Conformal Data-driven Control of Stochastic Multi-Agent Systems under Collaborative Signal Temporal Logic Specifications

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

***Abstract*— We study the control of stochastic linear multi-agent systems (MAS) under additive stochastic noise and collaborative signal temporal logic (STL) specifications to be satisfied with a desired probability. Given available disturbance datasets, we leverage conformal prediction (CP) to address the underlying chance-constrained multi-agent STL synthesis problem in a distribution-free manner. By introducing non-conformity scores as functions of prediction regions (PRs) of error trajectories, we develop an iterative PR-scaling and disturbance-feedback synthesis approach to bound training error trajectory samples. These bounds are calibrated using a separate dataset, providing probabilistic guarantees via CP. We then relax the stochastic problem by tightening the robustness functions using Lipschitz constants and the computed error bounds. To address scalability, we exploit the compositional structure of the multi-agent STL formula and propose a model-predictive-control-like algorithm, where agent-level problems are solved in a distributed fashion.**

I. Introduction

Multi-agent systems (MAS) arise when multiple agents collaborate to achieve global objectives, while signal temporal logic (STL) offers a formal framework to specify such objectives [1]. In stochastic settings, STL control synthesis often relies on chance constraints, which are computationally demanding and typically addressed via constraint tightening [2] or analytic techniques [3], but these can be conservative or intractable in non-Gaussian settings, limiting applicability to general MAS. We propose a data-driven control design for stochastic MAS under collaborative STL specifications, using conformal prediction (CP) to provide distribution-free probabilistic guarantees [4]. While CP has recently been explored in control and STL [5], mostly for single-agent systems, we focus on MAS under collaborative tasks. Given agent-level disturbance datasets, we iteratively train disturbance feedback controllers and prediction regions of aggregated error trajectories, which are then calibrated to ensure CP-based probabilistic guarantees, yielding tighter bounds than existing methods [2], [6]. Last, we relax the stochastic control problem via Lipschitz-based tightening of robustness functions and propose a distributed MPC-like algorithm exploiting the compositional STL structure to enhance scalability.

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.

II. Preliminaries and problem setup

Conformal Prediction: If $\mathcal{R}^{(0)},\dots,\mathcal{R}^{(k)}$ are i.i.d. random variables, then for any $\theta\in(0,1)$, we have

$$\Pr\left\{\mathcal{R}^{(0)}\leq Q_{1-\theta}\Big(\mathcal{R}^{(1)},\dots,\mathcal{R}^{(k)},\infty\Big)\right\}\geq 1-\theta, \quad (1)$$

where $Q_{1-\theta}\big(\mathcal{R}^{(1)},\dots,\mathcal{R}^{(k)},\infty\big)$ is the $(1-\theta)$th quantile of the empirical distribution $\{\mathcal{R}^{(1)},\dots,\mathcal{R}^{(k)},\infty\}$ [7].

Signal temporal logic: We consider the STL syntax

$$\varphi \coloneqq \top \mid \pi \mid \neg\phi \mid \phi_1\wedge\phi_2 \mid \phi_1 U_{[t_1,t_2]}\phi_2, \quad (2)$$

where $\pi:=(\mu(x)\geq 0)$ is a predicate, $\mu(x):\mathbb{R}^{n_x}\to\mathbb{R}$, 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. A scalar-valued function $\rho^{\phi}(\boldsymbol{x}(t))$ of a signal indicates how robustly a signal $\boldsymbol{x}(t)$ satisfies a formula ϕ. Specifically, $\rho^{\phi}(\boldsymbol{x}(t))\geq 0 \iff \boldsymbol{x}(t)\models\phi$.

Dynamics: The aggregate dynamics of $|\nu|$ agents are

$$x_\nu(t+1)=A_\nu x_\nu(t)+B_\nu u_\nu(t)+w_\nu(t). \quad (3)$$

STL specification: The MAS is subject to

$$\phi=\bigwedge_{\nu\in\mathcal{K}_\phi}\phi_\nu, \quad (4)$$

where ϕ_ν is a formula involving a clique of agents ν, with $1\leq|\nu|\leq M$, and $\mathcal{K}_\phi$ collects all these cliques induced by ϕ.

