuses on satisfying tasks within exact time windows with desired probabilities, which their work does not address.

2 PRELIMINARIES

Let Σ be a finite set. We denote the power set of Σ by 2^Σ . A finite or infinite sequence of elements from Σ is called a word over Σ . In this context, Σ is also called an alphabet. Let $k, i, j \in \mathbb{Z}_{\geq 0}$ with $i \leq j$. The k -th element of a word σ is denoted by σ_k , and the




Figure 1: Motivating example of our proposed framework. Primary tasks: repeatedly transporting resources from warehouses to processing stations and then to operation site within specified time windows, which must be satisfied with a high probability to prevent resource accumulation. Auxiliary tasks: monitoring traffic congestion to improve future routing and delivery times, or returning to stations after completing the primary tasks.

subword $\sigma_i, \dots, \sigma_j$ is denoted by $\sigma_{i,j}$. We denote the set of atomic propositions by AP .

2.1 Time Window Temporal Logic

Time Window Temporal Logic (TWTL) [38] is a language for expressing time-bounded specifications. A TWTL formula is defined over a set of atomic propositions AP as follows:

$$\phi ::= H^d s \mid H^d \neg s \mid \neg \phi_1 \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \mid \phi_1 \cdot \phi_2 \mid [\phi_1]^{[a,b]}.$$

Here, s represents either the constant “true” or an atomic proposition in AP ; ϕ_1 and ϕ_2 are TWTL formulas. The hold operator $H^d s$, with $d \in \mathbb{Z}_{\geq 0}$, specifies that $s \in AP$ should hold for d time units. The negation operator $\neg \phi_1$ specifies “do not satisfy the formula”. The conjunction operator $\phi_1 \wedge \phi_2$ and disjunction operator $\phi_1 \vee \phi_2$ specify “satisfy both formulas”, and “satisfy at least one formula”, respectively. The concatenation operator $\phi_1 \cdot \phi_2$ specifies that ϕ_1 must be satisfied first, and ϕ_2 must be satisfied immediately after. The within operator $[\phi]^{[a,b]}$, where $a, b \in \mathbb{Z}_{\geq 0}$ and $a \leq b$, restricts the satisfaction of ϕ to the time window $[a, b]$. The **time bound** of a TWTL formula ϕ , denoted as $\|(\phi)\|$, represents the maximum time allowed to satisfy ϕ . For a formal definition of the TWTL semantics and the time bound, we refer readers to [38].

Definition 2.1. (Deterministic Finite-State Automaton) A deterministic finite-state automaton (DFA) is a tuple $\mathcal{A} = (Q, q_0, \Sigma, \delta, F)$, where Q is a finite set of states, q_0 is the initial state, $\Sigma = 2^{AP}$ is the input alphabet, $\delta : Q \times \Sigma \rightarrow Q$ is the transition function, and F is the set of accepting states.

A finite input word $\sigma = \sigma_0, \sigma_1, \dots, \sigma_T$ over the alphabet 2^{AP} generates a trajectory $\mathbf{q} = q_0, q_1, \dots, q_{T+1}$ on the DFA, where q_0 is the initial state of the DFA and $q_{k+1} = \delta(q_k, \sigma_k)$ for all $k \geq 0$. A finite input word σ over Σ is accepted by a DFA if the corresponding trajectory \mathbf{q} ends in an accepting state of the DFA.

A TWTL formula ϕ can be translated into a DFA that either accepts or rejects an input word [38]. An input word σ is said to satisfy the corresponding TWTL formula ϕ if it is accepted by the DFA, denoted as $\sigma \models \phi$. For example, the TWTL formula $\phi = [H^1 P]^{[1,2]} \cdot [H^0 D]^{[0,2]}$, specifying a pickup (P) and hold for one time step within time window $[1, 2]$ followed by a delivery (D)

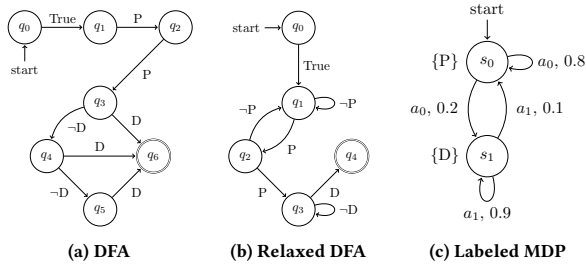


Figure 2: (a) DFA for TWTL formula $\phi = [H^1 P]^{[1,2]} \cdot [H^0 D]^{[0,2]}$. (b) DFA for the temporally relaxed TWTL formula $\phi(\tau) = [H^1 P]^{[1,2+\tau_1]} \cdot [H^0 D]^{[0,2+\tau_2]}$. (c) Example of a labeled MDP, where $S = \{s_0, s_1\}$, $A = \{a_0, a_1\}$, $AP = \{P, D\}$, $I(s_0) = \{P\}$ and $I(s_1) = \{D\}$.

within two time steps, is translated into a DFA shown in Fig. 2a. An input word $\sigma = \{P\}, \{P\}, \{P\}, \{D\}$ satisfies the TWTL formula, as it generates a trajectory $q = q_0, q_1, q_2, q_3, q_6$ that ends in the accepting state q_6 . This TWTL formula can also be translated into a temporally relaxed DFA (see Fig. 2b), which is more compact and facilitates more efficient computation and construction.

2.2 Markov Decision Process

Definition 2.2. (Labeled MDP) A labeled MDP is a tuple $\mathcal{M} = (S, A, \Delta, R, I)$, where S represents the state space, and A denotes the set of actions. The probabilistic transition function is given by $\Delta : S \times A \times S \rightarrow [0, 1]$, and the reward function is defined as $R : S \times A \rightarrow \mathbb{R}$. Additionally, $I : S \rightarrow 2^{AP}$ is a labeling function that maps each state to a set of atomic propositions.

Definition 2.3 (Policy). A policy over the MDP is a mapping $\pi : S \rightarrow P(A)$, where $P(A)$ represents the set of probability distributions over the actions set A . A policy is stationary if $\pi(\cdot|s) \in P(A)$ does not change over time, and it is deterministic if $\pi(\cdot|s)$ always selects the same action for a given state.

MDPs with labeling functions are widely used in RL with temporal logic [4, 6, 25] to allow verification of temporal logic satisfaction. Given a finite MDP trajectory $s = s_0, s_1, \dots, s_T$, the labeling function generates an output word $\mathbf{o} = o_0, o_1, \dots, o_T$ where $o_i = I(s_i)$. A subword of the output, $o_{t_1}, o_{t_1+1}, \dots, o_{t_2}$, is denoted as \mathbf{o}_{t_1, t_2} . Consider the labeled MDP example shown in Fig. 2c. An MDP trajectory $s = s_0, s_0, s_0, s_1$ generates an output word $\mathbf{o} = \{P\}, \{P\}, \{P\}, \{D\}$, which serves as an input word to the DFA to verify whether the robot satisfies the TWTL formula.

