Risk-Aware Real-Time Task Allocation for Stochastic Multi-Agent Systems under STL Specifications

M. H. W. Engelaar[1], Z. Zhang[1], E. E. Vlahakis[2], D.V. Dimarogonas[2], M. Lazar[1], and S. Haesaert[1]

***Abstract*— This paper addresses the control synthesis of heterogeneous stochastic linear multi-agent systems with real-time allocation of signal temporal logic (STL) specifications. Based on previous work, we decompose specifications into sub-specifications on the individual agent level. To leverage the efficiency of task allocation, a heuristic filter evaluates potential task allocation based on STL robustness, and subsequently, an auctioning algorithm determines the definitive allocation of specifications. Finally, a control strategy is synthesized for each agent-specification pair using tube-based model predictive control (MPC), ensuring provable probabilistic satisfaction. We demonstrate the efficacy of the proposed methods using a multi-shuttle scenario that highlights a promising extension to automated driving applications like vehicle routing.**

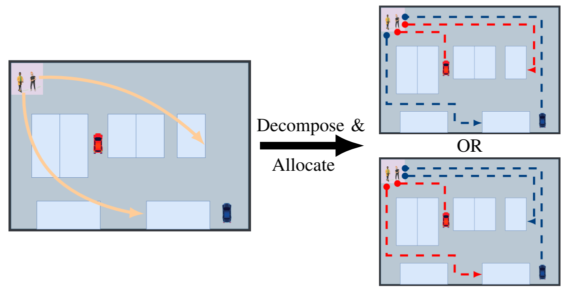

Fig. 1. Decomposing and allocating tasks for vehicle routing.

I. Introduction

Multi-agent systems are frequently employed in complex tasks that require close collaborations. Formal specifications, such as signal temporal logic (STL), are increasingly utilized to define complex collaborative task objectives [1]. The collaborative control of multi-agent systems satisfying STL specifications has been studied using mixed-integer programming [2], funnel functions [3], and control barrier functions [4]. Nevertheless, a critical challenge for the practical implementation of these methods is their reliability in uncertain and dynamic environments. For systems subject to stochastic disturbances, it is important to evaluate the impact of these disturbances on the task accomplishment of individual agents so that the associated risks are restricted [5].

Despite recent efforts, a notable gap exists in the literature regarding the control of stochastic multi-agent systems under real-time allocated STL specifications with probabilistic guarantees. Such systems have significant practical value, exemplified by scenarios such as vehicle routing [6]. For example, consider a scenario where two passengers require taxi rides, as illustrated in Fig. 1. This naturally renders a task allocation problem by assigning each passenger a taxi. Finding the most efficient assignment is a critical challenge.

Inspired by the example, this paper addresses 1) the problem of task allocation (TA) of STL specifications, which aims to determine an efficient agent-specification mapping in a multi-agent setting and 2) the control synthesis of agent-specification pairs assigned at runtime with probabilistic guarantees. TA necessitates decomposing specifications assigned to multiple collaborative agents into individual agent specifications, requiring the decomposition of STL specifications both among agents [7], [8] and across time [9]. Previous work yields a linear programming problem [10], efficiently solvable using auction algorithms [11]. With decomposed real-time specifications, control synthesis for stochastic multi-agent systems holds promise with the utilization of shrinking-horizon model predictive control (MPC) [12] to ensure bounded risks [13]. In previous work, probabilistic reachable tubes have been utilized to optimize risk bounds for both predefined specifications [14] and real-time specifications [15]. These methods underpin the risk-aware control synthesis of stochastic multi-agent systems with decomposed real-time specifications.

This paper proposes a novel approach to allocating STL specifications, involving subgroups of heterogeneous agents and synthesizing control strategies for individual agents in real-time. By decomposing newly assigned specifications into agent-level tasks, we employ a heuristic approach based on STL robustness [16] to assess agent-specification pairs and attack task assignment via an auctioning scheme formulated as a mixed integer linear program. Subsequently, a tube-based model predictive control strategy is employed where an optimization is solved at every time step at agent level. The solution of this optimization ensures the satisfaction of previous and newly allocated tasks with a desired risk level, while discarding tasks not adhering to the risk specification. We prove recursive feasibility of this MPC scheme, leading to an open-loop system with provable probabilistic guarantees and a closed-loop system for which probabilistic guarantees follow naturally. As trust is crucial in dynamic multi-agent systems, where agents and environment may often change [17], our contribution is proposing a novel concept of a trustworthy (correct-by-design) control algorithm for multi-agent systems with real-time task assignment.

This work is supported by the Dutch NWO Veni project CODEC, grant number 18244 and the European project SymAware, grant number 101070802.
[1]Department of Electrical Engineering (Control Systems Group), Eindhoven University of Technology, The Netherlands. Emails:{m.h.w.engelaar, z.zhang3, m.lazar, s.haesaert}@tue.nl. [2]Division of Decision and Control Systems, School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, Sweden. Email: {vlahakis, dimos}@kth.se

II. Preliminaries and Problem Statement

For a given probability measure $\mathbb{P}$ defined over a Borel measurable space $(\mathbb{X}, \mathcal{B}(\mathbb{X}))$, we denote the probability of an event $\mathcal{A} \in \mathcal{B}(\mathbb{X})$ as $\mathbb{P}(\mathcal{A})$. In this paper, we will work with Euclidean spaces and Borel measurability. Further details on measurability are omitted, and we refer the interested reader to [18]. Additionally, any half-space $H \subset \mathbb{R}^n$ is defined as $H := \{x \in \mathbb{R}^n \mid g^T x \geq b\}$ with $g \in \mathbb{R}^n$ and $b \in \mathbb{R}$. A polyhedron is the intersection of finitely many half-spaces, also denoted as $H := \{x \in \mathbb{R}^n \mid Gx \geq b\}$ with $G \in \mathbb{R}^{q \times n}$ and $b \in \mathbb{R}^q$. A polytope is a bounded polyhedron. The 2-norm is defined by $||x|| = \sqrt{x^T x}$ with $x \in \mathbb{R}^n$. The identity matrix is denoted by $I_n \in \mathbb{R}^{n \times n}$. The vector of all elements one and zero are denoted by $\mathbf{1}_n \in \mathbb{R}^n$ and $\mathbf{0}_n \in \mathbb{R}^n$, respectively. If the context is clear, they will be denoted by $\mathbf{1}$ and $\mathbf{0}$. For any finite set I, the cardinality is denoted by $|I|$. $\mathrm{Proj}_X(Y)$ denotes the projection of elements in Y onto X.

