IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

Multi-Agent Collaboration for PrSTL Specifications With Temporal Collective Counting Operators

Yicheng Quan, Yan Yang, Zhijie Liu, *Member, IEEE*, and Zhongjiao Shi

***Abstract*—We address the collaborative path planning problem for multi-agent systems with heterogeneous capabilities, subject to uncertainty and operating under complex task specifications. Conventional Probabilistic Signal Temporal Logic (PrSTL) frameworks exhibit significant limitations in describing multi-agent collaborative tasks with temporally cumulative properties. To address this challenge, we extend the PrSTL framework by introducing a Temporal Collective Counting Operator to characterize such spatio-temporal specifications. We then formulate the multi-agent collaborative planning problem under dynamics uncertainty as a Mixed-Integer Second-Order Cone Program. This formulation leverages PrSTL to specify tasks with cumulative temporal properties, while employing Polynomial Chaos Expansion to propagate uncertainty. Finally, we propose a constraint relaxation mechanism to address the conservatism introduced by formula transformations and probabilistic constraints' approximation.**

***Index Terms*—Formal methods in robotics and automation, planning under uncertainty, path planning for multiple mobile robots or agents.**

I. Introduction

NUMEROUS applications necessitate deploying multi-agent systems for collaborative missions whose complex spatio-temporal specifications are challenging to formalize mathematically [1]. Signal Temporal Logic (STL) provides a formal language to specify the complex spatio-temporal requirements of such collaborative missions concisely [2], [3], [4], [5], [6]. However, a core limitation of STL is that its standard temporal operators are suited for properties that must hold at some point or always over a time interval, not for describing the temporally cumulative properties essential to many collaborative tasks, such as requiring "a target region be visited a total of 10 times by agents with reconnaissance capabilities."

The challenge of specifying cumulative tasks is compounded by the inherent uncertainty in practical systems, motivating the development of various stochastic control methods. These range from stochastic tubes for handling bounded additive or multiplicative disturbances [7] and scenario-based approaches for parametric uncertainty [8] to various risk-based methods [9], many of which have been adapted for synthesis under temporal logic specifications [10], [11], [12], [13], [14], [15], [16]. Among these, Probabilistic Signal Temporal Logic (PrSTL) [11] is particularly relevant, as it integrates uncertainty directly into the logic to provide probabilistic guarantees, such as ensuring the probability of violating a safety constraint remains below a predefined threshold. While PrSTL effectively handles uncertainty, it inherits the limitations of STL's temporal operators, leaving the challenge of specifying cumulative properties in stochastic systems unaddressed.

While several formalisms address temporal accumulation in deterministic settings, such as the integral predicate for control synthesis [17] and Cumulative-Time Signal Temporal Logic (CT-STL) for runtime verification [18], extending these foundations to stochastic systems presents fundamental obstacles. A probabilistic integral predicate requires tracking cross-covariance terms, leading to intractable variance calculations. Similarly, augmenting the state with an auxiliary signal to track accumulation encounters the same tractability issues regarding variance propagation. At the same time, CT-STL's monitoring-oriented semantics lack a direct path to a tractable formulation for multi-agent synthesis. Existing formalisms such as state automata [19], Reward Machines [20], [21], and Control Barrier Functions [22] are also unsuitable, as they lack the probabilistic vocabulary to address temporally cumulative specifications under stochastic dynamics.

This paper addresses this challenge by extending the PrSTL framework to naturally handle tasks with cumulative temporal effects under dynamic uncertainty. To achieve this, our approach combines a novel operator for specification with Polynomial Chaos Expansion (PCE) for uncertainty propagation. As a spectral method, PCE represents stochastic uncertainty through a series of orthogonal polynomial basis functions, allowing for the propagation of probability distributions through the system dynamics. Its capability to yield analytical statistical moments facilitates computationally efficient implementations of probabilistic constraints [23], [24], [25], [26], [27]

The primary contributions of this work are threefold:

1) We introduce the Temporal Collective Counting Operator (TCCO), an extension to PrSTL whose logical structure is designed for tractable synthesis under uncertainty. We show that TCCO formalizes cumulative and collaborative properties for heterogeneous agents and can be represented as a Boolean

Received 6 July 2025; accepted 19 November 2025. Date of publication 8 December 2025; date of current version 11 December 2025. This article was recommended for publication by Associate Editor H. Tsukamoto and Editor S. J. Chung upon evaluation of the reviewers' comments. This work was supported by the National Natural Science Foundation of China under Grant 62522305, Grant 62473039, and Grant 62473041. *(Corresponding author: Zhongjiao Shi.)*
Yicheng Quan and Zhongjiao Shi are with the School of Aerospace Engineering, Beijing Institute of Technology, Beijing 100081, China (e-mail: quanyicheng0110@gmail.com; shizhongjiao@gmail.com).

Yan Yang and Zhijie Liu are with the School of Intelligence Science and Technology, University of Science and Technology Beijing, Beijing 100083, China (e-mail: yangyanxdu@gmail.com; liuzhijie@ustb.edu.cn).
Digital Object Identifier 10.1109/LRA.2025.3641092

2377-3766 © 2025 IEEE. All rights reserved, including rights for text and data mining, and training of artificial intelligence and similar technologies.
Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

©2026 IEEE

Authorized licensed use limited to: BEIJING INSTITUTE OF TECHNOLOGY. Downloaded on February 22,2026 at 17:01:03 UTC from IEEE Xplore. Restrictions apply.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

combination of atomic chance constraints, thus avoiding the intractable cross-covariance calculations inherent in extending prior methods like [17].

2) We leverage PCE for uncertainty propagation. Leveraging the properties of PCE, statistical moments of the stochastic variables are obtained, which subsequently allows for the transformation of atomic chance constraints into deterministic second-order cone constraints. Finally, a Mixed-Integer Second-Order Cone Program (MISOCP) is constructed for the generation of trajectories.

3) We propose a constraint relaxation mechanism to address the conservatism introduced by the approximations. This mechanism offers a trade-off between the rigor of the formal probabilistic guarantee and computational efficiency and feasibility. It ensures our synthesis framework is practically feasible.

II. Preliminaries

A. Agent Dynamics

The considered multi-agent system consists of N_a agents, indexed by $i \in \mathbb{I} := \{1, 2, \ldots, N_a\}$. Each agent $i \in \mathbb{I}$ follows the linear stochastic discrete-time dynamics:

$$x_i(k+1) = A_i(\theta)x_i(k) + B_i(\theta)u_i(k), \tag{1}$$

where $x_i(k) \in \mathbb{R}^{n_x}$ is the stochastic state of agent i at time step k and $u_i(k) \in \mathbb{R}^{n_u}$ is the deterministic control input. The stochastic dynamics of the system are characterized by $A_i(\theta)$ and $B_i(\theta)$, with both quantities being dependent on the random variables $\theta \in \mathbb{R}^{n_\theta}$. We focus on this parametric uncertainty to capture state-dependent physical variations, though the framework is applicable to additive disturbances as well. Here, θ consists of time-invariant and independently distributed scalar random variables θ_m. We assume these random variables are defined on a probability space, characterized by known cumulative distribution functions, and have finite second-order moments, which ensures they belong to the L^2 space of square-integrable functions. The expected value and variance of a stochastic variable ψ are denoted by $\mathbf{E}[\psi]$ and $\mathbf{Var}[\psi]$, respectively. The inner product for two random variables ψ_1 and ψ_2 is defined as $\langle \psi_1, \psi_2 \rangle := \mathbf{E}[\psi_1 \psi_2]$.

