Shrinking Horizon Model Predictive Control with Signal Temporal Logic Constraints under Stochastic Disturbances

Samira S. Farahani*, Rupak Majumdar*, Vinayak Prabhu*, Sadegh Esmaeil Zadeh Soudjani*

Abstract—We present Shrinking Horizon Model Predictive Control (SHMPC) for discrete-time linear systems with Signal Temporal Logic (STL) specification constraints under stochastic disturbances. The control objective is to maximize an optimization function under the restriction that a given STL specification is satisfied with high probability against stochastic uncertainties. We formulate a general solution, which does not require precise knowledge of the probability distributions of the (possibly dependent) stochastic disturbances; only the bounded support intervals of the density functions and moment intervals are used. For the specific case of disturbances that are independent and normally distributed, we optimize the controllers further by utilizing knowledge of the disturbance probability distributions. We show that in both cases, the control law can be obtained by solving optimization problems with linear constraints at each step. We experimentally demonstrate effectiveness of this approach by synthesizing a controller for an HVAC system.

I. INTRODUCTION

We consider the control synthesis problem for stochastic discrete-time linear systems under path constraints that are expressed as temporal logic specifications and are written in signal temporal logic (STL) [23]. Our aim is to obtain a controller that robustly satisfies desired temporal properties with high probability despite stochastic disturbances, while optimizing additional control objectives. With focus on temporal properties defined on a finite path segment, we use the model predictive control (MPC) scheme [3], [8], [20], [22] with a *shrinking horizon*: the horizon window is fixed and not shifted at each time step of the controller synthesis problem. We start with an initial prediction horizon dependent on the temporal logic constraints, compute the optimal control sequence for the horizon, apply the first step, observe the system evolution under the stochastic disturbance, and repeat the process (decreasing the prediction horizon by 1) until the end of the original time horizon.

Our proposed setting requires solving three technical challenges in the MPC framework.

First, in addition to optimizing control and state cost, the derived controller must ensure that the system evolution satisfies chance constraints arising from the STL specifications. Previous choices of control actions can impose temporal constraints on the rest of the path. We describe an algorithm that updates the temporal constraints based on previous actions.

Second, for some temporal constraints, we may require that the system satisfies the constraints *robustly*: small changes to the inputs should not invalidate the temporal constraint. To ensure robust satisfaction, we use a quantitative notion of robustness for STL [10]. We augment the control objective to maximize the expected robustness of an STL specification, in addition to minimizing control and state costs, under chance constraints. Unfortunately, the resulting optimization problem is not convex.

As a third contribution, we propose a tractable approximation method for the solution of the optimization problem. We conservatively approximate the chance constraints by linear inequalities. Second, we provide a tractable procedure to compute an upper bound for the expected value of the robustness function under these linear constraints.

Recently receding horizon control with STL constraints has been studied for a variety of domains [12], [24]. In these works, the disturbance is assumed to be deterministic but from a bounded polytope, and the worst-case MPC optimization problem is solved. The control synthesis for deterministic systems with probabilistic STL specifications is studied in [25] but only a fragment of STL is considered in order to obtain a convex optimization problem. Also, the receding horizon control has been applied to deterministic systems in the presence of perception uncertainty [17]. Additionally, chance-constrained MPC has been addressed in [26] for deterministic systems, in which the underlying probability space comes only from the measurement noise. Application of chance-constrained MPC to optimal control of drinking water networks is studied in [14].

In this paper, we assume that the the disturbance signal has an arbitrary probability distribution with bounded domain and that we only know the support and the first moment interval for each component of the disturbance signal. In order to solve the optimization problem more efficiently, we transform chance constraints into their equivalent (or approximate) linear constraints. To this end, we apply the technique presented by [4], to approximate the chance constraints with an upper bound. Also, the expected value of the robustness function can be approximated by the moment intervals of the disturbance signal, and can be computed without using numerical integration.

Furthermore, as an additional case in this study, we show that if the disturbance signal is normally distributed and hence, has no bounded support, instead of truncating the distribution to obtain a finite interval of support for random variables, we can use a different approach, which is based on the quantiles of the normally distributed random variables to replace the chance constraints by linear constraints. In this case, we also show that the expected value of the robustness function can be replaced by an upper bound using the methods presented in [13].

We empirically demonstrate the effectiveness of our approach by synthesizing a controller for a Heating, Ventilation

and Air Conditioning (HVAC) system. We compare our approach with open-loop optimal controller synthesis and with robust MPC [24], and show that our approach can lead to significant energy savings.

A. Notations

We use $\mathbb{R}$ for the set of reals and $\mathbb{N} := \{0,1,2,\ldots\}$ for the set of non-negative integers. The set $\mathbb{B} := \{\top, \bot\}$ indicates logical true and false. For a vector $v \in \mathbb{R}^s$, its components are denoted by v_k, $k \in \{1,\ldots,s\}$. For a sequence $\{v(t) \in \mathbb{R}^s, t \in \mathbb{N}\}$ and $t_1 < t_2$, we define $\tilde{v}(t_1 : t_2) := [v(t_1), v(t_1+1), \ldots, v(t_2-1)]$, In this paper, all random variables are denoted by capital letters and the deterministic variables are denoted by small letters. We also use small letter y to indicate observations of a random vector Y. For a sequence of random vectors $\{Y(t) \in \mathbb{R}^s, t \in \mathbb{N}\}$ and $t_1 \leq t < t_2$, we define $\bar{Y}(t_1 : t : t_2) := [y(t_1), y(t_1+1), \ldots, y(t), Y(t+1), \ldots, Y(t_2-1)]$, which is a matrix containing observations of the random variable up to time t augmented with its unobserved values after t. For a random variable $Y(t)$ denote the support interval by $I_{Y(t)}$ and the first moment[1] by $\mathbb{E}[Y(t)]$.

We consider operations on intervals according to interval arithmetic: for two arbitrary intervals $[a,b]$ and $[c,d]$, and constants $\lambda, \gamma \in \mathbb{R}$, we have $[a,b] + [c,d] = [a+c, b+d]$ and $\lambda \cdot [a,b] + \gamma = [\min(\lambda a, \lambda b) + \gamma, \max(\lambda a, \lambda b) + \gamma]$.

II. DISCRETE-TIME STOCHASTIC LINEAR SYSTEMS

In this paper, we consider systems in discrete-time that can be modeled by linear difference equations perturbed by stochastic disturbances. Depending on the probability distribution of the disturbance signal, we conduct our study for two cases: a) the disturbance signal has an arbitrary probability distribution with a bounded domain for which we only know the support and their first moment intervals; and b) the disturbance signal has a normal distribution. The first case can be extended to random variables with an unbounded support, such as normal or exponential random variables, by truncating their distributions. The specific form of the distribution in the second case enables us to perform a more precise analysis using properties of the normal distribution. Note that the support of a random variable X with values in $\mathbb{R}^n$ is defined as the set $\{x \in \mathbb{R}^N \,|\, \Pr_X[\mathscr{B}(x,r)] > 0, \forall r > 0\}$, where $\mathscr{B}(x,r)$ denotes the ball with center at x and radius r; alternatively, the support can be defined as the smallest closed set C such that $\Pr_X[C] = 1$.

Consider a (time-variant) discrete-time stochastic system modeled by the difference equation

$$X(t+1) = A(t)X(t) + B(t)u(t) + W(t), \ X(0) = x_0, \quad (1)$$

where $X(t) \in \mathbb{R}^n$ denotes the state of the system at time instant t, $u(t) \in \mathbb{R}^m$ denotes the control input at time instant t, and $W(t) \in \mathbb{R}^s$ is a vector of random variables, the components of which have either of the above mentioned probability distributions. The random vector $W(t)$ can be interpreted

as the process noise or an adversarial disturbance. Matrices $A(\cdot) \in \mathbb{R}^{n \times n}$, and $B(\cdot) \in \mathbb{R}^{n \times m}$ are possibly time-dependent appropriately defined system's matrices, and the initial state $X(0)$ is assumed to be known. We assume that $W(0),\ldots,W(t)$ are mutually independent random vectors for all time instants t. Note that, for any $t \in \mathbb{N}$, the state-space model (1) provides the following explicit form for $X(\tau)$, $\tau \geq t$, as a function of $X(t)$, input $u(\cdot)$, and the process noise $W(\cdot)$:

$$X(\tau) = \Phi(\tau,t)X(t) + \sum_{k=t}^{\tau-1} \Phi(\tau,k+1)\left(B(k)u(k) + W(k)\right), \quad (2)$$

where $\Phi(\cdot,\cdot)$ is the *state transition matrix* of model (1) defined as

$$\Phi(\tau,t) = \begin{cases} A(\tau-1)A(\tau-2)\ldots A(t) & \tau > t \geq 0 \\ \mathbb{I}_n & \tau = t \geq 0, \end{cases}$$

with $\mathbb{I}_n$ being the identity matrix.

For a fixed positive integer N, and a given time instant $t \in \mathbb{N}$, let $\tilde{u}(t : N) = [u(t), u(t+1), \ldots, u(N-1)]$ (matrix $\tilde{W}(t : N)$ is defined similarly). A *finite stochastic run* of system (1) for the time interval $[t : N]$ is defined as $\Xi(t : N) = X(t)X(t+1)\ldots X(N)$, which is a finite sequence of states satisfying (2). Since each state $X(\tau)$ depends on $X(t), \tilde{u}(t:N)$, and $\tilde{W}(t:N)$, we can rewrite $\Xi(t : N)$ in a more elaborative notation as $\Xi_N(X(t), \tilde{u}(t:N), \tilde{W}(t:N))$. Analogously, we define an *infinite stochastic run* $\Xi = X(t)X(t+1)X(t+2)\ldots$ as an infinite sequence of states. Stochastic runs will be used in Section III to define the system's specifications.