2.3 System Setup

We consider a multi-robot system with a set of N robots, denoted as $\{r_i\}_{i \in \{1, 2, \dots, N\}}$, and we refer to this set as $\{N\}$ for simplicity. Similarly, we define a set of K tasks, denoted by $\{t_k\}_{k \in \{1, 2, \dots, K\}}$, and refer to it as $\{K\}$. Each task $t_k \in \{K\}$ is defined by a tuple (ϕ_k, P_k) , where ϕ_k is a TWTL formula and P_k is the desired probability threshold for satisfying ϕ_k . The time bound for these tasks is $T = \|\phi_k\|$ time steps, and they must be fulfilled every T time steps over an infinite horizon or a long period. Each repetition of T time steps is called an **episode**. We introduce a virtual task $t_{K+1} = (\phi_{K+1}, P_{K+1})$, where ϕ_{K+1} is a null task and $P_{K+1} = 0$, to augment the task set. A

robot assigned to task t_{K+1} is considered unassigned to any TWTL task, allowing it to focus on its auxiliary objective. We denote this augmented task set as $\{K + 1\}$.

We assume that the robots in $\{N\}$ have decoupled dynamics and reward functions, and operate independently without interactions with others. Each robot $r_i \in \{N\}$ is modeled as a decoupled MDP, denoted by $\mathcal{M}^i = (S^i, A^i, \Delta^i, R^i, I^i)$. The collection of all robot MDPs $\{\mathcal{M}^i\}_{i \in \{N\}}$ is referred to as $\{\mathcal{M}\}$. The state transition probabilities Δ^i are unknown and account for uncertainties such as movement delays due to road conditions or potential robot failures. The reward function R^i encodes the robot's individual auxiliary objectives.

The output word of robot r_i over a time interval $[t_1, t_2]$ is denoted as \mathbf{o}_{t_1, t_2}^i . **Team output word**, the collection of output words from all robots $\{\mathbf{o}_{t_1, t_2}^i\}_{i \in \{N\}}$, is given by $\{\mathbf{o}_{t_1, t_2}\}$. Given a TWTL formula ϕ , we say $\{\mathbf{o}_{t_1, t_2}\} \models \phi$ if and only if there exists at least one robot r_i such that $\mathbf{o}_{t_1, t_2}^i \models \phi$. The probability that $\{\mathbf{o}_{t_1, t_2}\}$ satisfies a TWTL formula ϕ is denoted as $\Pr(\{\mathbf{o}_{t_1, t_2}\} \models \phi)$.

Guaranteeing probabilistic constraint satisfaction is unachievable without prior knowledge of the MDP transition probabilities. In this work, we assume that although the exact transition probabilities are unknown, the agent has partial knowledge: when taking action a at state s , the agent knows (i) which states have a non-zero probability of being reached and (ii) which states have a sufficiently high probability of being the next state. This partial knowledge is represented by ϵ^i , an upper bound on transition uncertainty, which is crucial for computing a lower bound on the probability of task satisfaction. We formalize this assumption as follows:

Assumption 1. Given an MDP \mathcal{M}^i for an arbitrary robot r_i , the exact transition probability Δ^i is unknown. However, a known parameter $\epsilon^i \in [0, 1]$ exists such that, for each state s and action a , (i) the set of states s' for which $\Delta^i(s, a, s') > 0$ and (ii) the set of states s'' for which $\Delta^i(s, a, s'') \geq 1 - \epsilon^i$ are non-empty and known.

Example 2.4. Consider a probabilistic transition model shown in Fig. 3a, where the action set is given by $A = \{N, NE, E, SE, S, SW, W, NW, Stay\}$. Under the 'Stay' action, robots remain stationary, while all other actions result in an intended transition (black arrows) with probability $1 - \epsilon$ and unintended transitions (red arrows) with a probability of ϵ . The exact values of ϵ are unknown to robots. However, each robot r_i is provided with an estimate ϵ^i , which overestimates the actual ϵ and reflects the robots' limited knowledge of their transition models and environmental disturbances, as stated in Assumption 1.

3 PROBLEM FORMULATION

Given a multi-robot system $\{N\}$ and a set of tasks $\{K + 1\}$, the objective is to find an optimal policy for each MDP in $\{\mathcal{M}\}$ to ensure that the team output word satisfies each TWTL formula ϕ_k with at least probability P_k while maximizing the sum of rewards. With unknown MDP transitions, this can be formulated as a constrained RL problem, with the probabilistic satisfaction of TWTL tasks encoded as constraints, and reward maximization as the objective. This problem presents multiple challenges, including coordinating robots under uncertain dynamics, ensuring probabilistic satisfaction of TWTL tasks, and maximizing individual rewards.

We propose a bi-level solution to address these challenges hierarchically. At the high level, we solve for multi-robot task allocation to ensure desired probabilistic satisfaction of the TWTL tasks. Specifically, we consider task allocation $\Pi : \{N\} \times \{K+1\} \rightarrow [0, 1]$, where $\Pi(i, k)$ represents the probability of robot r_i being assigned to task t_k . At the low level, each robot independently finds its own policy π_i^Π , to satisfy its assigned task while simultaneously maximizing its individual reward. We formally define the problem as below.

Problem 1. *Given a set of robots $\{N\}$, a set of TWTL tasks $\{K+1\}$, a set of MDP $\{\mathcal{M}\}$ subject to Assumption 1, and a discount factor γ , determine the optimal high-level task allocation and low-level policies*

$$\Pi^*, \pi_1^*, \dots, \pi_N^* = \arg \max_{\Pi, \pi_1^\Pi, \dots, \pi_N^\Pi} \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \sum_{i=1}^N R^i(s_t^i, a_t^i) \right], \quad (1)$$

such that, for every episode j ,

$$\Pr(\{\mathbf{o}_{jT, jT+T}\} \models \phi_k) \geq P_k, \quad \forall j \geq 0, (\phi_k, P_k) \in \{K+1\} \quad (2)$$

where $\{\mathbf{o}_{jT, jT+T}\}$ is the team output word in episode j .

Remark 1. *The reward functions can incorporate different auxiliary objectives into task allocations based on system requirements. (i) **Environment-driven vs. user-defined:** robots can learn rewards from environmental interactions (e.g., identifying high-traffic areas) or follow predefined incentives (e.g., encouraging timely returns to charging stations). (ii) **Robot-specific objectives:** reward functions can be designed to reflect the distinct roles of different robots (e.g. aerial robots prioritize large area coverage, while mobile robots remain in critical locations to improve future task efficiency). (iii) **Capability-aware rewards:** rewards encode robot capabilities, such as higher rewards for drones with superior sensors when monitoring key areas.*

4 PROPOSED SOLUTION

Our approach consists of a high-level coordinator for task allocation and a low-level task execution mechanism. The high-level coordinator considers each robot's probability of satisfying TWTL tasks from its current state and the expected rewards associated with specific tasks. It computes task assignments that maximize the total expected reward while ensuring that each task is completed with the desired probability threshold. At the low level, each robot maintains $K+1$ policies: K stationary policies, one for each TWTL task, and an additional policy learned over time to maximize individual rewards when not assigned to a TWTL task. Upon receiving an assignment from the high-level coordinator, each robot follows the corresponding policy for its designated task.

4.1 High-level Multi-robot Task Allocation

Before each episode, each robot provides the coordinator with its $P_{i,k}^\epsilon(p_i)$, the expected probability of satisfying TWTL formula ϕ_k , and $V_{i,k}(p_i)$, the expected reward under its current policy. Here, p_i denotes the current state of the robot. For simplicity, we refer to these quantities as $P_{i,k}^\epsilon$ and $V_{i,k}$ in the following sections.