We denote the stacked state vector and control input as $X(k) := [x_1(k)^T, x_2(k)^T, \ldots, x_{N_a}(k)^T]^T$ and $U(k) := [u_1(k)^T, u_2(k)^T, \ldots, u_{N_a}(k)^T]^T$, respectively. For a planning horizon of N, define $\boldsymbol{U}^N(k) := [U(k)^\top, U(k+1)^\top, \ldots, U(k+N-1)^\top]^\top$ as a control input vector at time k. A stochastic process initiated at time k is a sequence of random state vectors $X(k), X(k+1), \ldots$, which we denote by (Ξ, k). As the process depends on $X(k)$ and $\boldsymbol{U}^N(k)$, (Ξ, k) can be denoted in a more elaborative notation as $\Xi[X(0), \boldsymbol{U}^N]$. A specific realization of system trajectory is called a run, denoted by the infinite signal $\xi = Z(0)Z(1)\ldots$. In a run ξ, each $Z(k)$ is a specific realization of the random state vector $X(k)$. We use (Ξ_i, k) to refer to the stochastic process of a single agent i starting from time k.

B. Signal Temporal Logic

STL provides a formal language to specify properties over deterministic signals like the run ξ defined above [1]. In this paper, we consider a fragment of STL as:

$$\phi ::= \top | \mu | \neg\phi \,|\, \phi_1 \wedge \phi_2 |\, \phi_1 \vee \phi_2 \left| G_{[a,b]}\phi \right| F_{[a,b]}\phi \tag{2}$$

where $\top$ denotes the Boolean constant true; ϕ_1, ϕ_2 and ϕ denote STL formulas; $\neg$, $\wedge$ and $\vee$ are the Boolean negation, conjunction and disjunction operators; $G_{[a,b]}$ and $F_{[a,b]}$ are the "always" and "finally" operators with $a, b \in \mathbb{R}^+$ and $a \leq b$.

Each predicate μ is defined by the sign of corresponding predicate function $h : \mathbb{R}^{n_x} \to \mathbb{R}$ that maps a state to a real number, and we make the following assumption:

Assumption 1: The predicate function h for an predicate μ takes the form:

$$h(x_i(k)) = dx_{i,l}(k) + e, \tag{3}$$

where $x_{i,l}(k)$ is l-th state component of agent i at time step k, and d and e are deterministic scalars.

This assumption restricts predicates to half-spaces, which enables the tractable formulation of chance constraints using statistical moments. To evaluate the satisfaction of μ for a given run ξ at time k, we apply the function h to the determined state vector $Z(k)$ from that run. Given a run ξ, we say ξ satisfies ϕ at time k, denoted by $(\xi, k) \models \phi$, if the suffix sequence $Z(k)Z(k+1)\ldots$ satisfies ϕ.

C. Probabilistic Signal Temporal Logic

PrSTL extends STL to stochastic systems by evaluating the probability that a run generated by the stochastic process Ξ satisfies a given STL specification [11]. Given the stochastic process (Ξ, k) and an STL formula ϕ, we denote by $\mathcal{P}[(\Xi, k) \models \phi]$ the probability measure of the set of instantiations ξ of Ξ such that $(\xi, k) \models \phi$. For simplicity, we adopt the formula $\phi^{(\delta)}$ from [13], where $\phi^{(\delta)}$ denotes the requirement that ϕ is satisfied with a probability of at least δ. This formulation is a form of a chance constraint, which enforces that a certain condition must hold with a minimum probability:

$$(\Xi, k) \models \phi^{(\delta)} \Leftrightarrow \mathcal{P}[(\Xi, k) \models \phi] \geq \delta. \tag{4}$$

The syntax of PrSTL are defined in terms of temporal operators and Boolean operators over $\phi^{(\delta)}$ as follows:

$$\varphi ::= \phi^{(\delta)} \,|\varphi_1 \wedge \varphi_2 |\, \varphi_1 \vee \varphi_2 \left| G_{[a,b]}\varphi \right| F_{[a,b]}\varphi \tag{5}$$

where φ_1, φ_2 and φ denote PrSTL formulas, and ϕ is an STL formula given by (2). We denote the satisfaction of φ with respect to stochastic process (Ξ, k) by $(\Xi, k) \models \varphi$. The satisfaction definitions for PrSTL formulas involving Boolean and temporal operators are as follows:

$$\begin{aligned}
(\Xi, k) &\models \phi^{(\delta)} \Leftrightarrow \mathcal{P}[(\Xi, k) \models \phi] \geq \delta \\
(\Xi, k) &\models \varphi_1 \wedge \varphi_2 \Leftrightarrow (\Xi, k) \models \varphi_1 \wedge (\Xi, k) \models \varphi_2 \\
(\Xi, k) &\models \varphi_1 \vee \varphi_2 \Leftrightarrow (\Xi, k) \models \varphi_1 \vee (\Xi, k) \models \varphi_2 \\
(\Xi, k) &\models G_{[a,b]}\varphi \Leftrightarrow \forall k' \in [k+a, k+b], (\Xi, k') \models \varphi \\
(\Xi, k) &\models F_{[a,b]}\varphi \Leftrightarrow \exists k' \in [k+a, k+b], (\Xi, k') \models \varphi.
\end{aligned} \tag{6}$$

Authorized licensed use limited to: BEIJING INSTITUTE OF TECHNOLOGY. Downloaded on February 22,2026 at 17:01:03 UTC from IEEE Xplore. Restrictions apply.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

III. Problem Formulation

We denote $\mathbb{C}$ as the set of all unique capabilities available in the multi-agent system. Each agent i possesses a specific set of capabilities cap_i, where $cap_i \subseteq \mathbb{C}$. Let $\mathbb{D}$ be the finite set of predefined regions of interest, where each region $D \in \mathbb{D}$ is a subset of the workspace. Inspired by [17], we define the temporally cumulative collaborative task as follows.

Definition 1: A *temporally cumulative collaborative task* is a quadruple $S = (t^S, D^S, E^S, T^S) \in \mathbb{S}$, where $t^S = (t^S_{\min}, t^S_{\max})$ denotes the task time interval; $D^S \in \mathbb{D}$ denotes the region of interest; $E^S \subseteq \mathbb{C}$ denotes the set of required capabilities; $T^S \in \mathbb{R}^+$ denotes the minimum cumulative duration; and $\mathbb{S}$ denotes the set of all such tasks.

In plain English, a temporally cumulative collaborative task involves heterogeneous agents with diverse capabilities performing a specific mission over a period of time within the same region. We use the following example to demonstrate this type of task.

Example 1: A temporally cumulative collaborative task $S = ((10, 15), D, Irrigation, 5)$ requires a cumulative irrigation duration of at least 5 seconds within region D between $t = 10$ s and 15 s. Consider a system with three agents whose capabilities are $cap_1 = \{Irrigation\}$, $cap_2 = \{Irrigation, Fertilization\}$, and $cap_3 = \{Fertilization\}$. As the task requires Irrigation, only Agents 1 and 2 are qualified to contribute. This cumulative requirement can be met if Agent 1 alone spends 5 seconds in the region, or collaboratively if Agent 1 contributes 2 seconds and Agent 2 contributes 3 seconds, since their combined time satisfies the minimum.