III. SIGNAL TEMPORAL LOGIC

An infinite run of system (1) can be considered as a signal $\xi = x(0)x(1)x(2)\ldots$, which is a sequence of observed states. We consider Signal temporal logic (STL) formulas with bounded-time temporal operators defined recursively according to the grammar [23]

$$\varphi ::= \top \,|\, \pi \,|\, \neg\varphi \,|\, \varphi \wedge \psi \,|\, \varphi \mathscr{U}_{[a,b]} \psi$$

where $\top$ is the *true* predicate; π is a predicate whose truth value is determined by the sign of a function, i.e. $\pi = \{\alpha(x) \geq 0\}$ with $\alpha : \mathbb{R}^n \to \mathbb{R}$ being an affine function of state variables; ψ is an STL formula; $\neg$ and $\wedge$ indicate negation and conjunction of formulas; and $\mathscr{U}_{[a,b]}$ is the *until* operator with $a,b \in \mathbb{R}_{\geq 0}$. A run ξ satisfies φ at time t, denoted by $(\xi,t) \models \varphi$, if the sequence $x(t)x(t+1)\ldots$ satisfies φ. Accordingly, ξ satisfies φ, if $(\xi,0) \models \varphi$.

Semantics of STL formulas are defined as follows. Every run satisfies $\top$. The run ξ satisfies $\neg\varphi$ if it does not satisfy φ; it satisfies $\varphi \wedge \psi$ if both φ and ψ hold. For a run $\xi = x(0)x(1)x(2)\ldots$ and a predicate $\pi = \{\alpha(x) \geq 0\}$, we have $(\xi,t) \models \pi$ if $\alpha(x(t)) \geq 0$. Finally, $(\xi,t) \models \varphi \mathscr{U}_{[a,b]} \psi$ if φ holds at every time step starting from time t before ψ holds, and additionally ψ holds at some time instant between $a+t$ and $b+t$. Additionally, we derive the other standard operators as follows. *Disjunction* $\varphi \vee \psi := \neg(\neg\varphi \wedge \neg\psi)$, the *eventually* operator as $\Diamond_{[a,b]} \varphi := \top \mathscr{U}_{[a,b]} \varphi$, and the *always* operator as $\Box_{[a,b]} \varphi := \neg\Diamond_{[a,b]} \neg\varphi$.

---

[1] The expected value of a random variable X with support $\mathscr{D}$ and the cumulative distribution function F is defined as $\mathbb{E}[X] = \int_{\mathscr{D}} x \, dF(x)$. The expectation exists if the integral is well-defined and yields a finite value.

Thus $(\xi,t) \models \Diamond_{[a,b]} \varphi$ if φ holds at some time instant between $a+t$ and $b+t$ and $(\xi,t) \models \Box_{[a,b]} \varphi$ if φ holds at every time instant between $a+t$ and $b+t$.

Formula Horizon. The *horizon* of an STL formula φ is the smallest $n \in \mathbb{N}$ such that the following holds for all signals $\xi = x(0)x(1)x(2)\ldots$ and $\xi' = x'(0)x'(1)x'(2)\ldots$:

$$\text{If } x(t+i) = x'(t+i) \text{ for all } i \in \{0,\ldots,n\}$$
$$\text{Then } (\xi,t) \models \varphi \text{ iff } (\xi',t) \models \varphi.$$

Thus, in order to determine whether a signal ξ satisfies an STL formula φ, we can restrict our attention to the signal prefix $x(0),\ldots,x(\Delta)$ where Δ is the horizon of φ. This horizon can be upper-approximated by a bound, denoted by $\text{len}(\varphi)$, defined to be the maximum over the sums of all nested upper bounds on the temporal operators. Formally, $\text{len}(\varphi)$ is defined recursively as:

$$\varphi := \top \Rightarrow \text{len}(\varphi) = 0, \qquad \varphi := \pi \Rightarrow \text{len}(\varphi) = 0,$$
$$\varphi := \neg\varphi_1 \Rightarrow \text{len}(\varphi) = \text{len}(\varphi_1),$$
$$\varphi := \varphi_1 \wedge \varphi_2 \Rightarrow \text{len}(\varphi) = \max(\text{len}(\varphi_1),\text{len}(\varphi_2)),$$
$$\varphi := \varphi_1 \, \mathscr{U}_{[a,b]} \, \varphi_2 \Rightarrow \text{len}(\varphi) = b + \max(\text{len}(\varphi_1),\text{len}(\varphi_2)),$$

where φ_1, φ_2 and ψ are STL formulas. For example, for $\varphi = \Box_{[0,4]}\Diamond_{[3,6]} \pi$, we have $\text{len}(\varphi) = 4+6 = 10$. For a given STL formula φ, it is possible to verify that $\xi \models \varphi$ using only the finite run $x(0)x(1)\ldots x(N)$, where N is equal to $\text{len}(\varphi)$.

STL Robustness. In contrast to the above Boolean semantics, the quantitative semantics of STL [18] assigns to each formula φ a real-valued function ρ^φ of signal ξ and t such that $\rho^\varphi(\xi,t) > 0$ implies $(\xi,t) \models \varphi$. Robustness of a formula φ with respect to a run ξ at time t is defined recursively as

$$\rho^\top(\xi,t) = +\infty,$$
$$\rho^\pi(\xi,t) = \alpha(x(t)) \text{ with } \pi = \{\alpha(x) \geq 0\},$$
$$\rho^{\neg\varphi}(\xi,t) = -\rho^\varphi(\xi,t)$$
$$\rho^{\varphi \wedge \psi}(\xi,t) = \min(\rho^\varphi(\xi,t),\rho^\psi(\xi,t)),$$
$$\rho^{\varphi \, \mathscr{U}_{[a,b]} \, \psi}(\xi,t) = \max_{i \in [a,b]} \left(\min(\rho^\psi(\xi,t+i), \min_{j \in [0,i)} \rho^\varphi(\xi,t+j)) \right),$$

where $x(t)$ refers to signal ξ at time t. The robustness of the derived formula $\Diamond_{[a,b]} \varphi$ can be worked out to be $\rho^{\Diamond_{[a,b]} \varphi}(\xi,t) = \max_{i \in [a,b]} \rho^\varphi(\xi,t+i)$; and similarly for $\Box_{[a,b]} \varphi$ as $\rho^{\Box_{[a,b]} \varphi}(\xi,t) = \min_{i \in [a,b]} \rho^\varphi(\xi,t+i)$. The robustness of an arbitrary STL formula is computed recursively on the structure of the formula according to the above definition, by propagating the values of the functions associated with each operand using min and max operators.

STL Robustness for Stochastic Runs. With focus on stochastic runs $\Xi = X(0)X(1)X(2)\ldots$ and using the bound of a formula φ, the finite stochastic run $\Xi(t : t+N) = X(t)X(1)\ldots X(t+N)$ with $N = \text{len}(\varphi)$ is sufficient to study probabilistic properties of $(\Xi,t) \models \varphi$. Analogous to the definition of robustness for deterministic run, we can define *stochastic robustness* $\rho^\varphi(\Xi,t)$ of a formula φ with respect to the run Ξ for times t with the stochastic robustness being dependent on $\Xi(t : t+N)$ and φ.

Note that a general STL formula φ consists of several Boolean and/or temporal operators. Due to the system dynamics (1), the stochastic run $\Xi(t : t+N)$ and $\rho^\varphi(\Xi(t : t+N),t)$ are both functions of $\tilde{W}(t : t+N)$. Therefore, $\rho^\varphi(\Xi(t : t+N),t)$ is a random variable since affine operators, maximization and minimization are measurable functions.

The above definition of robustness implies that, for any formula φ and constant $\delta \in (0,1)$, a stochastic run $\Xi = X(0)X(1)X(2)\ldots$ satisfies φ with probability greater than or equal to $1 - \delta$, if the finite stochastic run $\Xi(0 : N) = X(0)X(1)\ldots X(N)$ with $N \geq \text{len}(\varphi)$ satisfies $\Pr[\rho^\varphi(\Xi(0 : N),0) > 0] \geq 1 - \delta$.

IV. CONTROL PROBLEM STATEMENT

For system (1) with a given initial state $X(0) = x_0$, the stochastic disturbance vector $W(t)$ with a given probability distribution, STL formulas φ and ψ, and some constant $N \geq \max(\text{len}(\varphi),\text{len}(\psi))$, the control problem can be defined as finding an optimal input sequence $\tilde{u}^*(0 : N) = [u^*(0),\ldots,u^*(N-1)]$, that minimizes the expected value of a given objective function $J(\tilde{X}(0 : N+1),\tilde{u}(0 : N))$ subject to constraints on states and input variables, where $\tilde{X}(0 : N+1) = [X(0),X(1),\ldots,X(N)]$. This optimization problem for the time interval $0 \leq t < N$ can be defined as

$$\min_{\tilde{u}(0:N)} \mathbb{E}\left[J(\tilde{X}(0 : N+1),\tilde{u}(0 : N)) \right] \quad \text{s.t.} \tag{3a}$$

$$X(t) = \Phi(t,0)x_0 + \sum_{k=0}^{t-1} \Phi(t,k+1)(B(k)u(k) + W(k)), \tag{3b}$$

$$\Pr\left[\Xi_N(x_0,\tilde{u}(0 : N),\tilde{W}(0 : N)) \models \varphi \right] \geq 1 - \delta, \tag{3c}$$

$$\tilde{u}(0 : N) \in U^N, \tag{3d}$$

where $\mathbb{E}[\cdot]$ denotes the expected value operator and the closed set $U^N \in \mathbb{R}^{mN}$ specifies the constraint set for the input variables. The chance constraints (3c) state that for a given $\delta \in (0,1)$, stochastic runs of the system should satisfy φ with a probability greater than or equal to $1 - \delta$. We consider the following objective function

$$J(\tilde{X}(0 : N+1),\tilde{u}(0 : N)) := J_{\text{robust}}(\tilde{X}(0 : N+1)) + J_{\text{in}}(\tilde{u}(0 : N)), \tag{4}$$

where the first term $J_{\text{robust}}(\tilde{X}(0 : N+1)) := -\rho^\psi(\tilde{X}(0 : N+1),0)$ represents the negative value of the robustness function on STL formula ψ at time 0 that needs to be minimized; and the second term $J_{\text{in}}(\tilde{u}(0 : N))$ reflects the cost on the input variables and can be defined as a linear or a quadratic function.