For now, we assume that $P_{i,k}^\epsilon$ and $V_{i,k}$ are known to the robots. The process of obtaining these values will be discussed in later sections. Given a set of tasks $\{K+1\}$, the high-level task assignment for robot $r_i \in \{N\}$ is given by the probability distribution $\{P_{i,k}\}_{k=1,2,\dots,K+1}$, where $\sum_{k=1}^{K+1} P_{i,k} = 1$. For $k = 1, 2, \dots, K$, $P_{i,k}$

denotes the probability of robot r_i being assigned to TWTL task t_k , while $P_{i,K+1}$ denotes the probability that robot r_i prioritizes maximizing its reward function instead of executing a TWTL task. Following the multi-robot task allocation taxonomy proposed by [7], our task allocation problem falls under ST-SR-IA, i.e. single-task robots (ST) execute single-robot tasks (SR), with instantaneous assignments (IA). Note that although the stochastic task assignments allow multiple robots to be assigned to the same task, we assume they execute the tasks in parallel rather than collaboratively. Therefore, tasks are classified as single-robot rather than multi-robot. The high-level task allocation is determined by solving the following optimization problem.

$$\arg \max_{P_{i,k}} \sum_{i=1}^N \sum_{k=1}^{K+1} P_{i,k} \cdot V_{i,k} \quad (3a)$$

$$\text{subject to } 1 - \prod_{i=1}^N (1 - P_{i,k} \cdot P_{i,k}^\epsilon) \geq P_k \quad \forall k = 1, 2, \dots, K \quad (3b)$$

$$\sum_{k=1}^{K+1} P_{i,k} = 1 \quad \forall i = 1, 2, \dots, N \quad (3c)$$

$$P_{i,k} \geq 0 \quad \forall i = 1, 2, \dots, N, \quad k = 1, 2, \dots, K+1 \quad (3d)$$

In problem (3), the objective function (3a) aims to maximize the sum of the expected rewards across all robots. Constraint (3b) ensures that the probability of at least one robot satisfying task t_k is not less than its desired probability P_k . Constraints (3c) and (3d) ensure that the task assignment $\{P_{i,k}\}$ define a valid probability distribution over tasks for each robot. Although the assignments are stochastic, each robot executes only one task at a time, sampled from the task assignment. The stochastic assignment generalizes the deterministic solutions commonly adopted in multi-agent task allocation, allowing better task allocations that optimize auxiliary rewards while ensuring TWTL satisfaction. The deterministic version can be achieved by restricting $\{P_{i,k}\}$ to either 0 or 1.

The task completion probability $P_{i,k}^\epsilon$ in (3b) is unknown due to the lack of knowledge about the transition probabilities in the MDPs. One solution is to substitute $P_{i,k}^\epsilon$ with its lower bound. The key question is whether solving problem (3) using lower bounds of $P_{i,k}^\epsilon$ still ensures satisfaction of the original constraint (3b). To answer this, we present the following proposition.

PROPOSITION 4.1. *Let $\lfloor P_{i,k}^\epsilon \rfloor$ be an arbitrary lower bound for $P_{i,k}^\epsilon$, that is, $0 \leq \lfloor P_{i,k}^\epsilon \rfloor \leq P_{i,k}^\epsilon \leq 1$. If a set of probabilities $\{P_{i,k}\}$ satisfies*

$$1 - \prod_{i=1}^N (1 - P_{i,k} \cdot \lfloor P_{i,k}^\epsilon \rfloor) \geq P_k \quad \forall k = 1, 2, \dots, K, \quad (4)$$

then $\{P_{i,k}\}$ also satisfies constraint (3b).

PROOF. Since $P_{i,k}^\epsilon \geq \lfloor P_{i,k}^\epsilon \rfloor$, it follows that:

$$\prod_{i=1}^N (1 - P_{i,k} \cdot P_{i,k}^\epsilon) \leq \prod_{i=1}^N (1 - P_{i,k} \cdot \lfloor P_{i,k}^\epsilon \rfloor). \quad (5)$$

$$\Rightarrow 1 - \prod_{i=1}^N (1 - P_{i,k} \cdot P_{i,k}^\epsilon) \geq 1 - \prod_{i=1}^N (1 - P_{i,k} \cdot \lfloor P_{i,k}^\epsilon \rfloor) \geq P_k. \quad (6)$$

□

This proposition guarantees that replacing $P_{i,k}^\epsilon$ with its lower bound in (3) preserves the original task satisfaction constraints (3b).

4.2 Low-level Single-Agent Task Execution

Upon receiving its task assignment $\{P_{i,k}\}_{k=1,2,\dots,K+1}$ from the high-level coordinator, robot r_i samples from this probability distribution to determine which task it will execute in the upcoming episode. The robot's policy during the episode depends on whether the selected task t_k is a TWTL task ($1 \leq k \leq K$), or the auxiliary task focused on maximizing the robot's individual reward ($k = K + 1$).

A. Single-Agent Policy for TWTL Tasks ($1 \leq k \leq K$).

In this case, the robot follows a stationary policy that maximizes the probability of satisfying the TWTL task t_k . We discuss the methodology for deriving this policy in the following section.

Definition 4.2. (Product MDP) Given a labeled MDP $\mathcal{M} = (S, A, \Delta, R, I)$ and an DFA $\mathcal{A} = (Q, q_0, \Sigma, \delta, F)$, a product MDP is a tuple $\mathcal{P} = \mathcal{M} \times \mathcal{A} = (S^\otimes, S_0^\otimes, A^\otimes, \Delta^\otimes, R^\otimes, F^\otimes)$, where $S^\otimes = S \times Q$ is a finite set of states; $S_0^\otimes = \{(s, \delta(q_0, l(s))) \mid \forall s \in S\}$ is the set of initial states; $A^\otimes = A$ is the set of actions; $\Delta^\otimes : S^\otimes \times A^\otimes \times S^\otimes \rightarrow [0, 1]$ is the probabilistic transition relation such that for any two states, $p = (s, q) \in S^\otimes$ and $p' = (s', q') \in S^\otimes$, and any action $a \in A^\otimes$, $\Delta^\otimes(p, a, p') = \Delta(s, a, s')$ if $\delta(q, l(s)) = q'$; $R^\otimes : S^\otimes \times A^\otimes \rightarrow \mathbb{R}$ is the reward function such that $R^\otimes(p, a) = R(s, a)$ for $p = (s, q) \in S^\otimes$ and $a \in A^\otimes$; $F^\otimes = (S \times F) \subseteq S^\otimes$ is the set of accepting states.

To maximize the probability of reaching F^\otimes from any state of the product MDP, we can select the action that is most likely to minimize the expected distance to any accepting states. Inspired by [1, 20], we define ϵ -stochastic transitions and distance-to- F^\otimes . With a slight abuse of notation, we use ϵ as a generalized parameter in the following definitions, where for each robot r_i , ϵ should be replaced with the robot-specific parameter ϵ^i , as stated in Assumption 1.

Definition 4.3 (ϵ -Stochastic Transitions). For a given product MDP and some $\epsilon \in [0, 1]$, any transition (p, a, p') such that $\Delta^\otimes(p, a, p') \geq 1 - \epsilon$ is defined as an ϵ -stochastic transition.

Definition 4.4 (Distance-To- F^\otimes). Given a product MDP, the distance from any state $p \in S^\otimes$ to the set of accepting states F^\otimes is

$$D^\epsilon(p) = \min_{p' \in F^\otimes} \text{dist}^\epsilon(p, p'), \quad (7)$$

where $\text{dist}^\epsilon(p, p')$ represents the minimum number of ϵ -stochastic transitions to move from p to another state p' .

