

Distributionally Robust Control for Chance-Constrained Signal Temporal Logic Specifications

Arash Bahari Kordabad, Eleftherios E. Vlahakis, Lars Lindemann, Dimos V. Dimarogonas, and Sadegh Soudjani

Abstract—We consider distributionally robust optimal control of stochastic linear systems under signal temporal logic (STL) chance constraints when the disturbance distribution is unknown. By assuming that the underlying predicate functions are Lipschitz continuous and the noise realizations are drawn from a distribution having a concentration of measure property, we first formulate the underlying chance-constrained control problem as stochastic programming with constraints on expectations and propose a solution using a distributionally robust approach based on the Wasserstein metric. We show that by choosing a proper Wasserstein radius, the original chance-constrained optimization can be satisfied with a user-defined confidence level. A numerical example illustrates the efficacy of the method.

I. INTRODUCTION

Control of stochastic systems under temporal logic finds application in a wide range of domains, including robotics, autonomous systems, and cyber-physical systems. The formal specification of system properties that can be formulated in a probabilistic setting, enabling a systematic approach to quantifying uncertainty and handling feasibility, lies at the core of this problem [1]. Signal temporal logic (STL) is a formal language that allows us to encode time-constrained tasks using both Boolean and quantitative semantics [2], [3]. When systems are subject to stochastic disturbances and STL specification, a typical probabilistic approach is to formulate the problem as a chance-constrained program (CCP) [4].

Most recent results in the probabilistic STL context focus on applying probability or risk measures to individual predicates. To address the impact of critical tail events when STL formulas are violated, [5] propose Risk Signal Temporal Logic, incorporating risk constraints over predicates while preserving Boolean and temporal operators. Authors in [6] introduce probabilistic signal temporal logic, allowing expression of uncertainty by incorporating random variables into predicates. Similarly, in [7], chance-constrained temporal logic formulates chance constraints as predicates to model perception uncertainty for autonomous vehicles. Stochastic

temporal logic introduced in [8] is similar in syntax to chance-constrained temporal logic but is designed for stochastic systems, where perturbations affect system dynamics rather than predicate coefficients.

Top-down approaches study STL probabilistic verification of stochastic systems considering chance constraints on the entire specification [9]. More closely related to our work is [10], in which the authors transform chance constraints into linear constraints using concentration of measure inequalities to provide a conservative approximation of the feasible domain. Due to the nonconvex feasible domains typically induced by CCPs, many studies focus on numerical methods to handle CCPs, such as randomized optimization where the original optimization is approximated by a scenario program (SP) by sampling the uncertainty space. SP approaches have been studied for convex [11], [12], and nonconvex CCPs [13].

Unlike stochastic optimization settings where the probability distribution is assumed to be known, distributionally robust optimization (DRO) addresses the lack of information on the probability distribution by considering the worst-case distribution within an ambiguity set. Various methods exist for constructing ambiguity sets, such as moment ambiguity [14], Kullback–Leibler divergence-based ball [15], and Wasserstein-based ball [16]. The Wasserstein ambiguity set represents a statistical ball within the space of probability distributions surrounding the empirical distribution, with its radius measured using Wasserstein distance. Wasserstein DRO offers a probabilistic guarantee based on finite samples within a tractable formulation [17] and has attracted significant attention recently [18].

In the DRO literature, several works have addressed CCP directly. An explicit reformulation for both individual and joint CCPs was presented in [19], where uncertainties are modeled as affine functions. The authors of [20] reformulated CCP as a conditional value-at-risk (CVaR) mixed-integer program for affine functions. A more general approach for CCP has been proposed in [21]. However, dealing with expectation-constrained programs (ECPs) preserves linearity and convexity, making it often computationally simpler and more straightforward compared to evaluating or approximating CCPs, particularly for non-standard distributions.

An SP approach has been proposed in [22] for solving a general CCP by transforming it into an ECP assuming that 1) the underlying distribution of the uncertain parameters satisfies a *concentration of measure* property and exhibits bounded variance, and 2) the constraint function is Lipschitz

Arash Bahari Kordabad and Sadegh Soudjani are with the Max Planck Institute for Software Systems, Kaiserslautern, Germany. E-mail: {arashbk, sadegh}@mpi-sws.org. Eleftherios E. Vlahakis and Dimos V. Dimarogonas are with the Division of Decision and Control Systems, KTH Royal Institute of Technology, Stockholm, Sweden. E-mail: {vlahakis, dimos}@kth.se. Lars Lindemann is with the Thomas Lord Department of Computer Science, Viterbi School of Engineering, University of Southern California, Los Angeles, USA. Email: llindema@usc.edu.

This research is supported by the following grants: EIC 101070802 and ERC 101089047.

continuous in the uncertainty parameters.

