2024 IEEE 63rd Conference on Decision and Control (CDC)
December 16-19, 2024. MiCo, Milan, Italy

Probabilistic Tube-based Control Synthesis of Stochastic Multi-Agent Systems under Signal Temporal Logic

Eleftherios E. Vlahakis[1], Lars Lindemann[2], Pantelis Sopasakis[3] and Dimos V. Dimarogonas[1]

***Abstract*— We consider the control design of stochastic discrete-time linear multi-agent systems (MASs) under a global signal temporal logic (STL) specification to be satisfied at a predefined probability. By decomposing the dynamics into deterministic and error components, we construct a probabilistic reachable tube (PRT) as the Cartesian product of reachable sets of the individual error systems driven by disturbances lying in confidence regions (CRs) with a fixed probability. By bounding the PRT probability with the specification probability, we tighten all state constraints induced by the STL specification by solving tractable optimization problems over segments of the PRT, and relax the underlying stochastic problem with a deterministic one. This approach reduces conservatism compared to tightening guided by the STL structure. Additionally, we propose a recursively feasible algorithm to attack the resulting problem by decomposing it into agent-level subproblems, which are solved iteratively according to a scheduling policy. We demonstrate our method on a ten-agent system, where existing approaches are impractical.**

I. INTRODUCTION

Multi-agent systems (MASs) can be found in many applications, such as robotics, autonomous vehicles, and cyber-physical systems. When these systems are stochastic, the formal specification of system properties can be formulated probabilistically. As the complexity in control synthesis from temporal logic under uncertainty grows with the dimensionality of the overall system, existing approaches typically focus on single-agent [1] or non-stochastic [2], [3] systems.

In this paper, we focus on signal temporal logic (STL) [4] to formally formulate and verify specifications for a wide range of MASs. STL employs predicates coupled with Boolean and temporal operators, allowing precise specification of complex spatio-temporal properties in a dynamical system. In a deterministic setting, it is possible to design sound and complete algorithms that guarantee STL satisfaction [5], based on the quantitative semantics of STL [6]. Here, we consider stochastic MASs and a stochastic optimal control problem, where the goal is to satisfy a multi-agent STL specification with a predefined probability.

To address stochasticity in the STL framework, the works in [7]–[9] propose risk constraints over predicates while preserving Boolean and temporal operators. Probabilistic STL in [10] allows one to express uncertainty by incorporating random variables into predicates, while [11] introduces chance-constrained temporal logic for modeling uncertainty. Similar approaches are found in [12], [13]. Top-down approaches imposing chance constraints on the entire specification are explored in [1], [14], [15]. Although important, these works focus on low-dimensional systems and lack guidance on extending to MASs [16]. A recent extension of [1] to stochastic MASs under STL [17] considers only a single joint task per agent and bounded distributions.

Here, we solve a stochastic optimal control problem of a discrete-time linear MAS subject to additive stochastic perturbations and a global STL specification permitting multiple individual and joint tasks per agent. First, we decompose the multi-agent dynamics into a deterministic system and an error closed-loop stochastic system, for which we construct a probabilistic reachable tube (PRT) [18] as the Cartesian product of reachable sets of individual error systems. These are driven by stochastic disturbances lying in confidence regions (CRs) with a fixed probability. By assuming independence among individual disturbances, we show that the PRT probability can be controlled by the product of probabilities selected for each individual CR and a union-bound argument applied over time. Thus, by lower bounding the PRT probability by the specification probability, we can tighten all state constraints induced by the STL specification by solving tractable optimization problems over segments of the PRT. For multi-agent STL specifications, this is a less conservative alternative to tightening approaches relying on the STL structure [1], [17]. An attainable feasible solution to the resulting deterministic problem can then be used to synthesize multi-agent trajectories that satisfy the STL specification with the desired confidence level. To the best of the authors' knowledge, this work is the first to address stochastic MASs under STL utilizing PRTs. Subsequently, to enhance scalability, we decompose the resulting deterministic problem into agent-level subproblems, which are solved iteratively according to a scheduling policy. We show that this iterative procedure is recursively feasible, ensures satisfaction of local tasks, and guarantees nondecreasing robustness for joint tasks.

The remainder of the paper is organized as follows. Preliminaries and the control problem setup are in Sec. II. The construction of PRTs, the constraint tightening and the distributed control synthesis, are in Sec. III. An illustrative numerical example is in Sec. IV, whereas concluding remarks are discussed in Sec. V.

This work was supported by the Swedish Research Council (VR), the Knut & Alice Wallenberg Foundation (KAW), the Horizon Europe Grant SymAware and the ERC Consolidator Grant LEAFHOUND. [1]Division of Decision and Control Systems, School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, 10044, Stockholm, Sweden. Email: {vlahakis,dimos}@kth.se. [2]Thomas Lord Department of Computer Science, Viterbi School of Engineering, University of Southern California, Los Angeles, 90089, CA, USA. Email: llindema@usc.edu. [3]School of Electronics, Electrical Engineering and Computer Science, Queen's University Belfast, Northern Ireland, UK. Email: p.sopasakis@qub.ac.uk

II. Problem setup

A. Notation and Preliminaries

The sets of real numbers and nonnegative integers are $\mathbb{R}$ and $\mathbb{N}$, respectively. Let $N \in \mathbb{N}$. Then, $\mathbb{N}_{[0,N]} = \{0, 1, \ldots, N\}$. Let $x_1, \ldots, x_n$ be vectors. Then, $x = (x_1, \ldots, x_n) = [x_1^\intercal \ \cdots \ x_n^\intercal]^\intercal$. We denote by $\boldsymbol{x}(a : b) = (x(a), \ldots, x(b))$ an aggregate vector consisting of $x(t)$, $t \in \mathbb{N}_{[a,b]}$, representing a trajectory. When it is clear from the context, we write $\boldsymbol{x}(t)$, omitting the endpoint. When $x(t)$, $t \in \mathbb{N}_{[a,b]}$, are random vectors, $\boldsymbol{x}(a : b) = (x(a), \ldots, x(b))$ is a random process. Let $x_i(t)$, for $t \in \mathbb{N}_{[0,N]}$ and $i \in \mathbb{N}_{[1,M]}$. Then, $\boldsymbol{x}(0 : N) = (x(0), \ldots, x(N))$ denotes an aggregate trajectory when $x(t) = (x_1(t), \ldots, x_M(t))$, $t \in \mathbb{N}_{[0,N]}$. The remainder of the division of a by b is $\mathrm{mod}(a, b)$. A random variable (vector) w following a distribution $\mathcal{D}_w$ is denoted as $w \sim \mathcal{D}_w$, the support of $\mathcal{D}_w$ is $\mathrm{supp}(\mathcal{D}_w)$, the expected value of w is $\mathbb{E}(w)$, and the variance (covariance matrix) of w is $\mathrm{Var}(w)$ ($\mathrm{Cov}(w)$). The probability of event Y is $\Pr\{Y\}$. The cardinality of a set $\mathcal{V}$ is $|\mathcal{V}|$. The Minkowski sum and the Pontryagin set difference of $S_1 \subseteq \mathbb{R}^n$ and $S_2 \subseteq \mathbb{R}^n$ are $S_1 \oplus S_2 = \{s_1 + s_2 \mid s_1 \in S_1,\ s_2 \in S_2\}$ and $S_1 \ominus S_2 = \{s_1 \in S_1 \mid s_1 + s_2 \in S_1, \forall s_2 \in S_2\}$, respectively.

Lemma 1 (Distributivity of Minkowski sum) *Let* $\mathcal{X}_i, \mathcal{Y}_i \subseteq \mathbb{R}^{n_i}$, $i \in \mathbb{N}_{[1,M]}$. *Then,* $(\mathcal{X}_1 \times \cdots \times \mathcal{X}_M) \oplus (\mathcal{Y}_1 \times \cdots \times \mathcal{Y}_M) = (\mathcal{X}_1 \oplus \mathcal{Y}_1) \times \cdots \times (\mathcal{X}_M \oplus \mathcal{Y}_M)$.

