

Robustly Complete Finite-State Abstractions for Control Synthesis of Stochastic Systems

YIMING MENG ^{id} AND JUN LIU ^{id} (Senior Member, IEEE)

(Formal Verification and Synthesis of Cyber-Physical Systems)

Department of Applied Mathematics, University of Waterloo, Waterloo, ON N2L 3G1, Canada

CORRESPONDING AUTHOR: YIMING MENG (e-mail: yiming.meng@uwaterloo.ca)

This work was supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC), in part by the Canada Research Chairs (CRC) program, and in part by the Ontario Early Researcher Awards (ERA) program.

ABSTRACT The essential step of abstraction-based control synthesis for nonlinear systems to satisfy a given specification is to obtain a finite-state abstraction of the original systems. The complexity of the abstraction is usually the dominating factor that determines the efficiency of the algorithm. For the control synthesis of discrete-time nonlinear stochastic systems modelled by nonlinear stochastic difference equations, recent literature has demonstrated the soundness of abstractions in preserving robust probabilistic satisfaction of ω -regular linear-time properties. However, unnecessary transitions exist within the abstractions, which are difficult to quantify, and the completeness of abstraction-based control synthesis in the stochastic setting remains an open theoretical question. In this article, we address this fundamental question from the topological view of metrizable space of probability measures, and propose constructive finite-state abstractions for control synthesis of probabilistic linear temporal specifications. Such abstractions are both sound and approximately complete. That is, given a concrete discrete-time stochastic system and an arbitrarily small \mathcal{L}^1 -perturbation of this system, there exists a family of finite-state controlled Markov chains that both abstracts the concrete system and is abstracted by the slightly perturbed system. In other words, given an arbitrarily small prescribed precision, an abstraction always exists to decide whether a control strategy exists for the concrete system to satisfy the probabilistic specification.

INDEX TERMS \mathcal{L}^1 -perturbation, abstraction, completeness, control synthesis, decidability, linear-time property, metrizable space of probability measures, robustness, stochastic systems, nonlinear systems.

I. INTRODUCTION

Abstraction-based formal synthesis relies on obtaining a finite-state abstraction (or symbolic model) of the original, and possibly nonlinear systems. Computational methods, such as graph-based model checking and automaton-guided controller synthesis, are then developed based on the abstraction to verify the system or synthesize controllers with respect to a temporal logic specification [1], [2], [3]. Abstractions enable autonomous decision making of physical systems to achieve more complex tasks, and received significant success in the past decade [2], [3], [4], [5], [6]. Regardless of heavy state-space discretization and complicated abstraction analysis, formal methods compute with guarantees a set of

initial states from which a controller exists to realize the given specification [2], [3], [7], [8].

Heuristically, abstractions use a finite-state automaton to solve the corresponding search problem at a cost of potentially including non-deterministic transitions in the automation. For non-stochastic control systems, both sound and approximately complete abstractions exist [3], [7], [8], [9], [10], [11]. This is in the sense that, the abstractions can not only include a sufficient number of transitions to design provably correct controllers, but also quantify the level of over-approximation allowed for a specified precision. Therefore, the completeness analysis theoretically removes the doubts of finding an abstraction-based approach once

a robust control strategy of a certain degree is supposed to exist with respect to a given specification, which makes any computational attempts not meaningless.

There is a recent surge of interest in studying formal methods for stochastic systems. We review some crucial results from the literature that are pertinent to the work presented in this article.

A. RELATED WORK

Probabilistic model checkers have been developed for discrete-time discrete-state fully observed Markov decision processes (MDP) and partially observed Markov decision processes (POMDP) [12], [13], [14], [15], [16], and have gained success in applications of control synthesis with probabilistic temporal logics [17], [18], [19], [20].

For continuous-state space, a major strategy is to approximate the transition kernels by some reference (control) stochastic matrices, known as finite-model approximations, to solve optimal control problem or control synthesis with respect to probabilistic temporal logics [21], [22]. Probabilistic reachability and safety related control synthesis can be resolved by relating the satisfaction probability to the corresponding value functions [23], [24], [25], [26], [27]. By necessarily imposing stability-like conditions, the problem can be reduced to solving the characterized dynamic programming problem using computable bounded-horizon counterparts [28], [29], [30], [31], [32].

For fully observed systems, other than the approximation schemes, a formal abstraction for stochastic systems provides an inclusion of all possible approximate transitions of the labelled processes, which eventually will preserve the probability of satisfaction in a proper sense. Bounded-Parameter Markov Decision Processes (BMDP) can naturally serve this purpose [33], [34]. A BMDP contains a family of finite-state MDPs with uncertain transitions given each action, and provide the upper and lower quasi-stochastic matrices as abstractions for the continuous-state controlled Markov systems. The authors of [34] developed algorithms based on [33], [35] to obtain the upper/lower bound of the satisfaction probability of fundamental formulas of probabilistic computation tree logic. The work in [36] formulated BMDP abstraction for bounded-linear temporal logic specifications. The most recent works [37], [38] developed a specification-guided refinement strategy on the partition of the state space and presented a synthesis procedure for finite-mode stochastic systems against any ω -regular specifications. Furthermore, the work [39] proved the soundness of abstractions for linear periodic event-triggered control systems, but in the sense that expected rewards computed along the abstraction's paths bound that of the concrete system.