Note that optimization problem (3) is an open-loop optimization problem and we cannot incorporate any information related to the process noise or the states of the system.

Remark 1: The above problem formulation enables us to distinguish the following two cases: we put the robustness of a formula in the objective function if the system is required to be robust with respect to satisfying the formula; we encode the formula in the probabilistic constraint if only satisfaction of the formula is important.

A. Model Predictive Control

To obtain a more well-behaved control input and to include the information about the disturbances, instead of solving the optimization problem (3), we apply *shrinking horizon model predictive control* (SHMPC), which can be summarized as follows: at time step one, we obtain a sequence of control inputs with length N (the prediction horizon) to optimize the cost function; then we only use the first component of the obtained control sequence to update the state of the system (or to observe the state in the case of having a stochastic disturbance); in the next time step, we fix the first component of the control sequence by the first component of the previously calculated optimal control sequence and hence, we only optimize for a control sequence of length $N-1$. As such, at each time step, the size of the control sequence decreases by 1. Note that in this approach, unlike the receding horizon approach, we do not shift the horizon at each time step and the end point of the prediction window is fixed. MPC allows us to incorporate the new information we obtain about the state variables and the disturbance signal, at each time step and hence, to improve the control performance comparing with the one of solving the open-loop optimization problem (3).

A natural choice for the prediction horizon N in this setting with STL constraints is to set it to be greater than or equal to the bound of the formula φ, i.e., $\text{len}(\varphi)$, which was defined in the previous section. This choice provides a conservative maximum trajectory length required to make a decision about the satisfiability of the formula.

Let $\bar{X}(0:t:N+1) = [x(0),\dots,x(t),X(t+1),\dots,X(N)]$ where $x(0),\dots,x(t)$ are the observed states up to time t and $X(\tau)$ is the random state variable at time $\tau > t$, also let $\bar{W}(0:t-1:N) = [w(0),\dots,w(t-1),W(t),W(t+1),\dots,W(N-1)]$ such that $w(0),\dots,w(t-1)$ are the noise realizations up to time $t-1$ and $W(\tau)$ are random vectors with given probability distributions at time $\tau \geq t$. Define $\bar{u}(0:t-1:N) = [u^*(0),\dots,u^*(t-1),u(t),\dots,u(N-1)]$ to be the vector of input variables such that $u^*(0),\dots,u^*(t-1)$ are the obtained optimal control inputs up to time $t-1$ and $u(t),\dots,u(N-1)$ are the input variables that need to be determined at time $t \geq 0$.

Given STL formula φ, observations of state variables $x(0),x(1),\dots,x(t)$, and designed control inputs $u^*(0),\dots,u^*(t-1)$ of system (1), the *stochastic SHMPC* optimization problem minimizes the expected value of the cost function

$$J(\bar{X}(0:t:N+1),\bar{u}(0:t-1:N)) = $$
$$J_{\text{robust}}(\bar{X}(0:t:N+1)) + J_{\text{in}}(\bar{u}(0:t-1:N)),$$

at each time instant $0 \leq t < N$, as follows

$$\min_{\bar{u}(t:N)} \mathbb{E}\left[J(\bar{X}(0:t:N+1),\bar{u}(0:t-1:N))\right] \quad \text{s.t.} \quad (5a)$$

$$X(\tau) = \Phi(\tau,t)x(t) + \sum_{k=t}^{\tau-1} \Phi(\tau,k+1)(B(k)u(k)+W(k)),$$

$$\text{for } t \leq \tau \leq N \quad (5b)$$

$$\Pr\left[\Xi_N(x_0,\bar{u}(0:t-1:N),\bar{W}(0:t-1:N)) \models \varphi\right] \geq 1-\delta_t \quad (5c)$$

$$\bar{u}(t:N) \in U^{N-t}, \quad (5d)$$

where the expected value $\mathbb{E}[\cdot]$ in (5a) is conditioned on observations $\tilde{X}(0:t+1) = [x(0),\dots,x(t)]$ and $\delta_t = \delta/N$ for all t. Optimization variables in (5) are the control inputs $\tilde{u}(t:N) = [u(t),\dots,u(N-1)]$. We indicate the argument of minimum by $\tilde{u}_{opt}(t:N) = [u_{opt}(t),\dots,u_{opt}(N-1)]$.

The complete procedure of obtaining an optimal control sequence using SHMPC is presented in Algorithm 1. Lines 3 to 8 of this algorithm specify the inputs and the parameters used in the algorithm and line 20 specifies the output. In line 10, the SHMPC optimization procedure starts for each time step $t \in [0,N-1]$. In line 11, we solve the optimization problem (5) to obtain an optimal control sequence for time instance t. In lines 12 to 16, we check whether the obtained solution satisfies the STL specifications or not; if yes, assign the first component of the obtained input sequence to $u^*(t)$, and if not, the optimization procedure will be terminated. Finally, in line 17, we apply $u^*(t)$ to the system (1) and observe the states at time instant t.

Algorithm 1

1: **procedure** CHANCE-CONSTRAINED STOCHASTIC SHMPC

2: *input*:

3:     STL formulas φ and ψ and a fixed $\delta \in (0,1)$

4: *parameters*:

5:     $N \geq \max(\text{len}(\varphi),\text{len}(\psi))$

6:     probability distributions of the process noise $\{W(t), t = 0,\dots,N-1\}$

7:     Initial state x_0

8:     $\delta_t = \delta/N$ for $t = 0,\dots,N-1$

9: *closed-loop optimization problem*:

10:     **for** $t = 0; t < N; t = t+1$ **do**

11:         Compute $\tilde{u}_{opt}(t:N) = [u_{opt}(t),\dots,u_{opt}(N-1)]$ by solving the optimization problem (5)

12:         **if** the solution of optimization problem (5) exists **then**

13:             Let $u^*(t) := u_{opt}(t)$;

14:         **else**

15:             Return *Infeasible Solution* and terminate the optimization procedure;

16:         **end if**

17:         Apply $u^*(t)$ to the system and observe the value of $X(t+1)$ as $x(t+1)$

18:     **end for**

19: *output*:

20:     $u^*(0:N) = [u^*(0),\dots,u^*(N-1)]$

21: **end procedure**

We show in the following theorem that in Algorithm 1, by using the shrinking horizon technique, the specific choice of δ_t, and keeping track of the control inputs and observed states, the closed-loop system satisfies the STL specification φ with probability greater than or equal to $1-\delta$.

Theorem 2: Given a constant $\delta \in (0,1)$ and an STL formula φ, if the optimization problems in Algorithm 1 are all feasible, the computed optimal control sequence $\tilde{u}^*(0:N) = [u^*(0),\dots,u^*(N-1)]$ ensures that the closed-loop satisfy φ with probability greater than or equal to $1-\delta$.

Proof: Considering the chance constraint (5c) the probability that a trajectory of the system violates φ at time step t is at most δ_t. This implies that the probability of violating φ in the time interval $t = 0, \ldots, N-1$ is at most $\sum_{t=0}^{N-1} \delta_t = \sum_{t=0}^{N-1} \delta/N = \delta$, which proves that the optimal control sequence $\tilde{u}^* = [u^*(0), \ldots, u^*(N-1)]$ obtained using Algorithm 1 results in trajectories that satisfy φ with probability greater than or equal to $1 - \delta$. $\blacksquare$

Note that in practice, if at each time step a feasible solution is not found, by using the previous control value, i.e., by setting $u^*(t) = u^*(t-1)$, we can give the controller a chance to retry in the next time step after observing the next state.

Remark 3: The choice of $\delta_t = \delta/N$ is completely arbitrary. In general, the positive constants δ_t can be picked freely with the condition that $\sum_{t=0}^{N-1} \delta_t = \delta$.

Computation of the solution of the optimization problem (5) requires addressing two main challenges: a) the objective function (5a) depends on the optimization variables $\tilde{u}(t:N)$ and on random variables $\tilde{W}(t:N)$, thus we have to compute the expected value as a function of these variables; and b) the feasible set of the optimization restricted by the chance constraint (5c) is in general difficult to characterize. We propose approximation methods in Sections V and VI to respectively address these two challenges.

V. APPROXIMATING THE OBJECTIVE FUNCTION

To solve the optimization problem (5), one needs to calculate the expected value of the objective function. One way to do this is via numerical integration methods [7]. However, numerical integration is in general both cumbersome and time-consuming. For example, the method of approximating the density function of the disturbance with piecewise polynomial functions defined on polyhedral sets [5], [19] suffers from scalability issues on top of the induced approximation error. Therefore, in this section, we discuss an efficient method that computes an upper bound for the expected value of the objective function and then, minimize this upper bound instead.

