

High-Level, Collaborative Task Planning Grammar and Execution for Heterogeneous Agents

Amy Fang
Cornell University
Ithaca, NY, USA
axf4@cornell.edu

Hadas Kress-Gazit
Cornell University
Ithaca, NY, USA
hadaskg@cornell.edu

ABSTRACT

We propose a new multi-agent task grammar to encode collaborative tasks for a team of heterogeneous agents that can have overlapping capabilities. The grammar allows users to specify the relationship between agents and parts of the task without providing explicit assignments or constraints on the number of agents required. We develop a method to automatically find a team of agents and synthesize correct-by-construction control with synchronization policies to satisfy the task. We demonstrate the scalability of our approach through simulation and compare our method to existing task grammars that encode multi-agent tasks.

KEYWORDS

Formal methods, multi-agent coordination, task planning, robotics

ACM Reference Format:

Amy Fang and Hadas Kress-Gazit. 2024. High-Level, Collaborative Task Planning Grammar and Execution for Heterogeneous Agents. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, Auckland, New Zealand, May 6 – 10, 2024, IFAAMAS, 9 pages.

1 INTRODUCTION

Agents working together to achieve common goals have a variety of applications, such as warehouse automation or disaster response. Multi-agent tasks have been defined in different ways in the scheduling and planning literature. For example, in multi-agent task allocation [8, 9, 12] and coalition formation [14, 22], each task is a single goal with an associated utility. Individual agents or agent teams then automatically assign themselves to a task based on some optimization metric. Swarm approaches [18, 21] consider emergent behavior of an agent collective as the task, for example, aggregation or shape formation.

Recently, formal methods, such as temporal logics for task specifications and correct-by-construction synthesis, have been used to solve different types of multi-agent planning tasks [2, 17, 20]. Tasks written in temporal logic, such as Linear Temporal Logic (LTL), allow users to capture complex tasks with temporal constraints. Existing work has extended LTL [15, 16] and Signal Temporal Logic [13] to encode tasks that require multiple agents.

In this paper, we consider tasks that a team of heterogeneous agents are required to collaboratively satisfy. For instance, consider

a precision agriculture scenario in which a farm contains agents with different on-board sensors to monitor crop health. The user may want to take a moisture measurement in one region, and then take a soil sample of a different region. Depending on the agents' sensors and sensing range, the agents may decide to collaborate to satisfy the task. For example, one agent may perform the entire task on its own if it has both a moisture sensor and an arm mechanism to pick up a soil sample and can move between the two regions. However, another possible solution is for two agents to team up so that one takes a moisture measurement and the other picks up the soil. Existing task grammars [13, 15, 16] capture tasks such as the above by providing explicit constraints on the types or number of agents for each part of the task, i.e. the task must explicitly encode whether it should be one agent, two agents, or either of these options. In this paper, we create a task grammar and associated control synthesis that removes the need to a priori decide on the number of agents necessary to accomplish a task, allowing users to focus solely on the actions required to achieve the task (e.g. "take a moisture measurement and then pick up a soil sample, irrespective of which or how many agents perform which actions").

Our task grammar has several unique aspects. First, this grammar enables the interleaving of agent actions, alleviating the need for explicit task decomposition in order to assign agents to parts of the task. Second, rather than providing explicit constraints on the types or number of agents for each part of the task, the task encodes, using the concept of *bindings* (inspired by [15]), the overall relationship between agent assignments and team behavior; we can require certain parts of the task to be satisfied by the same agent without assigning the exact agent or type of agent *a priori*. Lastly, the grammar allows users to make the distinction between the requirements "for all agents" and "at least one agent". Given these types of tasks, agents autonomously determine, based on their capabilities, which parts of the task they can and should do for the team to satisfy the task.

Tasks may require collaboration between different agents. Similar to [3, 11, 19], to ensure the actions are performed in the correct order, our framework takes the corresponding synchronization constraints into account while synthesizing agent behavior; agents must wait to execute the actions together. In our approach, execution of the synchronous behavior for each agent is decentralized; agents carry out their plan and communicate with one another when synchronization is necessary.

Depending on the task and the available agents, there might be different teams (i.e., subsets of the agent set) that can carry out the task; our algorithm for assigning a team and synthesizing behavior for the agents finds the largest team of agents that satisfies the task. This means that the team may have redundancies, i.e. agents can



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

be removed while still ensuring the overall task is satisfied. This is beneficial both for robustness and optimality; the user can choose a subset of the team (provided that all the required bindings are still assigned) to optimize different metrics, such as cost or overall number of agents.

Related work: One way to encode tasks is to first decompose them into independent sub-tasks and then allocate them to the agents. For example, [7, 17] address finite-horizon tasks for multi-agent teams. The authors first automatically decompose a global automaton representing the task into independent sub-tasks. To synthesize control policies, the authors build product automata for each heterogeneous agent. Each automaton is then sequentially linked using switch transitions to reduce state-space explosion in synthesizing parallel plans. In our prior work [6], we address infinite-horizon tasks that have already been decomposed into sub-tasks. Given a new task, we proposed a decentralized framework for agents to automatically update their behavior based on a new task and their existing tasks, allowing agents to interleave the tasks.

The works discussed above make the critical assumption that tasks are independent, i.e. agents do not collaborate with one another. One approach to including collaborative actions is to explicitly encode the agent assignments in the tasks. To synthesize agent control for these types of tasks, in [19], the authors construct a reduced product automaton in which the agents only synchronize when cooperative actions are required. The work in [10] proposes a sampling-based method that approximates the product automaton of the team by building trees incrementally while maintaining probabilistic completeness. In this paper, we consider the more general setting in which agents may need to collaborate with each other, but are not given explicit task assignments *a priori*.

Rather than providing predetermined task assignments, another approach for defining collaborative tasks is to capture information about the number and type of agents needed for parts of the specification. For example, [16] imposes constraints on the number of agents necessary in regions using counting LTL. [13] uses Capability Temporal Logic to encode both the number and capabilities necessary in certain abstracted locations in the environment and then formulates the problem as a MILP to find an optimal teaming strategy. The authors of [15] introduce the concept of induced propositions, where each atomic proposition not only encodes information about the number, type of agents, and target regions, but also has a connector that binds the truth of certain atomic propositions together. To synthesize behavior for the agents, they propose a hierarchical approach that first constructs the automaton representing the task and then decomposes the task into possible sub-tasks. The temporal order of these sub-tasks is captured using partially ordered sets and are used in the task allocation problem, which is formulated as a MILP.

Inspired by [15] and the concept of induced propositions, we create a task grammar that includes information about how the atomic propositions are related to one another, which represents the overall relationship between agents and task requirements. Unlike [15], which considers navigation tasks in which the same set of agents of a certain type may need to visit different regions, we generalize these tasks to any type of abstract action an agent may be able to perform. In addition, a key assumption we relax is that we do not require each agent to be only categorized as one type. As a

result, agents can have overlapping capabilities. To our knowledge, no other grammars have been proposed for these generalized types of multi-agent collaborative tasks.

Contributions: We propose a task description and control synthesis framework for heterogeneous agents to satisfy collaborative tasks. Specifically, we present a new, LTL-based task grammar for the formulation of collaborative tasks, and provide a framework to form a team of agents and synthesize control and synchronization policies to guarantee the team satisfies the task. We demonstrate our approach in simulated precision agriculture scenarios.

2 PRELIMINARIES

2.1 Linear Temporal Logic

LTL formulas are defined over a set of atomic propositions AP , where $\pi \in AP$ are Boolean variables [5]. We abstract agent actions as atomic propositions. For example, UV captures an agent taking UV measurement.