Here, we label transitions with sufficiently high probabilities ($\geq 1 - \epsilon$) as ϵ -stochastic, and define $D^\epsilon(p)$ as the minimum ϵ -stochastic transitions needed from state p to the set of accepting states. Next, we use $D^\epsilon(p)$ to define a policy that minimizes the ϵ -stochastic transitions required to reach an accepting state.

Definition 4.5 (π^ϵ Policy). Given a product MDP and $\epsilon \in [0, 1]$, $\pi^\epsilon : S^\otimes \rightarrow A$ is a stationary policy such that

$$\pi^\epsilon(p) = \arg \min_{a \in A} D_{\min}^\epsilon(p, a), \quad (8)$$

where $D_{\min}^\epsilon(p, a) = \min_{p' : \Delta^\otimes(p, a, p') \geq 1 - \epsilon} D^\epsilon(p')$, i.e., the minimum distance-to- F^\otimes among the states reachable from p under action a with probability of at least $1 - \epsilon$.

We summarize the procedure for synthesizing policy π^ϵ in Alg. 1. The inputs are the set of tasks and an MDP \mathcal{M}^i that adheres to Assumption 1 with parameter ϵ^i . First, a DFA is generated for each

TWTL formula (line 2). Then, for robot r_i and its corresponding MDP, a product MDP is constructed (line 3). The product MDP is used to calculate the Distance-To- F^\otimes for all states using ϵ^i -stochastic transitions (line 4). Finally, the $\pi_{i,k}^\epsilon$ policy is computed for this specific product MDP by selecting the action that minimizes the Distance-To- F^\otimes for each state (lines 5-6).

Algorithm 1: Off-line computation of π^ϵ policy

Input : A set of tasks $\{K + 1\}$; a MDP \mathcal{M}^i
Output : Stationary policies $\{\pi_{i,k}^\epsilon\}$; product MDPs $\{\mathcal{P}_{i,k}\}$

```

1 for  $t_k = (\phi_k, P_k) \in \{K + 1\}$ ,  $k \leq K$  do
2    $\mathcal{A}_k \leftarrow$  Create DFA of TWTL formula  $\phi_k$ ;
3    $\mathcal{P}_{i,k} = \mathcal{M}^i \times \mathcal{A}_k \leftarrow$  Create product MDP;
4   Calculate Distance-To- $F^\otimes$  for all states in  $\mathcal{P}_{i,k}$ ;
5   for  $p \in$  all states of  $\mathcal{P}_{i,k}$  do
6      $\pi_{i,k}^\epsilon(p) \leftarrow (8)$ 
```

Recall that we want to replace $P_{i,k}^\epsilon$ with its lower bound in problem (3). Now we will discuss methods for obtaining its lower bounds. By definition, $P_{i,k}^\epsilon$ is the probability that robot r_i , under stationary policy $\pi_{i,k}^\epsilon$, satisfies the TWTL formula ϕ_k . Equivalently, it is the probability of reaching the accepting states of the product MDP from an initial state under $\pi_{i,k}^\epsilon$. We adopt the method proposed in [20], which finds the lower bound of the probability of reaching accepting states within a time window, given the upper bounds of the transition uncertainties. We denote this lower bound as $\underline{P}_{i,k}^\epsilon$ and refer to it as the **static lower bound**. We define $\{\underline{P}_{i,k}^\epsilon\}$ as the set of static lower bounds. Alternatively, another approach is to use the partial information from Assumption 1 to estimate the upper and lower bounds of the transition probabilities in the MDP. These bounds can then be used with probabilistic model checkers such as PRISM [17] or the optimization-based method proposed in [18] to compute the lower bounds.

While the static lower bound can be computed offline using partial information ϵ^i without exact transition probabilities, it may be overly conservative when ϵ^i significantly overestimates the true uncertainty. To reduce conservativeness, we introduce a second lower bound that is refined in real time as robots interact with the environment. This data-driven method allows the lower bounds of $P_{i,k}^\epsilon$ to be adaptively refined for greater accuracy.

For robot r_i with an initial state p_i , the outcomes of the policy $\pi_{i,k}^\epsilon$ (either satisfying or violating the TWTL constraint) can be modeled as a Bernoulli distribution with the probability of success equal to $P_{i,k}^\epsilon(p_i)$. To estimate $P_{i,k}^\epsilon(p_i)$, we use the Wilson score interval [41], which computes a confidence interval $[P_{i,k}^{\epsilon-}(p_i), P_{i,k}^{\epsilon+}(p_i)]$ that bounds $P_{i,k}^\epsilon(p_i)$ with a specified confidence level, where

$$P_{i,k}^{\epsilon-}(p_i) = \frac{n_{i,k}^S(p_i) + \frac{1}{2}z^2}{n_{i,k}(p_i) + z^2} - \frac{z}{n_{i,k}(p_i) + z^2} \sqrt{\frac{n_{i,k}^S(p_i)n_{i,k}^F(p_i)}{n_{i,k}(p_i)} + \frac{z^2}{4}}. \quad (9)$$

Here, $n_{i,k}(p_i)$ represents the number of episodes in which robot r_i started at p_i and adopted $\pi_{i,k}^\epsilon$. The number of episodes that satisfied or violated the constraint are denoted by $n_{i,k}^S(p_i)$ and $n_{i,k}^F(p_i)$, respectively. Thus, $n_{i,k}(p_i) = n_{i,k}^F(p_i) + n_{i,k}^S(p_i)$. The value of z is

chosen based on the desired confidence level (e.g., 99% confidence level corresponds to $z = 2.58$). By selecting a sufficiently high value of z , we ensure that $P_{i,k}^\epsilon(p_i) \geq P_{i,k}^{\epsilon^-}(p_i)$ with high confidence.

In the initial episodes, we rely on the static lower bound $P_{i,k}^{\epsilon^-}$, as the confidence lower bound $P_{i,k}^{\epsilon^-}$ is initially far from the true probability $P_{i,k}^\epsilon$ due to the limited data available. After each episode, we update $P_{i,k}^\epsilon(p_i)$ for robot r_i and the task t_k it executed. As robots collect more data, $P_{i,k}^{\epsilon^-}$ becomes a more accurate estimate of $P_{i,k}^\epsilon$ and can gradually replace the static lower bound. The set of **confidence lower bounds** is denoted as $\{P_{i,k}^{\epsilon^-}\}$.

Finally, we discuss how to obtain $V_{i,k}$ for problem (3). As defined earlier, $V_{i,k}$ represents the expected reward for robot r_i under its current policy for task t_k and is therefore the state value function of the MDP under $\pi_{i,k}^\epsilon$, a stationary policy defined in Definition 4.5 and Alg. 1. We estimate $V_{i,k}$ using the tabular TD(0) method [34].

B. Single-Agent Policy for Reward Maximization ($k = K + 1$).

In this case, the robot is not assigned to any TWTL tasks and exclusively focuses on exploring the environment to maximize its reward function. In this work, we use the tabular Q-learning algorithm [40] to learn an optimal policy that maximizes the expected discounted reward. However, any RL algorithm can be used for this purpose. The Q-function is updated as follows:

$$Q_{i,k}(s_t^i, a_t^i) \leftarrow Q_{i,k}(s_t^i, a_t^i) + \alpha \left(r_{t+1}^i + \gamma \max_{a'} Q_{i,k}(s_{t+1}^i, a') - Q_{i,k}(s_t^i, a_t^i) \right), \quad (10)$$

where s_t^i , a_t^i , r_{t+1}^i , and γ denote the state, action, reward, and discount factor, respectively. At each state, the value function $V_{i,K+1}$ is the maximum Q-value over all available actions at that state.