Problem 1: Consider a multi-agent system comprising N_a agents, indexed by $i \in \mathbb{I}$, subject to uncertainties, temporally cumulative collaborative tasks $\mathbb{S}$ and PrSTL specification $\phi^{(\delta)}$. The overall capabilities of the system are represented by the set $\mathbb{C}$. The dynamics and the set of capabilities of each agent are given by (1) and cap_i, with the initial state $X(0)$ and a planning horizon of N. Find a deterministic control law $\boldsymbol{U}^N$ such that the resulting system trajectory satisfies $\phi^{(\delta)}$ and all tasks within the set of temporally cumulative collaborative tasks $\mathbb{S}$ with a probability greater than a predefined threshold σ by solving the following optimization problem:

$$\begin{aligned} &\min \|\boldsymbol{U}^N\|_2 \\ \text{s.t. } & x_i(k+1) = A_i(\theta)x_i(k) + B_i(\theta)u_i(k), i \in \mathbb{I}, \\ & cap_i \subseteq 2^{\mathbb{C}}, i \in \mathbb{I}, \\ & \Xi[X(0), \boldsymbol{U}^N] \models \phi^{(\delta)}, \\ & \mathcal{P}\left(\Xi[X(0), \boldsymbol{U}^N] \text{ satisfy } \mathbb{S}\right) \geq \sigma. \end{aligned} \tag{7}$$

IV. Control Synthesis

A. Temporal Collective Counting Operator

To incorporate the temporally cumulative collaborative property into the PrSTL framework, we introduce a TCCO to quantitatively measure the contribution of the multi-agent system towards satisfying the temporally cumulative collaborative tasks. Fig. 1 presents an overview of the proposed synthesis framework.

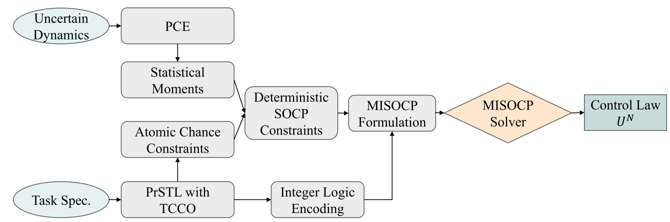

Fig. 1. Overview of the proposed synthesis framework.

Definition 2: A *temporal collective counting operator* corresponding to a temporally cumulative collaborative task $S = (t^S, D^S, E^S, T^S)$ is defined as

$$\mathcal{C}^{\tau = \left\lceil \frac{T^S}{\Delta t} \right\rceil}_{[t^{\mathcal{C}}_{\min}, t^{\mathcal{C}}_{\max}]} \phi_D\{\mathbb{I}_{\mathcal{C}}\}, \tag{8}$$

where $t^{\mathcal{C}}_{\min}, t^{\mathcal{C}}_{\max}$ denote the start and end times of the task, respectively; ϕ_D denotes an STL formula constructed from predicates using only Boolean operators, which corresponds to region D, and each predicate μ within ϕ_D is defined by the sign of a predicate function h; $\mathbb{I}_{\mathcal{C}}$ denotes the set of indices for agents whose capabilities match the task's requirements E^S; τ denotes the user-defined indicator that quantitatively determines the satisfaction of a temporally cumulative collaborative property, and Δt is the sampling time.

The satisfaction definition for the TCCO is:

$$\begin{aligned} &(\Xi, k) \models \mathcal{C}^{\tau = \left\lceil \frac{T^S}{\Delta t} \right\rceil}_{[a,b]} \phi_D\{\mathbb{I}_{\mathcal{C}}\} \\ &\Leftrightarrow \forall j \in \{1, 2, \ldots, \tau\}, \exists k'_j \in [k+a, k+b], \exists i_j \in \mathbb{I}_{\mathcal{C}}, \\ &\text{s.t. } \left((\Xi_{i_j}, k'_j) \models \phi_D\right) \wedge \left((i_{j_1} \neq i_{j_2}) \vee \left(k'_{j_1} \neq k'_{j_2}\right)\right). \end{aligned} \tag{9}$$

Essentially, the TCCO translates task requirements from a continuous duration requirement to a discrete counting one. Specifically, the satisfaction of the TCCO requires that the proposition ϕ_D holds for τ distinct instances. Specifically, If a TCCO is satisfied, there must exist a sequence of τ selections, indexed by $j \in \{1, \ldots, \tau\}$. For each index j, a time step $k'_j \in [k+a, k+b]$ and an agent $i_j \in \mathbb{I}_{\mathcal{C}}$ are chosen such that $(\Xi_{i_j}, k'_j) \models \phi_D$. Furthermore, a global uniqueness constraint is imposed on this sequence. For any two distinct indices $j_1 \neq j_2$, the corresponding selected pairs (i_{j_1}, k'_{j_1}) and (i_{j_2}, k'_{j_2}) must not be identical. This is enforced by the condition $(i_{j_1} \neq i_{j_2}) \vee (k'_{j_1} \neq k'_{j_2})$, which prevents satisfying the count requirement by repeatedly referencing the same agent's state at the exact same time step.

A key strength of our framework is its ability to integrate the TCCO with standard temporal operators into a unified specification. This integration is possible because a TCCO formula evaluates to a Boolean truth value over a given trajectory. As a result, it can be seamlessly composed with other temporal operators, such as "always" and "finally", using Boolean operators.

According to the satisfaction definition of the TCCO in (9), it follows that:

$$(\Xi, k) \models \mathcal{C}^{\tau = \left\lceil \frac{T^S}{\Delta t} \right\rceil}_{[a,b]} \phi_D\{\mathbb{I}_{\mathcal{C}}\}$$

Authorized licensed use limited to: BEIJING INSTITUTE OF TECHNOLOGY. Downloaded on February 22,2026 at 17:01:03 UTC from IEEE Xplore. Restrictions apply.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

Algorithm 1: Formulation of TCCO.

Input: Task $S = (t^S, D^S, E^S, T^S)$, sampling time Δt, agents' capabilities $\{cap_1, \ldots, cap_{N_a}\}$
Output: Operator $\mathcal{C}_{[t^{\mathcal{C}}_{\min}, t^{\mathcal{C}}_{\max}]}^{\tau = \left\lceil \frac{T^S}{\Delta t} \right\rceil} \phi_D\{\mathbb{I}_{\mathcal{C}}\}$

```
1:  Compute ⌈T^S/Δt⌉;
2:  (t^C_min, t^C_max) ← t^S;
3:  φ_D ← Boolean combination of predicates for region D;
4:  Initialize i ← 1, I_C = ∅;
5:  for i ≤ N_a do
6:    if cap_i ∩ E^S ≠ ∅ then
7:      I_C ← I_C ∪ {i};
8:    end if
9:    i ← i + 1;
10: end for
11: return C^{τ=⌈T^S/Δt⌉}_{[t^C_min, t^C_max]} φ_D{I_C};
```

$$\Leftrightarrow \bigvee_{\substack{(i_1, i_2, \ldots, i_\tau) \in \mathbb{I}_{\mathcal{C}}^{\tau} \\ (k'_1, k'_2, \ldots, k'_\tau) \in [k+a, k+b]^\tau}} \left[\left(\bigwedge_{j=1}^{\tau} (\Xi_{i_j}, k'_j) \models \phi \right) \right.$$

$$\left. \wedge \left(\bigwedge_{1 \le j_1 < j_2 \le \tau} \left((i_{j_1} \neq i_{j_2}) \vee (k'_{j_1} \neq k'_{j_2}) \right) \right) \right]. \tag{10}$$