Syntax: An LTL formula is defined as:

$$\varphi ::= \pi \mid \neg\varphi \mid \varphi \vee \varphi \mid \bigcirc\varphi \mid \varphi \mathcal{U} \varphi$$

where \neg (“not”) and \vee (“or”) are Boolean operators, and \bigcirc (“next”) and \mathcal{U} (“until”) are temporal operators. From these operators, we can define: conjunction $\varphi \wedge \varphi$, implication $\varphi \Rightarrow \varphi$, eventually $\diamond\varphi = \text{True } \mathcal{U} \varphi$, and always $\square\varphi = \neg\diamond\neg\varphi$.

Semantics: The semantics of an LTL formula φ are defined over an infinite trace $\sigma = \sigma(0)\sigma(1)\sigma(2)\dots$, where $\sigma(i)$ is the set of true AP at position i . We denote that σ satisfies LTL formula φ as $\sigma \models \varphi$.

Intuitively, $\diamond\varphi$ is satisfied if there exists a $\sigma(i)$ in which φ is true. $\square\varphi$ is satisfied if φ is true at every position in σ . To satisfy $\varphi_1 \mathcal{U} \varphi_2$, φ_1 must remain true until φ_2 becomes true. See [5] for the full semantics.

2.2 Büchi Automata

An LTL formula φ can be translated into a Nondeterministic Büchi Automaton that accepts infinite traces if and only if they satisfy φ . A Büchi automaton is a tuple $\mathcal{B} = (Z, z_0, \Sigma_{\mathcal{B}}, \delta_{\mathcal{B}}, F)$, where Z is the set of states, $z_0 \in Z$ is the initial state, $\Sigma_{\mathcal{B}}$ is the input alphabet, $\delta_{\mathcal{B}} : Z \times \Sigma_{\mathcal{B}} \times Z$ is the transition relation, and $F \subseteq Z$ is a set of accepting states. An infinite run of \mathcal{B} over a word $w = w_1w_2w_3\dots$, $w_i \in \Sigma_{\mathcal{B}}$ is an infinite sequence of states $z = z_0z_1z_2\dots$ such that $(z_{i-1}, w_i, z_i) \in \delta_{\mathcal{B}}$. A run is accepting if and only if $\text{Inf}(z) \cap F \neq \emptyset$, where $\text{Inf}(z)$ is the set of states that appear in z infinitely often [1].

2.3 Agent Model

Following [6], we create an abstract model for each agent based on its set of capabilities. A capability is a weighted transition system $\lambda = (S, s_0, AP, \Delta, L, W)$, where S is a finite set of states, $s_0 \in S$ is the initial state, AP is the set of atomic propositions, $\Delta \subseteq S \times S$ is a transition relation where for all $s \in S$, $\exists s' \in S$ such that $(s, s') \in \Delta$, $L : S \rightarrow 2^{AP}$ is the labeling function such that $L(s)$ is the set of propositions that are true in state s , and $W : \Delta \rightarrow \mathbb{R}_{\geq 0}$ is the cost function assigning a weight to each transition. Since we are considering a group of heterogeneous agents, agent j has its own set of k capabilities $\Lambda_j = \{\lambda_1, \dots, \lambda_k\}$.

An agent model A_j is the product of its capabilities: $A_j = \lambda_1 \times \dots \times \lambda_k$ such that $A_j = (S, s_0, AP_j, \gamma, L, W)$, where $S = S_1 \times \dots \times S_k$

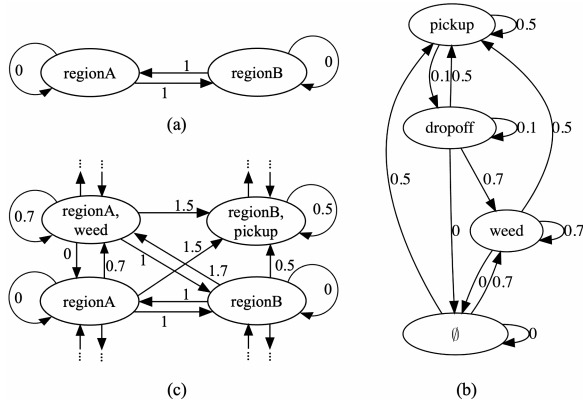


Figure 1: Agent partial model: (a) λ_{area} (b) λ_{arm} (c) A_{green}

is the set of states, $s_0 \in S$ is the initial state, $AP_j = \bigcup_{i=1}^k AP_i$ is the set of propositions, $\gamma \subseteq S \times S$ is the transition relation such that $(s, s') \in \gamma$, where $s = (s_1, \dots, s_k)$, $s' = (s'_1, \dots, s'_k)$, if and only if for all $i = \{1, \dots, k\}$, $(s_i, s'_i) \in \Delta_i$, $L : S \rightarrow 2^{AP_j}$ is the labeling function where $L(s) = \bigcup_{i=1}^k L_i(s_i)$, and $W : \gamma \rightarrow \mathbb{R}_{\geq 0}$ is the cost function that combines the costs of the capabilities. Fig. 1c depicts a snippet of an agent model where we treat the cost as additive. Fig. 1a represents the agent’s sensing area λ_{area} : the agent can orient its sensors to take measurements in different regions of a partitioned workspace (in this case, regions A and B). Fig. 1b represents the agent’s robot manipulator, which can pick up and drop off soil samples, as well as pull weeds.

3 TASK GRAMMAR - LTL $^\psi$

We define the task grammar LTL $^\psi$ that includes atomic propositions that abstract agent action, logical and temporal operators, as in LTL, and bindings that connect actions to specific agents; any action labeled with the same binding must be satisfied by the same agent(s) (the actual value of the binding is not important). We define a task recursively over LTL and binding formulas.

$$\psi := \rho \mid \psi_1 \vee \psi_2 \mid \psi_1 \wedge \psi_2 \quad (1)$$

$$\varphi := \pi \mid \neg\varphi \mid \varphi \vee \varphi \mid \bigcirc\varphi \mid \varphi \mathbf{U} \varphi \quad (2)$$

$$\varphi^\psi := \varphi^\psi \mid \neg(\varphi^\psi) \mid \varphi_1^{\psi_1} \wedge \varphi_2^{\psi_2} \mid \varphi_1^{\psi_1} \vee \varphi_2^{\psi_2} \mid \bigcirc\varphi^\psi \mid \varphi_1^{\psi_1} \mathbf{U} \varphi_2^{\psi_2} \mid \square\varphi^\psi \quad (3)$$

where ψ , the *binding formula*, is a Boolean formula excluding negation over $\rho \in AP_\psi$, and φ is an LTL formula. An LTL $^\psi$ formula consists of conjunction, disjunction, and temporal operators; we define eventually as $\diamond\varphi^\psi = \text{True} \mathbf{U} \varphi^\psi$. An example of an LTL $^\psi$ formula is shown in Eq. 4.

Semantics: The semantics of an LTL $^\psi$ formula φ^ψ are defined over σ and R ; $\sigma = \sigma_1\sigma_2\dots\sigma_n$ is the team trace where σ_j is agent j ’s trace, and $\forall i, \sigma(i) = \sigma_1(i)\sigma_2(i)\dots\sigma_n(i)$. $R = \{r_1, r_2, \dots, r_n\}$ is the set of binding assignments, where $r_j \in R$ is the set of AP_ψ that are assigned to agent j . Once a team is established, R is constant, i.e. an agent’s binding assignment does not change throughout the task execution. For example, $r_1 = \{2, 3\}$, $r_2 = \{1\}$ denotes that agent 1 is assigned bindings 2 and 3, and agent 2 is assigned binding 1.