A. Multi-agent Systems

In this paper, we consider a finite ordered set $\mathcal{I}$ of heterogeneous agents where each agent $i \in \mathcal{I}$ has linear time-invariant dynamics with additive noise, given by

$$x^i(k+1) = A^i x^i(k) + B^i u^i(k) + w^i(k), \tag{1}$$

where (A^i, B^i) is stabilizable. We impose that all agents have the same state and input dimensions. For the i-th agent, $x^i \in \mathbb{X}^i \subseteq \mathbb{R}^n$ is the state, $x^i(0) \in \mathbb{X}^i$ is an initial state, $u^i \in \mathbb{U}^i \subseteq \mathbb{R}^m$ is the input and $w^i \in \mathbb{R}^n$ is an independent, identically distributed (i.i.d.) noise disturbance with distribution $\mathcal{Q}_w^i$, i.e., $w^i(k) \sim \mathcal{Q}_w^i$, which can have infinite support. We will assume that the distribution has at least a known mean and variance, with the latter required to be strictly positive definite. Additionally, we impose that the distribution is central convex unimodal[1]. Finally, we define that any two disturbances $w^i(k)$ and $w^j(l)$ are independent.

For each agent $i \in \mathcal{I}$, we require there is a local controller. We define the local controllers as a sequence of policies

$$\boldsymbol{f}^i := \{f_0^i, f_1^i, \dots\},$$

such that $f_k^i : \mathbb{H}_k^i \to \mathbb{U}^i$ maps the history of states and inputs to inputs. Here $\mathbb{H}_k^i := (\mathbb{X}^i \times \mathbb{U}^i)^k \times \mathbb{X}^i$ with elements $\eta^i(k) := (x^i(0), u^i(0), \cdots, u^i(k-1), x^i(k))$. By implementing controller $\boldsymbol{f}^i$ upon the corresponding agent, we obtain the controlled form of (1) for which the control input satisfies the feedback law $u^i(k) = f_k^i(\eta^i(k))$ with $\eta^i(k) \in \mathbb{H}_k^i$. We indicate the input sequence of agent i by $\boldsymbol{u}^i := \{u^i(0), u^i(1), \dots\}$ and define its executions as sequences of states $\boldsymbol{x}^i := \{x^i(0), x^i(1), \dots\}$ referred to as signals.

We define the suffix, segment and signal element of any signal $\boldsymbol{x}^i := \{x^i(0), x^i(1), \ldots\}$, respectively, by $\boldsymbol{x}_k^i = \{x^i(k), x^i(k+1), \dots\}$, $\boldsymbol{x}_{[a,b]}^i := \{x^i(a), \dots, x^i(b)\}$ with $a \leq b$, and $\boldsymbol{x}^i(k) := \boldsymbol{x}_{[k,k]}^i$. We define the concatenation of $\boldsymbol{x}^i$, for all agents $i \in I \subseteq \mathcal{I}$, by a higher dimensional signal $\boldsymbol{x}^I$, where $\boldsymbol{x}^I(k) := [\boldsymbol{x}^{i_1}(k)^T, \cdots, \boldsymbol{x}^{i_l}(k)^T]^T$, $i_j \in I$, and $l = |I|$ is the number of elements in I. Any signal $\boldsymbol{x}^i$ can be interpreted as a realization of the probability distribution induced by implementing controller $\boldsymbol{f}^i$, denoted by $\boldsymbol{x}^i \sim \mathbb{P}_{\boldsymbol{f}^i}$. Similar realizations exist for segments, i.e., $\boldsymbol{x}_k^i \sim \mathbb{P}_{\boldsymbol{f}^i, k}$ and concatenations, i.e., $\boldsymbol{x}^I \sim \mathbb{P}_{\boldsymbol{f}^I}$, with $\boldsymbol{f}^I := \{\boldsymbol{f}^i \mid i \in I\}$.

B. Signal Temporal Logic & Probabilistic Satisfaction

To mathematically describe tasks, e.g., reaching a target within a given time frame, we consider specifications given by signal temporal logic (STL). Here, we assume that all STL specifications adhere to the *negation normal form* (NNF), see [1]. This assumption does not restrict the overall framework since [20] shows that every STL specification can be rewritten into the negation normal form.

STL consists of predicates μ^l for subgroups of l agents that are either true ($\top$) or false ($\bot$). Each predicate μ^l is described by a function $h : \mathbb{R}^{nl} \to \mathbb{R}^q$, $l \leq |\mathcal{I}|$, as follows

$$\mu^l := \begin{cases} \top \text{ if } \forall i \in \{1, \cdots, q\},\ h_i(x) \geq 0 \\ \bot \text{ if } \exists i \in \{1, \cdots, q\},\ h_i(x) < 0. \end{cases} \tag{2}$$

Here $h_i : \mathbb{R}^{nl} \to \mathbb{R}$ indicates the i-th element of function h. We will assume that all predicate functions are affine functions of the form $h(x) = Gx - b$ with $G \in \mathbb{R}^{q \times nl}$, $b \in \mathbb{R}^q$ and $||g_j|| = 1$ for all $j \in \{1, \cdots, q\}$, where g_j is the j-th row of G. Furthermore, we require that $H_{\mu^l} := \{x \in \mathbb{R}^{nl} \mid Gx \geq b\}$ describes a polytope. Let the set of all predicates be given by $\mathcal{P}$. The STL syntax will be given by

$$\phi ::= \top \mid \mu^l \mid \neg\mu^l \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \mid \square_{[a,b]}\phi \mid \phi_1\ U_{[a,b]}\ \phi_2,$$

where $\mu^l \in \mathcal{P}$, ϕ, ϕ_1 and ϕ_2 are STL formula, $a, b \in \mathbb{N}$, and $a \leq b$. The semantics are given next, where $\boldsymbol{x}_k^{\mathcal{I}} \vDash \phi$ denotes the satisfaction of ϕ verified over the suffix of signal $\boldsymbol{x}^{\mathcal{I}}$.