In this paper, we formulate a stochastic optimal control problem as a CCP, where the underlying system is subject to stochastic disturbances with unknown distribution and a set of STL specifications. We first assume that the disturbance realizations follow a concentration of measure property and that the underlying predicate functions involved in the STL constraints are Lipschitz continuous in the perturbation parameters. We build upon the results in [22] and transform the CCP to stochastic programming with expectation constraints. Since the exact distribution is assumed to be unknown, we propose a data-driven Wasserstein distributionally robust approach that guarantees STL satisfaction in a probabilistic sense.

The remainder of the paper is organized as follows. In Section II, we present the system setup, the STL formulation and the control synthesis problem we study. The construction of a stochastic program, its connection to the original CCP through the concentration of measure, and the distributionally robust solution of the stochastic program are in Section III. A numerical study is in Section IV, and concluding remarks are in Section V.

II. PROBLEM FORMULATION

A. Discrete-Time Stochastic Linear Systems

We consider systems in discrete time with state space $\mathcal{X} \subseteq \mathbb{R}^n$, input space $\mathcal{U} \subseteq \mathbb{R}^m$, and disturbance set $\mathcal{W} \subseteq \mathbb{R}^n$, that can be modeled by linear difference equations perturbed by stochastic disturbances:

$$x_{k+1} = Ax_k + Bu_k + w_k, \quad (1)$$

where $x_k \in \mathcal{X}$ denotes the state of the system at time instant k , $u_k \in \mathcal{U}$ denotes the control input at time instant k , and $w_k \in \mathcal{W}$ is a random vector that has an unknown probability distribution \mathcal{P} supported on \mathcal{W} . Matrices $A \in \mathbb{R}^{n \times n}$, and $B \in \mathbb{R}^{n \times m}$, and initial state x_0 are assumed to be known. In the view of (1), for any $k \in \mathbb{N}$, x_k is a function of x_0 , input sequence vector $u_{0:k} := [u_0^\top, \dots, u_{k-1}^\top]^\top$, and the process noise $w_{0:k} := [w_0^\top, \dots, w_{k-1}^\top]^\top$:

$$x_k = A^k x_0 + \sum_{i=0}^{k-1} A^{k-i-1} (Bu_i + w_i). \quad (2)$$

B. STL specifications

We consider signal temporal logic (STL) formulas defined recursively according to the grammar [23]:

$$\varphi ::= \mathcal{T} \mid \pi \mid \neg\varphi \mid \varphi \wedge \psi \mid \varphi U_{[a,b]} \psi,$$

where \mathcal{T} is the *true* predicate; π is a predicate whose truth value is determined by the sign of a predicate function of state variables, i.e. $\pi = \{\alpha(x) \geq 0\}$ with $\alpha: \mathbb{R}^n \rightarrow \mathbb{R}$; ψ is an STL formula; \neg and \wedge indicate negation and conjunction of formulas; and $U_{[a,b]}$ is the *until* operator with $a, b \in \mathbb{N}$. A finite run $\xi := \{x_0, x_1, x_2, \dots, x_N\}$ satisfies φ at time k , denoted by $(\xi, k) \models \varphi$ with the Boolean semantics of STL formulas defined as follows:

$$(\xi, k) \models \pi \quad \Leftrightarrow \quad \alpha(x_k) \geq 0,$$

$$\begin{aligned} (\xi, k) \models \neg\varphi &\quad \Leftrightarrow \quad \neg((\xi, k) \models \varphi), \\ (\xi, k) \models \varphi \wedge \psi &\quad \Leftrightarrow \quad (\xi, k) \models \varphi \wedge (\xi, k) \models \psi, \\ (\xi, k) \models \varphi U_{[a,b]} \psi &\quad \Leftrightarrow \quad \exists k' \in \{a, \dots, b\}, (\xi, k+k') \models \psi \\ &\quad \wedge \forall k'' \in \{k, \dots, k'\}, (\xi, k'') \models \varphi, \end{aligned}$$

Additionally, we derive the *disjunction* operator as $\varphi \vee \psi := \neg(\neg\varphi \wedge \neg\psi)$, the *eventually* operator as $\diamond_{[a,b]} \varphi := \mathcal{T}U_{[a,b]} \varphi$, and the *always* operator as $\square_{[a,b]} \varphi := \neg \diamond_{[a,b]} \neg\varphi$. Thus $(\xi, k) \models \diamond_{[a,b]} \varphi$ if φ holds at some time instant between $a+k$ and $b+k$ and $(\xi, k) \models \square_{[a,b]} \varphi$ if φ holds at every time instant between $a+k$ and $b+k$.