Given n agents and a set of binding propositions AP_ψ , we define the function $\zeta : \psi \rightarrow 2^{2^{AP_\psi}}$ such that $\zeta(\psi)$ is the set of all possible combinations of ρ that satisfy ψ . For example, $\zeta((1 \vee 2) \wedge 3) = \{\{1, 3\}, \{2, 3\}, \{1, 2, 3\}\}$.

The semantics of LTL $^\psi$ are:

- $(\sigma(i), R) \models \varphi^\psi$ iff $\exists K \in \zeta(\psi)$ s.t. $(K \subseteq \bigcup_{p=1}^n r_p)$ and $(\forall j$ s.t. $K \cap r_j \neq \emptyset, \sigma_j(i) \models \varphi)$
- $(\sigma(i), R) \models (\neg\varphi)^\psi$ iff $\exists K \in \zeta(\psi)$ s.t. $(K \subseteq \bigcup_{p=1}^n r_p)$ and $(\forall j$ s.t. $K \cap r_j \neq \emptyset, \sigma_j(i) \not\models \varphi)$
- $(\sigma(i), R) \models \neg(\varphi^\psi)$ iff $\exists K \in \zeta(\psi)$ s.t. $(K \subseteq \bigcup_{p=1}^n r_p)$ and $(\exists j$ s.t. $K \cap r_j \neq \emptyset, \sigma_j(i) \not\models \varphi)$
- $(\sigma(i), R) \models \varphi_1^{\psi_1} \wedge \varphi_2^{\psi_2}$ iff $(\sigma(i), R) \models \varphi_1^{\psi_1}$ and $(\sigma(i), R) \models \varphi_2^{\psi_2}$
- $(\sigma(i), R) \models \varphi_1^{\psi_1} \vee \varphi_2^{\psi_2}$ iff $(\sigma(i), R) \models \varphi_1^{\psi_1}$ or $(\sigma(i), R) \models \varphi_2^{\psi_2}$
- $(\sigma(i), R) \models \bigcirc\varphi^\psi$ iff $\sigma(i+1), R \models \varphi^\psi$
- $(\sigma(i), R) \models \varphi_1^{\psi_1} \mathbf{U} \varphi_2^{\psi_2}$ iff $\exists \ell \geq i$ s.t. $(\sigma(\ell), R) \models \varphi_2^{\psi_2}$ and $\forall i \leq k < \ell, (\sigma(k), R) \models \varphi_1^{\psi_1}$
- $(\sigma(i), R) \models \square\varphi^\psi$ iff $\forall \ell > i, (\sigma(\ell), R) \models \varphi^\psi$

Intuitively, the behavior of an agent team and their respective binding assignments satisfy φ^ψ if there exists a possible binding assignment in $\zeta(\psi)$ in which all the bindings are assigned to (at least one) agent, and the behavior of all agents with a relevant binding assignment satisfy φ . An agent can be assigned more than one binding, and a binding can be assigned to more than one agent.

Remark 1. For the sake of clarity in notation, $\neg\varphi^\psi$ is equivalent to $(\neg\varphi)^\psi$. For example, $\neg\text{pickup}^1 \triangleq (\neg\text{pickup})^1$.

Remark 2. Note the subtle but important difference between $(\neg\varphi)^\psi$ and $\neg(\varphi^\psi)$. Informally, the former requires all agents with binding assignments that satisfy ψ to satisfy $\neg\varphi$; the latter requires the formula φ^ψ to be violated, meaning that at least one agent’s trace violates φ , i.e. satisfies $\neg\varphi$.

Remark 3. Unique to LTL $^\psi$ is the ability to encode both tasks that include constraints on all agents or on at least one agent; “For all agents” is captured by φ^ψ ; “at least one agent” is encoded as $\neg(\neg\varphi)^\psi$, which captures “at least one agent assigned a binding in $K \in \zeta(\psi)$ satisfies φ ”. This allows for multiple agents to be assigned the same binding, but only one of those agents is necessary to satisfy φ . This can be particularly useful in tasks with safety constraints; for example, we can write $\neg(\neg\text{region}_A^1) \Rightarrow (\text{region}_A \wedge \text{visual})^2$, which says “if any agent assigned binding 1 is in region A, all agents assigned binding 2 must take a picture of the region.”

Example. Let $AP_\psi = \{1, 2, 3\}$, $AP_\varphi = \{\text{region}_A, \text{region}_B, \text{pickup}, \text{thermal}, \text{visual}, \text{moisture}, UV\}$, and $\varphi^\psi = \varphi_1^{\psi_1} \wedge \varphi_2^{\psi_2}$, where

$$\varphi_1^{\psi_1} = \diamond((\text{region}_B \wedge \text{moisture} \wedge UV)^{2\wedge 3} \wedge (\text{region}_A \wedge \text{pickup})^1) \quad (4a)$$

$$\varphi_2^{\psi_2} = \neg\text{pickup}^1 \mathbf{U} (\text{region}_A \wedge ((\text{thermal} \vee \text{visual}) \wedge \neg(\text{thermal} \wedge \text{visual})))^2 \quad (4b)$$

φ_1^ψ captures “Agent(s) assigned bindings 2 and 3 must take a moisture and UV measurement in region B at the same time that agent(s) assigned binding 1 picks up a soil sample in region A.” φ_2^ψ says “Before the soil sample is picked up, agent(s) assigned binding 2 must take a thermal or a visual image (but not both) of region A.”

Since multiple bindings can be assigned to the same agent, an agent can be assigned both bindings 2 and 3, provided it has the capabilities to satisfy the corresponding parts of the formula. In addition, depending on the final assignments, the agents may need to synchronize with one another to perform parts of the task. For example, to satisfy φ_1^ψ , agents assigned with any subset of bindings $\{1, 2, 3\}$ need to synchronize their respective actions.

4 CONTROL SYNTHESIS FOR LTL $^\psi$

Problem statement: Given n heterogeneous agents $A = \{A_1, \dots, A_n\}$ and a task φ^ψ in LTL $^\psi$, find a team of agents $A' \subseteq A$, their binding assignments $R_{A'}$, and synthesize behavior σ_j for each agent such that $(\sigma(0), R_{A'}) \models \varphi^\psi$. This behavior includes synchronization constraints for agents to satisfy the necessary collaborative actions. We assume that each agent is able to wait in any state (i.e. every state in the agent model has a self-transition).

Example. Consider a group of four agents $A = \{A_{green}, A_{blue}, A_{orange}, A_{pink}\}$ in a precision agriculture environment composed of 5 regions, as illustrated in Fig. 2. A_{orange} is a mobile robot manipulator, such as Harvest Automation’s HV-100, while the other agents are stationary with different onboard sensing capabilities. The set of all capabilities is $\Lambda = \{\lambda_{area_j}, \lambda_{motion}, \lambda_{arm}, \lambda_{UV}, \lambda_{moisture}, \lambda_{visual}, \lambda_{thermal}\}$, where $\forall j = \{green, blue, pink\}$, λ_{area_j} is agent j ’s sensing area model. The green agent can orient its arm to reach either region A or B. The blue agent can orient its sensors to see one of three regions, B, C, or D; in order to reorient its sensors from regions B to D, its sensing range must first pass through region C. Similarly, the pink agent can orient its sensors to see either region A, B, or C, and its sensing range must pass through region B to get from regions A to C. The orange agent’s ability to move between adjacent regions is represented by the capability λ_{motion} . Its sensing region is whichever region it is in. $AP_{arm} = \{pickup, dropoff, weed\}$ is an abstraction of a robot manipulator that represents different actions the arm can perform, such as picking up soil samples or pulling weeds. $AP_{UV}, AP_{moisture}, AP_{visual}, AP_{thermal}$ all contain a single proposition representing a agent’s ability to take UV measurements, soil moisture measurements, visual images, and thermal images, respectively. λ_{arm} has more states (see Fig. 1b). Each agent may have distinct cost functions corresponding to individual capabilities.