Algorithm 1 presents the procedure for converting a temporally cumulative collaborative task into the TCCO operator. The algorithm initially computes the value of $\left\lceil \frac{T^S}{\Delta t} \right\rceil$ (line 1). Then, the time interval $[t^{\mathcal{C}}_{\min}, t^{\mathcal{C}}_{\max}]$ is defined and region D^S is described by the conjunction of its associated predicate functions (lines 2-3). Finally, the set $\mathbb{I}_{\mathcal{C}}$ representing agents capable of completing task E^S is determined by checking for an overlap between capabilities and the requirements, and the operator $\mathcal{C}_{[t^{\mathcal{C}}_{\min}, t^{\mathcal{C}}_{\max}]}^{\tau = \left\lceil \frac{T^S}{\Delta t} \right\rceil} \phi_D\{\mathbb{I}_{\mathcal{C}}\}$ is then returned (lines 4-9).

Example 2: We continue with the scenario from Example 1. Assuming a sampling time Δt of 1s and the predicate functions corresponding to region D are given by ϕ_D, then the requirement to satisfy task S with probability at least σ is formalized as the following PrSTL specification:

$$\mathcal{P}\left(\Xi[X(0), \boldsymbol{U}^N] \models \mathcal{C}_{[10,15]}^{\tau=5} \phi_D\{1, 2\} \right) \ge \sigma. \tag{11}$$

Next, consider a specific trajectory over the time window $[10s, 15s]$. On this trajectory, Agent 1 satisfies ϕ_D at $t = 13$ s and 14 s, while Agent 2 satisfies ϕ_D at $t = 12$ s, 13 s, and 14 s. According to Definition 2, we count the number of distinct agent-time pairs that satisfy the predicate. This yields five such pairs: $(2, 12)$, $(1, 13)$, $(2, 13)$, $(1, 14)$, and $(2, 14)$. Since the total count of 5 meets the required threshold $\tau = 5$, the specification $\mathcal{C}_{[10,15]}^{\tau=5} \phi_D\{1, 2\}$ is satisfied for this trajectory.

We then incorporate the TCCOs into the PrSTL framework to specify tasks with temporally cumulative properties in the presence of uncertainties.

Theorem 1: Given a signal (Ξ, k), a probability threshold σ, and a TCCO specification $\mathcal{C}_{[a,b]}^{\tau = \left\lceil \frac{T^S}{\Delta t} \right\rceil} \phi\{\mathbb{I}_{\mathcal{C}}\}$, a sufficient condition for its satisfaction can be constructed as a Boolean combination of atomic chance constraints of the form $\mathcal{P}[h(x_i(k)) \ge 0] \ge 1 - \beta$. Satisfying this Boolean combination guarantees the satisfaction of the original PrSTL specification as:

$$\mathcal{P}\left((\Xi, k) \models \mathcal{C}_{[a,b]}^{\tau = \left\lceil \frac{T^S}{\Delta t} \right\rceil} \phi\{\mathbb{I}_{\mathcal{C}}\} \right) \ge \sigma. \tag{12}$$

Proof: We prove that satisfying the proposed Boolean combination of atomic chance constraints is a sufficient condition for the PrSTL specification in (12). The proof relies on a conservative approximation of the TCCO's probabilistic semantics. We only present the proof for the case of conjunctive propositions and the extension to other Boolean combinations is omitted for brevity.

Let the proposition ϕ be a conjunction of N_p predicates, $\phi = \bigwedge_{p=1}^{N_p} \mu_p$. If the specification is satisfied, there must exist at least one selection of τ distinct agent-time pairs, denoted $l^* = \{(i_1, k'_1), \ldots, (i_\tau, k'_\tau)\}$. For the TCCO to be satisfied by this selection, the proposition ϕ must hold for each pair (i_j, k'_j). This requires that for each of these τ pairs, all N_p predicates μ_p must hold. This results in a total of $N_p\tau$ atomic conditions that must be simultaneously satisfied.

For this specific selection l^*, let $E_{j,p}$ be the event that the predicate μ_p is satisfied for the j-th pair, i.e., $(\Xi_{i_j}, k'_j) \models \mu_p$. Our method asserts an atomic chance constraint on each of these events, such that the probability of the complementary event $\neg E_{j,p}$ is bounded by $\mathcal{P}(\neg E_{j,p}) \le \beta$. The failure of the entire selection l^* occurs if at least one of these $N_p\tau$ atomic conditions is not satisfied. By applying the Boole's inequality across all $N_p\tau$ events, we can bound the probability of this total failure:

$$\mathcal{P}\left(\bigvee_{j=1}^{\tau} \bigvee_{p=1}^{N_p} \neg E_{j,p} \right) \le \sum_{j=1}^{\tau} \sum_{p=1}^{N_p} \mathcal{P}(\neg E_{j,p}) \le N_p \tau \beta. \tag{13}$$

Consequently, the probability that the selection l^* is successful is lower-bounded by $1 - N_p\tau\beta$. The overall TCCO is satisfied if any valid selection is successful. Therefore, the probability of satisfying the TCCO specification is lower-bounded by the success probability of the specific selection l^* found by our method. By setting the violation probability for each atomic constraint as $\beta = (1 - \sigma)/(N_p\tau)$, we guarantee that:

$$\mathcal{P}\left((\Xi, k) \models \mathcal{C}_{[a,b]}^{\tau = \left\lceil \frac{T^S}{\Delta t} \right\rceil} \phi\{\mathbb{I}_{\mathcal{C}}\} \right) \ge 1 - N_p\tau\beta = \sigma. \tag{14}$$

This shows that satisfying the Boolean combination of $N_p\tau$ atomic chance constraints, each with a satisfaction probability of at least $1 - \beta$, is a sufficient condition. □

B. Deterministic Approximation of Chance Constraints

While the atomic chance constraints derived in Theorem 1 provide a formal guarantee, they are generally intractable for direct synthesis. We therefore employ the Chebyshev-Cantelli inequality to obtain a deterministic approximation.