STL Robustness. In contrast to the above Boolean semantics, the quantitative semantics (a.k.a. robustness function) of STL [24] assigns to each formula φ a real-valued function ρ^φ of signal ξ and k such that $\rho^\varphi > 0$ implies $(\xi, k) \models \varphi$. The robustness of a formula φ with respect to a run ξ at time k is defined recursively as

$$\begin{aligned} \rho^\top(\xi, k) &= +\infty \\ \rho^\mu(\xi, k) &= \alpha(x_k) \\ \rho^{\neg\phi}(\xi, k) &= -\rho^\phi(\xi, k) \\ \rho^{\phi \wedge \psi}(\xi, k) &= \min(\rho^\phi(\xi, k), \rho^\psi(\xi, k)) \\ \rho^{\phi U_{[a,b]} \psi}(\xi, k) &= \max_{k' \in \{a, \dots, b\}} \left(\min(\rho^\psi(\xi, k+k'), \right. \\ &\quad \left. \min_{k'' \in \{k, \dots, k'\}} \rho^\phi(\xi, k'')) \right). \end{aligned}$$

Therefore, the value of the robustness function $\rho^\phi(\xi, k)$ can be interpreted as how much the trajectory ξ satisfies a given STL formula ϕ . The robustness of the formulas $\diamond_{[a,b]} \varphi$ and $\square_{[a,b]} \varphi$ are

$$\begin{aligned} \rho^{\diamond_{[a,b]} \varphi}(\xi, k) &= \max_{k' \in \{a, \dots, b\}} \rho^\varphi(\xi, k+k'), \\ \rho^{\square_{[a,b]} \varphi}(\xi, k) &= \min_{k' \in \{a, \dots, b\}} \rho^\varphi(\xi, k+k'). \end{aligned}$$

As described above, the robustness function is commonly defined with its arguments being the system trajectory and the time index. However, for the scope of this study, it is more comprehensible to explicitly define the robustness as a function of the control input and the disturbance. In essence, as in (2), the system trajectory is determined by the initial state, input, and disturbance sequence. Therefore, we define the dynamic-dependent function ϱ^φ as follows:

$$\varrho^\varphi(u_{k:N}, w_{k:N}, x_k, k) := \rho^\varphi(\xi, k).$$

Moreover, at time $k=0$ for a given x_0 , we eliminate x_k and k from the argument of ϱ^φ and define

$$\varrho_0^\varphi(\mathbf{u}, \mathbf{w}) := \varrho^\varphi(\mathbf{u}, \mathbf{w}, x_0, 0),$$

where $\mathbf{u} := u_{0:N}$ and $\mathbf{w} := w_{0:N}$.

In the following, we make a regularity assumption on the STL predicate functions. This assumption allows us to establish the Lipschitz continuity of robustness functions, define a well-defined chance constraint for the robustness function, and utilize the concentration of measure property.

Assumption 1. We assume that the predicate functions are Lipschitz functions.

C. Chance-constrained optimization

In the following, we introduce the STL chance-constrained optimization. Our aim is to provide a control input \mathbf{u} such that the function ϱ_0^φ becomes positive, or specifically lower bounded by a pre-defined positive robustness level r_0 . Since this function is affected by the unknown disturbance \mathbf{w} , the purpose would be to satisfy the inequality

$$\varrho_0^\varphi(\mathbf{u}, \mathbf{w}) \geq r_0$$

with probability at least a given threshold. We denote the N -fold product of the probability measure \mathcal{P} by $P := \mathcal{P}^N$ supported on the Cartesian product space $\mathbb{W} := \mathcal{W}^N$, and we denote $\mathbb{U} := \mathcal{U}^N$. We then define the probability space $(\mathbb{W}, \mathcal{F}, P)$ and the following chance-constrained program (CCP):

$$\text{CCP} : \begin{cases} \min_{\mathbf{u} \in \mathbb{U}} \mathbb{E}_P [J(\mathbf{u}, \mathbf{w})], \\ \text{s.t. } P \{ \varrho_0^\varphi(\mathbf{u}, \mathbf{w}) \geq r_0 \} \geq 1 - \varepsilon, \end{cases} \quad (3)$$

where $\varepsilon \in (0, 1)$ is the constraint violation tolerance and J is a lower semi-continuous cost function. The aim of this optimization is to find an optimal control sequence that minimizes the expected value of a performance function J while ensuring that the STL constraint, with robustness level r_0 , is satisfied with a probability of at least $1 - \varepsilon$. Under Assumption 1 and for the linear system in (1), where the mapping from the disturbances to the system trajectory is continuous, the above optimization is well-defined and attains a solution if it is feasible [22]. This is due to the continuity of the robustness function under Assumption 1, which arises from the continuity of functions containing min and max operators with continuous arguments.

D. Concentration of measure property

In this paper, we make the following assumptions on the distribution of \mathbf{w} .

Assumption 2. (Light-tailed distribution). The distribution of random variable \mathbf{w} is a Light-tailed distribution. More specifically, there exists $a > 1$ such that $C := \mathbb{E}_P[\exp \|\mathbf{w}\|^a] < \infty$.