Disturbance: Sets $\mathcal{D}^{w_i}=\{\boldsymbol{w}_i^{(0)},\dots,\boldsymbol{w}_i^{(k)}\}$ of $k{+}1$ samples $\forall i\in\mathcal{V}$ are available, with $\boldsymbol{w}_i^{(\varsigma)}=(w_i^{(\varsigma)}(0),\dots,w_i^{(\varsigma)}(N{-}1))$.

Problem statement: Given $x(0)=x_0$, we wish to solve

$$\begin{aligned}&\underset{\substack{\boldsymbol{u}(0:N-1)\\ \boldsymbol{x}(0:N)}}{\text{Min.}}\ \mathbb{E}\left(\sum_{i=1}^{M}\Big(\sum_{t=0}^{N-1}(\ell_i(x_i(t),u_i(t)))+V_{f,i}(x_i(N))\Big)\right)\\ &\text{s.t. } x(t+1)=Ax(t)+Bu(t)+w(t),\ t\in\mathbb{N}_{[0,N)},\\ &\qquad \Pr\{\boldsymbol{x}_\nu(0:N)\models\phi_\nu,\ \forall\nu\in\mathcal{K}_\phi\}\geq 1-\theta, \end{aligned} \quad (5)$$

where $\boldsymbol{u}(0{:}N{-}1)$, $\boldsymbol{x}(0{:}N)$, are the opt. variables, with $u(t)=(u_1(t),\dots,u_M(t))$, and $x(t)=(x_1(t),\dots,x_M(t))$, resp., and ϕ is a multi-agent STL formula to be satisfied by $\boldsymbol{x}(0{:}N)$ with a probability $1-\theta$.

III. Summary of our approach

Decomposition of dynamics: Consider the feedback policy

$$u_i(t)=\sum_{k=0}^{t-1}\Gamma_i^{t,k}w_i(k)+v_i(t). \quad (6)$$

Then, the aggregate dynamics of the agents in ν can be decomposed into

$$z_\nu(t+1) = A_\nu z_\nu(t) + B_\nu v_\nu(t), \tag{7a}$$

$$e_\nu(t+1) = A_\nu e_\nu(t) + \sum_{k=0}^{t-1} \Gamma_\nu^{t,k} w_\nu(k) + w_\nu(t). \tag{7b}$$

Given disturbance feedback gains $\Gamma_\nu^{t,k}$, the systems in (7) can be analyzed independently.

Error trajectory samples: From disturbance samples $\boldsymbol{w}_i^{(\varsigma)}(0{:}N-1) \in \mathcal{D}^{w_i}$, we may construct error samples

$$\mathcal{D}^{e_i} = \{\boldsymbol{e}_i^{(0)}(1{:}N), \ldots, \boldsymbol{e}_i^{(k)}(1{:}N)\}, \tag{8a}$$

$$\boldsymbol{e}_i^{(\varsigma)}(1{:}N) = (\boldsymbol{A}_i + \boldsymbol{B}_i \boldsymbol{\Gamma}_i)\boldsymbol{w}_i^{(\varsigma)}(0{:}N-1),\ \varsigma \in \mathbb{N}_{[0,k]}. \tag{8b}$$

Training bounds for (7b) **and disturbance feedback:** Let nonconformity scores

$$E^{(\varsigma)}(C, \boldsymbol{\Gamma})) = \max_{\nu \in \mathcal{K}_\phi} \left(C_\nu \|\boldsymbol{e}_\nu^{(\varsigma)}(1:N)\| \right), \tag{9}$$

where $C = \{C_\nu\}_{\nu \in \mathcal{K}_\phi}$ and $\boldsymbol{\Gamma} = \{\boldsymbol{\Gamma}_1, \ldots, \boldsymbol{\Gamma}_M\}$. Synthesis of $\boldsymbol{\Gamma}$ and C is formulated as

$$\underset{C,\, \boldsymbol{\Gamma}}{\text{Minimize}}\ Q_{\hat{\theta}}\Big(E^{(k_1+1)}(C, \boldsymbol{\Gamma}), \ldots, E^{(k)}(C, \boldsymbol{\Gamma})\Big) \tag{10a}$$

$$\text{subject to } 0 \le C_\nu \le 1,\ \nu \in \mathcal{K}_\phi,\ \sum_{\nu \in \mathcal{K}_\phi} C_\nu = 1. \tag{10b}$$

Calibration: Using fresh datasets $\{\boldsymbol{e}_i^{(1)}, \ldots, \boldsymbol{e}_i^{(k_1)}\}$, and optimal C^*, $\boldsymbol{\Gamma}^*$, one can use (1) to obtain guarantees

$$\Pr\left\{\boldsymbol{e}_\nu^{(0)}(1:N) \in \mathbb{B}(q/C_\nu^*),\ \forall \nu \in \mathcal{K}_\phi\right\} \ge 1 - \theta, \tag{11}$$

by computing $q = Q_{1-\theta}\big(E^{(1)}(C^*, \boldsymbol{\Gamma}^*), \ldots, E^{(k_1)}(C^*, \boldsymbol{\Gamma}^*), \infty\big)$.