We discuss computation of such upper bounds for both cases of process noise with arbitrary probability distribution and with normal distribution in Sections V-A and V-B, respectively. For this purpose, we first provide a canonical form for the robustness function of a STL formula ψ, which is the min-max or max-min of random variables. This result is inspired by [9], in which the authors provide such canonical forms for max-min-plus-scaling functions.

Theorem 4: For a given STL formula ψ, the robustness function $\rho^\psi(\Xi(0:N), 0)$, and hence the function $J_{\text{robust}}(\bar{X}(0:t:N))$, can be written into a max-min canonical form

$$J_{\text{robust}}(\bar{X}(0:t:N)) = \max_{i \in \{1,\ldots,L\}} \min_{j \in \{1,\ldots,m_i\}} \left\{ \eta_{ij} + \lambda_{ij} \bar{W}(0:t:N) \right\},$$
$$(6)$$

and into a min-max canonical form

$$J_{\text{robust}}(\bar{X}(0:t:N)) = \min_{i \in \{1,\ldots,K\}} \max_{j \in \{1,\ldots,n_i\}} \left\{ \zeta_{ij} + \gamma_{ij} \bar{W}(0:t:N) \right\},$$
$$(7)$$

for some integers $K, L, n_1, \ldots, n_K, m_1, \ldots, m_L$, where λ_{ij} and γ_{ij} are weighting vectors and η_{ij} and ζ_{ij} are affine functions of $\bar{u}(0:t:N)$ and x_0.

Proof: The proof is inductive on the structure of ψ and uses the explicit form of the states in (2) utilizing the identities $-\max(f_1, f_2) = \min(-f_1, -f_2)$ and

$$\min\left(\max(f_1, f_2), \max(g_1, g_2)\right) =$$
$$\max\left(\min(f_1, g_1), \min(f_1, g_2), \min(f_2, g_1), \min(f_2, g_2)\right).$$

for functions f_1, f_2, g_1, and g_2. $\blacksquare$

A. Arbitrary probability distributions with bounded support

Suppose the elements of the stochastic vector $W(t)$, i.e., $W_k(t)$, $k \in \{1, \ldots, n\}$ have arbitrary probability distribution with known bounded support $I_{W_k(t)} = [a_k, b_k]$ and its first moment $\mathbb{E}[W_k(t)]$ belongs to the interval $\mathbb{M}_{W_k(t)} = [c_k, d_k]$, with known quantities $a_k, b_k, c_k, d_k \in \mathbb{R}$. Under this assumption, the explicit form of $X(\cdot)$ in (2) implies that, for the observed value of $X(t)$ as $x(t)$, $X(\tau)$ is a random variable with the following interval of support and the first moment interval

$$I_{X(\tau)} = [\bar{a}_\tau + \bar{C}_\tau, \bar{b}_\tau + \bar{C}_\tau], \ \ \mathbb{M}_{X(\tau)} = [\bar{c}_\tau + \bar{C}_\tau, \bar{d}_\tau + \bar{C}_\tau] \quad (8)$$

where $\bar{C}_\tau = \Phi(\tau, t)x(t) + \sum_{k=t}^{\tau-1} \Phi(\tau, k+1)B(k)u(k)$, and $\bar{a}_\tau, \bar{b}_\tau, \bar{c}_\tau$ and $\bar{d}_\tau$ are weighted sum of a_k, b_k, c_k, d_k, $k \in \mathbb{N}$, obtained by using interval arithmetics mentioned in Section I-A.

The objective function in (5) can be written as $\mathbb{E}[J_{\text{robust}}(\bar{X}(0:t:N+1))] + J_{\text{in}}(\bar{u}(0:t-1:N)))$ and that $J_{\text{robust}}(\bar{X}(0:t:N+1)) = -\rho^\psi(\bar{X}(0:t:N+1), 0)$. Recall that $\bar{X}(0:t:N+1) = [x(0), \ldots, x(t), X(t+1), \ldots, X(N)]$ with observed states $x(0), \ldots, x(t)$ of system (1) and random states $X(\tau)$, $\tau > t$. The following theorem shows how we can compute an upper bound for $\mathbb{E}[J_{\text{robust}}(\bar{X}(0:t:N+1))]$ based on the canonical form of J_{robust}.

Theorem 5: For a given STL formula ψ, $\mathbb{E}[J_{\text{robust}}(\bar{X}(0:t:N+1))]$ can be upper bounded by

$$\max_{i \in \{1,\ldots,L\}} \min_{j \in \{1,\ldots,m_i\}} (\hat{d}_{ij} + \eta_{ij}),$$

where the constants η_{ij}, $i \in \{1, \ldots, L\}, j \in \{1, \ldots, m_i\}$, are affine functions of $\bar{u}(0:t-1:N)$ and $x(0)$, and $\hat{d}_{ij}$ are weighted sum of $w(0), \ldots, w(t-1)$ and c_k, d_k for $k = t, \ldots, N-1$.

Proof: With focus on the canonical form (6), let $Y_{ij} = \eta_{ij} + \lambda_{ij} \bar{W}(0:t:N)$. Considering the support and moment interval of the components of $W(\tau), \tau = t, \ldots, N-1$, each random variable Y_{ij} has the following support and moment interval (similar to (8))

$$I_{Y_{ij}} = [\hat{a}_{ij} + \eta_{ij}, \hat{b}_{ij} + \eta_{ij}], \ \ \mathbb{M}_{Y_{ij}} = [\hat{c}_{ij} + \eta_{ij}, \hat{d}_{ij} + \eta_{ij}] \quad (9)$$

where the constants $\hat{a}_{ij}, \hat{b}_{ij}, \hat{c}_{ij}, \hat{d}_{ij}$, $i \in \{1, \ldots, L\}, j \in \{1, \ldots, m_i\}$, are weighted sum of $w(0), \ldots, w(t-1)$ and a_k, b_k, c_k, d_k for $k = t, \ldots, N-1$, using interval arithmetic (cf. Section I-A). Accordingly, J_{robust} is a random variable with the following support and moment intervals,

$$I_{J_{\text{robust}}} =$$

$$\mathbb{M}_{J_{\text{robust}}} = \begin{array}{l} \left[\displaystyle\max_{i\in\{1,\dots,L\}}\min_{j\in\{1,\dots,m_i\}}(\hat{a}_{ij}+\eta_{ij}), \max_{i\in\{1,\dots,L\}}\min_{j\in\{1,\dots,m_i\}}(\hat{b}_{ij}+\eta_{ij}\right] \\ \\ \left[\displaystyle\max_{i\in\{1,\dots,L\}}\min_{j\in\{1,\dots,m_i\}}(\hat{c}_{ij}+\eta_{ij}), \max_{i\in\{1,\dots,L\}}\min_{j\in\{1,\dots,m_i\}}(\hat{d}_{ij}+\eta_{ij})\right]. \end{array}$$
$$(10)$$

Hence, as we are minimizing the cost function in (5), we can utilize the upper bound $\max_{i\in\{1,\dots,L\}}\min_{j\in\{1,\dots,m_i\}}(\hat{d}_{ij}+\eta_{ij})$ for $\mathbb{E}[J_{\text{robust}}(\bar{X}(0:t:N+1))]$. ∎

Note that the approximation methodology of Theorem 5 is applicable also to the min-max canonical form (7).

By replacing the expected objective function by its upper bound given in Theorem 5, and by replacing the probabilistic constraints by their equivalent linear approximation (as is discussed in Section VI), the optimization problem (5) can be then recast as a mixed integer linear programming (MILP) problem, which can be solved using the available MILP solvers [2], [21].

B. Normal distribution

The upper bound on the objective function provided in the previous section does not apply to process noises with unbounded support, but knowing the distribution of the process noise provides more information about the statistics of the runs of the system. In this section we address process noises with normal distribution separately due the their wide use in engineering applications.

Suppose that for any $t\in\mathbb{N}$, $W(t)$ is normally distributed with mean $\mathbb{E}[W(t)]=0$ and covariance matrix $\Sigma_{W(t)}$, i.e., $W(t)\sim\mathcal{N}(0,\Sigma_{W(t)})$. The explicit form of $X(\tau)$ in (2) and the fact that normal distribution is closed under affine transformations result in normal distribution for $X(\tau)$, $\tau\in\mathbb{N}$. Its expected value and covariance matrix with an observed value $x(t)$ of $X(t)$ are

$$\mu_\tau = \Phi(\tau,t)x(t) + \sum_{k=t}^{\tau-1}\Phi(\tau,k+1)B(k)u(k) \text{ and}$$

$$\Sigma_\tau = \sum_{k=t}^{\tau-1}\Phi(\tau,k+1)\Sigma_{W(k)}\Phi(\tau,k+1)^T,$$

respectively, for $\tau\geq t\geq 0$.

In this section we use the canonical representation of $J_{\text{robust}}(\bar{X}(0:t:N+1))$ in Theorem 4, which states that J_{robust} (for fixed $\bar{u}(0:t:N)$ and x_0) can be written in either of the forms

$$\max_{i\in\{1,\dots,L\}}\min_{j\in\{1,\dots,m_i\}}Y_{ij} \quad\text{or}\quad \min_{i\in\{1,\dots,K\}}\max_{j\in\{1,\dots,n_i\}}Y_{ij} \quad (11)$$

with $Y_{ij}=\eta_{ij}+\lambda_{ij}\bar{W}(0:t-1:N)$ being affine functions of the process noise, thus normally distributed random variables (similar to $X(\tau)$ explained above). With focus on these canonical representations for J_{robust} we employ Proposition 6 to show how to approximate $\mathbb{E}[J_{\text{robust}}]$ using higher order moments of $W(t)\sim\mathcal{N}(0,\Sigma)$. This proposition, also used in [13], follows due to the relation between the infinity norm and the p-norm of a vector and Jensen's inequality.