Assumption 2 holds for many distributions, e.g., multivariate normal distribution, exponential distribution, log-normal distribution, and all distributions with bounded support.

Assumption 3. (Concentration of Measure). There exists a monotonically decreasing function $h : \mathbb{R}^{\geq 0} \rightarrow [0, 1]$ such that

$$P \{ |f(\mathbf{w}) - \mathbb{E}[f(\mathbf{w})]| \leq t \} \geq 1 - h(t), \quad \forall t \geq 0, \quad (4)$$

holds for any Lipschitz continuous function $f : \mathbb{W} \rightarrow \mathbb{R}$ with Lipschitz constant 1.

We recall that a function $f : \mathbb{W} \rightarrow \mathbb{R}$ is a Lipschitz continuous function if there exists $L \geq 0$ such that for any two vectors $\mathbf{w}_1, \mathbf{w}_2 \in \mathbb{W}$, $\frac{|f(\mathbf{w}_1) - f(\mathbf{w}_2)|}{d_{\mathbb{W}}(\mathbf{w}_1, \mathbf{w}_2)} \leq L$ holds, where $d_{\mathbb{W}}$ denotes a metric on the set \mathbb{W} . Constant L is referred

to as the Lipschitz constant. Throughout this paper, we use the 2-norm on \mathbb{W} to calculate the Lipschitz constants, given by $d_{\mathbb{W}}(\mathbf{w}_1, \mathbf{w}_2) = \sqrt{\langle \mathbf{w}_1 - \mathbf{w}_2, \mathbf{w}_1 - \mathbf{w}_2 \rangle}$. Note that, by employing alternative metrics, the general results of the paper remain valid, and changing the metric only affects the values of the Lipschitz constant.

Assumption 2 will be utilized in deriving a data-driven solution to (3) in Section III. We use Assumption 3 to construct a stochastic program for finding a (possibly sub-optimal) solution of CCP (3). Note that Assumption 3 also holds for many distributions. Examples of different distributions with the concentration of measure property and corresponding h functions can be found in [22]. For instance, the standard multi-variate Gaussian distribution satisfies (4) with $h(t) = \min\{2e^{-2t^2/\pi^2}, 1\}$ [25], [26]. Note that if we substitute $h(\cdot)$ with another monotonically decreasing function $\bar{h}(\cdot)$, and $\bar{h}(\cdot)$ satisfies $\bar{h}(\cdot) \geq h(\cdot)$, then the inequality (4) remains valid when using $\bar{h}(\cdot)$.

In the following, we provide a lemma and a theorem that enable us to use the concentration of measure property in the context of STL and robustness function.

Lemma 1. For any two Lipschitz functions $f_1 : X \rightarrow \mathbb{R}$ and $f_2 : X \rightarrow \mathbb{R}$, $\max(f_1, f_2)$ and $\min(f_1, f_2)$ are Lipschitz functions with $L := \max(L_1, L_2)$, where L_i is the Lipschitz constant of f_i , $i \in \{1, 2\}$.

Proof. The proof is given in the appendix of [27]. ■

Note that Lemma 1 can be extended for the cases where we have more than two functions inside the min or max operators. In this case, one can readily verify that $\min(f_1, \dots, f_n)$ and $\max(f_1, \dots, f_n)$ are Lipschitz functions with constant $\max\{L_1, \dots, L_n\}$ when f_i is a Lipschitz function with constant L_i for all $i \in \{1, \dots, n\}$. Moreover, the result holds for the combination of min and max with any number of operators. For instance, $\min(f_1, \dots, \max(f_j, \dots, f_l), \dots, f_n)$ with $1 \leq j \leq l \leq n$, is a Lipschitz function with constant $\max\{L_1, \dots, L_n\}$. The following theorem uses Lemma 1 and shows that the robustness of an STL specification for linear systems is a Lipschitz function. Moreover, it provides the corresponding Lipschitz constant.

Theorem 1. For any STL specification φ with Lipschitz atomic predicates, $\varrho_0^\varphi(\mathbf{u}, \mathbf{w})$ is Lipschitz continuous with respect to \mathbf{w} for the system defined in (1). The Lipschitz constant will be $L_\varphi = L_1 L_2$, where L_1 is the maximum Lipschitz constant of atomic predicates appearing in φ . More specifically, consider the STL formula φ consists of \mathcal{J} subformula with atomic predicate α_j , $j \in \{1, \dots, \mathcal{J}\}$, then:

$$L_1 := \max_{j \in \{1, \dots, \mathcal{J}\}} L_{\alpha_j},$$

where L_{α_j} is the Lipschitz constant of α_j for $j \in \{1, \dots, \mathcal{J}\}$. Constant L_2 is the maximum Lipschitz constant of x_k , the state at time $k \in \{1, 2, \dots, N\}$, with respect to \mathbf{w} , which is bounded by:

$$L_2 = \sqrt{\sum_{i=0}^{N-1} \|A^i\|^2}, \quad (5)$$

where $\|A^i\|$ is the induced 2-norm of the matrix A^i and N is the length of the sequence \mathbf{w} .