Relaxation: The original problem in (5) can be relaxed as

$$\begin{aligned} \underset{\boldsymbol{v}(0),\, \boldsymbol{z}(0)}{\text{Minimize}}\ & \sum_{i=1}^{M} \left(\sum_{t=0}^{N-1} (\ell_i(z_i(t), v_i(t))) + V_{f,i}(z_i(N)) \right) \\ \text{subject to } & z(t+1) = Az(t) + Bv(t),\ t \in \mathbb{N}_{[0,N)}, \\ & \rho^{\phi_\nu}(\boldsymbol{z}_\nu(0:N)) \ge L_{\phi_\nu} \frac{q}{C_\nu^*},\ \nu \in \mathcal{K}_\phi, \end{aligned} \tag{12}$$

where $z(0) = x(0)$, and L_{ϕ_ν} is the Lipschitz constant of the robustness function $\rho^{\phi_\nu}(\boldsymbol{z}_\nu(0{:}N) + \boldsymbol{e}_\nu(0{:}N))$ wrt $\boldsymbol{e}_\nu(0{:}N)$.

Distributed STL synthesis: Let $\mathcal{T}_i = \{\nu \in \mathcal{K}_\phi \mid \nu \ni i\}$, and $\hat{\phi} = \bigwedge_{i \in \mathcal{V}} \hat{\phi}_i$, where $\hat{\phi}_i = \bigwedge_{\nu_i \in \mathcal{T}_i} \phi_{\nu_i}$. In the following,

$$\varrho^{\phi_{\nu_i}}(\boldsymbol{z}_{\nu_i}(0{:}N)) = \rho^{\phi_{\nu_i}}(\boldsymbol{z}_{\nu_i}(0{:}N)) - L_{\phi_{\nu_i}} \frac{q}{C_{\nu_i}},\ \nu_i \in \mathcal{T}_i. \tag{13}$$

We introduce the following problems for the ith agent

$$P_i^0 := \underset{\boldsymbol{v}_i^0, \boldsymbol{z}_i^0}{\text{Minimize}}\ \mathcal{L}_i(\boldsymbol{z}_i^0, \boldsymbol{v}_i^0) \text{ subject to} \tag{14a}$$

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

$$\varrho^{\phi_i}(\boldsymbol{z}_i^0) \ge 0, \text{ with } z_i^0(0) = x_i(0), \tag{14c}$$

$$P_i^t := \underset{\boldsymbol{v}_i^t, \boldsymbol{z}_i^t}{\text{Minimize}}\ \mathcal{L}_i(\boldsymbol{z}_i^t, \boldsymbol{v}_i^t) - \Omega_i \mu_{\nu_t}^t \text{ subject to} \tag{15a}$$

$$z_i^t(k+1) = A_i z_i^t(k) + B_i v_i^t(k),\ k \in \mathbb{N}_{[t,N)}, \tag{15b}$$

$$\varrho^{\phi_i}(\boldsymbol{z}_i^t) \ge 0, \text{ with } z_i^t(t) = x_i(t), \tag{15c}$$

$$\varrho^{\phi_{\nu_t}}(\boldsymbol{z}_{\nu_t}^t) \ge \mu_{\nu_t}^t,\ \nu_t = \underset{\nu \in \mathcal{T}_i,\ |\nu|>1}{\operatorname{argmin}} \{\varrho^{\phi_\nu}(\boldsymbol{z}_\nu^{t-1})\}, \tag{15d}$$

$$\mu_{\nu_t}^t \ge \min\left(0, \varrho^{\phi_{\nu_t}}(\boldsymbol{z}_{\nu_t}^{t-1})\right), \tag{15e}$$

$$\varrho^{\phi_\nu}(\boldsymbol{z}_\nu^t) \ge \min\left(0, \varrho^{\phi_\nu}(\boldsymbol{z}_\nu^{t-1})\right), \forall \nu \in \mathcal{T}_i \backslash \{\nu_t, i\} \tag{15f}$$