Proof: It holds that $(\mathcal{X}_1 \times \ldots \times \mathcal{X}_M) \oplus (\mathcal{Y}_1 \times \ldots \times \mathcal{Y}_M) = \{(x_1, \ldots, x_M) + (y_1, \ldots y_M) \mid \forall x_i \in \mathcal{X}_i,\ i \in \mathbb{N}_{[1,M]}, \text{and } \forall y_i \in \mathcal{Y}_i,\ i \in \mathbb{N}_{[1,M]}\} = \{(x_1 + y_1, \ldots, x_M + y_M) \mid \forall x_i \in \mathcal{X}_i, y_i \in \mathcal{Y}_i,\ i \in \mathbb{N}_{[1,M]}\} = (\mathcal{X}_1 \oplus \mathcal{Y}_1) \times \ldots \times (\mathcal{X}_M \oplus \mathcal{Y}_M)$. ■

We consider STL formulas with standard syntax

$$\varphi \coloneqq \top \mid \pi \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid \phi_1 U_{[t_1,t_2]} \phi_2, \quad (1)$$

where $\pi := (\mu(x) \geq 0)$ is a predicate, $\mu(x) := a^\intercal x + b$ is an affine predicate function, with $a \in \mathbb{R}^{n_x}$, $x \in \mathbb{R}^{n_x}$, and $b \in \mathbb{R}$, and ϕ, ϕ_1, and ϕ_2 are STL formulas, which are built recursively using predicates π, logical operators $\neg$ and $\wedge$, and the *until* temporal operator U, with $[t_1, t_2] \equiv \mathbb{N}_{[t_1,t_2]}$. We omit $\vee$ (*or*), $\Diamond$ (*eventually*) and $\Box$ (*always*) operators from (1) and the sequel, as these may be defined by (1), e.g., $\phi_1 \vee \phi_2 = \neg(\neg\phi_1 \wedge \neg\phi_2)$, $\Diamond_{[t_1,t_2]}\phi = \top U_{[t_1,t_2]}\phi$, and $\Box_{[t_1,t_2]}\phi = \neg\Diamond_{[t_1,t_2]}\neg\phi$.

Let π be a predicate and ϕ an STL formula. We write $\pi \in \phi$ to indicate that π is part of the formula ϕ. We denote by $\boldsymbol{x}(t) \models \phi$, $t \in \mathbb{N}$, the satisfaction of ϕ, verified over $\boldsymbol{x}(t) = (x(t), x(t+1), \ldots)$. The validity of a formula ϕ can be verified using Boolean semantics: $\boldsymbol{x}(t) \models \pi \Leftrightarrow \mu(x(t)) \geq 0$, $\boldsymbol{x}(t) \models \neg\phi \Leftrightarrow \neg(\boldsymbol{x}(t) \models \phi)$, $\boldsymbol{x}(t) \models \phi_1 \wedge \phi_2 \Leftrightarrow \boldsymbol{x}(t) \models \phi_1 \wedge \boldsymbol{x}(t) \models \phi_2$, $\boldsymbol{x}(t) \models \phi_1 U_{[a,b]} \phi_2 \Leftrightarrow \exists \tau \in t \oplus \mathbb{N}_{[a,b]}$, s.t. $\boldsymbol{x}(\tau) \models \phi_2 \wedge \forall \tau' \in \mathbb{N}_{[t,\tau]}, \boldsymbol{x}(\tau') \models \phi_1$. Based on the Boolean semantics, the horizon of a formula is recursively defined as [4]: $N^\pi = 0$, $N^{\neg\phi} = N^\phi$, $N^{\phi_1 \wedge \phi_2} = \max(N^{\phi_1}, N^{\phi_2})$, $N^{\phi_1\ U_{[a,b]} \phi_2} = b + \max(N^{\phi_1}, N^{\phi_2})$.

STL is endowed with quantitative semantics [6]: A scalar-valued function $\rho^\phi : \mathbb{R}^n \times \cdots \times \mathbb{R}^n \to \mathbb{R}$ of a signal indicates how robustly a signal $\boldsymbol{x}(t)$ satisfies a formula ϕ. The robustness function is defined recursively as follows: $\rho^\pi(\boldsymbol{x}(t)) = \mu(x(t))$, $\rho^{\neg\phi}(\boldsymbol{x}(t)) = -\rho^\phi(\boldsymbol{x}(t))$, $\rho^{\phi_1 \wedge \phi_2}(\boldsymbol{x}(t)) = \min(\rho^{\phi_1}(\boldsymbol{x}(t)), \rho^{\phi_2}(\boldsymbol{x}(t)))$, and $\rho^{\phi_1 U_{[a,b]} \phi_2}(\boldsymbol{x}(t)) = \max_{\tau \in t \oplus \mathbb{N}_{[a,b]}} (\min(Y_1(\tau), Y_2(\tau')))$, with $Y_1(\tau) = \rho^{\phi_1}(\boldsymbol{x}(\tau))$, $Y_2(\tau') = \min_{\tau' \in \mathbb{N}_{[t,\tau]}} \rho^{\phi_2}(\boldsymbol{x}(\tau'))$, π being a predicate, and ϕ, ϕ_1, and ϕ_2 being STL formulas.

Definition 1 *Let* $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ *be an undirected graph containing no self-loops, with node set* $\mathcal{V}$*, cardinality* $M = |\mathcal{V}|$*, and edge set* $\mathcal{E}$*. Let also* $\mathcal{V}' \subseteq \mathcal{V}$*, with* $|\mathcal{V}'| > 1$*, and define* $\mathcal{E}_{\mathcal{V}'} \subseteq \mathcal{E}$ *as the set of edges attached to nodes* $\mathcal{V}'$*. Then,* $\mathcal{G}' = (\mathcal{V}', \mathcal{E}_{\mathcal{V}'})$ *is a clique [19], i.e., a complete subgraph of* $\mathcal{G}$*, if* $\mathcal{E}_{\mathcal{V}'}$ *contains all possible edges between nodes* $\mathcal{V}'$*. The set of cliques of* $\mathcal{G}$ *is defined as* $\mathcal{K} = \{\nu \subseteq \mathcal{V} \mid (\nu, \mathcal{E}_\nu)$ *is a complete subgraph of* $\mathcal{G}\}$.

Consider a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with clique set $\mathcal{K}$, a clique $\nu \in \mathcal{K}$, with $\nu = (i_1, \ldots, i_{|\nu|})$, and vectors $x_{i_j}(t)$, $j \in \mathbb{N}_{[1,|\nu|]}$, with $t \in \mathbb{N}$, Then, $x_\nu(t) = (x_{i_1}(t), \ldots, x_{i_{|\nu|}}(t))$ is an aggregate vector. We denote by $\boldsymbol{x}_\nu(t) \models \phi_\nu$ the validity of an STL formula defined over the aggregate trajectory $\boldsymbol{x}_\nu(t) = (x_\nu(t), x_\nu(t+1), \ldots)$. If $\pi_\nu \in \phi_\nu$, $\pi_\nu := (\mu_\nu(x_\nu) \geq 0)$, where $\mu_\nu(x_\nu)$ is an affine predicate function of x_ν, with $x_\nu = (x_{i_1}, \ldots, x_{i_{|\nu|}})$.

B. Multi-agent system

1) Dynamics: We consider M agents with dynamics

$$x_i(t+1) = A_i x_i(t) + B_i u_i(t) + w_i(t), \quad (2)$$

where $x_i(t) \in \mathcal{X}_i \subseteq \mathbb{R}^{n_i}$, $u_i \in \mathcal{U}_i \subseteq \mathbb{R}^{m_i}$, and $w_i(t) \in \mathcal{W}_i \subseteq \mathbb{R}^{n_i}$ are the state, input and disturbance vectors, respectively, the initial condition, $x_i(0)$, is known, (A_i, B_i) is a stabilizable pair, with $A_i \in \mathbb{R}^{n_i \times n_i}$, $B_i \in \mathbb{R}^{n_i \times m_i}$, $i \in \mathbb{N}_{[1,M]}$, and $t \in \mathbb{N}$. By collecting individual state, input, and disturbance vectors, as $x(t) = (x_1(t), \ldots, x_M(t)) \in \mathcal{X} \subseteq \mathbb{R}^n$, $u(t) = (u_1(t), \ldots, u_M(t)) \in \mathcal{U} \subseteq \mathbb{R}^m$, and $w(t) = (w_1(t), \ldots, w_M(t)) \in \mathcal{W} \subseteq \mathbb{R}^n$, respectively, we write the dynamics of the entire MAS as

$$x(t+1) = Ax(t) + Bu(t) + w(t), \quad (3)$$

where $A = \mathrm{diag}(A_1, \ldots, A_M)$, $B = \mathrm{diag}(B_1, \ldots, B_M)$, and the state, input, and disturbance sets are $\mathcal{X} = \mathcal{X}_1 \times \cdots \times \mathcal{X}_M$, $\mathcal{U} = \mathcal{U}_1 \times \cdots \times \mathcal{U}_M$, and $\mathcal{W} = \mathcal{W}_1 \times \cdots \times \mathcal{W}_M$, respectively.

2) Disturbance: We assume that the uncertain sequence $\boldsymbol{w}_i(0) = (w_i(0), w_i(1), \ldots)$, with $i \in \mathbb{N}_{[1,M]}$, is an independent and identically distributed random process, and that $w_i(t)$ is a random vector with mean $\mathbb{E}(w_i(t)) = 0$ and positive definite covariance matrix $\mathrm{Cov}(w_i(t)) = Q_i$, which is known, for all $t \in \mathbb{N}$. We also assume that $w_i(t)$, $i \in \mathbb{N}_{[1,M]}$, are independent for $t \in \mathbb{N}$. We denote by $\mathcal{D}_{w_i}$ the distribution of the disturbance $w_i(t) \in \mathcal{W}_i$, where $\mathcal{W}_i$ is its support, which may be unbounded. We also may write that $w(t) \sim \mathcal{D}_w$, with $\mathrm{supp}(\mathcal{D}_w) = \mathcal{W} = \mathcal{W}_1 \times \cdots \times \mathcal{W}_M$, $\mathbb{E}(w(t)) = 0$, $\mathrm{Cov}(w(t)) = \mathrm{diag}(Q_1, \ldots, Q_M) = Q$.