Authorized licensed use limited to: BEIJING INSTITUTE OF TECHNOLOGY. Downloaded on February 22,2026 at 17:01:03 UTC from IEEE Xplore. Restrictions apply.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

TABLE I
INDEX NOTATION FOR PCE

Index	Description
i	Agent index ($i \in \{1, \ldots, N_a\}$)
k	Discrete time step index ($k \in \{0, \ldots, N\}$)
l	State component index ($l \in \{1, \ldots, n_x\}$)
j	PCE basis function index ($j \in \{0, \ldots, L\}$)
g, m	Row and column indices for system matrices A_i, B_i

Lemma 1 ([28]): Given the chance constraint:

$$\mathcal{P}(dx + e \geq 0) \geq 1 - \beta, \tag{15}$$

where x is a scalar random variable with known mean $\mathbf{E}[x]$ and variance $\mathbf{Var}[x]$, d and e are deterministic scalars, and β is the violation probability, it can be conservatively approximated as a second-order cone constraint as follows:

$$d\mathbf{E}[x] + e - |d|\sqrt{\mathbf{Var}[x]}\sqrt{\frac{1-\beta}{\beta}} \geq 0. \tag{16}$$

Applying this lemma requires the statistical moments of the stochastic state variables. To this end, we employ PCE, a framework that approximates a stochastic variable as a series of orthogonal polynomials [24]. We expand all stochastic quantities in the system dynamics (1), including the states $x_i(k)$ and the parameter matrices $A_i(\theta)$ and $B_i(\theta)$:

$$x_{i,l}(k) \approx \hat{x}_{i,l}(k,\theta) = \sum_{j=0}^{L} \bar{x}_{i,l,j,k}\Phi_j(\theta) = \boldsymbol{x}_{i,l}^T(k)\boldsymbol{\Phi}(\theta), \tag{17}$$

$$A_{i,g,m}(\theta) \approx \hat{A}_{i,g,m}(\theta) = \sum_{j=0}^{L} \bar{A}_{i,g,m,j}\Phi_j(\theta), \tag{18}$$

$$B_{i,g,m}(\theta) \approx \hat{B}_{i,g,m}(\theta) = \sum_{j=0}^{L} \bar{B}_{i,g,m,j}\Phi_j(\theta). \tag{19}$$

The indices are summarized in Table I for clarity.

$\hat{x}_{i,l}$ is the PCE approximation of the true random state $x_{i,l}$. The core of PCE is to shift focus from the intractable random variable $x_{i,l}$ to the deterministic vector of its PCE coefficients, $\boldsymbol{x}_{i,l}(k) = [\bar{x}_{i,l,0,k}, \ldots, \bar{x}_{i,l,L,k}]^T$. The number of expansion terms $L = \frac{(n_\theta + c)!}{n_\theta! c!} - 1$ depends on the number of random variables n_θ and the polynomial order c. The coefficients $\bar{A}_{i,g,m,j}$ and $\bar{B}_{i,g,m,j}$ for the system matrices are pre-computed offline via Galerkin projection:

$$\bar{A}_{i,g,m,j} = \frac{\langle A_{i,g,m}(\theta), \Phi_j(\theta)\rangle}{\langle \Phi_j(\theta)^2\rangle}, \tag{20}$$

$$\bar{B}_{i,g,m,j} = \frac{\langle B_{i,g,m}(\theta), \Phi_j(\theta)\rangle}{\langle \Phi_j(\theta)^2\rangle}, \tag{21}$$

where $\langle \Phi_j(\theta)^2\rangle = \mathbf{E}[\Phi_j(\theta)^2]$. Our central goal is to derive the evolution dynamics for the vector of PCE coefficients. To achieve this, we substitute the PCE approximations (17)-(19) back into the original stochastic system dynamics (1):

$$\hat{x}_{i,l}(k+1,\theta) = \sum_{m=1}^{n_x}\sum_{p=0}^{L}\sum_{q=0}^{L} \bar{A}_{i,l,m,p}\bar{x}_{i,m,q,k}\Phi_p(\theta)\Phi_q(\theta)$$

$$+ \sum_{m=1}^{n_u}\sum_{p=0}^{L} \bar{B}_{i,l,m,p}u_{i,m,k}\Phi_p(\theta). \tag{22}$$

The crucial step to obtain a deterministic system is to apply the Galerkin projection to (22) (i.e., taking the inner product with each basis function Φ_j). This projection effectively isolates the dynamics of the coefficients, yielding a deterministic time-invariant system that governs the evolution of the entire PCE coefficient vector for agent i, which we denote as $\boldsymbol{x}_i(k) = [\boldsymbol{x}_{i,1}^T(k), \ldots, \boldsymbol{x}_{i,n_x}^T(k)]^T$:

$$\boldsymbol{x}_i(k+1) = \boldsymbol{A}_i\boldsymbol{x}_i(k) + \boldsymbol{B}_i u_i(k). \tag{23}$$

The expanded matrices $\boldsymbol{A}_i$ and $\boldsymbol{B}_i$ are constant and can be computed offline (see [25], [26] for details). This equation confirms that the PCE coefficient vector $\boldsymbol{x}_i(k+1)$ is uniquely determined by its current value $\boldsymbol{x}_i(k)$, the control input $u_i(k)$, and the pre-computed system matrices. The coefficients are therefore not independent but evolve according to this deterministic system.

A key advantage of PCE is that the moments of the approximated state $\hat{x}_{i,l}(k)$ can be computed directly from its coefficients in $\boldsymbol{x}_i(k)$ due to the orthogonality of the basis functions [23]:

$$\mathbf{E}[x_{i,l}(k)] \approx \mathbf{E}[\hat{x}_{i,l}(k,\theta)] = \bar{x}_{i,l,0,k},$$

$$\mathbf{Var}[x_{i,l}(k)] \approx \mathbf{Var}[\hat{x}_{i,l}(k,\theta)] = \sum_{j=1}^{L} \bar{x}_{i,l,j,k}^2\langle\Phi_j(\theta)^2\rangle. \tag{24}$$

Substituting these expressions into Lemma 1 allows us to transform each atomic chance constraint into a deterministic second-order cone constraint on the PCE coefficients:

$$\mathcal{P}((\Xi_i, k) \models \mu) \geq 1 - \beta \iff \mathcal{P}(dx_{i,l}(k) + e \geq 0) \geq 1 - \beta$$

$$\Longleftarrow d\bar{x}_{i,l,0,k} + e - |d|\sqrt{\sum\nolimits_{j=1}^{L} \bar{x}_{i,l,j,k}^2\langle\Phi_j(\theta)^2\rangle}\sqrt{\frac{1-\beta}{\beta}} \geq 0. \tag{25}$$

Remark 1: We consider the control input to be a deterministic, open-loop decision variable within the optimization horizon, rather than a state-dependent feedback policy. Therefore, we only apply PCE to the stochastic states and system parameters.

C. MISOCP Encoding

We formulate Problem 1 as an MISOCP. The encoding follows a hierarchical structure: we first introduce binary variables to indicate the satisfaction of atomic probabilistic predicates. These binary variables then serve as building blocks to represent the satisfaction of complex spatial regions, which are governed by the TCCO's temporal constraints. The encodings for Boolean operators and other temporal operators follow standard MILP formulations [29].

1) Predicate: Given an predicate μ^i defined by $h(x_i(k)) \geq 0$, we introduce a binary variable $z_k^{\mu^i}$ such that $z_k^{\mu^i} = 1$ if and only if $\mathcal{P}(h(x_i(k)) \geq 0) \geq 1 - \beta$. This logical equivalence is enforced by the following constraints:

$$d\bar{x}_{i,l,0,k} + e - \kappa_\beta|d|\sqrt{\sum\nolimits_{j=1}^{L} \bar{x}_{i,l,j,k}^2\langle\Phi_j(\theta)^2\rangle} \leq M z_k^{\mu^i} - \epsilon,$$

Authorized licensed use limited to: BEIJING INSTITUTE OF TECHNOLOGY. Downloaded on February 22,2026 at 17:01:03 UTC from IEEE Xplore. Restrictions apply.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

$$d\bar{x}_{i,l,0,k}+e-\kappa_\beta|d|\sqrt{\sum_{j=1}^{L}\bar{x}_{i,l,j,k}^2\langle\Phi_j(\theta)^2\rangle}\geq M(z_k^{\mu^i}-1)+\epsilon, \tag{26}$$

where $\kappa_\beta=\sqrt{(1-\beta)/\beta}$; M is a sufficiently large positive constant; ϵ is a sufficiently small positive constant.