where $\Omega_i \gg 0$, $z_i^t(\tau)$ denotes the prediction of $x_i(\tau)$ carried out at time t, $\varrho^{\phi_\nu}(\boldsymbol{z}_\nu^t)$ is the robustness function of the formula ϕ_ν, $\nu \in \mathcal{T}_i$, evaluated over the trajectory $\boldsymbol{z}_\nu^t$, and

$$\boldsymbol{z}_\nu^t = (x_\nu(0), \ldots, x_\nu(t-1), z_\nu^t(t), \ldots, z_\nu^t(N)), \tag{16}$$

$$\boldsymbol{v}_i^t = (v_i(0), \ldots, v_i(t-1), v_i^t(t), \ldots, v_i^t(N-1)), \tag{17}$$

$$\mathcal{L}_i(\boldsymbol{z}_i^t, \boldsymbol{v}_i^t) = \sum_{k=0}^{N-1} \ell_i(z_i^t(k), v_i^t(k)) + V_{f,i}(z_i^t(N)). \tag{18}$$

Let also $\boldsymbol{z}_i^t(x_i(t), \boldsymbol{v}_i^t) = (x_i(0), \ldots, x_i(t), z_i^t(t+1), \ldots, z_i^t(N))$ denote a trajectory where the last $N-t$ nominal states are generated by the last $N-t$ inputs of $\boldsymbol{v}_i^t$ starting from $x_i(t)$. Alg. 1 summarizes the proposed distributed STL control strategy. Complete version of this work is available in [8].

Algorithm 1 Distributed STL control of agent-i

for t in $1:N$ **do**
  Compute $r_i^t = \min_{\nu_i \in \mathcal{T}_i} \left(\varrho^{\phi_{\nu_i}}(\boldsymbol{z}_{\nu_i}^{t-1})\right)$
  Measure $x_i(t)$ and $w_i(t-1)$
  Construct $\boldsymbol{z}_i^t(x_i(t), \boldsymbol{v}_i^{t-1})$
  Communicate r_i^t, $\boldsymbol{z}_i^t(x_i(t), \boldsymbol{v}_i^{t-1})$ to $j \in \nu_i$, $\nu_i \in \mathcal{T}_i$
  Receive r_j^t, $\boldsymbol{z}_j^t(x_i(t), \boldsymbol{v}_i^{t-1})$ from $j \in \nu_i$, $\nu_i \in \mathcal{T}_i$
  if $r_i^t < r_j^t$ for all $j \in \nu_i$, $\nu_i \in \mathcal{T}_i$ **then**
    Solve P_i^t and store $(\boldsymbol{v}_i^t, \boldsymbol{z}_i^t)$
  else
    Update $\boldsymbol{v}_i^t \leftarrow \boldsymbol{v}_i^{t-1}$ and $\boldsymbol{z}_i^t \leftarrow \boldsymbol{z}_i^t(x_i(t), \boldsymbol{v}_i^{t-1})$
  Apply $u_i(t) = \sum_{k=0}^{t-1} \Gamma_i^{t,k} w_i(k) + v_i^t(t)$

References

[1] 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.

[2] 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.

[3] 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.

[4] V. Vovk, A. Gammerman, and G. Shafer, *Algorithmic Learning in a Random World*. Springer Science & Business Media, 2005.

[5] L. Lindemann, Y. Zhao, X. Yu, G. J. Pappas, and J. V. Deshmukh, "Formal verification and control with conformal prediction," 2024. [Online]. Available: https://arxiv.org/abs/2409.00536

[6] E. E. Vlahakis, L. Lindemann, P. Sopasakis, and D. V. Dimarogonas, "Probabilistic tube-based control synthesis of stochastic multi-agent systems under signal temporal logic," in *2024 IEEE 63rd Conference on Decision and Control (CDC)*, 2024, pp. 1586–1592.

[7] R. J. Tibshirani, R. F. Barber, E. J. Candès, and A. Ramdas, "Conformal prediction under covariate shift," in *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[8] E. E. Vlahakis, L. Lindemann, and D. V. Dimarogonas, "Conformal data-driven control of stochastic multi-agent systems under collaborative signal temporal logic specifications," 2025. [Online]. Available: https://arxiv.org/abs/2504.04615