3) STL specification: Let $\mathcal{V} = \{1, \dots, M\}$ be the set collecting the indices of all agents. The MAS is subject to a conjunctive STL formula ϕ with syntax as in (1), where each conjunct is either a local subformula ϕ_i involving agent $i \in \mathcal{V}$, or a joint subformula ϕ_ν involving a subset of agents $\nu \subseteq \mathcal{V}$, where ν is a clique. By collecting all cliques ν in $\mathcal{K}_\phi$, the global STL task is

$$\phi = \bigwedge_{i \in \mathcal{V}} \phi_i \wedge \bigwedge_{\nu \in \mathcal{K}_\phi} \phi_\nu. \tag{4}$$

The structure of ϕ in (4) induces an interaction graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is the set of nodes, and $\mathcal{E} = \{(\nu_i, \nu_j) \mid \nu_i, \nu_j \in \nu,\ i \neq j,\ \nu \in \mathcal{K}_\phi\}$ is the set of edges. Let $\pi := (\mu(y) \geq 0)$ be a predicate in ϕ, where $\mu(y) = a^\intercal y + b$, with $a, y \in \mathbb{R}^{n_y}$ and $b \in \mathbb{R}$. The vector $y \in \mathbb{R}^{n_y}$ represents either an individual state vector, $x_i \in \mathbb{R}^{n_i}$, $i \in \mathbb{N}_{[1,M]}$, or an aggregate vector, $x_\nu \in \mathbb{R}^{n_\nu}$, collecting the states of agents in the clique $\nu \in \mathcal{K}_\phi$. Formula ϕ can specify tasks between subsets of agents, by representing their entirety as cliques.

C. Problem statement

We wish to solve the stochastic optimal control problem

$$\underset{\substack{\boldsymbol{u}(0),\\ \boldsymbol{x}(0)}}{\text{Min.}}\ \mathbb{E}\left[\sum_{i=1}^{M}\Big(\sum_{t=0}^{N-1}(\ell_i(x_i(t), u_i(t))) + V_{f,i}(x_i(N))\Big)\right] \tag{5a}$$

$$\text{s.t. } x(t+1) = Ax(t) + Bu(t) + w(t),\ t \in \mathbb{N}_{[0,N)}, \tag{5b}$$

$$\Pr\{\boldsymbol{x}(0) \models \phi\} \geq \theta, \text{ with } x(0) = x_0, \tag{5c}$$

where $\ell_i : \mathbb{R}^{n_i} \times \mathbb{R}^{m_i} \to \mathbb{R}$, $V_{f,i} : \mathbb{R}^{n_i} \to \mathbb{R}$, the optimization variables are $\boldsymbol{u}(0) = (u(0), \dots, u(N-1))$, $\boldsymbol{x}(0) = (x(0), \dots, x(N))$, ϕ is a multi-agent STL formula, with structure as in (4) and syntax as in (1), to be satisfied by $\boldsymbol{x}(0)$ with a probability $\theta \in (0,1)$, x_0 is a known initial condition of the MAS, and N is the horizon of ϕ. Solving the problem directly is challenging due to the probabilistic constraint, the expectation operator in the cost function, and uncertain dynamics. To handle complexity, especially for a large number of agents and complex ϕ, we relax it with a deterministic problem, which subsequently, we decompose into smaller agent-level subproblems. Additionally, we make the following assumption.

Assumption 1 *For $x(0) = x_0$ and given $\theta \in (0,1)$, Problem* (5) *is feasible.*

III. Main results

A. Error dynamics and construction of probabilistic tubes

Due to the linearity in (2), the state of each agent can be decomposed into a deterministic part, $z_i(t)$, and an error, $e_i(t)$, i.e., $x_i(t) = z_i(t) + e_i(t)$. Consider the causal control law $u_i(t) = K_i e_i(t) + v_i(t)$, where $K_i \in \mathbb{R}^{m_i \times n_i}$ is a stabilizing state-feedback gain for the pair (A_i, B_i). Then,

$$z_i(t+1) = A_i z_i(t) + B_i v_i(t), \tag{6a}$$

$$e_i(t+1) = \bar{A}_i e_i(t) + w_i(t), \tag{6b}$$

where $z_i(0) = x_i(0)$, $e_i(0) = 0$, and $\bar{A}_i = A_i + B_i K_i$. The above choice of $u_i(t)$ will allow us to control the size of the reachable sets of (6b), in light of the probabilistic constraint in (5c). Define now the aggregate vectors $z(t) = (z_1(t), \dots, z_M(t))$, $e(t) = (e_1(t), \dots, e_M(t))$, and $v(t) = (v_1(t), \dots, v_M(t))$, the block-diagonal state-feedback gain $K = \text{diag}(K_1, \dots, K_M) \in \mathbb{R}^{m \times n}$, and the block-diagonal closed-loop matrix $\bar{A} = \text{diag}(\bar{A}_1, \dots, \bar{A}_M)$. Then, we decompose (3) into

$$z(t+1) = Az(t) + Bv(t), \tag{7a}$$

$$e(t+1) = \bar{A}e(t) + w(t). \tag{7b}$$

Given a particular state feedback gain K, the error system (7b) can be analyzed independently of (7a). As a closed-loop system driven by the random vector $w(t)$, we can predict its trajectory $\boldsymbol{e}(0) = (e(0), \dots, e(N))$, with $e(0) = 0$, by calculating its reachable sets probabilistically. Probabilistic reachable sets and tubes for system (7b) are defined next.

Definition 2 *A set $E(t) \subseteq \mathbb{R}^n$, $t \in \mathbb{N}_{[0,N]}$, is called a t-step probabilistic reachable set (t-PRS) for* (7b) *at probability level $\hat{\theta}_t \in [0,1]$, if $\Pr\{e(t) \in E(t) \mid e(0) = 0\} \geq \hat{\theta}_t$.*

It is worth noting that for a probability level $\hat{\theta}_t$ a t-PRS $E(t)$, $t \in \mathbb{N}_{[0,N]}$, for (7b) is not unique.

Definition 3 *Let $\boldsymbol{e}(0) = (e(0), \dots, e(N))$ be a trajectory of* (7b)*. Then, $\boldsymbol{E} \subseteq \mathbb{R}^n \times \cdots \times \mathbb{R}^n$ is called a probabilistic reachable tube (PRT) for* (7b) *at probability level $\Theta \in [0,1]$, if $\Pr\{\boldsymbol{e}(0) \in \boldsymbol{E}\} \geq \Theta$.*

Definition 4 *Let $w \sim \mathcal{D}_w$, with $\text{supp}(\mathcal{D}_w) = \mathcal{W}$. We call $\mathscr{E}_\theta(\mathcal{D}_w) \subseteq \mathcal{W}$ a confidence region (CR) for $w \in \mathcal{W}$ at probability level θ, if $\Pr\{w \in \mathscr{E}_\theta(\mathcal{D}_w)\} \geq \theta$.*

CRs for $w_i(t)$, $i \in \mathbb{N}_{[1,M]}$, can be approximated via Monte Carlo methods or computed analytically using concentration inequalities depending on the properties of $\mathcal{D}_w$. Here, since $\mathbb{E}(w_i(t)) = 0$ and $\text{Cov}(w_i(t)) = Q_i > 0$, for $t \in \mathbb{N}$ and $i \in \mathbb{N}_{[1,M]}$, we construct ellipsoidal CRs at probability level θ_i as $\mathscr{E}_{\theta_i}(\mathcal{D}_{w_i}) = \{w_i \in \mathbb{R}^{n_i} \mid w_i^\intercal Q_i^{-1} w_i \leq n_i/\theta_i\}$, by the multivariate Chebyshev's inequality. Next, we construct a CR for the aggregate random vector $w(t) = (w_1(t), \dots, w_M(t))$.