Proof. The proof is given in the appendix of [27]. \blacksquare

Theorem 1 enables us to use the results of [22] in the context of STL by providing an explicit formula for the Lipschitz constant. The next section details the proposed solution for the CCP in (3).

III. SOLUTION APPROACH

In the following, we use Assumption 3 for the Lipschitz continuous function $\varrho_0^\varphi(\mathbf{u}, \mathbf{w})$ (cf. Theorem 1) and provide an under approximation for CCP (3).

Theorem 2. *Under Assumption 3, the feasible domain of the CCP (3) includes the feasible domain of the following expectation-constrained program (ECP):*

$$\text{ECP} : \begin{cases} \min_{\mathbf{u} \in \mathbb{U}} \mathbb{E}_P [J(\mathbf{u}, \mathbf{w})], \\ \text{s.t. } \mathbb{E}_P [\varrho_0^\varphi(\mathbf{u}, \mathbf{w})] - L_\varphi h^{-1}(\varepsilon) \geq r_0, \end{cases} \quad (6)$$

where h is given by the class of distribution and L_φ is the Lipschitz constant of $\varrho_0^\varphi(\mathbf{u}, \mathbf{w})$ with respect to \mathbf{w} given in Theorem 1.

Proof. The proof is given in the appendix of [27]. \blacksquare

Note that the feasible domain of CCP (3) has been under-approximated by the tighter feasible domain of ECP (6) with constraints on the expectation and tightening function h^{-1} . Problem (6) can then be solved via: 1) computing the expectations $\mathbb{E}_P[\cdot]$, 2) utilizing knowledge of the family of distributions or an upper bound on the function h related to the concentration of measure property, and 3) computing the robustness Lipschitz constant L_φ .

In determining 2) and 3), knowledge of the exact distribution P is not necessary. However, to compute expectations in 1) one requires P . For instance, sample averaging is an often employed technique to approximate expectations numerically. Nonetheless, this approach necessitates a sufficiently large sample size to ensure the accuracy of the empirical expectation compared to the exact one.

In the following, we aim to solve the ECP (6) for the unknown distribution P with respect to the worst-case distribution in an ambiguity set using a Wasserstein distributionally robust approach and provide a finite sample guarantee with respect to the exact ECP (6). More specifically, the distributionally robust version of (6) can be written as the following distributionally robust program (DRP):

$$\text{DRP} : \begin{cases} \min_{\mathbf{u} \in \mathbb{U}} \sup_{Q \in \mathbb{Q}} \mathbb{E}_Q [J(\mathbf{u}, \mathbf{w})], \\ \text{s.t. } \inf_{Q \in \mathbb{Q}} \mathbb{E}_Q [\varrho_0^\varphi(\mathbf{u}, \mathbf{w})] - L_\varphi h^{-1}(\varepsilon) \geq r_0, \end{cases} \quad (7)$$

where \mathbb{Q} is an ambiguity set, defining a set of all distributions around an empirical distribution \hat{Q} that could contain the true distribution P with high confidence. In this paper, we use the Wasserstein metric $W : Q(\mathbb{W}) \times Q(\mathbb{W}) \rightarrow \mathbb{R}_{\geq 0}$ to define the ambiguity ball \mathbb{Q} as

$$\mathbb{Q} := \{Q \in Q(\mathbb{W}) \mid W(Q, \hat{Q}) \leq r\}, \quad (8)$$

where $Q(\mathbb{W})$ denotes the set of Borel probability measures on the support \mathbb{W} and $r \geq 0$ is the radius of the Wasserstein ball. For any two distributions $Q^1, Q^2 \in Q(\mathbb{W})$, the Wasserstein metric W is defined as follows:

$$W(Q^1, Q^2) := \min_{\kappa \in Q(\mathbb{W}^2)} \left\{ \int_{\mathbb{W}^2} \|\mathbf{w}_1 - \mathbf{w}_2\| d\kappa(\mathbf{w}_1, \mathbf{w}_2) \mid \Pi^j \kappa = Q^j, j = 1, 2 \right\}, \quad (9)$$

where $\Pi^j \kappa$ denotes the j^{th} marginal of the joint distribution κ for $j = 1, 2$. Note that, the sampling-based reformulation in (7) stems from the need to make decisions under uncertainty about the true distribution P that governs the random variable \mathbf{w} . Since the true distribution P is unknown, we rely on a finite number of i.i.d. samples $\{\mathbf{w}^i\}_{i=1}^M$ to infer information about P . These samples provide an empirical approximation \hat{Q} that can be constructed as follows:

$$\hat{Q} = \frac{1}{M} \sum_{i=1}^M \delta_{\mathbf{w}^i}, \quad (10)$$

where $\delta_{\mathbf{w}^i}$ is the Dirac measure concentrated at \mathbf{w}^i . Consider that, because \hat{Q} is constructed from a limited sample set, it may not perfectly capture the true distribution P . To account for this uncertainty, we introduced the ambiguity set \mathbb{Q} in (8), which includes all distributions that are close to the empirical distribution \hat{Q} within a Wasserstein ball of radius r . The parameter r reflects the confidence level: a larger r increases the probability that \mathbb{Q} contains the true distribution P , thereby providing a more robust solution to the optimization problem.