Proposition 6: Consider random variables Z_i for $i\in\{1,\dots,s\}$ and let p be an even integer. Then

$$\mathbb{E}[\max(Z_1,\dots,Z_s)]\leq\mathbb{E}[\max(|Z_1|,\dots,|Z_s|)]$$

$$\leq\mathbb{E}\left[((Z_1)^p+\dots+(Z_s)^p)^{1/p}\right]$$

$$\leq\left(\sum_{i=1}^s\mathbb{E}[(Z_i)^p]\right)^{1/p}.$$

Founded on Proposition 6, next theorem shows how we can upper bound $\mathbb{E}[J_{\text{robust}}]$ using the higher order moments of Y_{ij}.

Theorem 7: Considering the canonical forms in (11) for J_{robust} as a function of random variables Y_{ij}, $\mathbb{E}[J_{\text{robust}}]$ can be upper bounded by

$$\mathbb{E}\left[\max_{i\in\{1,\dots,L\}}\min_{j\in\{1,\dots,m_i\}}Y_{ij}\right]\leq\left(\sum_{i=1}^L\sum_{j=1}^{m_i}\mathbb{E}[Y_{ij}^p]\right)^{1/p}, \quad (12)$$

$$\mathbb{E}\left[\min_{i\in\{1,\dots,K\}}\max_{j\in\{1,\dots,n_i\}}Y_{ij}\right]\leq\min_{i\in\{1,\dots,K\}}\left(\sum_{j=1}^{n_i}\mathbb{E}\left[Y_{ij}^p\right]\right)^{1/p}. \quad (13)$$

Proof: For random variables Y_{ij}, $i\in\{1,\dots,L\}$, $j\in\{1,\dots,m_i\}$, and for a positive even integer p, the following inequality holds,

$$\mathbb{E}\left[\max_{i\in\{1,\dots,L\}}\min_{j\in\{1,\dots,m_i\}}Y_{ij}\right]\overset{(i)}{\leq}\left(\sum_{i=1}^L\mathbb{E}\left[\min_{j\in\{1,\dots,m_i\}}Y_{ij}\right]^p\right)^{1/p}$$

$$\overset{(ii)}{=}\left(\sum_{i=1}^L\mathbb{E}\left[-\max_{j\in\{1,\dots,m_i\}}-Y_{ij}\right]^p\right)^{1/p}$$

$$\overset{(iii)}{\leq}\left(\sum_{i=1}^L\sum_{j=1}^{m_i}\mathbb{E}[Y_{ij}^p]\right)^{1/p},$$

where in (i) we used the upper bound obtained in Proposition 6; in (ii) we used the fact that $\min_{k\in\{1,\dots,r\}}(\alpha_k)=-\max_{k\in\{1,\dots,r\}}(-\alpha_k)$; In (iii) we use again the inequality in Proposition 6. Moreover, for $i\in\{1,\dots,K\}$, $j\in\{1,\dots,n_i\}$, the following inequality holds,

$$\mathbb{E}\left[\min_{i\in\{1,\dots,K\}}\max_{j\in\{1,\dots,n_i\}}Y_{ij}\right]\overset{(i)}{\leq}\min_{i\in\{1,\dots,K\}}\mathbb{E}\left[\max_{j\in\{1,\dots,n_i\}}Y_{ij}\right]$$

$$\overset{(ii)}{\leq}\min_{i\in\{1,\dots,K\}}\left(\sum_{j=1}^{n_i}\mathbb{E}\left[Y_{ij}^p\right]\right)^{1/p},$$

where we apply Jensen's inequality to the concave function $\min(\cdot)$ to get (i). The inequality of Proposition 6 gives (ii). ∎

Note that random variables Y_{ij} are normally distributed in both (12) and (13). Higher order moments of normally distributed random variables can be computed analytically in a closed form as a function of the first two moments, i.e., using its mean and variance. More specifically, for a normally distributed random variable Z with mean μ and variance σ^2, the p-th moment has a closed form as

$$\mathbb{E}[Z^p]=\sigma^p i^{-p}H_p(i\mu/\sigma) \quad (14)$$

where i is the imaginary unit and

$$H_p(z)=p!\sum_{l=0}^{p/2}\frac{(-1)^l z^{p-2l}}{2^l l!(p-2l)!} \quad (15)$$

is the p-th Hermite polynomial [1, Chapter 22 and 26]. We use (14) to compute higher order moments of normal random

variables with p being even integers. Note that the right-hand side of (14) is in fact real because $H_p(z)$ contains only even powers of z when p is even.

In the next section we discuss how to cope with the second challenge of characterizing the feasible set of the optimization restricted by the chance constraint (5c).

VI. UNDER APPROXIMATION OF CHANCE CONSTRAINTS

In this section, we discuss methods for computing conservative lower approximations of the chance constraints in (5c) as linear constraints. For the sake of compact notation, we indicate the stochastic run $\Xi(0:N) = X(0)X(1)\ldots X(N)$ only by Ξ_N without declaring its dependency on the state, input, and disturbance variables. Recall the chance constraint (5c) as $\Pr[(\Xi_N,t) \models \varphi] \geq 1 - \delta_t$. In order to transform this constraint to linear inequalities, we first show in the following theorem, that this constraint can be transformed into similar inequalities but φ being an atomic predicate. Then in Sections VI-A and VI-B, we discuss how to transform the resulting constraints with atomic predicates into linear inequalities for the cases of arbitrary random variables with known bounded support and moment interval and of normally distributed random variables.

Theorem 8: for any formula φ and a constant $\vartheta \in (0,1)$, constraints of the forms

$$\Pr[(\Xi_N,t) \models \varphi] \geq \vartheta \text{ and } \Pr[(\Xi_N,t) \models \varphi] \leq \vartheta \quad (16)$$

can be transformed into similar constraints with φ being an atomic predicate using the structure of φ.

Proof: The proof is inductive on the structure of the formula φ as discussed in the following three cases.

Case I: $\varphi = \neg\varphi_1$ we have the following equivalences

$$\Pr[(\Xi_N,t) \models \neg\varphi_1] \geq \vartheta \Leftrightarrow \Pr[(\Xi_N,t) \nvDash \varphi_1] \geq \vartheta$$
$$\Leftrightarrow \Pr[(\Xi_N,t) \models \varphi_1] \leq 1 - \vartheta,$$
$$\Pr[(\Xi_N,t) \models \neg\varphi_1] \leq \vartheta \Leftrightarrow \Pr[(\Xi_N,t) \nvDash \varphi_1] \leq \vartheta$$
$$\Leftrightarrow \Pr[(\Xi_N,t) \models \varphi_1] \geq 1 - \vartheta.$$

Case II: $\varphi = \varphi_1 \wedge \varphi_2$ we obtain the following inequalities by using the fact that for possibly joint events $\mathscr{A}$ and $\mathscr{B}$, it holds that $\Pr[\mathscr{A} \wedge \mathscr{B}] \geq \vartheta \Leftrightarrow \Pr[\neg\mathscr{A} \vee \neg\mathscr{B}] \leq 1 - \vartheta$ and $\Pr(\mathscr{A} \vee \mathscr{B}) \leq \Pr[\mathscr{A}] + \Pr[\mathscr{B}]$.

$$\Pr[(\Xi_N,t) \models \varphi_1 \wedge \varphi_2] \geq \vartheta$$
$$\Leftrightarrow \Pr[(\Xi_N,t) \models \varphi_1 \wedge (\Xi_N,t) \models \varphi_2] \geq \vartheta$$
$$\Leftrightarrow \Pr[(\Xi_N,t) \nvDash \varphi_1 \vee (\Xi_N,t) \nvDash \varphi_2] \leq 1 - \vartheta$$
$$\Leftarrow \Pr[(\Xi_N,t) \nvDash \varphi_1] + \Pr[(\Xi_N,t) \nvDash \varphi_2] \leq 1 - \vartheta$$
$$\Leftarrow \Pr[(\Xi_N,t) \nvDash \varphi_i] \leq \frac{1-\vartheta}{2} \quad i = 1,2. \quad (17)$$

Note that in the last line of (17), we assume that the probability of the two events are upper bounded by the same value, i.e., $(1-\vartheta)/2$. However, this can be replaced by any two other probabilities δ_1 and δ_2 such that $\delta_1 + \delta_2 = 1 - \vartheta$. Now consider the second possibility:

$$\Pr[(\Xi_N,t) \models \varphi_1 \wedge \varphi_2] \leq \vartheta$$
$$\Leftrightarrow \Pr[(\Xi_N,t) \models \neg\varphi_1 \vee \neg\varphi_2] \geq 1 - \vartheta$$

$$\Leftrightarrow \Pr[(\Xi_N,t) \models \neg\varphi_1 \vee (\varphi_1 \wedge \neg\varphi_2)] \geq 1 - \vartheta$$
$$\Leftrightarrow \Pr[(\Xi_N,t) \models \neg\varphi_1] + \Pr[(\Xi_N,t) \models \varphi_1 \wedge \neg\varphi_2] \geq 1 - \vartheta, \quad (18)$$

where the last line of (18) is due to the fact that the events are disjoint. Assuming that the probabilities of these two events are lower bounded by the same values, i.e., $(1-\vartheta)/2$, we have the inequalities

$$\Pr[(\Xi_N,t) \models \neg\varphi_1] \geq (1-\vartheta)/2, \quad (19)$$

$$\Pr[(\Xi_N,t) \models \varphi_1 \wedge \neg\varphi_2] \geq \frac{1-\vartheta}{2}, \quad (20)$$

which are in the form of inequalities discussed previously. Note that Equations (17) to (19) discuss the case of having conjunction of two STL formulas. The results can be easily extended to conjunction of n STL formulas by replacing $(1-\vartheta)/2$ with $(1-\vartheta)/n$.