Lemma 2 *Let $\mathscr{E}_{\theta_i}(\mathcal{D}_{w_i})$ be a CR for $w_i(t) \in \mathcal{W}_i$, $i \in \mathbb{N}_{[1,M]}$, at probability level θ_i, for all $t \in \mathbb{N}$. Then, $\mathscr{E}_{\hat{\theta}}(\mathcal{D}_w) = \mathscr{E}_{\theta_1}(\mathcal{D}_{w_1}) \times \cdots \times \mathscr{E}_{\theta_M}(\mathcal{D}_{w_M})$, is a CR of $w(t) \in \mathcal{W}$ at probability level $\hat{\theta}$ for all $t \in \mathbb{N}$, where $\hat{\theta} \geq \Pi_{i=1}^M \theta_i$.*

Proof: Without loss of generality let $M = 2$. Then, $\Pr\{w(t) \in \mathscr{E}_{\hat{\theta}}(\mathcal{D}_w)\} \geq \Pr\{(w_1(t) \in \mathscr{E}_{\theta_1}(\mathcal{D}_{w_1})) \wedge (w_2(t) \in \mathscr{E}_{\theta_2}(\mathcal{D}_{w_2}))\} = \Pr\{w_1(t) \in \mathscr{E}_{\theta_1}(\mathcal{D}_{w_1})\}\Pr\{w_2(t) \in \mathscr{E}_{\theta_2}(\mathcal{D}_{w_2})\} \geq \theta_1\theta_2$, which is true due to independence of $w_1(t)$, $w_2(t)$, for all $t \in \mathbb{N}$. ■

Based on the CR construction of the disturbance $w(t)$, we construct t-PRSs at certain probability levels for the multi-agent error system (7b) as follows.

Proposition 1 *Let $\mathscr{E}_{\hat{\theta}}(\mathcal{D}_w) = \mathscr{E}_{\theta_1}(\mathcal{D}_{w_1}) \times \cdots \times \mathscr{E}_{\theta_M}(\mathcal{D}_{w_M})$ be a CR for $w(t)$, where $\mathscr{E}_{\theta_i}(\mathcal{D}_{w_i})$ is an ellipsoidal CR for $w_i(t)$ at probability level θ_i, $i \in \mathbb{N}_{[1,M]}$. Then, the sets $E(t) \subseteq \mathbb{R}^n$, $t \in \mathbb{N}_{[0,N]}$, which are recursively defined as $E(t+1) = \bar{A}E(t) \oplus \mathscr{E}_{\hat{\theta}}(\mathcal{D}_w)$, with $E(0) = \{0\} \times \cdots \times \{0\}$, are t-PRSs for* (7b) *at probability level $\hat{\theta} \geq \Pi_{i=1}^{M}\theta_i$, and $E(t) = E_1(t) \times \cdots \times E_M(t)$, $t \in \mathbb{N}_{[0,N]}$, where $E_i(t)$ is a t-PRS for* (6b) *at probability level θ_i, with $i \in \mathbb{N}_{[1,M]}$.*

Proof: Since $E(0) = \{0\} \times \cdots \times \{0\}$, we may write $E(0) = E_1(0)\times\cdots\times E_M(0)$, with $E_i(0) = \{0\}$, $i \in \mathbb{N}_{[1,M]}$, from which we compute $E(1) = \bar{A}E(0) \oplus \mathscr{E}_{\hat{\theta}}(\mathcal{D}_w) = (\text{diag}(\bar{A}_1, \ldots, \bar{A}_M)E_1(0)\times\cdots\times E_M(0))\oplus(\mathscr{E}_{\theta_1}(\mathcal{D}_{w_1})\times\cdots\times \mathscr{E}_{\theta_M}(\mathcal{D}_{w_M})) = (\bar{A}_1E_1(0)\times\cdots\times\bar{A}_ME_M(0))\oplus(\mathscr{E}_{\theta_1}(\mathcal{D}_{w_1})\times \cdots\times\mathscr{E}_{\theta_M}(\mathcal{D}_{w_M}))$, which from Lemma 1 results in $E(1) = \bar{A}_1E_1(0)\oplus\mathscr{E}_{\theta_1}(\mathcal{D}_{w_1}) \times \cdots \times \bar{A}_ME_M(0)\oplus\mathscr{E}_{\theta_M}(\mathcal{D}_{w_M})$, that is, $E(1) = E_1(1) \times \cdots \times E_M(1)$, with $E_i(1) = \bar{A}_iE_i(0) \oplus \mathscr{E}_{\theta_i}(\mathcal{D}_{w_i})$, $i \in \mathbb{N}_{[1,M]}$. Following the recursion, one can show that $E_i(t+1) = \bar{A}_iE_i(t)\oplus\mathscr{E}_{\theta_i}(\mathcal{D}_{w_i})$, for $t \in \mathbb{N}_{[0,N-1]}$, and $E(t) = E_1(t) \times \cdots \times E_M(t)$, for $t \in \mathbb{N}_{[0,N]}$.

Let now $\mathcal{D}_{e_i(t)}$ be the distribution of $e_i(t)$, and $\mathscr{E}_{\theta_i}(\mathcal{D}_{e_i(t)})$ be a CR for $e_i(t)$ at probability θ_i. Since, $\mathscr{E}_{\theta_i}(\mathcal{D}_{e_i(0)}) \subseteq E_i(0) = \{0\}$, we have $\bar{A}_i\mathscr{E}_{\theta_i}(\mathcal{D}_{e(0)}) \oplus \mathscr{E}_{\theta_i}(\mathcal{D}_{w_i}) \subseteq \bar{A}_iE_i(0) \oplus \mathscr{E}_{\theta_i}(\mathcal{D}_{w_i}) = E_i(1)$, so $\mathscr{E}_{\theta_i}(\mathcal{D}_{e_i(1)}) \subseteq E_i(1)$, as $\mathscr{E}_{\theta_i}(\mathcal{D}_{e_i(t+1)}) \subseteq \bar{A}_i\mathscr{E}_{\theta_i}(\mathcal{D}_{e_i(t)}) \oplus \mathscr{E}_{\theta_i}(\mathcal{D}_{w_i})$ for all $t \in \mathbb{N}$, since $\mathscr{E}_{\theta_i}(\mathcal{D}_{w_i})$ is an ellipsoidal region by [18, Cor. 4]. Inductively we show that $\mathscr{E}_{\theta_i}(\mathcal{D}_{e_i(t)}) \subseteq E_i(t)$, $i \in \mathbb{N}_{[1,M]}$, for all $t \in \mathbb{N}$. The latter implies that $\Pr\{e_i(t) \in E_i(t)\} \geq \theta_i$, $i \in \mathbb{N}_{[1,M]}$, from which we have $\Pr\{e(t) \in E(t)\} = \Pr\{(e_1(t) \in E_1(t))\wedge\cdots\wedge(e_M(t) \in E_M(t))\} = \Pr\{(e_1(t) \in E_1(t))\}\cdots\Pr\{(e_M(t) \in E_M(t))\} \geq \theta_1\theta_2\cdots\theta_M$, which follows from the independence of $E_i(t)$, $i \in \mathbb{N}_{[1,M]}$. ■

Prop. 1 leads to the following PRT result.

Theorem 1 *Let $\mathscr{E}_{\hat{\theta}}(\mathcal{D}_w) = \mathscr{E}_{\theta_1}(\mathcal{D}_{w_1}) \times \cdots \times \mathscr{E}_{\theta_M}(\mathcal{D}_{w_M})$ be a CR for $w(t)$, and define t-PRSs, $E(t)$, $t \in \mathbb{N}_{[0,N]}$ for system* (7b) *at probability level $\hat{\theta} \geq \Pi_{i=1}^{M}\theta_i$, where θ_i is the confidence level of the region $\mathscr{E}_{\theta_i}(\mathcal{D}_{w_i})$, $i \in \mathbb{N}_{[1,M]}$, as in Prop. 1. Then, i) $\boldsymbol{E}_i = E_i(0) \times \cdots \times E_i(N)$ is a PRT for* (6b) *at probability level $\Theta_i \geq 1 - N(1 - \theta_i)$, where $E_i(t)$, $t \in \mathbb{N}_{[0,N]}$, is a t-PRS for* (6b) *at probability level θ_i. ii) $\boldsymbol{E} = E(0) \times \cdots \times E(N)$ is a PRT for* (7b) *at probability level $\Theta = \Pi_{i=1}^{M}\Theta_i$. iii) Let $e_\nu(t + 1) = \bar{A}_\nu e_\nu(t) + w_\nu(t)$ be the aggregate system collecting individual error systems from the clique $\nu \in \mathcal{K}_\phi$, where $\nu = (i_1, \ldots, i_{|\nu|})$, $e_\nu(t) = (e_{i_1}(t), \ldots, e_{i_{|\nu|}}(t))$, $w_\nu(t) = (w_{i_1}(t), \ldots, w_{i_{|\nu|}}(t))$, and $\bar{A}_\nu = \text{diag}(\bar{A}_{i_1}, \ldots, \bar{A}_{i_{|\nu|}})$, and let $E_\nu(t) = E_{i_1}(t) \times \cdots \times E_{i_{|\nu|}}(t)$, $t \in \mathbb{N}_{[0,N]}$, be its t-PRSs, with $E_{i_j}(t)$, being t-PRS for* (6b) *at probability level θ_{i_j}, with $j \in \mathbb{N}_{[1,|\nu|]}$. Then, $\boldsymbol{E}_\nu = E_\nu(0) \times \cdots \times E_\nu(N)$ is a PRT at probability level $\Theta = \Pi_{j=1}^{|\nu|}\Theta_{i_j}$.*