2) Region Proposition: The proposition ϕ that defines a region is a Boolean combination of predicates. A corresponding binary variable $z_k^{\phi^i}=1$ is introduced to indicate that agent i satisfies ϕ at time k.

3) Temporal Collective Counting Operator: This encoding enforces that the temporally cumulative property of formula $\mathcal{C}_{[a,b]}^{\tau=\left\lceil\frac{T^S}{\Delta t}\right\rceil}\phi\{\mathbb{I}_\mathcal{C}\}$ is satisfied over the time interval $[a,b]$. We introduce a binary variable $z_k^\mathcal{C}$ to represent the satisfaction of this TCCO formula at time k, which is enforced by:

$$\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}-\tau z_k^\mathcal{C}\geq 0,$$
$$\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}-(\zeta-\tau+1)\,z_k^\mathcal{C}\leq\tau-1, \tag{27}$$

where $\zeta=|\mathbb{I}_\mathcal{C}|(\lfloor\frac{b-a}{\Delta t}\rfloor+1)$ denotes the total number of these binary variables.

Theorem 2: The TCCO $\mathcal{C}_{[a,b]}^{\tau=\left\lceil\frac{T^S}{\Delta t}\right\rceil}\phi\{\mathbb{I}_\mathcal{C}\}$ is satisfied if and only if the corresponding binary variable $z_k^\mathcal{C}$ governed by the constraints in (27) is equal to 1.

Proof: According to (10), the TCCO serves to regulate designated agents' residence duration with respect to target region. The temporally cumulative property is imparted by controlling the quantity of binary variables that take the value of 1 as:

$$\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}\geq\tau. \tag{28}$$

Hence, we need only demonstrate the equivalence between (27) and the assertion that $z_k^\mathcal{C}=1$ iff (28) is satisfied.

The proof commences with the establishment of the sufficient condition. Assuming constraints (27) are satisfied, it is required to demonstrate the subsequent two propositions: (i) if $\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}\geq\tau$, then $z_k^\mathcal{C}=1$, and (ii) if $\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}<\tau$, then $z_k^\mathcal{C}=0$.

(i) We employ the method of proof by contradiction. Assuming $z_k^\mathcal{C}=0$, then we have:

$$\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}-0\leq\tau-1\Rightarrow\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}\leq\tau-1. \tag{29}$$

This constitutes a contradiction with the premise $\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}\geq\tau$. Given that $z_k^\mathcal{C}$ is a binary variable, it necessarily follows that $z_k^\mathcal{C}=1$.

(ii) We employ the method of proof by contradiction. Assuming $z_k^\mathcal{C}=1$, then we have:

$$\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}\geq\tau\cdot 1\Rightarrow\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}\geq\tau. \tag{30}$$

This constitutes a contradiction with the premise $\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}<\tau$. Given that $z_k^\mathcal{C}$ is a binary variable, it necessarily follows that $z_k^\mathcal{C}=0$.

(i) and (ii) jointly establish the sufficiency of the proposition. We now proceed to demonstrate the necessity. Assume that $z_k^\mathcal{C}=1$ if and only if (28) is satisfied, it is required to demonstrate the subsequent two propositions: (A) if $\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}\geq\tau$ and $z_k^\mathcal{C}=1$, then (27) is satisfied, and (B) if $\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}<\tau$ and $z_k^\mathcal{C}=0$, then (27) is satisfied.

(A) We first verify whether the first constraint is satisfied:

$$\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}-\tau z_k^\mathcal{C}\geq 0\Rightarrow\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}-\tau\geq 0. \tag{31}$$

Given that $\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}\geq\tau$, the first constraint is thereby satisfied.

Then verify the second constraint:

$$\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}-\gamma z_k^\mathcal{C}\leq\tau-1\Rightarrow\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}\leq\zeta, \tag{32}$$

where $\gamma=(\zeta-\tau+1)$. Given that $z_{k'}^{\phi^i}\in\{0,1\}$, the second constraint is thereby satisfied.

(B) We first verify whether the first constraint is satisfied:

$$\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}-\tau z_k^\mathcal{C}\geq 0\Rightarrow\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}\geq 0. \tag{33}$$

Given that $z_{k'}^{\phi^i}\geq 0$, the first constraint is thereby satisfied.

Then verify the second constraint:

$$\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}-\gamma z_k^\mathcal{C}\leq\tau-1\Rightarrow\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}\leq\tau-1. \tag{34}$$

Given that $\sum_{i\in\mathbb{I}_\mathcal{C}}\sum_{k'=k+a}^{k+b}z_{k'}^{\phi^i}<\tau$ and $z_{k'}^{\phi^i}\in\{0,1\}$, the second constraint is thereby satisfied.

(A) and (B) jointly establish the necessity of the proposition. □

4) Encoding for Compound PrSTL-TCCO Specifications: A key feature of our framework is its ability to integrate TCCOs with standard PrSTL specifications. This integration extends to the optimization level, where a composite specification is systematically converted into MISOCP constraints by recursively encoding its logical structure. Crucially, the method for encoding a TCCO is consistent with standard formulations for conventional temporal logic operators.

D. Constraint Relaxation Mechanism

While PCE introduces a controllable truncation error in the moment calculations, the dominant source of conservatism in our method stems from the Boole's and Chebyshev-Cantelli inequalities. The resulting safety margin typically far exceeds the PCE approximation error, rendering the approach practically sound. However, this conservatism can also lead to infeasibility or excessive computational cost. We therefore introduce a constraint relaxation mechanism to improve solution efficiency. While this sacrifices the strict formal guarantee to gain feasibility, the inherent conservatism of our formulation often means

Authorized licensed use limited to: BEIJING INSTITUTE OF TECHNOLOGY. Downloaded on February 22,2026 at 17:01:03 UTC from IEEE Xplore. Restrictions apply.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

TABLE II
CAPABILITIES OF AGENTS

	Agent 1,2,3	Agent 4,5,6	Agent 7,8
Capabilities	Irrigation Fertilization	Fertilization Pest control	Irrigation Fertilization Pest control

the controller can still satisfy the specification with a high probability.

The constraint relaxation mechanism reduces conservatism by modifying the violation probability of predicates. To handle cases where the initial problem formulation is too restrictive, an iterative approach is employed. If the solver reports that it cannot find a solution or fails to produce a solution within a predefined time threshold T, the attempt is halted and the violation probability β is relaxed. This is achieved by increasing β via the update rule $\beta \leftarrow \min(\beta k, \beta_{\max})$, using a relaxation coefficient $k > 1$ and a saturation limit $\beta_{\max} < 1$ to maintain a valid probability measure. The process repeats with relaxed constraints until a solution is found. A final attempt is made with $\beta = \beta_{\max}$ and no time limit. If this also fails, the synthesis process terminates and reports infeasibility.

V. Case Study

We evaluate our method through multi-agent simulations across two scenarios, each with configurations of 3, 4, and 6 agents. For each configuration, 500 Monte Carlo simulations are run. All problems are solved using Gurobi on a desktop with an Intel Core i5 13600 K (5.1 GHz) and 32 GB of RAM.