Case III: $\varphi = \varphi_1 \mathscr{U}_{[a,b]} \varphi_2$ The satisfaction $(\Xi_N,t) \models \varphi_1 \mathscr{U}_{[a,b]} \varphi_2$ is equivalent to $\bigvee_{j=t+a}^{t+b} \psi_j$ with disjoint events

$$\psi_j = \bigwedge_{i=t}^{t+a-1} (\Xi_N,i) \models \varphi_1 \bigwedge_{i=a+t}^{j-1} (\Xi_N,i) \models (\varphi_1 \wedge \neg\varphi_2) \wedge (\Xi_N,j) \models \varphi_2.$$

Thus $\Pr[(\Xi_N,t) \models \varphi_1 \mathscr{U}_{[a,b]} \varphi_2] \geq \vartheta$ is equivalent to $\sum_{j=t+a}^{t+b} \Pr[\psi_j] \geq \vartheta$. Assuming the probabilities of events are lower bounded by the same values, we have $\Pr[\psi_j] \geq \vartheta/(b-a+1)$ for $j = a+t,\ldots,b+t$, which again can be reduced as in Case II.

The second possible probabilistic constraint in Case III can be obtained as

$$\Pr[(\Xi_N,t) \models \varphi_1 \mathscr{U}_{[a,b]} \varphi_2] \leq \vartheta \Leftrightarrow \Pr\left[\bigvee_{j=a+t}^{b+t} \psi_j\right] \leq \vartheta$$
$$\Leftrightarrow \sum_{j=t+a}^{t+b} \Pr[\psi_j] \leq \vartheta$$
$$\Leftrightarrow \Pr[\psi_j] \geq \vartheta/(b-a+1), \quad (21)$$

which can be again reduced as in Case II. Here also, we used the fact that ψ_j consists of disjoint events and we assume that he probabilities of events are lower bounded by the same value, i.e., by $\vartheta/(b-a+1)$, for $j = a+t,\ldots,b+t$. ∎

So far we have shown how to inductively reduce the chance constraint (5c) to inequalities of the form (16) with atomic predicates. In the rest of this section we discuss their corresponding linear inequalities for the two types of probability distributions considered in this paper.

A. Arbitrary probability distributions with bounded support

To transform the chance constraints into linear constraints in the case of having random variables with arbitrary probability distributions, we apply an approximation method based on the upper bound proposed by [4]. Let $Z_1,\ldots,Z_n$ be random variables with interval of bounded support $[a_i,b_i]$ and let $\mathbb{E}[Z_1],\ldots,\mathbb{E}[Z_n]$ denote their expected values belonging to the moment intervals $\mathbb{M}_i$ for $i = 1,\ldots,n$. Define $Z = \sum_{i=1}^n Z_i$ and

$\mathbb{E}(Z) = \sum_{i=1}^{n} \mathbb{E}[Z_i]$. Using Chernoff-Hoeffding inequality, the following upper bound exists for any $\varsigma \geq 0$ [16]

$$\Pr[Z - \mathbb{E}[Z] \leq -\varsigma] \leq \exp\left(\frac{-\varsigma^2}{\nu \sum_{i=1}^{n}(b_i - a_i)^2}\right). \quad (22)$$

where $\nu > 0$ is a constant. If $Z_1, \ldots, Z_n$ are dependent, then the inequality applies with a constant $\nu = \chi(\hat{G})/2$, where $\hat{G}$ denotes the indirected dependency graph of $Z_1, \ldots, Z_n$ and $\chi(\hat{G})$ is the chromatic number of the graph $\hat{G}$ defined as the minimum number of colors required to color $\hat{G}$. For the independent case, $\chi(\hat{G}) = 1$. The expression for the right tail probability is derived identically. For more details, the reader is referred to [4].

Consider the chance constraints (16) with $\varphi = \{\alpha \geq 0\}$. Since α is an affine function of random state variables, it is a random variable itself with the following interval of support and moment interval

$$\begin{aligned} I_{\alpha(X(t))} &= [\tilde{a}_t + \tilde{C}_t, \tilde{b}_t + \tilde{C}_t] \\ \mathbb{M}_{\alpha(X(t))} &= [\tilde{c}_t + \tilde{C}_t, \tilde{d}_t + \tilde{C}_t] \end{aligned} \quad (23)$$

where for $t = 0, \ldots, N$, we have $\tilde{a}_t, \tilde{b}_t, \tilde{c}_t$ and $\tilde{d}_t$ are weighted sum of $\bar{a}_t, \bar{b}_t, \bar{c}_t, \bar{d}_t$ related to the interval of support and moment interval of random variables $X(t)$ (cf. (8)), and $\tilde{C}_t$ is a linear expression of input variables.

Let $\varsigma = \mathbb{E}[\alpha(X(t))]$; we can directly use (22) as

$$\begin{aligned} \Pr[(\Xi_N, t) \models \pi^\alpha] &\geq 1 - \delta_t \Leftrightarrow \Pr[\alpha(X(t)) > 0] \geq 1 - \delta_t \\ &\Leftrightarrow \Pr[\alpha(X(t)) \leq 0] \leq \delta_t \\ &\Leftarrow \exp\left(\frac{-\varsigma^2}{\nu \sum_{t=1}^{N}(\tilde{b}_t - \tilde{a}_t)^2}\right) \leq \delta_t \\ &\Leftrightarrow \frac{-\varsigma^2}{\nu \sum_{t=1}^{N}(\tilde{b}_t - \tilde{a}_t)^2} \leq \log(\delta_t) \\ &\Leftrightarrow -\varsigma^2 \leq \nu \log(\delta_t) \sum_{t=1}^{N}(\tilde{b}_t - \tilde{a}_t)^2 \\ &\Leftarrow \varsigma \geq \sqrt{-\nu \log(\delta_t) \sum_{t=1}^{N}(\tilde{b}_t - \tilde{a}_t)^2} \end{aligned} \quad (24)$$

Note that since $\delta_t \in (0,1)$, we have $\log(\delta_t) < 0$; hence, by multiplying both sides of the inequality by -1 in line 5 of (24), the expression $-\log(\delta_t) \cdot \sum_{t=1}^{N}(\tilde{b}_t - \tilde{a}_t)^2$ becomes a positive number, and hence, its square root is a real number. Note also that the last inequality is due to the fact that $\varsigma \geq 0$. Hence, we can replace ς in the last inequality of (24) by the lower bound of its moment interval in (23), i.e., with $\tilde{c}_t + \tilde{C}_t$, which is a linear expression in the input variables.

Consequently, in this case, the chance constraint in (5) can be replaced by

$$\tilde{c}_t + \tilde{C}_t \geq \sqrt{-\nu \log(\delta_t) \cdot \sum_{t=1}^{N}(\tilde{b}_t - \tilde{a}_t)^2}. \quad (25)$$

For the second type of probabilistic inequality (cf. (16)), we can again use (22) for the right tail probability; hence we have

$$\begin{aligned} \Pr[(\Xi_N, t) \models \pi^\alpha] &\leq 1 - \delta_t \\ &\Leftarrow \Pr[\alpha(X(t)) \geq 0] \leq 1 - \delta_t \end{aligned}$$

$$\Leftarrow \exp\left(\frac{-\varsigma^2}{\nu \sum_{t=1}^{N}(\tilde{b}_t - \tilde{a}_t)^2}\right) \leq 1 - \delta_t, \quad (26)$$

and then following the same steps as in (24), we obtain the same linear expression for the chance constant as in (25) by only replacing δ_t by $1 - \delta_t$ in the related expressions.

B. Normal distribution

To transform the chance constraints into linear constraints in the case of having normally distributed random variables, we use the quantile of the normal distribution. By definition, for a normally distributed random variable x with mean μ and standard deviation σ,

$$\Pr[x \leq b] \leq \delta_t \Leftrightarrow F^{-1}(\delta_t) \geq b \Leftrightarrow \mu + \sigma\phi^{-1}(\delta_t) \geq b \quad (27)$$
$$\Pr[x \leq b] \geq \delta_t \Leftrightarrow F^{-1}(\delta_t) \leq b \Leftrightarrow \mu + \sigma\phi^{-1}(\delta_t) \leq b \quad (28)$$

where F^{-1} denotes the inverse of the cumulative distribution function or the quantile function and ϕ^{-1} is the inverse of the error function of a normally distributed random variable.

Recall the chance constraints (16) with $\varphi = \{\alpha \geq 0\}$. Since α is an affine function of normally distributed state variables, it is also normally distributed with appropriately defined mean μ_t and variance σ_t^2. Hence, we can directly use (27) and (28) as

$$\begin{aligned} \Pr[(\Xi_N, t) \models \pi^\alpha] \geq 1 - \delta_t &\Leftrightarrow \Pr[\alpha(X(t)) > 0] \geq 1 - \delta_t \\ &\Leftrightarrow \Pr[\alpha(X(t)) \leq 0] \leq \delta_t \\ &\Leftrightarrow F^{-1}(\delta_t) \geq 0 \\ &\Leftrightarrow \mu_t + \sigma_t\phi^{-1}(\delta_t) \geq 0, \quad (29) \end{aligned}$$

$$\begin{aligned} \Pr[(\Xi_N, t) \models \pi^\alpha] \leq 1 - \delta_t &\Leftrightarrow \Pr[\alpha(X(t)) > 0] \leq 1 - \delta_t \\ &\Leftrightarrow \Pr[\alpha(X(t)) \leq 0] \geq \delta_t \\ &\Leftrightarrow F^{-1}(\delta_t) \leq 0 \\ &\Leftrightarrow \mu_t + \sigma_t\phi^{-1}(\delta_t) \leq 0. \quad (30) \end{aligned}$$

Therefore, the chance constraint can be replaced by the equivalent linear constraint (29) or (30), depending on the type of constraint we have.