Definition 1: The STL semantics, for a specification jointly assigned to a set of agents, are recursively given by:

$$\begin{aligned}
&\boldsymbol{x}_k^{\mathcal{I}} \vDash \mu^l && \iff \exists I \subseteq \mathcal{I}:\ |I| = l,\ h(\boldsymbol{x}^I(k)) \geq \mathbf{0} \\
&\boldsymbol{x}_k^{\mathcal{I}} \vDash \neg\mu^l && \iff \nexists I \subseteq \mathcal{I}:\ |I| = l,\ h(\boldsymbol{x}^I(k)) \geq \mathbf{0} \\
&\boldsymbol{x}_k^{\mathcal{I}} \vDash \phi_1 \wedge \phi_2 && \iff \boldsymbol{x}_k^{\mathcal{I}} \vDash \phi_1 \text{ and } \boldsymbol{x}_k^{\mathcal{I}} \vDash \phi_2 \\
&\boldsymbol{x}_k^{\mathcal{I}} \vDash \phi_1 \vee \phi_2 && \iff \boldsymbol{x}_k^{\mathcal{I}} \vDash \phi_1 \text{ or } \boldsymbol{x}_k^{\mathcal{I}} \vDash \phi_2 \\
&\boldsymbol{x}_k^{\mathcal{I}} \vDash \square_{[a,b]}\ \phi && \iff \forall k' \in [k+a, k+b] : \boldsymbol{x}_{k'}^{\mathcal{I}} \vDash \phi \\
&\boldsymbol{x}_k^{\mathcal{I}} \vDash \phi_1\ U_{[a,b]}\ \phi_2 && \iff \exists k_1 \in [k+a, k+b] : \boldsymbol{x}_{k_1}^{\mathcal{I}} \vDash \phi_2 \\
&&& \quad \text{and } \forall k_2 \in [k, k_1],\ \boldsymbol{x}_{k_2}^{\mathcal{I}} \vDash \phi_1.
\end{aligned}$$

Here, we introduced a slight abuse of the notation as $[a, b]$ is used to describe the set of all natural numbers within the real interval $[a, b]$. Note that the definition is invariant of the actual indices of the l agents. Additional operators can be derived such as the eventually-operator $\lozenge_{[a,b]}\phi := \top U_{[a,b]}\phi$.

Since all agents behave stochastically, each STL specification can only be satisfied probabilistically. Should a specification be assigned at time k, the probability can be determined based on the state measurement $x^{\mathcal{I}}(k)$, the controller $\boldsymbol{f}^{\mathcal{I}}$, and the system dynamics (1). Accordingly, we consider the probability that suffix fragment $\boldsymbol{x}_k^{\mathcal{I}}$ satisfies specification ϕ, given state $x^{\mathcal{I}}(k)$. We denote this by

$$\mathbb{P}_{\boldsymbol{f}}(\phi, k) := \mathbb{P}_{\boldsymbol{f}^{\mathcal{I}}, k}(\boldsymbol{x}_k^{\mathcal{I}} \vDash \phi \mid x^{\mathcal{I}}(k)). \tag{3}$$

[1] $\mathcal{Q}_w^i$ is in the closed convex hull of all uniform distributions on symmetric compact convex bodies in $\mathbb{R}^n$ (c.f. [19, Def. 3.1]).

For each specification ϕ, we require that the maximal allowable risk $r_{\phi,\max}$ of not satisfying specification ϕ is given. In the remainder, we will refer to this as the maximal risk. Note that for $|\mathcal{I}| = 1$, the above definitions of STL and probability will become the standard ones as seen in [15].

C. Problem Statement & Approach

Consider each agent $i \in \mathcal{I}$ has dynamics (1) with noise distribution $\mathcal{Q}_w^i$ and state update $x^i(k)$, at each time step $k \in \mathbb{N}$. We assume that a new specification may be provided at any time. The objective is to update the suffix controller $\boldsymbol{f}_k^{\mathcal{I}}$ at each time step $k \in \mathbb{N}$. Here, any newly provided specification is either accepted or rejected based on the maximal risks of the newly provided specification and previously accepted specifications. An illustrative example can be found in the extended report [21] to further showcase the paper's objective. We consider the following problem statement.

Problem Statement. Given a sequence of STL specifications $\boldsymbol{\phi} := \{\phi_1, \cdots, \phi_t\}$ assigned, respectively, at time instances $\{k_1, \cdots, k_t\}$ with $k_i < k_j$ for $i < j$, with maximal risk $r_{\phi_i,\max}$, develop a method for *updating*, at each time step $k \in \mathbb{N}$, suffix $\boldsymbol{f}_k^{\mathcal{I}}$ such that $\mathbb{P}_{\boldsymbol{f}}(\phi_i, k_i) \geq 1 - r_{\phi_i,\max}$ if $k_i \leq k$ and ϕ_i is a previously accepted specification, where ϕ_i with $k_i = k$ is allowed to be rejected.

Approach. First, a newly provided STL specification and maximal risk are decomposed onto the single-agent level (Section III-B). Subsequently, a filter determines which agent-specification pairs are favorable based on STL robustness and heuristics (Section III-C). Afterwards, for each remaining agent-specification pair, a potential control strategy is determined together with the local risk value (Section III-D). Based on the local risk value for each feasible agent-specification pair, auctioning determines the definitive assignment (Section III-E). Should no valid assignment exist, the specification is rejected. Finally, the control strategy for each agent is updated and implemented (Section III-F).