The agent capabilities and label on the initial state are:

$\Lambda_{green} = \{\lambda_{area_1}, \lambda_{arm}\}, L(s_0) = \{region_B\}$
 $\Lambda_{blue} = \{\lambda_{area_2}, \lambda_{moisture}, \lambda_{UV}\}, L(s_0) = \{region_D\}$
 $\Lambda_{orange} = \{\lambda_{motion}, \lambda_{moisture}, \lambda_{UV}, \lambda_{arm}\}, L(s_0) = \{region_E\}$
 $\Lambda_{pink} = \{\lambda_{area_4}, \lambda_{thermal}, \lambda_{visual}, \lambda_{moisture}, \lambda_{UV}\}, L(s_0) = \{region_C\}$
 The team receives the task (Eq. 4) and must determine a teaming assignment and behavior to satisfy the task. During execution, the agents must also synchronize with each other when necessary.

5 APPROACH

To find a teaming assignment and synthesize the corresponding synchronization and control, we first automatically generate a Büchi




Figure 2: Agriculture environment and initial agent states. The green, blue, and pink agents are stationary; the orientation of their sensors are indicated by the colored boxes.

automaton \mathcal{B} for the task φ^ψ (Sec. 5.1). Each agent A_j then constructs a product automaton $\mathcal{G}_j = A_j \times \mathcal{B}$ (Sec. 5.2). For each binding $\rho \in AP_\psi$, it checks whether or not it can perform the task associated with that binding by finding a path to an accepting cycle in \mathcal{G}_j . Each agent creates a copy of the Büchi automaton \mathcal{B}_j pruned to remove any unreachable transitions and stores information about which combinations of binding assignments it can do.

For parts of the task that require collaboration (e.g., when a transition calls for actions with bindings $\{1, 2\}$ and $r_{green} = \{1, 2\}, r_{blue} = \{2\}$), we need agents to synchronize. Thus, we synthesize behavior that allows for parallel execution while also guaranteeing that the team’s overall behavior satisfies the global specification.

To find a team of agents that can satisfy the task and their assignments, we need to guarantee that 1) every binding is assigned to at least one agent and 2) the agents synchronize for the collaborative portions of the task. To do so, we first run a depth-first search (DFS) to find a path through the \mathcal{B} to an accepting cycle in which there exists a team of agents such that for every transition in the path, every proposition in AP_ψ is assigned to at least one agent (Sec. 5.4). Each agent then synthesizes behavior to satisfy this path and communicates to other agents when synchronization is necessary.

5.1 Büchi Automaton for an LTL $^\psi$ Formula

When constructing a Büchi automaton for an LTL $^\psi$ specification, we automatically rewrite the specification such that the binding propositions are only over individual atomic proposition $\pi \in AP_\varphi$ (i.e. the formula is composed of π^ρ). For instance, the formula $(\neg pickup \mathcal{U} region_A)^{1 \vee 2}$ is rewritten as $(\neg pickup^1 \mathcal{U} region_A^1) \vee (\neg pickup^2 \mathcal{U} region_A^2)$.

In our running example, we rewrite the formula in Eq. 4a as

$$\diamond (region_B^2 \wedge moisture^2 \wedge UV^2) \wedge region_B^3 \wedge moisture^3 \wedge UV^3 \wedge region_A^1 \wedge pickup^1 \quad (5)$$

Remark 4. In rewriting the specification, negation follows bindings in the order of operations. For example, $\neg pickup^{1 \wedge 2} = \neg pickup^1 \wedge \neg pickup^2$, and $\neg (pickup^{1 \wedge 2}) = \neg (pickup^1 \wedge pickup^2) = \neg (pickup^1) \vee \neg (pickup^2)$.

From AP_φ and AP_ψ , we define the set of propositions AP_φ^ψ , where $\forall \pi \in AP_\varphi$ and $\forall \rho \in AP_\psi, \pi^\rho \in AP_\varphi^\psi$. Given AP_φ^ψ , we automatically translate the specification into a Büchi automaton using Spot [4].

To facilitate control synthesis, we transform any transitions in the Büchi automaton labeled with disjunctive formulas into disjunctive normal form (DNF). We then replace the transition labeled with a DNF formula containing ℓ conjunctive clauses with ℓ transitions between the same states, each labeled with a different conjunction of the original label.

In general, when creating a Büchi automaton from an LTL formula φ , $w \in \Sigma_{\mathcal{B}}$ are Boolean formulas over AP_{φ} , the atomic propositions that appear in φ , as seen in Fig. 3. In the following, for creating the product automaton, we use an equivalent representation, where $\Sigma_{\mathcal{B}} = 2^{AP_{\varphi}^{\psi}} \times 2^{AP_{\varphi}^{\psi}}$ and $w = (\sigma_T, \sigma_F) \in \Sigma_{\mathcal{B}}$ contains the set of propositions that must be true, σ_T , and the set of propositions that must be false, σ_F , for the Boolean formula over a transition to evaluate to True. These sets are unique in our case since each transition is labeled with a conjunctive clause (i.e. no disjunction). Note that $\sigma_T \cap \sigma_F = \emptyset$ and $\sigma_T \cup \sigma_F \subseteq AP_{\varphi}$: propositions that do not appear in w can have any truth value.

Given a Büchi automaton for an LTL ^{ψ} specification \mathcal{B} , we define the following functions:

Definition 1 (Binding Function). $\mathfrak{B} : \Sigma_{\mathcal{B}} \rightarrow 2^{AP_{\psi}}$ such that for $\sigma = (\sigma_T, \sigma_F) \in \Sigma_{\mathcal{B}}$, $\mathfrak{B}(\sigma) \subseteq AP_{\psi}$ is the set $\{\rho \in AP_{\psi} \mid \exists \pi^{\rho} \in \sigma_T \cup \sigma_F\}$. Intuitively, it is the set of bindings that appear in label σ of a Büchi transition.

Definition 2 (Capability Function). $\mathfrak{C} : \Sigma_{\mathcal{B}} \times AP_{\psi} \rightarrow 2^{AP_{\varphi}} \times 2^{AP_{\varphi}}$ such that for $\sigma = (\sigma_T, \sigma_F) \in \Sigma_{\mathcal{B}}$, $\rho \in AP_{\psi}$, $\mathfrak{C}(\sigma, \rho) = (C_T, C_F)$, where $C_T = \{\pi \in AP_{\varphi} \mid \exists \pi^{\rho} \in \sigma_T\}$ and $C_F = \{\pi \in AP_{\varphi} \mid \exists \pi^{\rho} \in \sigma_F\}$. Here, C_T and C_F are the sets of action propositions that are True/False and appear with binding ρ in label σ of a Büchi transition.

5.2 Agent Behavior for an LTL ^{ψ} Specification

To synthesize behavior for an agent, we find an accepting trace in its product automaton $\mathcal{G}_j = A_j \times \mathcal{B}$, where $A_j = (S, s_0, AP_j, \gamma, L, W)$ is the agent model, and $\mathcal{B} = (Z, z_0, \Sigma_{\mathcal{B}}, \delta_{\mathcal{B}}, F)$ is the Büchi automaton.