VII. EXPERIMENTAL RESULTS

We now apply our controller synthesis approach to the room temperature control in a building. The details of the thermal model can be found in [15], [24], and is briefly explained here for clarity. The temperature of room r is denoted by T_r and the wall and the temperature of the wall between the room and its surrounding j (e.g. other rooms) are denoted by w_j and T_{w_j}, respectively. Dynamics of the temperature of wall w_j and room r can be written as [15]

$$C_j^w \frac{dT_{w_j}}{dt} = \sum_{k \in \mathcal{N}_{w_j}} \frac{T_{r,k} - T_{w_j}}{R_{j,k}} + r_j \alpha_j A_{w_j} Q_{\text{rad}j} \quad (31)$$

$$\begin{aligned} C_j^r \frac{dT_r}{dt} = &\sum_{k \in \mathcal{N}_r} \frac{T_k - T_r}{R_k} + \dot{m}_r c_a (T_s - T_r) \\ &+ w_i \tau_w A_{\text{win}} Q_{\text{rad}} + \dot{Q}_{\text{int}} \end{aligned} \quad (32)$$

where C_j^w, α_j and A_{w_j} are heat capacity, a radiative heat absorption coefficient, and the area of w_j, respectively. The

total thermal resistance between the centerline of wall j and the side of the wall on which node k is located is denoted by R_{j_k}. The radiative heat flux density on w_j is denoted by $Q_{\mathrm{rad}j}$, the set of all neighboring nodes to w_j is denoted by $\mathcal{N}_{w_j}$, and r_j is a wall identifier, which equals 0 for internal walls and 1 for peripheral walls, where j is the outside node. The temperature, heat capacity and air mass flow into room r are denoted by T_r, C_j^r and $\dot{m}_r$, respectively; c_a is the specific heat capacity of air, and T_s is the temperature of the supply air to room r, w is a window identifier, which equals 0 if none of the walls surrounding room r have windows, and 1 if at least one of them does, τ_w is the transmissivity of the glass of the window in room r, A_{win} is the total area of the windows on walls surrounding room r, Q_{rad} is the radiative heat flux density per unit area radiated to room r, and $\dot{Q}_{\mathrm{int}r}$ is the internal heat generation in room r. Finally, $\mathcal{N}_r$ is the set of neighboring room nodes for room r. Further details on this thermal model can be found in [15].

As such, the heat transfer equations for each wall and room r is in the form of nonlinear differential equation. After linearization and time-discretization, the model of the system becomes in the form of dynamical equation

$$X(t+1) = AX(t) + Bu(t) + W(t),$$

where $X \in \mathbb{R}^n$ is the state vector representing the temperature of the walls and the rooms and $u \in \mathbb{R}^m$ is the input vector representing the air mass flow rate and discharge air temperature of conditioned air into each thermal zone. Matrices A and B are obtained by time discretization of dynamics of the thermal model (31)-(32) with a sampling time of $t_s = 30$ minutes. The disturbance $W(\cdot)$ aggregates various *unmodeled* dynamics and the uncertainty in physical variables such as the outside temperature, internal heat generation and radiative heat flux density. The statistics of $W(\cdot)$ can be estimated using historical data [15].

In this example, we only control the temperature of one room and include the temperature of the neighboring rooms as part of the disturbance signal $W(t)$. We also assume that there is a reference for the disturbance signal, denoted by $w_r(t)$, and the reference is perturbed by independent and identically distributed random vectors $e(t) \sim \mathcal{N}(0, \mathbb{I}_n)$, i.e., the disturbance is $W(t) = w_r(t) + e(t)$, which is normally distributed with mean $\mu_t = w_r(t)$ and identity covariance matrix $\Sigma = \mathbb{I}_n$.

In contrast to [24], which considers deterministic disturbances from a bounded set, we consider stochastic disturbances and we maximize the robustness of satisfaction of the STL specifications in the presence of such disturbance. Accordingly, we handle chance constraints and include the expected value of the robustness function in the objective function.

We consider a signal $\mathrm{occ} : \mathbb{N} \to \{-1, 1\}$ representing the room occupancy; $\mathrm{occ}(t) = 1$ if the room is occupied at time t and $\mathrm{occ}(t) = -1$ otherwise. This signal is assumed to be known for the entire simulation period. The MPC prediction horizon N is chosen to be 24, representing 12 hours monitoring of the room temperature. We select $\delta = 0.1$ so that the obtained control input provides confidence level of 90% on

the satisfaction of the desired behavior. We are interested in keeping the room temperature above a reference temperature T_r when the room is occupied; thus the specification is

$$\psi = \square_{[0,N]} \left(\mathrm{occ}(t) = 1 \rightarrow X(t) > T_r \right).$$

At each time instant $0 \leq t < N$, the optimization problem (5) obtains an optimal control input $\tilde{u}_{opt}(t : N) = [u_{opt}(t), \ldots, u_{opt}(N-1)]$ that minimizes

$$\mathbb{E}[-\rho^\psi(\bar{X}(0 : t : N), 0)] + \sum_{k=t}^{N-1} ||u(k)||_1,$$

where the robustness function is defined as

$$\rho^\psi(\bar{X}(0 : t : N), 0) = \min\{X(\tau) - T_r \mid \tau \in [0, N], \mathrm{occ}(\tau) > 0\}.$$

The chance constraint (5c) is defined with the same specification $\varphi = \psi$. We approximate $\mathbb{E}[-\rho^\psi(\tilde{X}(0 : t : N), 0)]$ using the upper bound (13) and transform the chance constraint (5c) into linear inequalities using the approach of Section VI. We also assume that inputs are bounded, i.e., for each $0 \leq t < N$, we have $0 \leq u(t) \leq 380$.

The simulations are done using MATLAB R2014b on a 2.6 GHz Intel Core i5 processor and the optimizations are solved using *fmincon* solver in MATLAB. we perform $n_s = 200$ simulations in order to check the satisfiability of the STL specifications with a probability greater than or equal to 0.9. Figure 1 shows the results of these 200 simulations. The top plot shows the occupancy signal and the middle plot illustrates the average, minimum, and maximum of the obtained room temperatures over 12 hours as well as the room reference temperature T_r in Fahrenheit. The controller ensures that the room temperature goes above the reference temperature T_r once the occupancy signal is 1 and stays there as long as the room is occupied. The minimum and maximum bounds on the room temperature shows that the specifications have never been violated in these 200 simulations. The bottom plot shows the average, minimum, and maximum of the air flow rate in $\left[\frac{\mathrm{ft}^3}{\mathrm{min}}\right]$, which indicates that the input constraint is not violated.

Note that all these $n_s = 200$ runs result in feasible solutions, which gives a confidence bound on the feasibility of the original problem as follows. Since all the n_s runs of the simulation are feasible, we can claim that the original problem is also feasible with probability at least $(\beta/2)^{1/n_s}$, $\beta \in (0, 1)$, with confidence level $1 - \beta$ [6]; hence, having 200 runs being all feasible, the optimization problem (5) is also feasible with probability 0.98 with confidence level 0.95.

To further illustrate the performance of the proposed method, we compare our SHMPC approach with the robust MPC (RMPC) approach of [24], in which the disturbance belongs to a bounded polyhedral set. Note that RMPC approach is not directly applicable to unbounded uncertainties. Therefore, in the optimization procedure, we truncate a normally distributed disturbance in the 2σ interval such that $e(t) \in [-1, 1]$. Further, we solve the RMPC optimization problem using Monte Carlo sampling.

The total fan energy consumption is proportional to the cubic of the airflow. Table I shows the total fan energy consumption and the computation time for the three approaches.
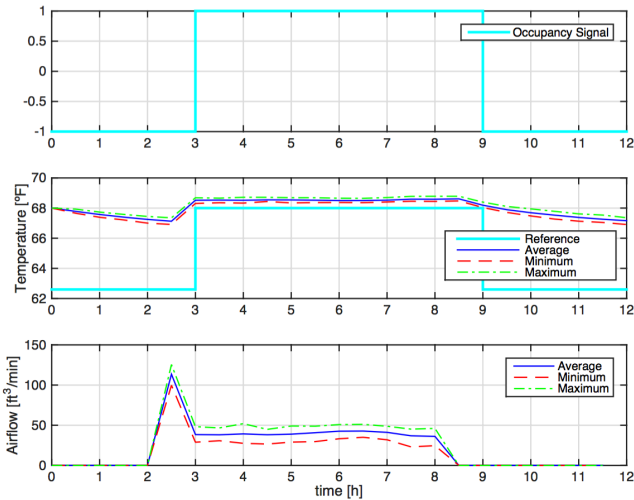
Fig. 1. Controlling the room temperature using SHMPC in the presence of normally distributed disturbance and STL constraints.

TABLE I
COMPARISON OF THE STATISTICS OF THE FAN ENERGY CONSUMPTION USING DIFFERENT CONTROL APPROACHES.