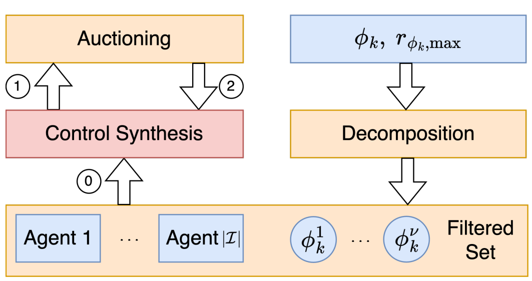

Fig. 2. Illustration of the approach at time k. (0) Determine for each agent-specification pair a control strategy and the local risk value. (1) Auctioning offers from agent-specification pairs. (2) Decision after auctioning.

III. Risk-Aware Multi-Agent Task Allocation and Control Synthesis

To tackle the problem above, we will first consider agent dynamic decomposition, specification decomposition and heuristic agent-specification filtering. Second, we establish control synthesis and local risk value computation for any individual agent-specification pair. Third, we utilize auctioning, based on the local risk values, to obtain a definitive agent-specification assignment. This section concludes with an implementation algorithm and theoretical analysis of the overall method.

A. Decomposing Agent Dynamics

We decompose the dynamics (1), given by

$$x^i(k+1) = A^i x^i(k) + B^i u^i(k) + w^i(k),$$

into a nominal and an error part [22]. The nominal dynamics, denoted as z^i, contain no stochasticity, and the (stochastic) error dynamics, denoted as e^i, are autonomous. This yields

$$z^i(k+1) = A^i z^i(k) + B^i v^i(k), \tag{4a}$$
$$e^i(k+1) = A_K^i e^i(k) + w^i(k), \tag{4b}$$
$$\text{with } x^i(k) = z^i(k) + e^i(k), \tag{4c}$$
$$u^i(k) = v^i(k) + K^i e^i(k). \tag{4d}$$

Here $A_K^i = A^i + B^i K^i$ and K^i is a stabilizing feedback gain meant to keep the error e^i small. Throughout this section, we omit superscripts indicating agent indices when appropriate.

B. Specification Decomposition

Let ϕ be a newly assigned STL specification with maximal risk $r_{\phi,\max}$, involving a distinct number of agents, $\nu \leq |\mathcal{I}|$. The objective is to decompose the specification ϕ into agent-level sub-specifications, ϕ^j, $j \in \mathcal{I}$, and allocate individual maximal risks, $r_{\phi^j,\max}$, in accordance to $r_{\phi,\max}$. First, we define $\operatorname{dec}(\phi, \nu) = \wedge_{j=1}^{\nu} \phi^j$ as the decomposition of ϕ into a conjunctive form, with each conjunct representing a sub-specification, ϕ^j, the satisfaction of which depends on a single agent j, and ν being the number of agents needed to satisfy ϕ. Second, we assign a maximal risk to ϕ^j, following a union-bound argument, which requires $\sum_{j=1}^{\nu} r_{\phi^j,\max} \geq r_{\phi,\max}$. In this paper, we select a uniform allocation of risks across sub-specifications, obtaining $r_{\phi^j,\max} = \frac{r_{\phi,\max}}{\nu}$, for all agents $j \in \{1, \ldots, \nu\}$.

The decomposition $\operatorname{dec}(\phi, \nu)$ acts as a transformation on ϕ returning a conjunctive formula, the satisfaction of which implies satisfaction of ϕ. However, the converse is not necessarily true, implying a lack of completeness. There are a few methods [7], [8], [23] for splitting STL specifications into sub-specifications. In this paper, given a specification ϕ, we convert Boolean and temporal operators to conjunctive forms based on the heuristics in [23] and decompose predicates following a similar approach to [7]. Therein, predicates are defined on general convex sets, whereas negated forms are not studied. Here, we focus on polytopic settings and allow negations. For our setting, we assume the following.

Assumption 1: Let ϕ be an STL specification assigned at time k. There exists decomposition $\operatorname{dec}(\phi, \nu) = \wedge_{j=1}^{\nu} \phi^j$, where ϕ^j indicates a task involving a single agent.

In view of existing decomposition approaches, where sub-specifications are assigned to subteams of agents, Assumption 1 is restrictive. However, in the setting of polytopic predicates, we can produce agent-level sub-specifications by projecting the associated predicates onto the agents' state space. Before explaining this, we define the following.

Definition 2: Let $H := \{x \mid h(x) \geq \mathbf{0}\}$ be a polytope in $\mathbb{R}^n$. $B_{\rm in}(H) := \{x \mid \underline{b}_j^{\rm in} \leq x_j \leq \overline{b}_j^{\rm in},\ j = 1, \ldots, n\}$ is an orthotope inscribed in H, i.e., $B_{\rm in}(H) \subseteq H$, and $B_{\rm out}(H) := \{x \mid \underline{b}_j^{\rm out} \leq x_j \leq \overline{b}_j^{\rm out},\ j = 1, \ldots, n\}$ is an orthotope circumscribing H, i.e., $B_{\rm out}(H) \supseteq H$.

The following lemma shows that any (negated) predicate can be decomposed into a conjunction of agent-level sub-predicates, of which satisfaction implies the satisfaction of the original (negated) predicate. The decomposition of the Boolean and temporal operators is described in [23].

Lemma 3: Let $\mu := (h(x) \geq \mathbf{0})$ be a predicate, $H := \{x \mid h(x) \geq \mathbf{0}\}$ be a polytope, with $x = [x_1^\intercal \ \ldots \ x_\nu^\intercal]^\intercal$, $x_i \in \mathbb{X}^i$, $i \in \{1, \ldots, \nu\}$, and $B_{\rm in}(H) \subseteq H$ and $B_{\rm out}(H) \supseteq H$ be the associated orthotopes inscribed in and circumscribing H, respectively. Then, it holds that (i) $\operatorname{dec}(\mu, \nu) = \wedge_{i=1}^{\nu} \mu_{\rm in}^i$ satisfies μ, with $\mu_{\rm in}^i := (x_i \in \operatorname{Proj}_{\mathbb{X}^i}(B_{\rm in}(H))$, and (ii) $\operatorname{dec}(\neg\mu, \nu) = \wedge_{i=1}^{\nu} \neg\mu_{\rm out}^i$ satisfies $\neg\mu$, with $\mu_{\rm out}^i := (x_i \in \operatorname{Proj}_{\mathbb{X}^i}(B_{\rm out}(H))$.