Although the DRP (7) overcomes knowledge of the exact distribution, it is tricky to solve in general since it contains decision variables in the continuous probability measure space. It is desirable to derive an (approximated) solution of (7) based on the finite samples \mathbf{w}^i , ensuring both feasibility and performance guarantees.

It is important to note that most of the data-driven DRO literature focuses on providing data-driven optimizations that are equivalent to the original problem in convex optimization [17], [19], [20]. In the field of stochastic optimization, many guarantees, such as those for approximating expectations from data, are also typically established when the function within the expectation is convex [28].

However, since the robustness functions of STL are generally non-convex, it is necessary to develop an equivalent optimization approach that does not rely on this convexity assumption. Inspired by [17] and results developed in [29], the following theorem offers a data-driven equivalent solution to the DRP (7) for the Wasserstein ambiguity set, defined in (8), and with a finite number of samples that eliminates the decision variables in the probability measure space. We can guarantee that any feasible solution obtained from the proposed optimization is a feasible solution to the DRP (7) and, with a predefined confidence level, is a feasible solution to the main ECP (6). Moreover, the obtained control input minimizes an upper bound on the objective in (6) with high confidence.

Theorem 3. (Data-driven solution of the DRP) Consider the following optimization:

$$\inf_{\mathbf{u} \in \mathcal{U}, \lambda_1, \lambda_2, y_1^i, y_2^i} \lambda_1 r + \frac{1}{M} \sum_{i=1}^M y_1^i, \quad (11a)$$

$$\text{s.t. } \sup_{\mathbf{w} \in \mathbb{W}} [J(\mathbf{u}, \mathbf{w}) - \lambda_1 \|\mathbf{w} - \mathbf{w}^i\|] \leq y_1^i, \forall i \leq M, \quad (11b)$$

$$\sup_{\mathbf{w} \in \mathbb{W}} [-\varrho_0^\varphi(\mathbf{u}, \mathbf{w}) - \lambda_2 \|\mathbf{w} - \mathbf{w}^i\|] \leq y_2^i, \forall i \leq M, \quad (11c)$$

$$\lambda_1 \geq 0, \quad \lambda_2 \geq 0, \quad (11d)$$

$$-\lambda_2 r - \frac{1}{M} \sum_{i=1}^M y_2^i - L_\varphi h^{-1}(\varepsilon) \geq r_0, \quad (11e)$$

with the optimal solution and value of the objective function denoted by $\hat{\mathbf{u}}$ and \hat{J} , respectively. Based on the Wasserstein ambiguity set \mathbb{Q} defined in (8)-(10), we have:

- **Relation to the DRP (7):** Optimization (11) is equivalent to (7).
- **Relation to the ECP (6):** By choosing a proper Wasserstein radius r , the following statements hold with the probability of at least $1 - \beta$ for a user-specified confidence level $\beta \in (0, 1)$:
 - 1) Any feasible solution to (11) is a feasible solution to (6).
 - 2) The cost function in (6), evaluated for the input $\hat{\mathbf{u}}$, is upper-bounded by \hat{J} . More specifically, the following out-of-sample performance guarantee holds:

$$\mathbb{E}_P [J(\hat{\mathbf{u}}, \mathbf{w})] \leq \hat{J}. \quad (12)$$

Proof. The proof is given in the appendix of [27]. ■

Note that the term *data-driven*, as used in e.g., [17], [19], [21], refers to optimization that utilizes a set of samples to construct the DRP. In the control literature, this term has also been used in contexts where the system matrices are unknown. However, this is beyond the scope of our paper, which assumes that the system matrices are known in advance.

Remark 1. The min and max operators utilized in defining robust semantic ϱ_0^φ in Section II-B are not smooth. Numerical solvers commonly encounter difficulties when these operators appear in the objective function or constraints. Inspired by [30], we opt for smooth under-approximations for these operators, as follows:

$$\min([a_1, \dots, a_m]^\top) \approx -\frac{1}{C} \log \left(\sum_{i=1}^m \exp(-Ca_i) \right),$$

$$\max([a_1, \dots, a_m]^\top) \approx \frac{\sum_{i=1}^m a_i \exp(Ca_i)}{\sum_{i=1}^m \exp(Ca_i)},$$

where C is a positive constant. It is noteworthy that these approximations under-approximate the exact min and max operators. Consequently, the robust semantics derived from these approximations are not greater than the original robust semantics. Hence, fulfilling the approximated robust semantics ensures the satisfaction of the original semantics directly. Additionally, as demonstrated in [30], for a sufficiently large

C , the approximated robust semantics converge to the original semantics with the exact min and max operators.