The dynamics of agent $i \in \mathbb{I} = \{1, 2, 3, 4, 5, 6, 7, 8\}$ are given as follows:

$$x_i(k+1) = \begin{bmatrix} I_2 & \Delta t \theta_1 I_2 \\ 0 & I_2 \end{bmatrix} x_i(k) + \begin{bmatrix} \frac{(\Delta t)^2}{2}\theta_2 I_2 \\ \Delta t \theta_2 I_2 \end{bmatrix} u_i(k) \quad (35)$$

where the parametric uncertainties are assumed to follow a normal distribution ($\theta_1 \sim \mathcal{N}(1, 0.01^2)$ and $\theta_2 \sim \mathcal{N}(1, 0.01^2)$). The capabilities of each agent are presented in Table II. The state $x_i = [l_{x_i}, l_{y_i}, v_{x_i}, v_{y_i}]^\top \in \mathbb{R}^4$ denotes the position of agent i and its velocity along the X and Y axes and $u_i = [a_{x_i}, a_{y_i}]^\top$ denotes the commanded accelerations, subject to the constraint that $\|u_i\|_2 \leq 0.2\ m/s^2$. The sampling time Δt is set to 1 s. For the constraint relaxation mechanism, we use a relaxation coefficient of $k = 10$, a solver time threshold of T = 15 s, and a saturation limit of $\beta_{\max} = 0.5$.

Scenario 1 assesses the framework on complex collaborative missions with temporally cumulative requirements defined by multiple TCCO specifications. Each agent is initialized at a uniformly random position within the initial area. An example of a resulting trajectory from a successful run is shown in Fig. 2(a). The task specifications for this scenario are as follows:

- *Irrigation in Region A:* $\varphi_1 := \mathcal{C}^{\tau=4}_{[4,11]} \phi_A\{1, 2, 3, 7, 8\}$
- *Fertilization in Region B:* $\varphi_2 := \mathcal{C}^{\tau=8}_{[0,14]} \phi_B\{1, 2, \dots, 8\}$
- *Pest control in Region B:* $\varphi_3 := \mathcal{C}^{\tau=5}_{[0,9]} \phi_B\{4, 5, 6, 7, 8\}$
- *Pest control in Region C:* $\varphi_4 := \mathcal{C}^{\tau=8}_{[0,14]} \phi_C\{4, 5, 6, 7, 8\}$
- *Avoidance of Region D:* $\varphi_5 := G_{[0,14]} \neg \phi_D$

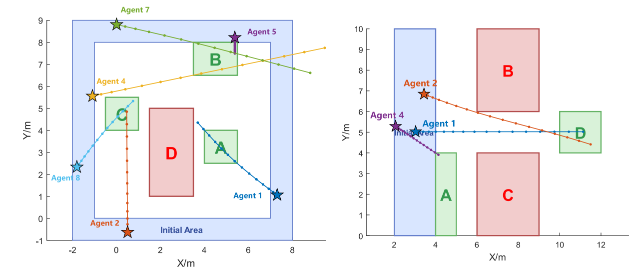

(a) (b)

Fig. 2. Example trajectories of the multi-agent system with PCE. (a) Shows a run from a Monte Carlo simulation for Scenario 1, and (b) shows a run for Scenario 2.

TABLE III
PERFORMANCE COMPARISON FOR SCENARIO 1

Agent Configuration		3 Agents (1,4,7)	4 Agents (1,4,5,7)	6 Agents (1,2,4,5,7,8)
MILP	Solver Time (s)	2.03 ± 1.04	5.58 ± 2.82	38.34 ± 25.98
	Success Rate	0%	0%	0%
	Infeasibility Rate	65.6%	23.2%	12.0%
	Failure Rate	34.4%	76.8%	88.0%
Proposed Method	Solver Time (s)	9.70 ± 3.23	20.47 ± 10.60	83.61 ± 56.99
	Success Rate	30.2%	67.4%	87.8%
	Infeasibility Rate	68.4%	31.4%	10.8%
	Failure Rate	1.4%	1.2%	1.4%

The global specification is the conjunction of these individual tasks: $\varphi_{\text{total}} := \varphi_1^{(0.9)} \wedge \varphi_2^{(0.9)} \wedge \varphi_3^{(0.9)} \wedge \varphi_4^{(0.9)} \wedge \varphi_5^{(0.9)}$.

We evaluate performance using four metrics. Solver Time is the time reported by the solver, and Infeasibility Rate is the percentage of runs where no feasible solution is found. For the remaining runs, Success and Failure Rates are the percentages of trajectories that respectively satisfy or fail the specification due to uncertainty. As presented in Table III, where indices in parentheses denote the specific indices of agents as defined in Table II, the standard MILP approach is unreliable due to a high failure rate. In contrast, our proposed method consistently generates robust solutions with near-zero failures, validating PCE's effectiveness despite longer solver times. Regarding scalability, while adding agents improves MILP feasibility, high failure rates render it impractical. In contrast, our method leverages the expanded collaborative possibilities, trading increased computation time for a significantly higher success rate. Furthermore, the large standard deviation in solver time for the 6-agent scenario indicates that solver performance is highly sensitive to the agents' initial configurations.

Scenario 2 evaluates the integration of TCCO with standard PrSTL specifications. Each agent is initialized at a uniformly random position within the initial area. An example of a resulting trajectory from a successful run is shown in Fig. 2(b). The mission is formally specified as follows:

- *Reaching Region A:* $\varphi_1 := F_{[0,18]} \phi_A$
- *Avoidance of Region B:* $\varphi_2 := G_{[0,18]} \neg \phi_B$
- *Avoidance of Region C:* $\varphi_3 := G_{[0,18]} \neg \phi_C$

Authorized licensed use limited to: BEIJING INSTITUTE OF TECHNOLOGY. Downloaded on February 22,2026 at 17:01:03 UTC from IEEE Xplore. Restrictions apply.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

TABLE IV
PERFORMANCE COMPARISON FOR SCENARIO 2

Agent Configuration		3 Agents (1,2,4)	4 Agents (1,2,4,5)	6 Agents (1,2,3,4,5,6)
MILP	Solver Time (s)	2.28 ± 0.60	9.84 ± 6.04	48.99 ± 26.87
	Success Rate	0%	0%	0%
	Infeasibility Rate	0.4%	0%	0%
	Failure Rate	99.6%	100%	100%
Proposed Method	Solver Time (s)	13.67 ± 5.33	25.93 ± 9.42	98.11 ± 62.56
	Success Rate	99.8%	100.0%	100.0%
	Infeasibility Rate	0.2%	0%	0%
	Failure Rate	0%	0%	0%

- *Fertilization in Region D:* $\varphi_4 := \mathcal{C}^{\tau=5}_{[3,18]}\phi_D\{1,2,\ldots,6\}$

The global specification is the conjunction of these individual tasks: $\varphi_{\text{total}} := \varphi_1^{(0.9)} \wedge \varphi_2^{(0.9)} \wedge \varphi_3^{(0.9)} \wedge \varphi_4^{(0.9)}$.