Proof: (i) Let $\mu_{\rm in} := (x \in B_{\rm in}(H))$. By construction, it holds that $B_{\rm in}(H) \equiv \operatorname{Proj}_{\mathbb{X}^1}(B_{\rm in}(H)) \times \cdots \times \operatorname{Proj}_{\mathbb{X}^\nu}(B_{\rm in}(H))$, implying that $\mu_{\rm in} \equiv \wedge_{i=1}^{\nu} \mu_{\rm in}^i$. Since, $B_{\rm in}(H) \subseteq H$, $\mu_{\rm in}$ satisfies μ, which leads to $\wedge_{i=1}^{\nu} \mu_{\rm in}^i$ satisfying μ as needed. (ii) Let $\mu_{\rm out} := (x \in B_{\rm out}(H))$. Similarly to (i), it holds that $\mu_{\rm out} \equiv \wedge_{i=1}^{\nu} \mu_{\rm out}^i$, since $B_{\rm out}(H) \equiv \operatorname{Proj}_{\mathbb{X}^1}(B_{\rm out}(H)) \times \cdots \times \operatorname{Proj}_{\mathbb{X}^\nu}(B_{\rm out}(H))$, by construction. Then, we have $\neg\mu_{\rm out} \equiv \vee_{i=1}^{\nu} \neg\mu_{\rm out}^i$. Since $\wedge_{i=1}^{\nu} \neg\mu_{\rm out}^i$ satisfies $\vee_{i=1}^{\nu} \neg\mu_{\rm out}^i$ and $\neg\mu_{\rm out}$ satisfies $\neg\mu$, we have that $\wedge_{i=1}^{\nu} \neg\mu_{\rm out}^i$ satisfies $\neg\mu$, completing the proof. ∎

An illustrative example of the specification decomposition method can be found in the extended report [21].

C. Heuristic Agent-Specification Filtering

STL is equipped with robustness metrics for assessing the robust satisfaction of a formula [16]. A scalar-valued function $\rho^\phi(\boldsymbol{x}, k)$ of a signal $\boldsymbol{x}$ and time k indicates how robustly a signal $\boldsymbol{x}$ satisfies a formula ϕ at time k. The robustness function is defined recursively in [16, Definition 3].

Let ϕ_k be a newly assigned STL specification at time k, and assume that $\operatorname{dec}(\phi_k, \nu_k) = \wedge_{j=1}^{\nu_k} \phi_k^j$. In the following, we require that the total number of agents, denoted by $|\mathcal{I}|$, satisfies $|\mathcal{I}| > \nu_k$, otherwise, the following filtering procedure is ignored. Although, ϕ_k is split into ν_k agent-level sub-specifications, the assignment of ϕ_k^j to a specific agent $i \in \mathcal{I}$ is not determined by $\operatorname{dec}(\phi_k, \nu_k)$. In fact, each sub-specification, ϕ_k^j, may be assigned to any agent $i \in \mathcal{I}$ resulting in $\nu_k|\mathcal{I}|$ agent-specification pairs. To reduce the number of potential pairs, we utilize the following heuristic approach: Based on the quantitative semantics of STL [16], we compute the robustness function of ϕ_k^j over the ith agent's trajectory. This metric indicates roughly the distance between the trajectory of agent i and the space of signals that satisfy ϕ_k^j [24], thereby facilitating a preliminary agent-specification assessment. We approximate the ith agent's trajectory by $\boldsymbol{z}^{i,k-1} = \{z_0^{i,k-1}, z_1^{i,k-1}, \ldots\}$, $i \in \mathcal{I}$, that is, the signal associated with the (deterministic) nominal dynamics of agent i, (4a), computed at time $k-1$, which is available at time k. Specifically, for each $j \in \{1, \ldots, \nu_k\}$, and $i \in \mathcal{I}$, we efficiently calculate $\rho^{\phi_k^j}(\boldsymbol{z}^{i,k-1}, k)$, which requires negligible computational cost. Subsequently, for each ϕ_k^j, we collect the largest ν_k values of $\rho^{\phi_k^j}(\boldsymbol{z}^{i,k-1}, k)$ among the $|\mathcal{I}|$ options, resulting in $(\nu_k)^2$ agent-specification pairs. Hence, should $v_k \ll |\mathcal{I}|$, the filtering significantly reduces potential pairs.

D. Control Synthesis and Local Risk Value

In this subsection, we compute the local risk value and synthesize a controller for a given agent-specification pair utilizing previous work on risk-aware MPC for single-agent systems [15]. We refer the reader to [15] for more details.

We assume a new specification ϕ is provided at time k together with maximal risk $r_{\phi,\max}$. For any specifications ψ, we denote with k_ψ the time at which the specification was assigned. We denote with $\mathcal{P}^k$ the set of all previously accepted specifications ψ with $k_\psi < k$, together with the newly assigned specification ϕ, i.e., $\phi \in \mathcal{P}^k$. The objective is to synthesize a controller $\boldsymbol{f}$ such that $\varphi := \wedge_{\psi\in\mathcal{P}^k} \Diamond_{[k_\psi,k_\psi]}\psi$ is satisfied in accordance to the maximal risk of each conjunction element, i.e., $\mathbb{P}_{\boldsymbol{f}}(\psi, k_\psi) \geq 1 - r_{\psi,\max},\ \forall\psi \in \mathcal{P}^k$. To accomplish the objective, we make the following assumption.

Assumption 2: Distribution $\mathcal{Q}_w$ has zero mean and $A_K\Sigma_\infty A_K^T + \operatorname{var}(\mathcal{Q}_w) = \Sigma_\infty$ has solution $\Sigma_\infty = I_n$

Note that the above assumption is not restrictive as a simple coordinate transformation $y = \Sigma_\infty^{-\frac{1}{2}} x$ will ensure the second part of the assumption is satisfied. We call the transformed dynamics the *normalized dynamics*, and in the remainder, we maintain our notations as if the dynamics are normalized.