Since the set of propositions of A_j may not be equivalent to the set of propositions of \mathcal{B} , we borrow from the definition of the product automaton in [6]. We first define the following function:

Definition 3 (Binding Assignment Function). Let $q = (s, z)$, $q' = (s', z')$, $\sigma = (\sigma_T, \sigma_F) \in \Sigma_{\mathcal{B}}$. Then $\mathfrak{R}(q, \sigma, q') = \{r \in 2^{AP_{\psi}} \setminus \emptyset \mid \forall \rho \in r, (C_T, C_F) = \mathfrak{C}(\sigma, \rho), \bigcup_{\rho \in r} C_T \subseteq L(s') \text{ and } \bigcup_{\rho \in r} C_F \cap L(s') = \emptyset\}$.

Intuitively, \mathfrak{R} outputs all possible combinations of binding propositions that the agent can be assigned for a transition (q, σ, q') . An agent can be assigned ρ if and only if the agent's next state s' is labeled with all the action and motion propositions $\pi \in AP_{\varphi}$ that appear in σ_T as π^{ρ} , and all the propositions $\pi \in AP_{\varphi}$ that appear in σ_F as π^{ρ} are not part of the state label (i.e. the agent is not performing that action). If a proposition π^{ρ} is in σ_F and π is not in AP_j (e.g. $scan^1 \in \sigma_F$ and the agent does not have λ_{scan}), the agent may be assigned ρ . Note that r may include any binding propositions that are not in σ , since there are no actions required by those bindings in that transition. For example, if $\sigma = (\{scan^1\}, \{pickup^2\})$ and $AP_{\psi} = \{1, 2, 3\}$, then $\{3\}$ will be in the set $\mathfrak{R}(q, \sigma, q')$ for all q, q' .

Given A_j and \mathcal{B} , we define the product automaton $\mathcal{G}_j = A_j \times \mathcal{B}$: **Definition 4** (Product Automaton). The product automaton $\mathcal{G}_j = (Q, q_0, AP_j, \delta_{\mathcal{G}}, L_{\mathcal{G}}, W_{\mathcal{G}}, F_{\mathcal{G}})$, where

- $Q = S \times Z$ is a finite set of states
- $q_0 = (s_0, z_0) \in Q$ is the initial state

- $\delta_{\mathcal{G}} \subseteq Q \times Q$ is the transition relation, where for $q = (s, z)$ and $q' = (s', z')$, $(q, q') \in \delta_{\mathcal{G}}$ if and only if $(s, s') \in \gamma$ and $\exists \sigma \in \Sigma_{\mathcal{B}}$ such that $(z, \sigma, z') \in \delta_{\mathcal{B}}$ and $\mathfrak{R}(q, \sigma, q') \neq \emptyset$
- $L_{\mathcal{G}}$ is the labeling function s.t. for $q = (s, z)$, $L_{\mathcal{G}}(q) = L(s) \subseteq AP_j$
- $W_{\mathcal{G}} : \delta_{\mathcal{G}} \rightarrow \mathbb{R}_{\geq 0}$ is the cost function s.t. for $(q, q') \in \delta_{\mathcal{G}}$, $q = (s, z)$, $q' = (s', z')$, $W_{\mathcal{G}}((q, q')) = W((s, s'))$
- $F_{\mathcal{G}} = S \times F$ is the set of accepting states

Example. Fig. 4 depicts a small portion of \mathcal{G}_{green} ; for the self-transition in \mathcal{B} that is labeled with $\sigma = (\emptyset, \{pickup^1, region_A^2\})$ (labeled as $e1$ in Fig. 3), and for states in A_{green} where $L(s_1) = \{region_B\}$, $L(s_2) = \{region_A\}$, $L(s_3) = \{region_A, pickup\}$, then the possible binding assignments are $\mathfrak{R}((s_1, 1), \sigma, (s_1, 2)) = 2^{\{1, 2, 3\}} \setminus \emptyset$ and $\mathfrak{R}((s_1, 1), \sigma, (s_2, 2)) = \{\{1\}, \{3\}, \{1, 3\}\}$. When the agent is in s_3 , it cannot be assigned either bindings 1 or 2, but since no propositions appear with binding 3 in σ , $\mathfrak{R}((s_1, 1), \sigma, (s_3, 2)) = \{\{3\}\}$.

5.3 Finding Possible Individual Agent Bindings

To construct a team, we first reason about each agent and the sets of bindings it can perform. For example, for a formula $region_A^1 \wedge region_B^2$, an agent may be assigned $r_j = \{1\}$ or $r_j = \{2\}$ but not $r_j = \{1, 2\}$, since it cannot be in two regions at the same time.

To find the set of possible binding assignments $R_j \subseteq 2^{AP_{\psi}}$, we search for an accepting trace in \mathcal{G}_j for every binding assignment $r_j \in 2^{AP_{\psi}}$. We start from the full set of bindings $r_j = AP_{\psi}$. Given an assignment r_j to check, we find an accepting trace in \mathcal{G}_j such that for all transitions (q, q') in the trace, $r_j \subseteq \mathfrak{R}(q, \sigma, q')$. This ensures that the agent can satisfy its binding assignment for the entirety of its execution (i.e. r_j does not change). Since every subset of a binding assignment r_j is itself a possible binding assignment, if the agent can be assigned all $m = |AP_{\psi}|$ bindings, then we know it can also be assigned every possible subset of m . If not, we check the $\binom{m}{m-1}$ combinations, and continue iterating until we have determined the agent's ability to perform every combination of the m bindings.

Once an agent determines its possible binding assignments R_j , it creates the Büchi automaton \mathcal{B}_j by removing any transition in \mathcal{B} that cannot be traversed by any assignment in R_j . In our example (Fig. 3), each agent can be assigned at least one binding over every transition in \mathcal{B} . Thus, $\forall j \in \{green, blue, orange, pink\}, \mathcal{B}_j = \mathcal{B}$.

5.4 Agent Team Assignment

A team of agents can perform the task if 1) all the bindings are assigned, with each agent maintaining the same binding assignment for the entirety of the task, and 2) the agents satisfy synchronization requirements. For a viable team, the agents' control follows the same path in the Büchi automaton \mathcal{B} to an accepting cycle. We perform DFS over \mathcal{B} to find an accepting trace (Alg. 1), where each tuple in *stack* contains the current edge (z, σ, z') , the current team of agents $R_{\hat{A}}$, and the path traversed so far $\beta_{\hat{A}}$.

We initialize the team with all agents A_j and all possible binding assignments R_j , and each path $\beta_{\hat{A}}$ starts from state z_0 of \mathcal{B} . When checking a transition (z, σ, z') , we remove any agent j if $\forall ((s, z), (s', z')) \in \delta_{\mathcal{G}_j}$, there are no possible binding assignments it can satisfy. This is done by checking each agent's pruned Büchi automaton \mathcal{B}_j in `UPDATE_TEAM` (line 8). We want the agent's behavior to satisfy not only the current transition, but also the entire path