Proof: i) From Prop. 1, we have that $E(t) = E_1(t)\times\cdots\times E_M(t)$, where $E_i(t)$ is a t-PRS for (6b) at probability level θ_i. Let $\boldsymbol{e}_i(0) = (e_i(0), \ldots, e_i(N))$ be a trajectory of (6b). Then, $\Pr\{\boldsymbol{e}_i(0) \in \boldsymbol{E}_i\} = \Pr\{(e_i(0) \in E_i(0)) \wedge \cdots \wedge (e_i(N) \in E_i(N))\} = 1 - \Pr\{(e_i(0) \notin E_i(0)) \vee \cdots \vee (e_i(N) \notin E_i(N))\} \geq 1 - \sum_{t=0}^{N}\Pr\{e_i(t) \notin E_i(t)\} = 1 - N(1 - \theta_i)$, where we use Boole's inequality, that $E_i(t)$, $t \in \mathbb{N}_{[0,N]}$, is a t-PRS at probability level θ_i, and $\Pr(e_i(0) \notin E_i(0)\} = 0$. ii) It holds that $\Pr\{\boldsymbol{e}(0) \in \boldsymbol{E}\} = \Pr\{((e_1(0) \in E_1(0)) \wedge \cdots \wedge (e_M(0) \in E_M(0))) \wedge \cdots \wedge ((e_1(N) \in E_1(N)) \wedge \cdots \wedge (e_M(N) \in E_M(N)))\} = \Pr\{((e_1(0) \in E_1(0)) \wedge \cdots \wedge (e_1(N) \in E_1(N))) \wedge \cdots \wedge ((e_M(0) \in E_M(0)) \wedge \cdots \wedge (e_M(N) \in E_M(N)))\}$, which is $\Pr\{(\boldsymbol{e}_1(0) \in \boldsymbol{E}_1) \wedge \cdots \wedge (\boldsymbol{e}_M(0) \in \boldsymbol{E}_M)\} = \Pr\{\boldsymbol{e}_1(0) \in \boldsymbol{E}_1\}\cdots\Pr\{\boldsymbol{e}_M(0) \in \boldsymbol{E}_M\} = \Pi_{i=1}^{M}\Theta_i$, by the independence of the PRTs $\boldsymbol{E}_i$, $i \in \mathbb{N}_{[1,M]}$. iii) By setting $M = |\nu|$ the result follows from Prop. 1 and item ii) herein. ■

Remark 1 *Note that our PRT construction reduces conservatism for a large number of agents, while utilizing the union-bound argument only over time. This may require conservative choices for the probability levels, θ_i, for the CRs of $w_i(t)$, $i \in \mathbb{N}_{[1,M]}$, for large horizons. To construct, e.g., a PRT for* (7b) *at probability level Θ, one may select uniform probability levels for the CRs as $\theta_i \geq 1 - \frac{1-\Theta^{\frac{1}{M}}}{N}$, where $\theta_i \to 1$ for large N, regardless of Θ.*

B. Constraint tightening

We aim to design a trajectory for the deterministic system (7a) that satisfies an STL formula derived from ϕ, incorporating *tighter* predicates. The following proposition underpins this approach.

Proposition 2 *Let $\boldsymbol{x}(0) = \boldsymbol{z}(0) + \boldsymbol{e}(0)$, with $\boldsymbol{x}(0) = (x(0), \ldots, x(N))$, $\boldsymbol{z}(0) = (z(0), \ldots, z(N))$ and $\boldsymbol{e}(0) = (e(0), \ldots, e(N))$. Suppose that $\Pr\{\boldsymbol{e}(0) \in \boldsymbol{E}\} \geq \theta$, for some $\boldsymbol{E} = E(0) \times \cdots \times E(N)$, with $E(t) \subseteq \mathbb{R}^n$, $t \in \mathbb{N}$. If $\boldsymbol{z}(0) + \boldsymbol{e}(0) \models \phi$ for all $\boldsymbol{e}(0) \in \boldsymbol{E}$, then $\Pr\{\boldsymbol{x}(0) \models \phi\} \geq \theta$.*

Proof: Define events $Y_x := \boldsymbol{x}(0) \models \phi$, $Y_e := \boldsymbol{e}(0) \in \boldsymbol{E}$, and $Y_e' := \boldsymbol{e}(0) \notin \boldsymbol{E}$. From the law of total probability, we have $\Pr\{Y_x\} = \Pr\{Y_x|Y_e\}\Pr\{Y_e\} + \Pr\{Y_x|Y_e'\}\Pr\{Y_e'\} \geq \theta$, since by assumption, $\Pr\{Y_x|Y_e\} = 1$ and $\Pr\{Y_e\} \geq \theta$, and $\Pr\{Y_x|Y_e'\}\Pr\{Y_e'\} \geq 0$. ■

Let $\boldsymbol{E}$ be a PRT for (7b) at probability level Θ. Next, we construct a formula ψ such that $\boldsymbol{z}(0) \models \psi$ implies that $\boldsymbol{z}(0)+\boldsymbol{e}(0) \models \phi$, for all $\boldsymbol{e}(0) \in \boldsymbol{E}$, that is, $\Pr\{\boldsymbol{x}(0) \models \phi\} \geq \theta$, by Prop. 2. Formula ψ has identical Boolean and temporal operators with ϕ in (4), and retains its multi-agent structure:

$$\psi = \bigwedge_{i\in\mathcal{V}} \psi_i \wedge \bigwedge_{\nu\in\mathcal{K}_\phi} \psi_\nu. \tag{8}$$

Let $\pi := (a^\intercal y + b \geq 0)$ be a predicate, with $a, y \in \mathbb{R}^{n_y}$, $b \in \mathbb{R}$. We denote by $\tau(\pi)$ the *tighter* version of π, where $\tau(\pi) \in \psi$ if $\pi \in \phi$, and $\neg\tau(\pi) \in \psi$ if $\neg\pi \in \phi$, with

$$\tau(\pi) := (a^\intercal y + b + \min_{g\in G} a^\intercal g \geq 0), \text{ if } \tau(\pi) \in \psi, \tag{9a}$$

$$\tau(\pi) := (a^\intercal y + b + \max_{g\in G} a^\intercal g \geq 0), \text{if } \neg\tau(\pi) \in \psi. \tag{9b}$$

Here, $G = \bigcup_{t=1}^{N} E_i(t)$ if $y = x_i(t)$, for some $i \in \mathcal{V}$, or $G = \bigcup_{t=1}^{N} E_\nu(t)$ if $y = x_\nu(t)$ for some $\nu \in \mathcal{K}_\phi$. We remark that the optimizations in (9) are tractable since the domain G is the union of finitely many convex and compact sets by the construction of t-PRSs, $E_i(t)$, $E_\nu(t)$, $t \in \mathbb{N}_{[1,N]}$, in Prop. 1 and Thm. 1. Practically, the tightening in (9) can be retrieved by solving a convex optimization by taking the convex hull of G, or, better, by solving N convex optimization problems, one for every $E_i(t)$, $E_\nu(t)$, $t \in \mathbb{N}_{[1,N]}$, and selecting the worst-case (minimum for (9a) and maximum for (9b)) solution among them. We are now ready to state the following result.

Theorem 2 *Let ψ be the STL formula resulting from ϕ according to* (8)-(9)*, and assume that the deterministic problem:*

$$\underset{\substack{\boldsymbol{v}(0),\\ \boldsymbol{z}(0)}}{\text{Minim.}} \sum_{i=1}^{M} \left(\sum_{t=0}^{N-1} (\ell_i(z_i(t), v_i(t))) + V_{f,i}(z_i(N)) \right) \quad (10a)$$

$$\text{s.t. } z(t+1) = Az(t) + Bv(t),\ t \in \mathbb{N}_{[0,N)}, \quad (10b)$$

$$\boldsymbol{z}(0) \models \psi, \text{ with } z(0) = x_0, \quad (10c)$$

has a feasible solution $\boldsymbol{v}(0) = (v(0), \ldots, v(N-1))$, with $v(t) \in \mathcal{U} \ominus KE(t)$, $t \in \mathbb{N}_{[0,N-1]}$, where K is a stabilising gain for (A, B) in (7b)*, and $E(t) = E_1(t) \times \cdots \times E_M(t)$, $E_i(t)$, $i \in \mathbb{N}_{[1,M]}$, being t-PRS for* (6b)*, at probability level θ_i, such that $\Pi_{i=1}^{M} (1 - N(1-\theta_i)) \geq \theta$. Let $\boldsymbol{e}(0) = (e(0), \ldots, e(N))$ be a trajectory of* (7b)*. Then, $\boldsymbol{u}(0) = \text{diag}(K, \ldots, K)\boldsymbol{e}(0) + \boldsymbol{v}(0)$ is a feasible solution for* (5)*.*