Note that in transitioning from the CCP (3) to the ECP (6), we can assume that the robustness function ϱ_0^φ is evaluated using the exact min and max operators, ensuring the validity of the Lipschitz constant L_φ as obtained from Theorem 1. We then substitute the expectation of the exact robustness with the expectation of the under-approximated robustness in the constraint of (6). Therefore, the results of the paper, and particularly the Lipschitz constant L_φ , remain valid when employing these under-approximations of the exact min and max operators.

IV. CASE STUDY

We consider the following two-dimensional stochastic dynamics:

$$x_{k+1} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} x_k + \begin{bmatrix} 0.5 \\ 1 \end{bmatrix} u_k + w_k,$$

where $u_k \in \mathcal{U} := [-1, 1]$ and $w_k \in \mathcal{W} = \mathbb{R}^2$ with Gaussian distribution with unknown mean and covariance. Using (2) and assuming $x_0 = [-8, 0]^\top$, the aim is to satisfy a safety constraint $[0, 1]x_k \leq 0.75$ for the whole bounded horizon $0 \leq k \leq N = 15$ and reaching the region $x_k^\top T x_k \leq 1$, with $T = \text{diag}(\frac{1}{4}, \frac{1}{25})$, sometime at $k \in [0, 15]$ with probability at least 0.9 while optimizing the following quadratic cost:

$$J(\mathbf{u}, \mathbf{w}) = 10x_N^\top x_N + \sum_{k=0}^{N-1} (10x_k^\top x_k + u_k^2),$$

where x_k is obtained from (2). The STL formula, described above, can be expressed as $\varphi = \diamond_{[0,15]} \pi_1 \wedge \square_{[0,15]} \pi_2$, where π_1 and π_2 are predicates with corresponding predicate functions $\alpha_1(x) = 1 - x^\top T x$ and $\alpha_2(x) = 0.75 - [0, 1]x$. The robustness function ϱ_0^φ can be written as follows:

$$\varrho_0^\varphi(\mathbf{u}, \mathbf{w}) = \min \left\{ \max_{k \in \{0, \dots, 15\}} \alpha_1(x_k), \min_{k \in \{0, \dots, 15\}} \alpha_2(x_k) \right\}.$$

As explained in Remark 1, we have chosen the smoothing constant C as 100, 10, and 10 for the inner minimization, the maximization, and the outer minimization in ϱ_0^φ , respectively.

Figure 1 shows the system trajectories using the proposed DRP approach for the Wasserstein radius $r = 10^{-3}$. As it can be seen, the trajectories have greater distance with the bound $x_2 = 0.75$ compared to the sample averaging method and the STL specification is satisfied for all trajectories. We have employed 10 times more sampling for the sample averaging method compared to the DRP method.

V. CONCLUSIONS

We have shown how to optimize control sequences for stochastic linear systems to satisfy signal temporal logic (STL) specifications probabilistically when the underlying predicate functions are Lipschitz continuous, and the disturbance distribution is unknown but attains a concentration of measure property. These assumptions allow us to reformulate the control problem as a chance-constrained program (CCP) and present an efficient two-step solution. First, leveraging the




Fig. 1. System trajectories for different realizations for deterministic system (black), ECP solution using the sample average approximation (green), and the proposed DRP solution (blue).

concentration of measure property, we transform the CCP into an expectation-based optimization problem. To account for unknown distributions, we proceed to the second step, where we tackle a distributionally robust optimization problem, which considers all distributions around the empirical one using an ambiguity set based on the Wasserstein metric. In the future, we plan to extend the method to multi-agent systems and nonlinear dynamics.