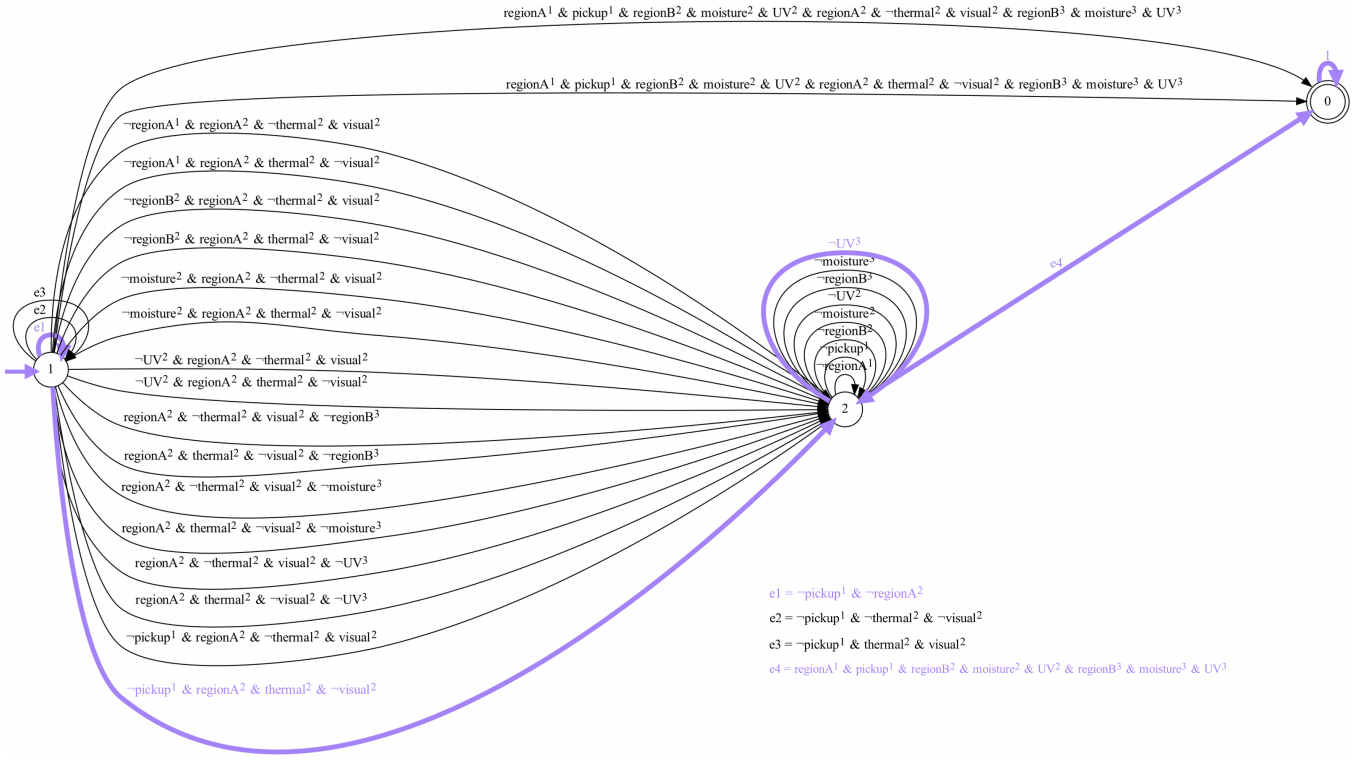


Figure 3: \mathcal{B} for ϕ^ψ (Eq. 4). The purple transitions illustrate a possible accepting trace.

with a consistent binding assignment. Thus, we update possible bindings (UPDATE_BINDINGS, lines 9-14).

To guarantee the overall team behavior, we need to ensure agents are able to “wait in a state” before they synchronize, as they may reach states at different times. This means that each state in the trace must have a corresponding self-transition. Thus, for every (z, σ, z') that we add to the path in which $z \neq z'$, the next edge to traverse must be a self-transition from z' to itself; the same holds vice-versa. In line 21, we check if the current transition is self-looping or not, and add subsequent transitions into the stack accordingly. If there is no self-transition on z' (i.e. $(z', \sigma, z') \notin \delta_{\mathcal{B}}$), then we do not consider z' to be valid and do not add it to the path.

Once we find a valid path to an accepting cycle, we parse it into β , the path without self-transitions, and δ_{self} , which contains the corresponding self-transition for each state in the path. Fig. 3 shows

a valid path in \mathcal{B} for the example in Sec. 4 and the corresponding team assignment $\hat{A} = \{A_{green}, A_{blue}, A_{orange}, A_{pink}\}$ and bindings $r_{green} = \{1\}, r_{blue} = \{3\}, r_{orange} = \{1\}, r_{pink} = \{2, 3\}$. Note that we find a valid path rather than a globally optimal one. However, the algorithm is complete; it will find a feasible path if one exists.

5.5 Synthesis and Execution of Control and Synchronization Policies

Given an accepting trace β through \mathcal{B} and the corresponding self-transitions δ_{self} that are valid for all agents in $R_{\hat{A}}$, we synthesize control and synchronization for each agent such that the overall team execution satisfies β (Alg. 2). For each transition (z, σ, z') in β , we find \bar{R} , which contains the binding assignments of all agents that require synchronization at state z' . Agent j participates in the synchronization step if r_j contains a binding ρ that is required by σ and is not the only agent assigned bindings from σ (line 3).

Subsequently, agent j finds an accepting trace in \mathcal{G}_j that reaches z' with minimum cost, following self-transitions stored in δ_{self} if necessary. As it executes this behavior, it communicates with other agents the tuple p , which contains 1) its ID, 2) the state z' it is currently going to, and 3) if it is ready for synchronization (line 8). If no synchronization is required (line 3), the agent can simply execute the behavior. Otherwise, to guarantee that the behavior does not violate the requirements of the task, the agent executes the synthesized behavior up until the penultimate state, z_{wait} .

When the agent reaches z_{wait} , it signals to other agents that it is ready for synchronization. Since all agents know the overall




Figure 4: A small portion of \mathcal{G}_{green}

Algorithm 1: Find Accepting Trace for Agent Team

```

Input :  $A = \{A_1, A_2, \dots, A_n\}$ ,  $R = \{R_1, R_2, \dots, R_n\}$ ,  $\mathcal{B}$ ,
          $\{\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_n\}$ 
Output:  $\beta$ ,  $\delta_{self}$ ,  $\hat{A} \subseteq A$ ,  $R_{\hat{A}}$ 
1  $stack = \emptyset$ ,  $visited = \emptyset$ 
2 for  $e \in \{(z, \sigma, z') \in \delta_{\mathcal{G}} \mid z = z_0\}$  do
3    $stack = stack \cup \{(e, R, [e])\}$ 
4 while  $stack \neq \emptyset$  do
5    $((z, \sigma, z'), R_{\hat{A}}, \beta_{\hat{A}}) = stack.pop()$ 
6   if  $(z, \sigma, z') \notin visited$  then
7      $visited = visited \cup (z, \sigma, z')$ 
8      $R_{\hat{A}} = UPDATE\_TEAM((z, \sigma, z'), \{\mathcal{B}_1, \dots, \mathcal{B}_n\})$ 
9     for  $R_j \in R_{\hat{A}}$  do
10       $R'_j = UPDATE\_BINDINGS(R_j, (z, \sigma, z'))$ 
11      if  $R'_j = \emptyset$  then
12         $R_{\hat{A}} = R_{\hat{A}} \setminus R_j$ 
13      else
14         $R_{\hat{A}} = (R_{\hat{A}} \setminus R_j) \cup R'_j$ 
15      if  $\bigcup_j (R_j \in R_{\hat{A}}) = AP_{\psi}$  then
16        if  $z' \in F$  then
17           $\beta, \delta_{self} = PARSE\_PATH(\beta_{\hat{A}})$ 
18          return  $\beta, \delta_{self}, R_{\hat{A}}$ 
19           $E = \{(z', \sigma', z'') \in \delta_{\mathcal{G}}\}$ 
20          for  $(z', \sigma', z'') \in E$  do
21            if  $(z = z' \text{ and } z' \neq z'')$  or
22               $(z \neq z' \text{ and } z' = z'')$  then
                 $stack = stack \cup$ 
                 $\{(z', \sigma', z''), R_{\hat{A}}, [\beta_{\hat{A}}(z', \sigma', z'')]\}$ 

```

teaming assignment, the agent continues to wait in state z_{wait} until it receives a signal that all other agents in \bar{R} are ready (line 13). These agents then move to the next state in the behavior simultaneously. Agent j continues synthesizing behavior through β until synchronization is necessary again, and this process is repeated.