Proof: By Thm. 1, $\boldsymbol{E} = E(0) \times \cdots \times E(N)$ is a PRT for (7b) at probability level $\Theta \geq \theta$, that is, $\Pr\{\boldsymbol{e}(0) \in \boldsymbol{E}\} \geq \theta$. Let $\boldsymbol{z}(0) \models \psi$ be a trajectory resulting from the input trajectory $\boldsymbol{v}(0)$ starting from $z(0) = x_0$. By the tightening in (9), we have that for all $\pi \in \phi$ and $\tau(\pi) \in \psi$, if $\boldsymbol{z}(0) \models \tau(\pi)$, then $\boldsymbol{z}(0) + \boldsymbol{e}(0) \models \pi\ \forall \boldsymbol{e}(0) \in \boldsymbol{E}$, and for all $\neg\pi' \in \phi$ and $\neg\tau(\pi') \in \psi$, if $\boldsymbol{z}(0) \models \neg\tau(\pi')$, then $\boldsymbol{z}(0) + \boldsymbol{e}(0) \models \neg\pi'\ \forall \boldsymbol{e}(0) \in \boldsymbol{E}$. Since ϕ and ψ differ only in predicates, it follows that if $\boldsymbol{z}(0) \models \psi$, then $\boldsymbol{z}(0) + \boldsymbol{e}(0) \models \phi\ \forall \boldsymbol{e}(0) \in \boldsymbol{E}$. Since $\boldsymbol{u}(0) = \text{diag}(K, \ldots, K)\boldsymbol{e}(0) + \boldsymbol{v}(0)$ is a feasible input trajectory for (3) $\forall \boldsymbol{e}(0) \in \boldsymbol{E}$, the resulting state trajectory of (3), $\boldsymbol{x}(0) = \boldsymbol{z}(0) + \boldsymbol{e}(0)$, ensures that $\boldsymbol{x}(0) \models \phi\ \forall \boldsymbol{e}(0) \in \boldsymbol{E}$. The result follows by Prop. 2 since $\Pr\{\boldsymbol{e}(0) \in \boldsymbol{E}\} \geq \theta$. ∎

The gain K affects the feasible domain of (10) and the volume of $\boldsymbol{E}$. Its construction will be addressed in future work. By selecting $\|\cdot\|_1$-based costs, problem (10) can be formulated as a mixed-integer linear program (MILP) [5]. Next, we decompose (10) into individual agent-level problems to address its complexity.

C. Distributed control synthesis

We propose an iterative procedure that handles the complexity of (10). First, we assume the following.

Assumption 2 *The optimization* (10) *has a feasible solution $\boldsymbol{v}(0) = (v(0), \ldots, v(N-1))$, $\boldsymbol{z}(0) = (z(0), \ldots, z(N))$, where $v(t) \in \mathcal{U} \ominus KE(t)$, $t \in \mathbb{N}_{[0,N)}$, with the gain K and the t-PRSs $E(t)$, $t \in \mathbb{N}_{[0,N]}$, being as in Thm. 2.*

1) Decomposition of STL formula ψ: For a node i participating in at least one clique, i.e., $i \in \nu$, with $\nu \in \mathcal{K}_\phi$, we define $\mathcal{T}_i$ by the set of cliques containing i excluding i, i.e.,

$$\mathcal{T}_i = \{\nu \setminus i : \nu \in \text{cl}(i)\}, \quad (11)$$

where $\text{cl}(i) = \{\nu \in \mathcal{K}_\phi,\ \nu \ni i\}$ is the set of cliques that contain i. Let $j \in \mathcal{T}_i$, with $j = (i_1, \ldots, i_{|j|})$. Let a trajectory $\boldsymbol{z}_{ij}(0) = (z_{ij}(0), \ldots, z_{ij}(N))$, where $z_{ij}(t) = (z_{i_1}(t), \ldots, z_i(t), \ldots, z_{i_{|j|}}(t))$, with $t \in \mathbb{N}_{[0,N]}$, and the order $i_1 < \ldots < i < \ldots < i_{|j|}$ being specified by the lexicographic ordering of the node set $\mathcal{V} = \mathbb{N}_{[1,M]}$. Using (11), an equivalent formula to the *tighter* formula (8)-(9), ψ, is defined as $\hat{\psi} = \bigwedge_{i \in \mathcal{V}} \hat{\psi}_i$, where

$$\hat{\psi}_i = \psi_i \wedge \bigwedge_{j \in \mathcal{T}_i} \psi_{ij}. \quad (12)$$

2) Iterative algorithm: For simplicity, we drop the time argument and introduce an iteration index as a superscript in the trajectory notation, e.g., $\boldsymbol{z}_i^k$ ($\boldsymbol{z}_{ij}^k$) indicates a trajectory $\boldsymbol{z}_i(0)$ ($\boldsymbol{z}_{ij}(0)$) that is retrieved at the kth iteration of the following procedure. To initialize the procedure, we generate initial guesses on the agents' trajectories by solving

$$\underset{\boldsymbol{v}_i^0, \boldsymbol{z}_i^0}{\text{Minimize}} \sum_{t=0}^{N-1} (\ell_i(z_i^0(t), v_i^0(t))) + V_{f,i}(z_i^0(N)) \quad (13a)$$

$$\text{s.t. } z_i^0(t+1) = A_i z_i^0(t) + B_i v_i^0(t),\ t \in \mathbb{N}_{[0,N)}, \quad (13b)$$

$$\boldsymbol{z}_i^0 \models \psi_i, \text{ with } z_i^0(0) = x_{0,i}, \quad (13c)$$

at $k = 0$ for $i \in \mathbb{N}_{[1,M]}$. After solving problem (13), which is feasible by Assumption 2, at iteration $k \geq 1$, only a subset of agents, denoted by $\mathcal{O}_k \subset \mathcal{V}$, are allowed to update their input sequences by performing an optimization. The remaining agents retrieve their input sequences from the previous iteration $k-1 \geq 0$. Roughly, the set $\mathcal{O}_k \subset \mathcal{V}$ is constructed so that any combination of its elements does not belong to a clique $\nu \in \mathcal{K}_\phi$. Due to space limitations, we simply construct $\mathcal{O}_k$ as a singleton, which only affects the number of agents' trajectories that can be optimized in parallel per iteration. For the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, with $\mathcal{V} = \mathbb{N}_{[1,M]}$, $\mathcal{O}_k = \text{mod}(\mathcal{O}_{k-1}, M) + 1$, for $k > 1$, with $\mathcal{O}_1 = 1$. We refer readers to [20, Sec. V.B] for a more efficient construction of $\mathcal{O}_k$ enabling parallel computations at each iteration (see Sec. IV for a numerical example).

At the kth iteration, with $k \geq 1$, if $i \notin \mathcal{O}_k$, then, $\boldsymbol{z}_i^k = \boldsymbol{z}_i^{k-1}$ and $\boldsymbol{v}_i^k = \boldsymbol{v}_i^{k-1}$. Otherwise, the input sequence of the ith agent is updated by solving

$$\underset{\boldsymbol{v}_i^k, \boldsymbol{z}_i^k}{\text{Minim.}} \sum_{t=0}^{N-1} (\ell_i(z_i^k(t), v_i^k(t))) + V_{f,i}(z_i^k(N)) - \mu_{ij_k}^k \quad (14a)$$

$$\text{s.t. } z_i^k(t+1) = A_i z_i^k(t) + B_i v_i^k(t),\ t \in \mathbb{N}_{[0,N)}, \quad (14b)$$

$$\boldsymbol{z}_i^k \models \psi_i, \text{ with } z_i(0) = x_{0,i}, \quad (14c)$$

$$\rho^{\psi_{ij_k}}(\boldsymbol{z}_{ij_k}^k) \geq \mu_{ij_k}^k,\ j_k = \underset{j \in \mathcal{T}_i}{\text{argmin}}\{\rho^{\psi_{ij}}(\boldsymbol{z}_{ij}^{k-1})\}, \quad (14d)$$

$$\mu_{ij_k}^k \geq \min\left(0, \rho^{\psi_{ij_k}}(\boldsymbol{z}_{ij_k}^{k-1})\right), \quad (14e)$$

$$\rho^{\psi_{ij}}(\boldsymbol{z}_{ij}^k) \geq \min\left(0, \rho^{\psi_{ij}}(\boldsymbol{z}_{ij}^{k-1})\right),\ \forall j \in \mathcal{T}_i \setminus j_k, \quad (14f)$$

where $\rho^{\psi_{ij}}(\boldsymbol{z}_{ij}^k)$ is the robustness function of the formula ψ_{ij} evaluated over the trajectory $\boldsymbol{z}_{ij}^k$. Agent-i, with $i \in \mathcal{O}_k$, by solving (14), retrieves an input sequence that guarantees 1) the satisfaction of the individual task ψ_i (see constraint (14c)), 2) the improvement of the most violating (or least robust) joint task ψ_{ij_k} (see constraints (14d)-(14e)), and 3) either improvement on or non-violation of the remaining joint tasks (see constraint (14f)). The inclusion of the $\min$ operator in the constraints (14e)-(14f) relaxes the satisfaction of joint tasks that have already been found to be satisfiable in previous iterations. This allows the algorithm to emphasize the satisfaction of joint tasks with the smallest robustness function. The algorithm may terminate if it exceeds a maximum number of iterations, denoted as $k_{\max}$ and defined by the designer, yielding a minimally violating solution. Alternatively, termination occurs when verifying the satisfiability of all joint tasks, i.e., when $\mu_{ij}^k \geq 0$ for all $i \in \mathcal{V}$, $j \in \mathcal{T}_i$, and some $k \leq k_{\max}$, returning a feasible solution to (10). The overall iterative procedure is summarized in Alg. 1, the integrity of which relies on the following result.