C. Combined Single-Agent Policy.

We summarize the overall single-agent policy in Alg. 2. The algorithm begins by constructing the product MDPs and computing π^ϵ policies (line 2). Then it initializes the variables for the Wilson score lower bound (line 5), computes the static lower bound (line 6), initializes a value function table (line 7) and a Q-table (line 8), and its corresponding policy (line 9). Upon receiving task allocation $\{P_{i,k}\}$ from the high-level coordinator, a task is sampled from this probability distribution (line 11), and the robot's initial product MDP state is set accordingly (lines 12-13). Depending on the selected task, the robot either (i) select actions from policy $\pi_{i,k}^\epsilon$ and update the value function $V_{i,k}$ (lines 15-19) if executing a TWTL task ($k \leq K$) or (ii) learn a policy to maximize rewards via Q-learning (lines 20-25), if in an exploration phase ($k = K + 1$) or if the TWTL constraint is met before the episode ends. After executing the task, the robot updates the success or failure counts and the confidence lower bound (line 26).

Function `stats()` is used to provide the expected rewards $\{V_{i,k}\}$, lower bounds $\{P_{i,k}^\epsilon\}$ and $\{P_{i,k}^{\epsilon^-}\}$ to the high-level coordinator (lines 33-39). The set $\{P_{i,k}^{\epsilon^-}\}$ initially consists of static lower bounds $P_{i,k}^{\epsilon^-}$, which are gradually replaced by the confidence lower bounds $P_{i,k}^{\epsilon^-}$ as more data is collected (line 37). In line 37, N is a user-defined data count threshold for switching from static to confidence lower bounds. We refer to $\{P_{i,k}^{\epsilon^-}\}$ as a set of **adaptive lower bounds**.

Algorithm 2: Single-agent Task Execution

Input : A set of tasks $\{K + 1\}$; a set of policies $\pi_{i,k}^\epsilon$; MDP \mathcal{M}^i for robot r_i ; episode length T ; initial MDP state s_0

```

1 Initialization:
2  $\{\pi_{i,k}^\epsilon\}, \{\mathcal{P}_{i,k}\} \leftarrow \text{Alg. 1}$ 
3 for  $\mathcal{P}_{i,k} \in \{\mathcal{P}_{i,k}\}_{k=1,2,\dots,K}$  do
4   for  $p \in \text{all initial states of } \mathcal{P}_{i,k}$  do
5      $n_{i,k}(p) \leftarrow 0, n_{i,k}^S(p) \leftarrow 0, n_{i,k}^F(p) \leftarrow 0$ 
6      $P_{i,k}^{\epsilon^-} \leftarrow \text{static lower bound (proposed in [20])}$ 
7    $V_{i,k}(p) \leftarrow 0$  for  $p \in \text{all states of } \mathcal{P}_{i,k}$ 
8  $Q(s, a) \leftarrow 0$  for  $(s, a) \in \text{all state-action pairs of } \mathcal{M}^i$ 
9  $\pi(s) \leftarrow \arg \max_a Q(s, a)$ 
10 Function execute( $\{P_{i,k}\}_{k=1,2,\dots,K+1}$ ):
11    $k \leftarrow \text{Sampled from } \{P_{i,k}\}_{k=1,2,\dots,K+1}$ 
12    $p \leftarrow \bar{p} \in \text{initial states of } \mathcal{P}_{i,k} \text{ s.t. mdp\_state}(\bar{p}) = s_0$ 
13    $p_0 \leftarrow p, s \leftarrow s_0$ 
14   for  $t = 1 : T$  do
15     if not constraint_satisfied and  $k \leq K$  then
16        $a \leftarrow \text{Select action from policy } \pi_{i,k}^\epsilon(p)$ 
17       Take action  $a$ , observe  $p'$  (next state) and  $r$  (reward)
18        $V_{i,k}(p) \leftarrow V_{i,k}(p) + \alpha (r + \gamma V_{i,k}(p') - V_{i,k}(p))$ 
19        $p \leftarrow p', s \leftarrow \text{mdp\_state}(p')$ 
20     else if constraint_satisfied or  $k = K + 1$  then
21        $a \leftarrow \text{Select action from } \pi(s) \text{ via } \epsilon\text{-greedy}$ 
22       Take action  $a$ , observe  $s'$  (next state) and  $r$  (reward)
23        $Q(s, a) \leftarrow Q(s, a) + \alpha (r + \gamma \max_{a'} Q(s', a') - Q(s, a))$ 
24       Update policy  $\pi(s)$ 
25        $s \leftarrow s'$ 
26   update( $p_0, n_{i,k}(p_0), n_{i,k}^S(p_0), n_{i,k}^F(p_0), \text{constraint\_satisfied}$ )
27    $s_0 \leftarrow s$ 
28 Function update( $p_0, n_{i,k}(p_0), n_{i,k}^S(p_0), n_{i,k}^F(p_0), \text{satisfied}$ ):
29    $n_{i,k}(p_0) \leftarrow n_{i,k}(p_0) + 1$ 
30   if satisfied then  $n_{i,k}^S(p_0) \leftarrow n_{i,k}^S(p_0) + 1$ 
31   else  $n_{i,k}^F(p_0) \leftarrow n_{i,k}^F(p_0) + 1$ 
32    $P_{i,k}^{\epsilon^-}(p_0) \leftarrow \text{equation (9)}$ 
33 Function stats():
34   for  $k = 1 : K$  do
35      $p \leftarrow \bar{p} \in \text{initial states of } \mathcal{P}_{i,k} \text{ s.t. mdp\_state}(\bar{p}) = s_0$ 
36      $V_{i,k} \leftarrow V_{i,k}(p), P_{i,k}^\epsilon \leftarrow P_{i,k}^\epsilon(p)$ 
37      $P_{i,k}^{\epsilon^-} \leftarrow P_{i,k}^{\epsilon^-}$  if  $n_{i,k}(p) \geq N$  else  $P_{i,k}^{\epsilon^-} \leftarrow P_{i,k}^{\epsilon^-}$ 
38    $V_{i,K+1} \leftarrow \max_a Q(s_0, a), P_{i,K+1}^\epsilon \leftarrow 0, P_{i,K+1}^{\epsilon^-} \leftarrow 0$ 
39   return  $\{V_{i,k}\}_{i=1,2,\dots,K+1}, \{P_{i,k}^\epsilon\}_{i=1,2,\dots,K+1}, \{P_{i,k}^{\epsilon^-}\}_{i=1,2,\dots,K+1}$ 

```

Remark 2. The single-agent policies provide movement guidance for lower-level task execution (e.g., determining which grid cells to visit next, as shown in Figure 1). However, these policies do not explicitly handle low-level motion planning or collision avoidance. To ensure safe and efficient navigation during the actual execution, real-time collision avoidance algorithms such as [26, 36, 37] and motion planning algorithms [16, 28] can be integrated.

4.3 Bi-level Task Allocation and Policy Learning

Assumption 2. We assume that, after substituting $P_{i,k}^\epsilon$ with the static lower bound $P_{i,k}^{\epsilon^-}$ in the optimization problem (3), a feasible solution $\{P_{i,k}\}$ exists for the resulting optimization problem.