6 RESULTS AND DISCUSSION

Fig. 5 shows the final step of the synchronized behavior of the agents for the example in Section 4, where $\hat{A} = \{A_{green}, A_{blue}, A_{orange}, A_{pink}\}$ with binding assignments $r_{green} = \{1\}$, $r_{blue} = \{3\}$, $r_{orange} = \{1\}$, $r_{pink} = \{2, 3\}$. A simulation of the full behavior is shown in the accompanying video¹.

Optimizing teams: Our synthesis algorithm can be seen as a greatest fixpoint computation, where we start with the full set of agents and remove those that cannot contribute to the task. As a result, the team may have redundancies, i.e. agents can be removed while still ensuring the overall task will be completed; this may be beneficial for robustness. Furthermore, we can choose a sub-team to optimize different metrics, as long as the agent bindings assignments still cover all the required bindings. For example, minimizing the number of bindings per agent could result in $\hat{A} = \{A_{green}, A_{blue}, A_{pink}\}$,

Algorithm 2: Synthesize an Agent's Behavior

```

Input :  $\mathcal{G}_j, r_j, R_{\hat{A}}, \beta, \delta_{self}$ 
1 for  $(z, \sigma, z') \in \beta$  do
2    $b_j = FIND\_BEHAVIOR(\mathcal{G}_j, r_j, (z, \sigma, z'), \delta_{self})$ 
3    $\bar{R} = \{r_k \in R_{\hat{A}} \mid r_k \cap \mathfrak{B}(\sigma) \neq \emptyset\}$  if  $r_j \notin \bar{R}$  or  $\bar{R} = \{r_j\}$ 
4     then
5        $p = ()$ 
6       EXECUTE( $b_j, p$ )
7     else
8        $p = (j, z', 0)$ ,  $\ell = length(b_j)$ 
9       EXECUTE( $b_j[1 : \ell - 1], p$ )
10       $z_{wait} = b_j[\ell - 1]$ ,  $P = \{j\}$ 
11      while  $\bigcup_{i \in P} (r_i \in \bar{R}) \neq \mathfrak{B}(\sigma)$  do
12         $p = (j, z', 1)$ 
13        EXECUTE( $z_{wait}, p$ )
14         $P = j \cup \{k \mid (k, z', 1) \in RECEIVE()\}$ 
15      EXECUTE( $b_j[\ell]$ )

```

$r_{green} = \{1\}$, $r_{blue} = \{3\}$, $r_{pink} = \{2\}$; minimizing the number of agents results in $\hat{A} = \{A_{green}, A_{pink}\}$, $r_{green} = \{1\}$, $r_{pink} = \{2, 3\}$.

To illustrate other possible metrics, we consider a set of 20 agents and create a team for the specification in Eq. 4. Their final binding assignments and costs are shown in Table 1. Minimizing cost results in a team $\hat{A} = \{A_7, A_{11}\}$. Minimizing cost while requiring each binding to be assigned to two agents results in $\hat{A} = \{A_4, A_7, A_{11}, A_{16}\}$.




Figure 5: The final step in the synchronized behavior of the agent team with their corresponding actions.

Computational complexity: The control synthesis algorithm (Alg. 2) is agnostic to the number of agents, since each agent determines its own possible binding assignments and behavior. For the team assignment (Alg. 1), we store the current team and possible binding assignments as we build an accepting trace. Thus, it has both a space and time complexity of $O(|E| * 2^m * n)$, where $|E|$ is the number of edges in \mathcal{B} , m is the number of bindings, and n is the number of agents.

Fig. 6a shows the computation time of the synthesis framework (Sec. 5.2 – 5.4) for simulated agent teams in which we vary the number of agents from 3 to 20, running 30 simulations for each set of agents and randomizing their capabilities. The task for each simulation is the example in Eq. 4. We also ran simulations in which

¹<https://youtu.be/Qx14xGza3Bg>

Agent	r_j	cost	Agent	r_j	cost	Agent	r_j	cost	Agent	r_j	cost	Agent	r_j	cost
1	1	1.2	5	3	2.75	9	3	2.6	13	3	2.0	17	2,3	3.275
2	3	1.0	6	1	0.95	10	1	2.8	14	1	1.2	18	3	2.55
3	1	1.2	7	1	0.65	11	2,3	0.9	15	3	1.1	19	1	1.9
4	2,3	1.3	8	1	1.0	12	2,3	1.825	16	1	0.775	20	2,3	2.35

Table 1: Example teaming assignment with 20 robots




Figure 6: Computation time as the number of agents (a) and bindings (b) increases. The error bars show min/max values.

we increase the number of bindings from 3 to 10 and randomized the capabilities of 4 agents (Fig. 6b). The variance in computation time is a result of the randomized agent capabilities, which affects the computation time of possible binding assignments (Sec. 5.3). All simulations ran on a 2.5 GHz quad-core Intel Core i7 CPU.

Task expressivity: We compare LTL^ψ with other temporal logic approaches to encode collaborative heterogeneous multi-agent tasks.

Standard LTL: One approach is to use LTL to express the task by enumerating all possible assignments in the specification. In our example, Eq. 4a would be rewritten as:

$$\begin{aligned} \varphi_1^\psi = & (\diamond((region_B^g \wedge mois^g \wedge UV^g) \wedge (region_A^b \wedge pickup^b))) \\ & \vee (\diamond((region_B^g \wedge mois^g \wedge UV^g) \wedge (region_A^o \wedge pickup^o))) \vee \dots \end{aligned}$$

where each agent has its own unique set of AP, denoted here by each proposition’s superscript. As a result, the number of propositions increases exponentially with the number of agents. The task complexity also increases, as the specification must include all possible agent assignments. Another drawback of using LTL for such tasks is that the specification is not generalizable to any number of agents; it must be rewritten when the set of agents change.

LTL $^\chi$: In [15], tasks are written in LTL $^\chi$, where proposition $\pi_{i,j}^{k,\chi}$ is true if at least i agents of type j are in region k with binding

χ . We can express φ_1^ψ (Eq. 4a) of our example as $\diamond(\pi_{1,mois}^{regionB,2} \wedge \pi_{1,UV}^{regionB,2} \wedge \pi_{1,mois}^{regionB,3} \wedge \pi_{1,UV}^{regionB,3} \wedge \pi_{1,arm}^{regionA,1})$. The truth value of $\pi_{i,j}^{k,\chi}$ is not dependent on any particular action an agent might take. LTL $^\chi$ can be extended to action propositions, but since an agent can only be categorized as one type, each type of agent must have non-overlapping capabilities (here, we have written the LTL $^\chi$ formula such that each type of agent only has one capability). In addition, φ_2^ψ (Eq. 4b) cannot be written in LTL $^\chi$ because the negation defined in our grammar cannot be expressed in LTL $^\chi$. On the other hand, the negative proposition $\neg\pi_{i,j}^{k,\chi}$ from [15] is equivalent to “less than i agents of type j are in region k ”, which our logic cannot encode.