The results for Scenario 2 are shown in Table IV. Here, the cumulative requirement is low, making safety the primary challenge. Consequently, increasing the number of agents does not significantly improve the success rate, which is already high even for a 3-agent team. Our method achieves a far superior success rate over the MILP approach due to its explicit handling of uncertainty, despite substantially longer solver times. In terms of scalability, the computational cost of our method still increases significantly with system size.

VI. Conclusion

This letter extends the PrSTL framework to handle temporally cumulative collaborative tasks for uncertain stochastic systems. We introduce the TCCO to represent these tasks. By employing PCE for uncertainty propagation, the problem is formulated as a MISOCP, with a proposed constraint relaxation mechanism to reduce conservatism. While computationally intractable for large-scale systems, future work will focus on real-time implementation by developing a distributed control architecture to reduce the computational load.

References

[1] X. Yin, B. Gao, and X. Yu, "Formal synthesis of controllers for safety-critical autonomous systems: Developments and challenges," *Annu. Rev. Control*, vol. 57, 2024, Art. no. 100940.

[2] K. Leahy et al., "Scalable and robust algorithms for task-based coordination from high-level specifications (scratches)," *IEEE Trans. Robot.*, vol. 38, no. 4, pp. 2516–2535, Aug. 2022.

[3] D. Sun, J. Chen, S. Mitra, and C. Fan, "Multi-agent motion planning from signal temporal logic specifications," *IEEE Robot. Automat. Lett.*, vol. 7, no. 2, pp. 3451–3458, Apr. 2022.

[4] L. Li, Z. Chen, H. Wang, and Z. Kan, "Task allocation of heterogeneous robots under temporal logic specifications with inter-task constraints and variable capabilities," *IEEE Trans. Autom. Sci. Eng.*, vol. 22, pp. 14030–14047, 2025.

[5] Z. Chen and Z. Kan, "Real-time reactive task allocation and planning of large heterogeneous multi-robot systems with temporal logic specifications," *Int. J. Robot. Res.*, vol. 44, no. 4, pp. 640–664, 2025.

[6] Y. Takayama, K. Hashimoto, and T. Ohtsuka, "STLCCP: Efficient convex optimization-based framework for signal temporal logic specifications," *IEEE Trans. Autom. Control*, vol. 70, no. 9, pp. 6064–6079, Sep. 2025.

[7] M. Cannon, B. Kouvaritakis, and X. Wu, "Probabilistic constrained MPC for multiplicative and additive stochastic uncertainty," *IEEE Trans. Autom. Control*, vol. 54, no. 7, pp. 1626–1632, Jul. 2009.

[8] D. Bernardini and A. Bemporad, "Stabilizing model predictive control of stochastic constrained linear systems," *IEEE Trans. Autom. Control*, vol. 57, no. 6, pp. 1468–1480, Jun. 2012.

[9] O. A. Hafez, G. D. Arana, and M. Spenko, "Integrity risk-based model predictive control for mobile robots," in *Proc. Int. Conf. Robot. Automat.*, 2019, pp. 5793–5799.

[10] Y. Yao, J. Sun, and Y. Zhang, "Model predictive control of stochastic hybrid systems with signal temporal logic constraints," *Automatica*, vol. 173, 2025, Art. no. 112038.

[11] D. Sadigh and A. Kapoor, "Safe control under uncertainty with probabilistic signal temporal logic," in *Proc. Robot. Sci. Syst.*, 2016.

[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," *IEEE Trans. Autom. Control*, vol. 64, no. 8, pp. 3324–3331, Aug. 2019.

[13] T. Yang, Y. Zou, S. Li, and Y. Yang, "Distributed model predictive control for probabilistic signal temporal logic specifications," *IEEE Trans. Autom. Sci. Eng.*, vol. 21, no. 4, pp. 6239–6249, Oct. 2024.

[14] S. Safaoui, L. Lindemann, D. V. Dimarogonas, I. Shames, and T. H. Summers, "Control design for risk-based signal temporal logic specifications," *IEEE Contr. Syst. Lett.*, vol. 4, no. 4, pp. 1000–1005, Oct. 2020.

[15] M. Vahs, C. Pek, and J. Tumova, "Risk-aware spatio-temporal logic planning in Gaussian belief spaces," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2023, pp. 7879–7885.

[16] L. Lindemann, G. J. Pappas, and D. V. Dimarogonas, "Reactive and risk-aware control for signal temporal logic," *IEEE Trans. Autom. Control*, vol. 67, no. 10, pp. 5262–5277, Oct. 2022.

[17] A. T. Buyukkocak, D. Aksaray, and Y. Yazıcıoğlu, "Planning of heterogeneous multi-agent systems under signal temporal logic specifications with integral predicates," *IEEE Robot. Automat. Lett.*, vol. 6, no. 2, pp. 1375–1382, Apr. 2021.

[18] H. Chen et al., "Cumulative-time signal temporal logic," *ACM Trans. Embedded Comput. Syst.*, vol. 24, pp. 1–23, 2025.

[19] B. Yalcinkaya, N. Lauffer, M. Vazquez-Chanlatte, and S. A. Seshia, "Provably correct automata embeddings for optimal automata-conditioned reinforcement learning," 2025, *arXiv:2503.05042*.

[20] C. Neary, Z. Xu, B. Wu, and U. Topcu, "Reward machines for cooperative multi-agent reinforcement learning," 2020, *arXiv:2007.01962*.

[21] A. Shah, N. Lauffer, T. Chen, N. Pitta, and S. A. Seshia, "Learning symbolic task decompositions for multi-agent teams," 2025, *arXiv:arXiv:2502.13376*.

[22] L. Lindemann and D. V. Dimarogonas, *Formal Methods for Multi-Agent Feedback Control Systems*. Cambridge, MA, USA: MIT Press, 2025.

[23] O. Le Maître and O. M. Knio, *Spectral Methods for Uncertainty Quantification: With Applications to Computational Fluid Dynamics*. Berlin, Germany: Springer, 2010.

[24] D. Landgraf, A. Völz, F. Berkel, K. Schmidt, T. Specker, and K. Graichen, "Probabilistic prediction methods for nonlinear systems with application to stochastic model predictive control," *Annu. Rev. Control*, vol. 56, 2023, Art. no. 100905.

[25] J. A. Paulson, A. Mesbah, S. Streif, R. Findeisen, and R. D. Braatz, "Fast stochastic model predictive control of high-dimensional systems," in *Proc. 53rd IEEE Conf. Decis. Control*, 2014, pp. 2802–2809.

[26] S. Lucia, P. Zometa, M. Kögel, and R. Findeisen, "Efficient stochastic model predictive control based on polynomial chaos expansions for embedded applications," in *Proc. 54th IEEE Conf. Decis. Control*, 2015, pp. 3006–3012.

[27] J. R. Fisher, *Stability Analysis and Control of Stochastic Dynamic Systems Using Polynomial Chaos*. College Station, TX, USA: Texas A&M Univ., 2008.

[28] G. C. Calafiore and L. E. Ghaoui, "On distributionally robust chance-constrained linear programs," *J. Optim. Theory Appl.*, vol. 130, pp. 1–22, 2006.

[29] 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 *Proc. 53rd IEEE Conf. Decis. Control*, 2014, pp. 81–87.

Authorized licensed use limited to: BEIJING INSTITUTE OF TECHNOLOGY. Downloaded on February 22,2026 at 17:01:03 UTC from IEEE Xplore. Restrictions apply.