We are interested in implementing a tube-based MPC algorithm that at each time instance $k \in [0, N]$ recomputes the optimal nominal trajectory $\boldsymbol{z}_{[k,N]}$, the optimal nominal input $\boldsymbol{v}_{[k,N-1]}$, and the optimal risk levels $\boldsymbol{r}_{[k,N]}$. For this, we assume given $x(k)$ together with $\boldsymbol{z}_{[0,k]}$, $\boldsymbol{v}_{[0,k-1]}$, $\boldsymbol{\rho}_{[0,k]}$ and $\boldsymbol{r}_{[0,k]}$ computed at time $k-1$. Furthermore, for all $\psi \in \mathcal{P}^k$, we let $\mathcal{H}_\psi^{k_\psi} \subseteq [0, N]$ without loss of generality. Here, the active horizon $\mathcal{H}_\phi^k$ is the set of all time instances needed to evaluate $\boldsymbol{x}_k \vDash \phi$ (see [15, Definition 4]). It should be noted that any active horizon consists of only finite elements.

Let us consider the following tube-based MPC problem

$$\min_{\Xi} \ J(\boldsymbol{z}_{[0,N]}, \boldsymbol{v}_{[0,N-1]}, \boldsymbol{r}_{[0,N]}), \tag{5a}$$
$$\text{s.t.} \ \ z(k) = x(k) \tag{5b}$$
$$z(\tau+1) = Az(\tau) + Bv(\tau), \forall\tau \in \{k, \cdots, N-1\}, \tag{5c}$$
$$G(j)z(j) + MF(j)s(k) \geq b(M, j) + \rho(j)\mathbf{1}, \tag{5d}$$
$$Cs(k) \geq d,\ \epsilon \leq \rho(j) \leq \tfrac{M}{2},\ \forall j \in \{0, \cdots, N\}, \tag{5e}$$
$$n = \rho(j)^2 r(j), \qquad \forall j \in \{0, \cdots, N\}, \tag{5f}$$
$$r_{\psi,\max} \geq \textstyle\sum_{j\in\mathcal{H}_\psi^{k_\psi}\setminus\{k_\psi\}} r(j), \forall\psi \in \mathcal{P}^k, \tag{5g}$$

where $\Xi := \{\boldsymbol{z}_{[k,N]}, \boldsymbol{v}_{[k,N-1]}, \boldsymbol{r}_{[k+1,N]}, \boldsymbol{\rho}_{[k+1,N]}, s(k)\}$, J is either a linear or quadratic cost function, the constraints (5d)-(5e) are obtained from $\varphi = \wedge_{\psi\in\mathcal{P}^k} \Diamond_{[k_\psi,k_\psi]}\psi$ and the constraints (5f)-(5g) are obtained from [15, Thm. 4]. From the above tube-based MPC, an (updated) control strategy is obtained by $\boldsymbol{v}_{[k,N-1]}$, and a local risk value $r_{\phi,mpc}$ is obtained by the right-hand side of (5g), i.e.,

$$r_{\phi,mpc} = \sum_{j\in\mathcal{H}_\phi^k\setminus\{k\}} r(j). \tag{6}$$

A solution to (5), by [15, Thm. 4], ensures that for all $\psi \in \mathcal{P}^k$, $\mathbb{P}_{\boldsymbol{f}}(\psi, k_\psi) \geq 1 - r_{\psi,\max}$. Should there exist no solution to problem (5), the newly provided agent-specification pair is deemed infeasible. In such a case, we will update the control strategy based only on measurement $x(k)$, as per constraint (5b). This can be achieved by considering the MPC problem (5) with $\mathcal{P}^k = \mathcal{P}^k \setminus \{\phi\}$. However, should the measurement also fail to produce a controller, the previously computed control strategy will be implemented instead.

E. Task Allocation Algorithm

Let ϕ be a specification which can be decomposed into ν sub-specifications. After filtering agent-specification pairs and attempting optimization problem (5) with the remainder, a list $\mathbb{L} \subset \mathcal{I} \times \{1, \cdots, \nu\}$ of feasible agent-specification pairs is generated. Each feasible agent-specification pair has a local risk value (6) denoted by $r^i_{\phi_j,mpc}$, with $i \in \mathcal{I}$ and $j \in \{1, \cdots, \nu\}$. Our objective is to auction sub-specifications to agents based on the local risk values, obtaining a definitive assignment in the process. Here, each agent can only receive one specification, and all specifications must be allocated.

To determine an agent-specification assignment, we solve the following binary optimization problem given by

$$\min_{\boldsymbol{\lambda}} \ \sum_{(i,j)\in\mathbb{L}} \lambda_{i,j} r^i_{\phi_j,mpc}, \tag{7a}$$

$$\text{s.t.} \ \sum_{i\in\mathcal{I}} \lambda_{i,j} = 1, \ \forall j \in \{1, \cdots, \nu\} \tag{7b}$$

$$\sum_{j\in\{1,\cdots,\nu\}} \lambda_{i,j} \leq 1, \ \forall i \in \mathcal{I}, \tag{7c}$$

$$\lambda_{i,j} = 0, \qquad \forall (i,j) \notin \mathbb{L}, \tag{7d}$$

$$\lambda_{i,j} \in \{0,1\}, \ \forall (i,j) \in \mathbb{L}, \tag{7e}$$

where $\boldsymbol{\lambda} = \{\lambda_{i,j} \mid (i,j) \in \mathbb{L}\}$. Here, constraint (7b) entails that each specification must be assigned to exactly one agent; constraint (7c) entails that each agent can at most receive one specification; and $\lambda_{i,j} = 1$ implies agent i and sub-specification j are chosen as an agent-specification pair. The cost minimizes the sum of the local risk values, thereby ensuring the best probability satisfaction of specification ϕ. Should no solution be found, specification ϕ will be rejected.