REFERENCES

- [1] S. Haesaert, P. Nilsson, C. Vasile, R. Thakker, A. Agha-mohammadi, A. Ames, and R. Murray, "Temporal logic control of POMDPs via label-based stochastic simulation relations," *IFAC-PapersOnLine*, vol. 51, no. 16, pp. 271–276, 2018.
- [2] A. Donzé and O. Maler, "Robust satisfaction of temporal logic over real-valued signals," in *International Conference on Formal Modeling and Analysis of Timed Systems*. Springer, 2010, pp. 92–106.
- [3] O. Maler and D. Nickovic, "Monitoring temporal properties of continuous signals," in *International Symposium on Formal Techniques in Real-Time and Fault-Tolerant Systems*. Springer, 2004, pp. 152–166.
- [4] S. S. Farahani, R. Majumdar, V. S. Prabhu, and S. E. Z. Soudjani, "Shrinking horizon model predictive control with chance-constrained signal temporal logic specifications," in *2017 American Control Conference (ACC)*. IEEE, 2017, pp. 1740–1746.
- [5] L. Lindemann, G. J. Pappas, and D. V. Dimarogonas, "Control barrier functions for nonholonomic systems under risk signal temporal logic specifications," in *Proceedings of the IEEE Conference on Decision and Control*, 2020, pp. 1422–1428.
- [6] D. Sadigh and A. Kapoor, "Safe control under uncertainty with probabilistic signal temporal logic," in *Proceedings of Robotics: Science and Systems*, jun 2016.
- [7] 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.
- [8] 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.
- [9] G. Scher, S. Sadraddini, and H. Kress-Gazit, "Robustness-based synthesis for stochastic systems under signal temporal logic tasks," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 1269–1275.

- [10] 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.
- [11] M. C. Campi and S. Garatti, "A Sampling-and-Discarding Approach to Chance-Constrained Optimization: Feasibility and Optimality," *Journal of Optimization Theory and Applications*, vol. 148, no. 2, pp. 257–280, 2011.
- [12] G. C. Calafiore, "Random convex programs," *SIAM Journal on Optimization*, vol. 20, no. 6, pp. 3427–3464, 2010.
- [13] S. Garatti and M. C. Campi, "Non-convex scenario optimization," *Mathematical Programming*, pp. 1–52, 2024.
- [14] E. Delage and Y. Ye, "Distributionally robust optimization under moment uncertainty with application to data-driven problems," *Operations research*, vol. 58, no. 3, pp. 595–612, 2010.
- [15] Z. Hu and L. J. Hong, "Kullback-Leibler divergence constrained distributionally robust optimization," *Available at Optimization Online*, pp. 1695–1724, 2013.
- [16] G. Pflug and D. Wozabal, "Ambiguity in portfolio selection," *Quantitative Finance*, vol. 7, no. 4, pp. 435–442, 2007.
- [17] P. Mohajerin Esfahani and D. Kuhn, "Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations," *Mathematical Programming*, vol. 171, no. 1, pp. 115–166, 2018.
- [18] A. B. Kordabad, R. Wisniewski, and S. Gros, "Safe reinforcement learning using Wasserstein distributionally robust MPC and chance constraint," *IEEE Access*, vol. 10, pp. 130 058–130 067, 2022.
- [19] Z. Chen, D. Kuhn, and W. Wiesemann, "Data-driven chance constrained programs over Wasserstein balls," *Operations Research*, vol. 72, no. 1, pp. 410–424, 2024.
- [20] W. Xie, "On distributionally robust chance constrained programs with Wasserstein distance," *Mathematical Programming*, vol. 186, no. 1, pp. 115–155, 2021.
- [21] A. R. Hota, A. Cherukuri, and J. Lygeros, "Data-driven chance constrained optimization under Wasserstein ambiguity sets," in *2019 American Control Conference (ACC)*. IEEE, 2019, pp. 1501–1506.
- [22] S. Soudjani and R. Majumdar, "Concentration of measure for chance-constrained optimization," *IFAC-PapersOnLine*, vol. 51, no. 16, pp. 277–282, 2018, 6th IFAC Conference on Analysis and Design of Hybrid Systems ADHS 2018.
- [23] O. Maler and D. Nickovic, "Monitoring temporal properties of continuous signals," in *FORMATS/FTRTFT*, ser. LNCS 3253. Springer, 2004.
- [24] X. Jin, A. Donzé, J. V. Deshmukh, and S. A. Seshia, "Mining requirements from closed-loop control models," *IEEE Trans. on CAD of Integrated Circuits and Systems*, vol. 34, no. 11, pp. 1704–1717, 2015.
- [25] A. Barvinok, "Measure concentration in optimization," *Mathematical Programming*, vol. 79, pp. 33–53, 1997.
- [26] G. Pisier, *The volume of convex bodies and Banach space geometry*. Cambridge University Press, 1999, vol. 94.
- [27] A. B. Kordabad, E. E. Vlahakis, L. Lindemann, D. V. Dimarogonas, and S. Soudjani, "Distributionally robust control for chance-constrained signal temporal logic specifications," 2024. [Online]. Available: <https://arxiv.org/abs/2409.03855>
- [28] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on stochastic programming: modeling and theory*. SIAM, 2021.
- [29] R. Gao and A. Kleywegt, "Distributionally robust stochastic optimization with Wasserstein distance," *Mathematics of Operations Research*, vol. 48, no. 2, pp. 603–655, 2023.
- [30] Y. Gilpin, V. Kurtz, and H. Lin, "A smooth robustness measure of signal temporal logic for symbolic control," *IEEE Control Systems Letters*, vol. 5, no. 1, pp. 241–246, 2020.