Theorem 3 *At each iteration $k \geq 1$, the optimization problem* (14) *is feasible for all $i \in \mathcal{O}_k$.*

Proof: Let $k = 1$. The lower bounds in (14d)-(14f) are defined over trajectories, $\boldsymbol{z}_i^0$, $i {\in} \mathcal{O}_1$, obtained by solving (13) at $k{=}0$. Thus $\boldsymbol{z}_i^0$ satisfies the constraints in (14d)-(14f) for all $i {\in} \mathcal{O}_1 \subset \mathcal{V}$. Moreover, the constraint in (14c) is satisfied by $\boldsymbol{z}_i^0$, since (13) is feasible. Hence, $\boldsymbol{u}_i^1 = \boldsymbol{u}_i^0$ is a feasible solution of (14) at iteration $k = 1$. Now, let $k{>}1$. The lower bounds in (14d)-(14f) are defined over trajectories, $\boldsymbol{z}_i^{k-1}$, obtained by the solutions $\boldsymbol{u}_i^{k-1}$, $i {\in} \mathcal{O}_k$, at iteration $k-1$. Thus, $\boldsymbol{z}_i^{k-1}$, $i {\in} \mathcal{O}_k$, satisfy the constraints in (14d)-(14f). Additionally, the constraint in (14c) is satisfied by $\boldsymbol{z}_i^{k-1}$, $i \in \mathcal{O}_k$, since it is retrieved by $\boldsymbol{u}_i^{k-1}$, which is obtained by solving (14) or (13) at some iteration $\kappa \leq k-1$. Thus, $\boldsymbol{u}_i^k {=} \boldsymbol{u}_i^{k-1}$ is a feasible solution of (14) for all $k \geq 1$. ∎

Algorithm 1 Iterative procedure for solving (10)

1: Compute $\mathcal{T}_i$ (11) and construct $\hat{\psi}_i$ (12), for $i \in \mathbb{N}_{[1,M]}$
2: Solve (13) and store $\boldsymbol{v}_i^0$, $\boldsymbol{z}_i^0$, for $i \in \mathbb{N}_{[1,M]}$
3: Construct $\mathcal{O}_k$, for $k \in \mathbb{N}_{[1,k_{\max}]}$
4: **for** k in $1:k_{\max}$ **do**
5:     **for** i in $1:M$ **do**
6:         **if** $i \in \mathcal{O}_k$, solve (14), and store $(\boldsymbol{v}_i^k, \boldsymbol{z}_i^k)$
7:         **if** $i \notin \mathcal{O}_k$, update $\boldsymbol{v}_i^k \leftarrow \boldsymbol{v}_i^{k-1}$ and $\boldsymbol{z}_i^k \leftarrow \boldsymbol{z}_i^{k-1}$
8:     Construct $(\boldsymbol{v}(0), \boldsymbol{z}(0))$ from $(\boldsymbol{v}_i^k, \boldsymbol{z}_i^k)$, $i \in \mathbb{N}_{[1,M]}$
9:     **if** $\rho^{\psi}(\boldsymbol{z}(0)) \geq 0$ **go to** 10
10: **return** $(\boldsymbol{v}(0), \boldsymbol{z}(0))$

IV. EXAMPLE

We consider ten agents with aggregate dynamics given by $x(t+1) = x(t) + u(t) + w(t)$, where $x(t) = (x_1(t), ..., x_{10}(t)) \in \mathbb{R}^{20}$, $u(t) = (u_1(t), ..., u_{10}(t)) \in \mathbb{R}^{20}$, and $w(t) = (w_1(t), ..., w_{10}(t)) \in \mathbb{R}^{20}$. States $x_i(t) \in \mathcal{X}$, where $\mathcal{X}$ is the workspace confined by the dashed border in Fig. 2. Individual inputs are constrained by $\|u_i(t)\|_\infty \leq 0.8$, and disturbances, $w_i(t)$, are Gaussian random vectors, independent time- and agent-wise, with zero mean and covariance, $Q_i = 0.05 I_2$, for all $t \in \mathbb{N}$ and $i \in \mathbb{N}_{[1,10]}$. The MAS is assigned a specification $\phi = \bigwedge_{i=1}^{10} \phi_i \wedge \bigwedge_{\nu \in \mathcal{K}_\phi} \phi_\nu$, with horizon $N = 100$, where $\mathcal{K}_\phi$ is the set of cliques shown in Fig. 1, and ϕ_i, ϕ_ν, are tasks assigned to agent i, and the agents in $\nu \in \mathcal{K}_\phi$, respectively. In (5), we select $\ell_i(x_i(t), u_i(t)) = \|u_i(t)\|_1$, $V_{f,i}(x_i(100)) = 0$, for all $i \in \mathbb{N}_{[1,10]}$, and set $\theta = 0.70$.

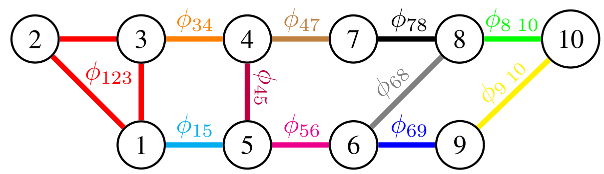

Fig. 1: The ten agents of the MAS and the cliques $\mathcal{K}_\phi$ in ϕ.

Let $\phi_i = \left(\square_{[0,100]}(\varphi_i^{\mathcal{X}} \wedge \neg\varphi_i^{O_1} \wedge \neg\varphi_i^{O_2} \wedge \neg\varphi_i^{O_3})\right) \wedge \left(\Diamond_{[10,50]}\varphi_i^{T_i}\right) \wedge \left(\Diamond_{[70,100]}\varphi_i^{G_i}\right)$ be an individual task, which requires agent-i, starting from $x_i(0)$ to pass through T_i and G_i within the intervals $\mathbb{N}_{[10,50]}$ and $\mathbb{N}_{[70,100]}$, respectively, while always staying within $\mathcal{X}$ and avoiding O_1, O_2, O_3. Regions T_i, G_i, $i \in \mathbb{N}_{[1,10]}$, and obstacles O_1, O_2, O_3, are in Fig. 2. Note that $x_i(t) \models \varphi_i^{\mathcal{Y}}$ if $x_i(t) \in \mathcal{Y}$, $\mathcal{Y} = \{\mathcal{X}, O_1, O_2, O_3, T_1, \ldots, T_{10}, G_1, \ldots, G_{10}\}$, for $t \in \mathbb{N}_{[0,100]}$.

Let $\phi_\nu = \Diamond_{[0,100]}\left(\|C_\nu x_\nu(t)\|_\infty \leq 1\right)$, where $C_\nu = \begin{bmatrix} I & -I \end{bmatrix}$ if $|\nu| = 2$ or $C_\nu = \begin{bmatrix} I & -I & 0 \\ 0 & I & -I \\ I & 0 & -I \end{bmatrix}$ if $|\nu| = 3$, be a joint task requiring agents in $\nu \in \mathcal{K}_\phi$ (see Fig. 1) to approach one another at least once within the horizon.