This assumption sets a necessary condition for the proposed algorithm. The static lower bound $P_{i,k}^{\epsilon^-}$ reflects the maximum achievable probability of constraint satisfaction given the robot's limited knowledge of transition dynamics. If a feasible task allocation $\{P_{i,k}\}$

does not exist, task satisfaction cannot be guaranteed at the desired probability thresholds, implying a need to either relax the TWTL task requirements or increase the number of robots.

We summarize the proposed bi-level framework in Alg. 3. Before each episode, the high-level coordinator collects task satisfaction probabilities and expected rewards from the robots (lines 2-3). The coordinator first attempts to solve the task allocation problem using the adaptive lower bounds (line 4). If no feasible solution is found, it falls back on the static lower bounds $\underline{P}_{i,k}^\epsilon$, as Assumption 2 ensures that a feasible solution exists (lines 5-6). Once the task assignment is determined, each robot independently executes its policy (Alg. 2) to complete the assigned task (lines 7-8).

Algorithm 3: Bi-level Multi-robot Coordination

Input: A set of robots $\{N\}$; number of episodes $N_{episode}$

```

1 for  $j = 1 : N_{episode}$  do
2   for  $r_i \in \{N\}$  do
3      $\{V_{i,k}\}, \{\underline{P}_{i,k}^\epsilon\}, \{\underline{P}_{i,k}^{\epsilon^-}\} \leftarrow r_i.stats() \text{ (Alg. 2)}$ 
4     try:  $\{P_{i,k}\} \leftarrow$  substitute  $\underline{P}_{i,k}^\epsilon$  with  $\underline{P}_{i,k}^{\epsilon^-}$  in Problem (3) and solve
5     if no feasible solution  $\{P_{i,k}\}$  found then
6        $\{P_{i,k}\} \leftarrow$  substitute  $\underline{P}_{i,k}^\epsilon$  with  $\underline{P}_{i,k}^\epsilon$  in Problem (3) and solve
7   for  $r_i \in \{N\}$  do
8      $r_i.execute(\{P_{i,k}\}_{k=1,2,\dots,K+1}) \text{ (Alg. 2)}$ 

```

Remark 3. In cases where some robots fail, re-planning can be achieved by solving a modified version of Problem (3). The reallocation process focuses on the affected TWTL tasks and redistributes them among available robots that are not committed to any TWTL tasks in the current episode. This re-planning scheme can improve robustness against failures while minimizing disruptions to ongoing tasks.

THEOREM 4.6. Given a set of robots $\{N\}$, a set of TWTL tasks $\{K+1\}$, and a set of MDPs $\{M\}$ subject to Assumption 1. Suppose Assumption 2 holds. Then, in every episode where Alg. 3 finds a feasible solution $\{P_{i,k}\}$, the probability of satisfying ϕ_k for each task t_k is guaranteed to be at least P_k with high confidence.

PROOF. Theorem 4.6 follows directly from Proposition 4.1, with the proof available in [19]. \square

5 SIMULATION

To validate the proposed algorithm, we conduct simulations based on the pickup and delivery example illustrated in Fig. 1, incorporating time window requirements. The environment depicted in Fig. 3b is used across all experiments. In Fig. 3b, the colored cells represent stations (S_1, S_2) where robots can idle, warehouses (W_1, W_2) for storing resources, resource processing units (P_1, P_2), and an operation site (O) where processed resources are delivered. The black zones indicate restricted areas that robots must avoid, while the blue zones represent water, which ground robots are prohibited from entering. The gray zones represent areas of interest requiring aerial monitoring, which is unknown to the robots. The arrows show two one-way bridges over the water.

We consider robots with transition models described in Example 2.4. The robot specifications are summarized in Table 1. Drones are incentivized to identify and monitor traffic congestions (grey




Figure 3: (a) Transitions (intended - black, unintended - red) under each action. (b) Illustration of the simulation environment.

Table 1: Robot Types, Rewards, and Uncertainties

Robot Type	Robot Index	Reward	ϵ	ϵ^i
Drone	1, 2	5 at grey cells, 0 at other cells	0.1	0.2
	3, 4	1 at grey cells, 0 at other cells		
Mobile Robot	5, 6	1 at S_2 , 0 at other cells	0.15	0.3
	7, 8			0.25

cells with initially unknown locations), where drones 1 and 2 get higher rewards due to their better sensing capability. Mobile robots are encouraged to return to Station S_2 as part of their auxiliary objectives. All robots have a conservative estimate ϵ^i of their actual unknown transition uncertainty ϵ .

Table 2 presents the task specifications. Tasks ϕ_1 and ϕ_2 require robots to pick up resources from warehouse W_2 , process them at P_2 or P_1 , respectively, and deliver them to O —all within 15 time steps, ensuring a 90% success rate. Tasks ϕ_3 and ϕ_4 involve transporting resources from a different warehouse W_1 , delivering them to P_1 , and to O , with different time window requirements. The lower desired probabilities for ϕ_3 and ϕ_4 provide robots with more flexibility to pursue their auxiliary task of environmental monitoring. While this relaxes the requirement for individual task satisfaction, the probability of completing at least one of these tasks remains high.

Table 2: Task Specifications as TWTL Formulas

Task	TWTL Formula	Desired Probability
ϕ_1	$[H^1 W_2]^{[0,15]} \cdot [H^1 P_2]^{[0,15]} \cdot [H^1 O]^{[0,15]}$	0.9
ϕ_2	$[H^1 W_2]^{[0,15]} \cdot [H^1 P_1]^{[0,15]} \cdot [H^1 O]^{[0,15]}$	0.9
ϕ_3, ϕ_4	$[H^1 W_1]^{[0,20]} \cdot [H^1 P_1]^{[0,10]} \cdot [H^1 O]^{[0,15]}$	0.7

We utilize the Robot Operating System (ROS 2) [23], where each robot is implemented as an independent ROS node. Robots communicate with the coordinator using ROS communication mechanisms, enabling realistic simulation of a distributed multi-robot system. Simulations are conducted over $N_{episode} = 2000$ episodes per iteration. The results presented are averaged over 20 iterations. Details of the parameters used are provided in [19].

Case 1. We evaluate the proposed algorithm under two conditions: using static lower bounds (by skipping line 4 in Alg. 3) and adaptive lower bounds (as described in Alg. 3). Our results show that solving the task allocation problem with static lower bounds tends to assign more robots to TWTL tasks ($\phi_1 - \phi_4$), compared to the adaptive lower bounds. This finding aligns with our expectations, as the static lower bound computes task satisfaction probability based on ϵ^i , which overestimates the actual uncertainty ϵ . In contrast, the adaptive lower bounds refine the estimate of task satisfaction probability as the robot interacts with the environment,

resulting in a less conservative lower bound. With the adaptive lower bounds, robots progressively gain confidence in TWTL satisfaction, allowing the high-level coordinator to adaptively improve the task allocation by assigning fewer robots to TWTL tasks.

Fig. 4a shows that average task satisfaction rates consistently exceed the desired probability, regardless of the lower bound selection. This confirms that the proposed framework successfully ensures probabilistic satisfaction of the TWTL tasks. Nevertheless, static lower bounds result in a conservatively higher satisfaction rate. Fig. 4b compares the TWTL satisfaction rate in the first 100 episodes versus 2000 episodes when using adaptive lower bounds. We observe that as the adaptive lower bounds refine over time, satisfaction rates converge toward the desired probabilities by reducing assignments to TWTL tasks. As a result, using the adaptive lower bounds yields higher rewards, as shown in Fig. 5a.