F. Implementation & Theoretical Analysis

To enact the proposed method, we consider Algorithm 1, which demonstrates multi-agent task allocation and control strategy updates for each agent at every time step. For a more detailed algorithm on single-agent control synthesis, we defer to [15] and the algorithm therein.

Let us consider the recursive feasibility of updating a control strategy for each agent at any time.

Theorem 4 (Recursive feasibility): If there exists a feasible control strategy $\boldsymbol{f}^{\mathcal{I},k}$, at time k, then there exists a feasible control strategy $\boldsymbol{f}^{\mathcal{I},k+1}$ at time $k+1$.

Proof: It is sufficient to prove that the existence of $\boldsymbol{f}^{i,k}$ implies the existence of $\boldsymbol{f}^{i,k+1}$ for each agent $i \in \mathcal{I}$. Notice that if an agent either accepts a new sub-specification or measurement only, an updated control strategy is obtained, respectively, as $\boldsymbol{f}^{i,k+1} = \boldsymbol{f}^{(i,j)}$ or $\boldsymbol{f}^{i,k+1} = \boldsymbol{f}^i_{meas}$. If neither is available, the previous control strategy is again implemented, $\boldsymbol{f}^{i,k+1} = \boldsymbol{f}^{i,k}$, completing the proof. ■

Algorithm 1 Task Allocation and Control Strategy Update

```
1: Given: Normalized system (1) with x^i(0), ∀i ∈ I
2: for k ∈ {0, ··· , N − 1} do
3:     if no new spec. then set L_opt = ∅ & go to line 15
4:     Task Allocation:
5:     (ψ, r_{ψ,max}) ← load new specification
6:     Decompose into (ψ^j, r_{ψ^j,max}), j ∈ {1, ··· , ν}
7:     Filter pairs to obtain L ⊂ I × {1, ··· , ν}
8:     for (i, j) ∈ L do
9:         Solve (5), get r^i_{ψ^j,mpc} and control strategy f^{(i,j)}
10:        if no solution then L = L \ {(i, j)}
11:    end for
12:    Solve (7) to get definitive assignment list L_opt
13:    Control Strategy Update:
14:    ∀(i, j) ∈ L_opt, set f^{i,k} = f^{(i,j)} ← update
15:    for i ∈ I \ Proj_I(L_opt) do
16:        Solve (5) to get control strategy f^i_meas
17:        if solution then set f^{i,k} = f^i_meas ← update
18:        else set f^{i,k} = f^{i,k−1} ← no update
19:    end for
20:    Implement f^{I,k} and measure x^i(k + 1), ∀i ∈ I
21: end for
```

For each specifications ϕ accepted at time k, by subgroup $I \subseteq \mathcal{I}$, we need that $\mathbb{P}_{\boldsymbol{f}}(\phi, k) \geq 1 - r_{\phi,\max}$. Here, it is sufficient to show that $\forall i \in I$, $\mathbb{P}_{\boldsymbol{f}}(\phi^i, k) \geq 1 - r_{\phi^i,\max}$. This, in relation to optimization problem (5), is discussed within open-loop and closed-loop implementation in our previous work [15]. In said paper, we established the validity of open-loop implementation and suggested that closed-loop implementation follows naturally from the convex unimodality of the disturbance.

Remark 1: Note that the local risk value $r_{\phi^j,mpc}$ compared to the average maximal risk $r_{\phi^j,\max}$ is conservative, $r_{\phi^j,mpc} \geq r_{\phi^j,\max}$, due to a union-bound argument used in proving [15, Thm. 4]. Also, the obtained task allocation may not always be optimal or feasible due to our heuristic choice of filtering, even if an allocation exists.

IV. Case Study

We validate our proposed method using a multi-shuttle routing case. In a tourist attraction spot (Fig. 3), four shuttle buses of differing speeds (A-D) should pick up tourists from gathering points (GP) and deliver them to an unloading point (ULP) while avoiding two buildings (B). The shuttles initiate from their respective terminals (T) and should return when required. Each task should be assigned to a most probable shuttle to fetch the tourists. Thus, new routing may be assigned anytime, rendering a dynamic task allocation problem. The technical details can be found in report [21].

After decomposing, filtering, control synthesis, and auctioning, we obtain the results in Fig. 3. The logging inform-

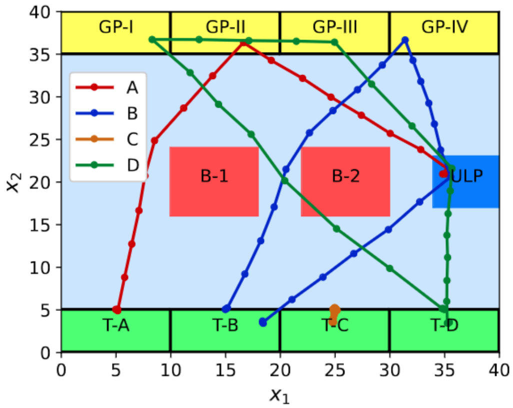

Fig. 3. The shuttle trajectories.

ation of the experiment indicates that: 1) at time $k = 2$, buses 'A', 'B', and 'D' are assigned with GP-II, GP-IV, and GP-I, respectively; 2) at time $k = 8$, buses 'A' and 'D' are assigned with GP-II and GP-III, respectively; 3) at time $k = 15$, buses 'B', 'C', and 'D' have accepted the return specification, while shuttle 'A' has rejected it. Bus 'C' is not assigned any task throughout the experiment and, as a result, stays at its terminal. Since buses 'A' and 'D' accept additional tasks at $k = 8$, our method allows agents to adjust their behaviours according to newly assigned tasks. As mentioned, shuttle 'A' has rejected the return task at time $k = 15$. This is due to the incapability to move long distances considering its small speed limit. Additionally, all buses are confined within the 'BOX' region and kept outside the buildings, implying the compatibility of our method to explicitly assigned specifications. Part of the trajectory of shuttle 'D' intersects with B-2 due to naive linear interpolation. In practice, nonlinear interpolation can guarantee critical safety, which may generate curved trajectories when navigating around the building.