Computational Methods	Fan energy consumption [kWh]	Average computation time [s]
Open-loop OC	1337.016	3.9277
RMPC	$\mu_1 = 12.2216$, $\sigma_1 = 0.045\mu_1$	33.4891
SHMPC	$\mu_2 = 2.5101$, $\sigma_2 = 0.104\mu_2$	19.3622

For RMPC and SHMPC, we report the average and standard deviation of total energy consumption using the sum of cubes of the optimal input sequences corresponding to the 200 simulations. Also, for these two approaches, we report the average computation time over the 200 simulations. Comparing statistics of these two approaches is essential because of the chance constraints in SHMPC and the Monte Carlo sampling based optimization in RMPC. The energy consumption using open-loop optimal control (OC) is very high, comparing to both RMPC and SHMPC. This is due to the fact that the open-loop strategy computes the solution of optimization problem (5) only once and hence, the computation time is smaller compared to the two other methods. As a result, the input sequence has an aggressive behavior to make sure that it reacts in time to the changes happening in the system. Since RMPC is more conservative compared to SHMPC, the average energy consumption is much higher for the RMPC controller compared to the SHMPC controller: the SHMPC controller achieves a 80% reduction of total energy consumption on average compared to RMPC.

VIII. CONCLUSIONS

In this paper, we presented shrinking horizon model predictive control (SHMPC) for stochastic discrete-time linear systems with signal temporal logic (STL) specifications. Our aim was to obtain an optimal control sequence that guarantees the satisfaction of STL specifications with a probability greater than a certain level. By assumption, the stochastic disturbance signal had an arbitrary probability distribution with a bounded support and the only available information related to this distribution is the intervals of support and the moment intervals of each component of the disturbance signal. Using an existing approximation technique, we showed that the chance constraints could be approximated by an upper bound, which resulted in having approximate linear constraints for the chance constraints. Moreover, in the case of having the state costs and/or the robustness function related to the degree of satisfaction of the specifications by the state trajectory, their expected value can be also approximated using the moment intervals of components of the disturbance signal. As an additional case, we further considered disturbances that are normally distributed and we showed that the chance constraints in this case can be replaced by the quantile expressions which are linear in the input variables. In the end, in an example, we applied the proposed method to control a HVAC system.

REFERENCES

[1] M. A. Abramowitz and I. Stegun. *Handbook of Mathematical Functions*. National Bureau of Standards, US Government Printing Office, Washington DC, 1964.

[2] A. Atamtürk and M. W. P. Savelsbergh. Integer-programming software systems. *Annals of Operations Research*, 140(1):67–124, November 2005.

[3] A. Bemporad, W. P. M. H. Heemels, and B. De Schutter. On hybrid systems and closed-loop MPC systems. *IEEE Transactions on Automatic Control*, 47(5):863–869, May 2002.

[4] O. Bouissou, E. Goubault, S. Putot, A. Chakarov, and S. Sankaranarayanan. Uncertainty propagation using probabilistic affine forms and concentration of measure inequalities. In *Proceedings of Tools and Algorithms for Construction and Analysis of Systems (TACAS)*, volume TBA of *Lecture Notes in Computer Science*, page TBA. Springer-Verlag, 2016.

[5] B. Büeler, A. Enge, and K. Fukuda. Exact volume computation for convex polytopes: A practical study. In G. Kalai and G.M. Ziegler, editors, *Polytopes – Combinatorics and Computation*, pages 131–154. Birkäuser Verlag, Basel, Switzerland, 2000.

[6] C. Clopper and E. S. Pearson. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, 26(4):404–413, 1934.

[7] P. J. Davis and P. Rabinowitz. *Methods of Numerical Integration*. Academic Press, New York, 2nd edition, 1984.

[8] B. De Schutter and T. van den Boom. Model predictive control for max-plus-linear discrete event systems. *Automatica*, 37(7):1049–1056, July 2001.

[9] B. De Schutter and T. J. J. van den Boom. MPC for continuous piecewise-affine systems. *Systems & Control Letters*, 52(3–4):179–192, July 2004.

[10] G. E. Fainekos and G. J. Pappas. Robustness of temporal logic specifications for continuous-time signals. *Theoretical Computer Science*, 410(42):4262–4291, 2009.

[11] S. S. Farahani, R. Majumdar, V.S. Prabhu, and S. Esmaeil Zadeh Soudjani. Shrinking horizon model predictive control with chance-constrained signal temporal logic specifications. In *To appear in the proceedings of American Control Conference (ACC) 2017*.

[12] S. S. Farahani, V. Raman, and R. M. Murray. Robust model predictive control for signal temporal logic synthesis. In *Proceedings of the IFAC Conference on Analysis and Design of Hybrid Systems*, Atlanta, Georgia, October 2015.

[13] S. S. Farahani, T. van den Boom, H. van der Weide, and B. De Schutter. An approximation method for computing the expected value of max-affine expressions. *European Journal of Control*, 27:17–27, January 2016.

[14] J. M. Grosso, C. Ocampo-Martinez, V. Puig, and B. Joseph. Chance-constrained model predictive control for drinking waternetworks. *Journal of process control*, 24:504–516, 2014.

[15] M. Maasoumy Haghighi. *Controlling Energy Efficient Buildings in the Context of Smart Grid: A Cyber Physical System Approach*. PhD thesis, University of California, Berkeley, Berkeley, California, December 2013.

[16] S. Janson. Large deviations for sums of partly dependent random variables. *Random Structures Algorithms*, 24(3):234–248, 2004.

[17] S. Jha and V. Raman. Automated synthesis of safe autonomous vehicle control under perception uncertainty. In *8th International Symposium of NASA Formal Methods*, Lecture Notes in Computer Science, pages 117–132, Minneapolis, Minnesota, June 2016.

[18] X. Jin, A. Donzé, J. V. Deshmukh, and S. A. Seshia. Mining requirements from closed-loop control models. In *Hybrid Systems: Computation and Control, HSCC 2013, April 8-11, 2013, Philadelphia, PA, USA*, pages 43–52, 2013.

[19] J. B. Lasserre. Integration on a convex polytope. *Proceedings of the American Mathematical Society*, 126(8):2433–2441, August 1998.

[20] M. Lazar, M. Heemels, S. Weiland, and A. Bemporad. Stability of hybrid model predictive control. *IEEE Transactions on Automatic Control*, 15(11):1813–1818, 2006.

[21] J. T. Linderoth and T. K. Ralphs. Noncommercial software for mixed-integer linear programming. In J. Karlof, editor, *Reinforcement Learning: State-Of-The-Art*, CRC Press Operations Research Series, pages 253–303. 2005.

[22] J. M. Maciejowski. *Predictive Control with Constraints*. Prentice Hall, Harlow, England, 2002.

[23] O. Maler and D. Nickovic. Monitoring temporal properties of continuous signals. In *In FORMATS/FTRTFT*, pages 152–166, 2004.

[24] V. Raman, A. Donzé, D. Sadigh, R. M. Murray, and S. A. Seshia. Reactive synthesis from signal temporal logic specifications. In *Hybrid Systems: Computation and Control, HSCC 2015, Seattle, WA, USA, April 14-16, 2015*, pages 239–248, 2015.

[25] D. Sadigh and A. Kapoor. Safe control under uncertainty with probabilistic signal temporal logic. In *Proceedings of Robotics: Science and Systems Conference*, AnnArbor, Michigan, June 2016.

[26] A. T. Schwarm and M. Nikolaou. Chance-constrained model predictive control. *AIChE Journal*, 45(8):1743–1752, 1999.

Vinayak Prabhu did his PhD in the department of Electrical Engineering and Computer Sciences at the University of California, Berkeley (USA) in 2008, and is currently a postdoctoral researcher at Max Planck Institute for Software Systems. His research interests are in modelling, analysis, verification, and control, of reactive, real-time, hybrid, and multi-agent systems.

Samira S. Farahani is a postdoctoral researcher at Max Planck Institute for Software Systems. Her research interests are in control synthesis for stochastic hybrid systems, temporal logic and formal methods, and discrete-event systems. Prior to this position, she was a postdoctoral researcher at CMS department, California Institute of Technology, and at Energy and Industry Department, Delft University of Technology. Dr. Farahani has received her PhD degree in Systems and Control and her MSc. degree in Applied Mathematics, both from Delft University of Technology, the Netherlands, in 2012 and 2008, respectively. She obtained her BSc. degree in Applied Mathematics from Sharif University of Technology, Iran, in 2005.

Sadegh Esmaeil Zadeh Soudjani is a postdoctoral researcher at the Max Planck Institute for Software Systems, Germany. His research interests are formal synthesis, abstraction, and verification of complex dynamical systems with application in Cyber-Physical Systems, particularly involving power and energy networks, smart grids, and transportation systems. Sadegh received the B.Sc. degrees in electrical engineering and pure mathematics, and the M.Sc. degree in control engineering from the University of Tehran, Tehran, Iran, in 2007 and 2009, respectively. He received the Ph.D. degree in Systems and Control in November 2014 from the Delft Center for Systems and Control at the Delft University of Technology, Delft, the Netherlands. His PhD thesis received the best thesis award from the Dutch Institute of Systems and Control. Before joining Max Planck Institute, he was a postdoctoral researcher at the department of Computer Science, University of Oxford, United Kingdom.

Rupak Majumdar is a Scientific Director at the Max Planck Institute for Software Systems. His research interests are in the verification and control of reactive, real-time, hybrid, and probabilistic systems, software verification and programming languages, logic, and automata theory. Dr. Majumdar received the President's Gold Medal from IIT, Kanpur, the Leon O. Chua award from UC Berkeley, an NSF CAREER award, a Sloan Foundation Fellowship, an ERC Synergy award, "Most Influential Paper" awards from PLDI and POPL, and several best paper awards (including from SIGBED, EAPLS, and SIGDA). He received the B.Tech. degree in Computer Science from the Indian Institute of Technology at Kanpur and the Ph.D. degree in Computer Science from the University of California at Berkeley.