Figure 4: TWTL Task satisfaction rate. (a): Satisfaction rate over 2000 episodes using static and adaptive lower bounds. (b) Satisfaction rate using adaptive lower bounds, over the first 100 episodes vs 2000 episodes. The dashed line represents the desired probability of each TWTL task. The bars show the mean values over the 20 iterations, with error bars depicting the standard deviation.

Our framework, which incorporates auxiliary reward functions alongside the TWTL constraint, enables adaptive and self-improving task allocation. In the early stages of each iteration, the system predominantly assigns TWTL tasks to the drones (Robots 1-4) as the mobile robots (Robots 5-8) initially have high transition uncertainty estimates (ϵ^i), resulting in lower task satisfaction probability estimates. As the simulation progresses, the robots continuously update their task satisfaction probabilities. Meanwhile, drones gradually discover high-reward areas, prompting the coordinator to adjust the task assignments. By the end of the simulation, the system converges to a reasonable division of labor, with TWTL tasks mainly assigned to mobile robots (Robots 5-8) and low-reward drones (Robots 3-4), while drones with high-quality sensors (Robots 1-2) remain unassigned to focus on monitoring the high-interest area. This adaptive behavior demonstrates the system’s ability to dynamically reallocate resources, balancing reward maximization with guaranteed satisfaction of TWTL constraints.

Case 2. In this case study, we evaluate the computation time of the proposed framework as the number of robots increases. The experimental setup is identical to Case 1, where tasks listed in Table 2 must be satisfied with their corresponding probability thresholds. We simulate systems with 8, 16, 24, 32, 40, and 48 robots, duplicating the robots defined in Table 1 by factors from 1 to 6. In this setup, we utilize the adaptive lower bound.

Fig. 5b presents the average computation time per episode for solving the task allocation problem (3), averaged over 2000 episodes.

The results show that, even with a large number of robots, the proposed framework solves the task allocation problem in a relatively short time. Notably, for 48 robots, the average computation time remains under 0.4 seconds, demonstrating the framework’s capability for efficient real-time applications.

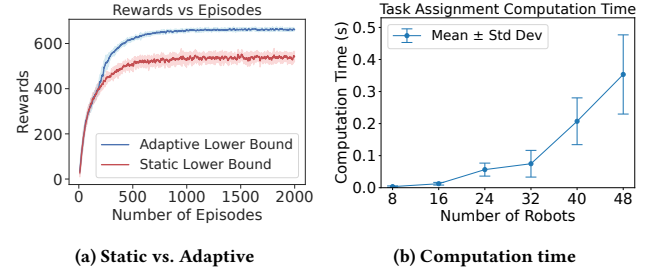


Figure 5: (a): Total rewards accumulated by all robots over the episodes. The reward represents the sum of individual rewards collected by each robot. (b): Computation time for solving task allocation (Alg. 3, line 4-6) with different number of robots.

In Alg. 1, each robot independently constructs the product MDPs and synthesizes the policies offline. In Alg. 2, each robot executes its task independently, without the need for communication with others. Thus, computation time for these processes is unaffected by the number of robots in the system. While our evaluation focused on varying the number of robots, increasing the number of tasks similarly impacts computation time by adding nonlinear constraints to the optimization problem. Analysis of this effect is included in the extended version of this work [19].

6 CONCLUSION

This paper presents a bi-level framework that integrates high-level task allocation with lower-level task execution and learning to ensure robots satisfy Time-Window Temporal Logic (TWTL) tasks with guaranteed probability thresholds. We introduce adaptive lower bounds on task completion probabilities, enabling robots to iteratively refine their probability estimates for more informed task allocation decisions. By incorporating reward functions as an auxiliary objective, the system can iteratively improve task allocation to maximize expected rewards while maintaining probabilistic task satisfaction. The flexibility to impose different reward functions on robots enables the system to achieve multiple objectives simultaneously, allowing different reward functions to guide user-customized allocation plans without compromising the primary objective of satisfying TWTL tasks.

We provide theoretical analysis and conduct comprehensive simulations to validate the framework. The results highlight its ability to ensure constraint satisfaction at desired probability thresholds under uncertainty. Furthermore, simulation results also show that the task allocation problem can be efficiently solved for large numbers of robots in real time, demonstrating the framework’s practical applicability in real-world scenarios.

ACKNOWLEDGMENTS

We sincerely thank Dr. Kevin Leahy for his insightful discussions and valuable suggestions on this paper’s problem formulation.