Remark 2: In this experiment, the computation of control inputs by solving problem (5) is distributed to each agent, leading to a computational complexity linear to the number of assigned specifications. In this sense, our proposed method allows for a comparable computational load of each agent to a conventional single-agent case, allowing it to be implemented in practical scenarios.

V. Conclusion

This paper considers real-time task allocation and control synthesis for stochastic heterogeneous linear multi-agent systems with probabilistic satisfaction of STL specifications. The proposed method relies on the dynamic allocation of decomposed specifications to individual agents using risk measures via auctioning. The risk measures are computed using tube-based MPC on the agent level during the control synthesis. These two approaches provide a novel perspective for coordination control of multi-agent systems in dynamic environments. In future work, we will improve the risk quantification for less conservative control solutions.

References

[1] V. Raman, A. Donzé, M. Maasoumy, R. M. Murray, A. Sangiovanni-Vincentelli, and S. A. Seshia, "Model predictive control with signal temporal logic specifications," in *Conference on Decision and Control*, 2014, pp. 81–87.

[2] D. Sun, J. Chen, S. Mitra, and C. Fan, "Multi-agent motion planning from signal temporal logic specifications," *Robotics and Automation Letters*, vol. 7, no. 2, pp. 3451–3458, 2022.

[3] S. Liu, A. Saoud, P. Jagtap, D. V. Dimarogonas, and M. Zamani, "Compositional synthesis of signal temporal logic tasks via assume-guarantee contracts," in *Conference on Decision and Control*, 2022, pp. 2184–2189.

[4] L. Lindemann and D. V. Dimarogonas, "Barrier function based collaborative control of multiple robots under signal temporal logic tasks," *Transactions on Control of Network Systems*, vol. 7, no. 4, pp. 1916–1928, 2020.

[5] T. Yang, Y. Zou, S. Li, and Y. Yang, "Distributed model predictive control for probabilistic signal temporal logic specifications," *Transactions on Automation Science and Engineering*, pp. 1–11, 2023.

[6] K. Braekers, K. Ramaekers, and I. Van Nieuwenhuyse, "The vehicle routing problem: State of the art classification and review," *Computers & industrial engineering*, vol. 99, pp. 300–313, 2016.

[7] M. Charitidou and D. V. Dimarogonas, "Signal temporal logic task decomposition via convex optimization," *Control Systems Letters*, vol. 6, pp. 1238–1243, 2021.

[8] K. Leahy, A. Jones, and C.-I. Vasile, "Fast decomposition of temporal logic specifications for heterogeneous teams," *Robotics and Automation Letters*, vol. 7, no. 2, pp. 2297–2304, 2022.

[9] Z. Zhang and S. Haesaert, "Modularized control synthesis for complex signal temporal logic specifications," in *Conference on Decision and Control*, 2023, pp. 7856–7861.

[10] H. Liu, P. Zhang, B. Hu, and P. Moore, "A novel approach to task assignment in a cooperative multi-agent design system," *Applied Intelligence*, vol. 43, pp. 162–175, 2015.

[11] D. P. Bertsekas, "Auction algorithms for network flow problems: A tutorial introduction," *Computational optimization and applications*, vol. 1, pp. 7–66, 1992.

[12] S. S. Farahani, R. Majumdar, V. S. Prabhu, and S. Soudjani, "Shrinking horizon model predictive control with signal temporal logic constraints under stochastic disturbances," *Transactions on Automatic Control*, vol. 64, no. 8, pp. 3324–3331, 2019.

[13] K. A. Mustafa, O. de Groot, X. Wang, J. Kober, and J. Alonso-Mora, "Probabilistic risk assessment for chance-constrained collision avoidance in uncertain dynamic environments," in *International Conference on Robotics and Automation*, 2023, pp. 3628–3634.

[14] M. H. W. Engelaar, S. Haesaert, and M. Lazar, "Stochastic model predictive control with dynamic chance constraints," in *International Conference on System Theory, Control and Computing*, 2023, pp. 356–361.

[15] M. H. W. Engelaar, Z. Zhang, M. Lazar, and S. Haesaert, "Risk-aware MPC for stochastic systems with runtime temporal logics," *arXiv preprint arXiv:2402.03165*, 2024.

[16] A. Donzé and O. Maler, "Robust satisfaction of temporal logic over real-valued signals," in *International Conference on Formal Modeling and Analysis of Timed Systems*, 2010, pp. 92–106.

[17] C. Burnett, T. J. Norman, and K. Sycara, "Trust decision-making in multi-agent systems," in *international joint conference on artificial intelligence*, 2011.

[18] D. Bertsekas and S. E. Shreve, *Stochastic optimal control: the discrete-time case*. Athena Scientific, 1996, vol. 5.

[19] S. Dharmadhikari and K. Jogdeo, "Multivariate unimodality," *The Annals of Statistics*, pp. 607–613, 1976.

[20] S. Sadraddini and C. Belta, "Robust temporal logic model predictive control," in *Conference on Communication, Control, and Computing (Allerton)*, 2015, pp. 772–779.

[21] M. H. W. Engelaar, Z. Zhang, E. E. Vlahakis, D. V. Dimarogonas, M. Lazar, and S. Haesaert, "Risk-aware real-time task allocation for stochastic multi-agent systems under STL specifications," *arXiv preprint arXiv:2404.02111*, 2024.

[22] L. Hewing and M. N. Zeilinger, "Stochastic model predictive control for linear systems using probabilistic reachable sets," in *Conference on Decision and Control*, 2018, pp. 5182–5188.

[23] K. Leahy, M. Mann, and C.-I. Vasile, "Rewrite-based decomposition of signal temporal logic specifications," in *NASA Formal Methods Symposium*, 2023, pp. 224–240.

[24] G. E. Fainekos and G. J. Pappas, "Robustness of temporal logic specifications for continuous-time signals," *Theoretical Computer Science*, vol. 410, no. 42, pp. 4262–4291, 2009.