To formulate the deterministic problem (10), we proceed as follows: First, we select closed-loop matrices $\bar{A}_i = I_2 + K_i$, with $K_i = -0.5 I_2$, $i \in \mathbb{N}_{[1,10]}$, and construct t-PRSs, $E_i(t)$, $t \in \mathbb{N}_{[0,100]}$, for (6b), by the recursion $E_i(t+1) = \bar{A}_i E_i(t) \oplus \mathscr{E}_{\theta_i}(\mathcal{D}_{w_i})$, with $E_i(0) = \{0\}$, at probability levels $\theta_i = 1 - \frac{1-0.7^{\frac{1}{10}}}{100} = 0.9996$, $i \in \mathbb{N}_{[1,10]}$, such that $\Pi_{i=1}^{10}(1 - 100(1-\theta_i)) \geq 0.7$, and $\mathscr{E}_{\theta_i}(\mathcal{D}_{w_i}) = \{w_i | w_i^\intercal Q_i^{-1} w_i \leq \chi_2^2(\theta_i)\}$, where χ_2^2 is the chi-squared distribution of degree 2. Last, given that $\boldsymbol{E} = E(0) \times \cdots \times E(100)$, with $E(t) = E_1(t) \times \cdots \times E_{10}(t)$, $t \in \mathbb{N}_{[0,100]}$, is a PRT for (7b) at probability level $\Theta = 0.7$, by Thm. 1, we perform the optimizations in (9), derive the *tighter* formula ψ as in (8), and formulate (10) as an MILP. We have attempted to solve (10) in a centralized manner using the GUROBI solver [21], which produces a solution after running for 2.5 hours, with its feasibility iteration limit set to a maximum of ten. To obtain a solution faster, we first decompose ψ according to (12), based on the sets $\mathcal{T}_1 = \{(2,3), 5\}$, $\mathcal{T}_2 = \{(1,3)\}$, $\mathcal{T}_3 = \{(1,2), 4\}$, $\mathcal{T}_4 = \{3,5\}$, $\mathcal{T}_5 = \{1,4,6\}$, $\mathcal{T}_6 = \{5,8,9\}$, $\mathcal{T}_7 = \{4,8\}$, $\mathcal{T}_8 = \{6,7,10\}$, $\mathcal{T}_9 = \{6,10\}$, and $\mathcal{T}_{10} = \{8,9\}$. By selecting sets $\mathcal{O}_k$, $k \geq 1$, as $\mathcal{O}_1 = \{1,4,6,10\}$, $\mathcal{O}_2 = \{3,5,7,9\}$, $\mathcal{O}_3 = \{8,9,2,5\}$, $\mathcal{O}_4 = \mathcal{O}_1$, $\mathcal{O}_5 = \mathcal{O}_2$, $\mathcal{O}_6 = \mathcal{O}_3$, $\mathcal{O}_7 = \mathcal{O}_1$, and so on, we run Alg. 1, which terminates in less than six minutes returning a multi-agent trajectory, illustrated in Fig. 2, that satisfies the global STL task ψ. Fig. 3 shows the computational overhead of Alg. 1 vs. the centralized solution for varying agent numbers, using a log scale to highlight the different runtime magnitudes. Note that the runtime of Alg. 1 can further be improved if agent-level subproblems, (13), (14), can be solved in parallel. By evaluating the robustness function of ϕ for numerous noisy realizations, we see that ϕ is violated in less than 30% of the time, verifying Thm. 2.

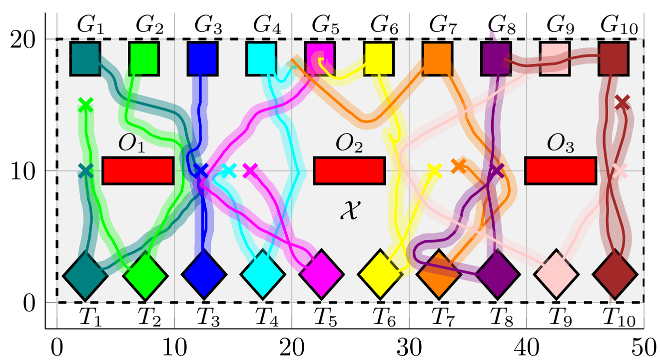

Fig. 2: Nominal trajectories (solid lines) by Alg. 1, with initial states marked by crosses. Tubes (transparent covers) around trajectories, at probability levels $\Theta_i = 0.965$.

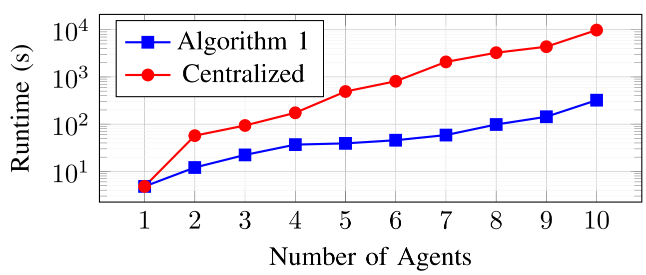

Fig. 3: Compute times (log scale) for solving (10) via Alg. 1 and a centralized approach for varying agent numbers.

V. Conclusion

We have considered stochastic linear multi-agent systems under STL specifications formulated probabilistically. Leveraging linearity, we construct a PRT at the specification probability level and relax the underlying stochastic control problem with a deterministic one with tighter constraints. Our PRT-based tightening reduces conservatism compared to approaches relying on the STL specification structure. To enhance scalability, we propose an algorithm, where the multi-agent problem is decomposed into agent-level subproblems that can be solved iteratively. Although our method fits large-scale MAS settings, the conservatism introduced by the construction of PRTs increases with the specification horizon. Future work will address this via efficient, data-driven approaches, avoiding union-bound arguments.

References

[1] 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 Transactions on Automatic Control*, vol. 64, no. 8, pp. 3324–3331, 2019.

[2] Z. Liu, B. Wu, J. Dai, and H. Lin, “Distributed communication-aware motion planning for multi-agent systems from STL and SpaTeL specifications,” in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, 2017, pp. 4452–4457.

[3] A. T. Buyukkocak, D. Aksaray, and Y. Yazıcıoğlu, “Planning of heterogeneous multi-agent systems under signal temporal logic specifications with integral predicates,” *IEEE Rob. and Autom. Letters*, vol. 6, no. 2, pp. 1375–1382, 2021.

[4] O. Maler and D. Nickovic, “Monitoring Temporal Properties of Continuous Signals,” in *Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems*, Y. Lakhnech and S. Yovine, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 152–166.

[5] V. Raman, M. Maasoumy, A. Donze, R. M. Murray, A. Sangiovanni-Vincentelli, and S. A. Seshia, “Model predictive control with signal temporal logic specifications,” in *Proceedings of the IEEE Conf. on Decis. and Cont.* IEEE, 2014, pp. 81–87.

[6] A. Donzé and O. Maler, “Robust Satisfaction of Temporal Logic over Real-Valued Signals,” in *8th International Conference on Formal Modeling and Analysis of Timed Systems, FORMATS, 2010*, Klosterneuburg, Austria, 2010, pp. 92–106.

[7] S. Safaoui, L. Lindemann, D. V. Dimarogonas, I. Shames, and T. H. Summers, “Control Design for Risk-Based Signal Temporal Logic Specifications,” *IEEE L-CSS*, vol. 4, no. 4, pp. 1000–1005, 2020.

[8] L. Lindemann, G. J. Pappas, and D. V. Dimarogonas, “Control barrier functions for nonholonomic systems under risk signal temporal logic specifications,” in *2020 59th IEEE Conference on Decision and Control (CDC)*, 2020, pp. 1422–1428.

[9] ——, “Reactive and Risk-Aware Control for Signal Temporal Logic,” *IEEE Trans. Autom. Control*, vol. 67, no. 10, pp. 5262–5277, 2022.

[10] D. Sadigh and A. Kapoor, “Safe Control under Uncertainty with Probabilistic Signal Temporal Logic,” in *Proceedings of Robotics: Science and Systems*, jun 2016.

[11] S. Jha, V. Raman, D. Sadigh, and S. A. Seshia, “Safe Autonomy Under Perception Uncertainty Using Chance-Constrained Temporal Logic,” *Journal of Automated Reasoning*, vol. 60, no. 1, pp. 43–62, 2018.

[12] J. Li, P. Nuzzo, A. Sangiovanni-Vincentelli, Y. Xi, and D. Li, “Stochastic contracts for cyber-physical system design under probabilistic requirements,” in *MEMOCODE 2017 - 15th ACM-IEEE International Conference on Formal Methods and Models for System Design*, 2017, pp. 5–14.

[13] P. Kyriakis, J. V. Deshmukh, and P. Bogdan, “Specification mining and robust design under uncertainty: A stochastic temporal logic approach,” *ACM Trans. on Emb. Comp. Syst.*, vol. 18, no. 5s, 2019.

[14] G. Scher, S. Sadraddini, R. Tedrake, and H. Kress-Gazit, “Elliptical slice sampling for probabilistic verification of stochastic systems with signal temporal logic specifications,” in *Proceedings of the 25th ACM International Conference on Hybrid Systems: Computation and Control*, ser. HSCC ’22. New York, NY, USA: Association for Computing Machinery, 2022.

[15] G. Scher, S. Sadraddini, and H. Kress-Gazit, “Robustness-based synthesis for stochastic systems under signal temporal logic tasks,” in *IEEE IROS*, 2022, pp. 1269–1275.

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

[17] T. Yang, Y. Zou, S. Li, and Y. Yang, “Distributed model predictive control for probabilistic signal temporal logic specifications,” *IEEE Trans. on Autom. Sc. and Eng.*, pp. 1–11, 2023.

[18] L. Hewing, A. Carron, K. P. Wabersich, and M. N. Zeilinger, “On a correspondence between probabilistic and robust invariant sets for linear systems,” in *2018 European Control Conference (ECC)*, 2018, pp. 1642–1647.

[19] J. Orlin, “Contentment in graph theory: Covering graphs with cliques,” *Indagationes Mathematicae*, vol. 80, no. 5, pp. 406–424, 1977.

[20] E. E. Vlahakis, L. Lindemann, and D. V. Dimarogonas, “Distributed sequential receding horizon control of multi-agent systems under recurring signal temporal logic,” in *2024 European Control Conference (ECC)*, 2024, pp. 305–310.

[21] “Gurobi optimizer,” https://www.gurobi.com, 2023.