REFERENCES

- [1] Derya Aksaray, Yasin Yazıcıoğlu, and Ahmet Semi Asarkaya. 2021. Probabilistically guaranteed satisfaction of temporal logic constraints during reinforcement learning. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 6531–6537.
- [2] Ali Tefvik Buyukkokac, Derya Aksaray, and Yasin Yazıcıoğlu. 2021. Planning of heterogeneous multi-agent systems under signal temporal logic specifications with integral predicates. *IEEE Robotics and Automation Letters* 6, 2 (2021), 1375–1382.
- [3] Mingyu Cai, Mohammadhosein Hasanbeig, Shaoping Xiao, Alessandro Abate, and Zhen Kan. 2021. Modular deep reinforcement learning for continuous motion planning with temporal logic. *IEEE Robotics and Automation Letters* 6, 4 (2021), 7973–7980.
- [4] Mingyu Cai, Shaoping Xiao, Baoluo Li, Zhiliang Li, and Zhen Kan. 2021. Reinforcement learning based temporal logic control with maximum probabilistic satisfaction. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 806–812.
- [5] Shushman Choudhury, Jayesh K Gupta, Mykel J Kochenderfer, Dorsa Sadigh, and Jeannette Bohg. 2022. Dynamic multi-robot task allocation under uncertainty and temporal constraints. *Autonomous Robots* 46, 1 (2022), 231–247.
- [6] Jie Fu and Ufuk Topcu. 2014. Probably approximately correct MDP learning and control with temporal logic constraints. *arXiv preprint arXiv:1404.7073* (2014).
- [7] Brian P Gerkey and Maja J Mataric. 2004. A formal analysis and taxonomy of task allocation in multi-robot systems. *The International journal of robotics research* 23, 9 (2004), 939–954.
- [8] Dhaval Gujarathi and Indranil Saha. 2022. MT*: Multi-robot path planning for temporal logic specifications. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 13692–13699.
- [9] Lewis Hammond, Alessandro Abate, Julian Gutierrez, and Michael Wooldridge. 2021. Multi-agent reinforcement learning with temporal logic specifications. *arXiv preprint arXiv:2102.00582* (2021).
- [10] Mohammadhosein Hasanbeig, Alessandro Abate, and Daniel Kroening. 2020. Cautious reinforcement learning with logical constraints. *arXiv preprint arXiv:2002.12156* (2020).
- [11] Michael Hyland and Hani S Mahmassani. 2018. Dynamic autonomous vehicle fleet operations: Optimization-based strategies to assign AVs to immediate traveler demand requests. *Transportation Research Part C: Emerging Technologies* 92 (2018), 278–297.
- [12] Yuqian Jiang, Suda Bharadwaj, Bo Wu, Rishi Shah, Ufuk Topcu, and Peter Stone. 2021. Temporal-logic-based reward shaping for continuing reinforcement learning tasks. In *Proceedings of the AAAI Conference on artificial Intelligence*, Vol. 35. 7995–8003.
- [13] Lars Johansmeier and Sami Haddadin. 2016. A hierarchical human-robot interaction-planning framework for task allocation in collaborative industrial assembly processes. *IEEE Robotics and Automation Letters* 2, 1 (2016), 41–48.
- [14] Yiannis Kantaros and Michael M Zavlanos. 2020. Stylus*: A temporal logic optimal control synthesis algorithm for large-scale multi-robot systems. *The International Journal of Robotics Research* 39, 7 (2020), 812–836.
- [15] Sertac Karaman and Emilio Frazzoli. 2011. Linear temporal logic vehicle routing with applications to multi-UAV mission planning. *International Journal of Robust and Nonlinear Control* 21, 12 (2011), 1372–1395.
- [16] Sertac Karaman and Emilio Frazzoli. 2011. Sampling-based algorithms for optimal motion planning. *The international journal of robotics research* 30, 7 (2011), 846–894.
- [17] M. Kwiatkowska, G. Norman, and D. Parker. 2011. PRISM 4.0: Verification of Probabilistic Real-time Systems. In *Proc. 23rd International Conference on Computer Aided Verification (CAV’11) (LNCS, Vol. 6806)*, G. Gopalakrishnan and S. Qadeer (Eds.). Springer, 585–591.
- [18] Xiaoshan Lin, Abbasali Koochakzadeh, Yasin Yazıcıoğlu, and Derya Aksaray. 2023. Reinforcement learning under probabilistic spatio-temporal constraints with time windows. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 8680–8686.
- [19] Xiaoshan Lin and Roberto Tron. 2025. Adaptive Bi-Level Multi-Robot Task Allocation and Learning under Uncertainty with Temporal Logic Constraints. *arXiv preprint arXiv* (2025).
- [20] Xiaoshan Lin, Sadık Bera Yüksel, Yasin Yazıcıoğlu, and Derya Aksaray. 2024. Probabilistic Satisfaction of Temporal Logic Constraints in Reinforcement Learning via Adaptive Policy-Switching. *arXiv preprint arXiv:2410.08022* (2024).
- [21] Wenliang Liu, Kevin Leahy, Zachary Serlin, and Calin Belta. 2023. CatNet: Learning Communication and Coordination Policies from CaTL+ Specifications. In *Learning for Dynamics and Control Conference*. PMLR, 705–717.
- [22] Xusheng Luo, Yiannis Kantaros, and Michael M Zavlanos. 2021. An abstraction-free method for multirobot temporal logic optimal control synthesis. *IEEE Transactions on Robotics* 37, 5 (2021), 1487–1507.
- [23] Steven Macenski, Tully Foote, Brian Gerkey, Chris Lalancette, and William Woodall. 2022. Robot Operating System 2: Design, architecture, and uses in the wild. *Science Robotics* 7, 66 (2022), eabm6074. <https://doi.org/10.1126/scirobotics.abm6074>
- [24] Devaprakash Muniraj, Kyriakos G Vamvoudakis, and Mazen Farhood. 2018. Enforcing signal temporal logic specifications in multi-agent adversarial environments: A deep q-learning approach. In *2018 IEEE Conference on Decision and Control (CDC)*. IEEE, 4141–4146.
- [25] Daniel Neider, Jean-Raphael Gaglione, Ivan Gavran, Ufuk Topcu, Bo Wu, and Zhe Xu. 2021. Advice-guided reinforcement learning in a non-Markovian environment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 9073–9080.
- [26] Dimitra Panagou. 2014. Motion planning and collision avoidance using navigation vector fields. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2513–2518.
- [27] Jinwoo Park, Andrew Messing, Harish Ravichandar, and Seth Hutchinson. 2023. Risk-Tolerant Task Allocation and Scheduling in Heterogeneous Multi-Robot Teams. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 5372–5379.
- [28] Christoph Rösmann, Wendelin Feiten, Thomas Wösch, Frank Hoffmann, and Torsten Bertram. 2012. Trajectory modification considering dynamic constraints of autonomous robots. In *ROBOTIK 2012; 7th German Conference on Robotics*. VDE, 1–6.
- [29] Philipp Schillinger, Mathias Bürger, and Dimos V Dimarogonas. 2018. Simultaneous task allocation and planning for temporal logic goals in heterogeneous multi-robot systems. *The international journal of robotics research* 37, 7 (2018), 818–838.
- [30] Philipp Schillinger, Mathias Bürger, and Dimos V Dimarogonas. 2019. Hierarchical ltl-task mdp for multi-agent coordination through auctioning and learning. *The international journal of robotics research* (2019).
- [31] Chuangchuan Sun, Xiao Li, and Calin Belta. 2020. Automata guided semi-decentralized multi-agent reinforcement learning. In *2020 American Control Conference (ACC)*. IEEE, 3900–3905.
- [32] Dawei Sun, Jingkai Chen, Sayan Mitra, and Chuchu Fan. 2022. Multi-agent motion planning from signal temporal logic specifications. *IEEE Robotics and Automation Letters* 7, 2 (2022), 3451–3458.
- [33] Elina Suslova and Pooyan Fazli. 2020. Multi-robot task allocation with time window and ordering constraints. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 6909–6916.
- [34] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*.
- [35] Jana Tumova and Dimos V Dimarogonas. 2016. Multi-agent planning under local LTL specifications and event-based synchronization. *Automatica* 70 (2016), 239–248.
- [36] Jur Van Den Berg, Stephen J Guy, Ming Lin, and Dinesh Manocha. 2011. Reciprocal n-body collision avoidance. In *Robotics Research: The 14th International Symposium ISRR*. Springer, 3–19.
- [37] Jur Van den Berg, Ming Lin, and Dinesh Manocha. 2008. Reciprocal velocity obstacles for real-time multi-agent navigation. In *2008 IEEE international conference on robotics and automation*. Ieee, 1928–1935.
- [38] Cristian-Ioan Vasile, Derya Aksaray, and Calin Belta. 2017. Time window temporal logic. *Theoretical Computer Science* 691 (2017), 27–54.
- [39] Jiangwei Wang, Shuo Yang, Ziyan An, Songyang Han, Zhili Zhang, Rahul Mangharam, Meiyi Ma, and Fei Miao. 2023. Multi-agent reinforcement learning guided by signal temporal logic specifications. *arXiv preprint arXiv:2306.06808* (2023).
- [40] Christopher JCH Watkins and Peter Dayan. 1992. Q-learning. *Machine learning* 8 (1992), 279–292.
- [41] Edwin B Wilson. 1927. Probable inference, the law of succession, and statistical inference. *J. Amer. Statist. Assoc.* 22, 158 (1927), 209–212.
- [42] Eric M Wolff, Ufuk Topcu, and Richard M Murray. 2014. Optimization-based trajectory generation with linear temporal logic specifications. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 5319–5325.
- [43] Zhi Yan, Nicolas Jouandeau, and Arab Ali-Chérif. 2012. Multi-robot heuristic goods transportation. In *2012 6th IEEE International Conference Intelligent Systems*. IEEE, 409–414.
- [44] Ningyuan Zhang, Wenliang Liu, and Calin Belta. 2022. Distributed control using reinforcement learning with temporal-logic-based reward shaping. In *Learning for Dynamics and Control Conference*. PMLR, 751–762.