Capability Temporal Logic (CaTL): Tasks in CaTL [13] are constructed over tasks $T = (d, \pi, cp_T)$, where d is a time duration, π is a region, $(c_i, m_i) \in cp_T$ denotes that at least m_i agents with capability c_i are required. Similar to our grammar, CaTL allows agents to have multiple capabilities, but each task must specify the number of agents required. Since it is an extension of STL, tasks provide timing requirements, which our logic cannot encode. However, it does not include the concept of binding assignments; in our example φ_1^ψ (Eq. 4a), CaTL cannot express that the same agent that took a UV measurement must also take a thermal image. Ignoring binding assignments and adding timing constraints, φ_1^ψ (Eq. 4a) can be rewritten in CaTL as $\diamond_{[0,10]}(T(0.1, region_B, \{(moisture, 2), (UV, 2)\}) \wedge T(0.5, region_A, \{(arm, 1)\}))$. Each capability in CaTL is represented as a sensor and therefore cannot include more complex capabilities, e.g. a robot arm that can perform several different actions. In addition, because CaTL requires the formula to be in positive normal form (i.e. no negation), we cannot express φ_2^ψ (Eq. 4b) in this grammar.

7 CONCLUSION

We define a new task grammar for heterogeneous teams of agents and develop a framework to automatically assign the task to a (sub)team of agents and synthesize correct-by-construction control policies to satisfy the task. We include synchronization constraints to guarantee that the agents perform the necessary collaborations.

In the future, we plan to demonstrate the approach on physical systems where we must ensure that the continuous execution satisfies all safety constraints. We will also explore different notions of optimality when finding a teaming plan, and increase expressivity of the grammar to allow for reactive tasks where agents modify their behavior at runtime in response to environment events.

ACKNOWLEDGMENTS

This work is supported by the the National Defense Science & Engineering Graduate (NDSEG) Fellowship Program.

REFERENCES

- [1] Christel Baier and Joost-Pieter Katoen. 2008. *Principles of Model Checking*. The MIT Press.
- [2] Ji Chen, Ruojia Sun, and Hadas Kress-Gazit. 2021. Distributed Control of Robotic Swarms from Reactive High-Level Specifications. In *2021 IEEE 17th International Conference on Automation Science and Engineering (CASE)*. 1247–1254. <https://doi.org/10.1109/CASE49439.2021.9551578>
- [3] Yushan Chen, Xu Chu Ding, and Calin Belta. 2011. Synthesis of distributed control and communication schemes from global LTL specifications. In *2011 50th IEEE Conference on Decision and Control and European Control Conference*. 2718–2723. <https://doi.org/10.1109/CDC.2011.6160740>
- [4] Alexandre Duret-Lutz, Alexandre Lewkowicz, Amaury Fauchille, Thibaud Michaud, Etienne Renault, and Laurent Xu. 2016. Spot 2.0 – a framework for LTL and ω -automata manipulation. In *Proceedings of the 14th International Symposium on Automated Technology for Verification and Analysis (ATVA'16) (Lecture Notes in Computer Science, Vol. 9938)*. Springer, 122–129. https://doi.org/10.1007/978-3-319-46520-3_8
- [5] E. Allen Emerson. 1990. Temporal and Modal Logic. In *Formal Models and Semantics*, JAN Van Leeuwen (Ed.). Elsevier, Amsterdam, 995–1072. <https://doi.org/10.1016/B978-0-444-88074-1.50021-4>
- [6] Amy Fang and Hadas Kress-Gazit. 2022. Automated Task Updates of Temporal Logic Specifications for Heterogeneous Robots. In *2022 International Conference on Robotics and Automation (ICRA)*. 4363–4369. <https://doi.org/10.1109/ICRA46639.2022.9812045>
- [7] Fatma Faruq, David Parker, Bruno Lacorda, and Nick Hawes. 2018. Simultaneous Task Allocation and Planning Under Uncertainty. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 3559–3564. <https://doi.org/10.1109/IROS.2018.8594404>
- [8] Brian P. Gerkey and Maja J. Matarić. 2004. A Formal Analysis and Taxonomy of Task Allocation in Multi-Robot Systems. *The International Journal of Robotics Research* 23, 9 (2004), 939–954. <https://doi.org/10.1177/0278364904045564> arXiv:<https://doi.org/10.1177/0278364904045564>
- [9] Xiao Jia and Max Q.-H. Meng. 2013. A survey and analysis of task allocation algorithms in multi-robot systems. In *2013 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. 2280–2285. <https://doi.org/10.1109/ROBIO.2013.6739809>
- [10] 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. <https://doi.org/10.1177/0278364920913922> arXiv:<https://doi.org/10.1177/0278364920913922>
- [11] Marius Kloetzer and Calin Belta. 2010. Automatic Deployment of Distributed Teams of Robots From Temporal Logic Motion Specifications. *IEEE Transactions on Robotics* 26, 1 (2010), 48–61. <https://doi.org/10.1109/TRO.2009.2035776>
- [12] G. Ayorkor Korsah, Anthony Stentz, and M. Bernardine Dias. 2013. A comprehensive taxonomy for multi-robot task allocation. *The International Journal of Robotics Research* 32, 12 (2013), 1495–1512. <https://doi.org/10.1177/0278364913496484> arXiv:<https://doi.org/10.1177/0278364913496484>
- [13] Kevin Leahy, Zachary Serlin, Cristian-Ioan Vasile, Andrew Schoer, Austin M. Jones, Roberto Tron, and Calin Belta. 2022. Scalable and Robust Algorithms for Task-Based Coordination From High-Level Specifications (ScrATChES). *IEEE Transactions on Robotics* 38, 4 (2022), 2516–2535. <https://doi.org/10.1109/TRO.2021.3130794>
- [14] Zhiyong Li, Bo Xu, Lei Yang, Jun Chen, and Kenli Li. 2009. Quantum Evolutionary Algorithm for Multi-Robot Coalition Formation. In *Proceedings of the First ACM/SIGEVO Summit on Genetic and Evolutionary Computation (Shanghai, China) (GEC '09)*. Association for Computing Machinery, New York, NY, USA, 295–302. <https://doi.org/10.1145/1543834.1543874>
- [15] Xusheng Luo and Michael M. Zavlanos. 2022. Temporal Logic Task Allocation in Heterogeneous Multirobot Systems. *IEEE Transactions on Robotics* (2022), 1–20. <https://doi.org/10.1109/TRO.2022.3181948>
- [16] Yunus Emre Sahin, Petter Nilsson, and Necmiye Ozay. 2017. Synchronous and asynchronous multi-agent coordination with cLTL+ constraints. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. 335–342. <https://doi.org/10.1109/CDC.2017.8263687>
- [17] 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. <https://doi.org/10.1177/0278364918774135> arXiv:<https://doi.org/10.1177/0278364918774135>
- [18] Thomas Schmickl, Christoph Möslinger, and Karl Crailsheim. 2006. Collective Perception in a Robot Swarm. 144–157. https://doi.org/10.1007/978-3-540-71541-2_10
- [19] Jana Tumova and Dimos V. Dimarogonas. 2016. Multi-agent planning under local LTL specifications and event-based synchronization. *Automatica* 70 (2016), 239–248. <https://doi.org/10.1016/j.automatica.2016.04.006>
- [20] Alphan Ulusoy, Stephen L. Smith, Xu Chu Ding, and Calin A. Belta. 2012. Robust multi-robot optimal path planning with temporal logic constraints. *2012 IEEE International Conference on Robotics and Automation* (2012), 4693–4698.
- [21] Hanlin Wang and Michael Rubenstein. 2020. Shape Formation in Homogeneous Swarms Using Local Task Swapping. *IEEE Transactions on Robotics* 36, 3 (2020), 597–612. <https://doi.org/10.1109/TRO.2020.2967656>
- [22] Bo Xu, Zhaofeng Yang, Yu Ge, and Zhiping Peng. 2015. Coalition Formation in Multi-agent Systems Based on Improved Particle Swarm Optimization Algorithm. *International Journal of Hybrid Information Technology* 8 (03 2015), 1–8. <https://doi.org/10.14257/ijhit.2015.8.3.01>