The recent research [40] proposed a notion of completeness for stochastic abstractions in verification of probabilistic ω -regular properties. That is, given a concrete discrete-time continuous-state Markov process X , and an arbitrarily small \mathcal{L}^1 -bounded perturbation of this system, there always exists

an IMC abstraction whose interval of satisfaction probability contains that of X , and meanwhile is contained by that of the slightly perturbed system. Instead of imposing the mix-monotone conditions [41] and the strong stability (ergodicity) assumptions [42] of the stochastic systems, the analysis in [40] is based on the topology of metrizable space of probability measures with only mild conditions. This methodology proves to be more effective than simply discussing the value of probabilities and enables us to demonstrate the approximate completeness of abstraction-based stochastic control synthesis.

B. CONTRIBUTIONS

In this article, we establish theoretical results on abstraction-based control of discrete-time nonlinear Markov systems, building on recent work [40] on formal verification. In brief:

- We define abstractions based on the weak topology of probability laws and propose the concept of robust completeness for controlled Markov systems.
- While it is often believed to be true that abstraction-based stochastic control synthesis is sound in the sense of Fact 1 of [38], [41], as well as similar statements in [34], [43], we provide the first formal proof of its soundness, to the best knowledge of the authors.
- We prove that robustly complete abstractions of fully observed controlled Markov systems (even with additional uncertainties) exist under a mild assumption, which demonstrates the decidability¹ of robust realization of probabilistic ω -regular temporal logic formulas.
- We improve upon the analysis in [40] by providing tighter inequalities that avoid unnecessarily refined partitions to guarantee the prescribed precision.
- We discuss the applicability of formal abstraction to partially observed controlled stochastic systems.

We would like to clarify that the results of this article differ from those presented in [40]. Specifically, [40] focuses solely on constructing robustly complete abstractions for control-free systems, which do not address the decidability of controlled Markov systems. In contrast, this current article tackles this issue by introducing an additional concept known as the 'implementation' relation (refer to Definition 4.3) between control policies of the original system and the abstraction. The construction of abstractions for controlled Markov systems in this article exhibits subtle but fundamental differences. We hence include explicit proofs in order to provide clearer guidelines and demonstrate our approach.

The rest of the article is organized as follows. Section II presents some preliminaries on probability spaces and controlled Markov systems. Section III provides the problem definition. Section IV presents the soundness of abstractions in verifying ω -regular linear-time properties for fully observed discrete-time controlled Markov systems. Section V

¹In this context, the decidability of an abstraction is to decide whether a control strategy exists for the concrete system to satisfy the probabilistic specification.

presents the constructive robust abstractions with soundness and approximate completeness guarantees. We discuss the applicability of the proposed method for partially observed discrete-time controlled Markov systems in Section VI. The article is concluded in Section VII.

C. CONVENTIONS FOR NOTATION

We denote by \prod the product of ordinary sets, spaces, or function values. Denote by \otimes the product of collections of sets, or sigma algebras, or measures. The n -times repeated product of any kind is denoted by $(\cdot)^n$ for simplification. Denote by $\pi_j : \prod_{i=0}^{\infty} (\cdot)_i \rightarrow (\cdot)_j$ the projection to the j th component. We denote the Borel σ -algebra of a set by $\mathcal{B}(\cdot)$ and the space of all probability measures on $\mathcal{B}(\cdot)$ by $\mathfrak{P}(\cdot)$.

For a set $A \subseteq \mathbb{R}^n$, \bar{A} denotes its closure, $\text{Int}(A)$ denotes its interior, and ∂A denotes its boundary. For two sets $A, B \subseteq \mathbb{R}^n$, the set difference is defined by $A \setminus B = \{x : x \in A, x \notin B\}$. Let $|\cdot|$ denote the infinity norm in \mathbb{R}^n and let $\mathbb{B} := \{x \in \mathbb{R}^n : |x| < 1\}$. Given a probability space $(\Omega, \mathcal{F}, \mathbf{P})$, we denote by $\|\cdot\|_1 := \mathbf{E}|\cdot|$ the \mathcal{L}^1 -norm for \mathbb{R}^n -valued random variables, and let $\mathcal{B} := \{X : \mathbb{R}^n\text{-valued random variable with } \|X\|_1 < 1\}$.

Given a matrix M , we denote by M_i its i th row and by M_{ij} its entry at the i th row and j th column.

D. A GUIDELINE OF TERMINOLOGY

To enhance clarity for readership, we would like to create a summary of terminologies. (1) Processes are referred to specifically when discussing stochastic processes individually. (2) A system generates a family of stochastic processes that will be investigated together. (3) Labelled processes/systems specifically pertain to situations where the topic is related to temporal logics and observations. (4) A model refers to a mechanism or tool that produces a system.

II. PRELIMINARIES

We consider $\mathbb{N} = \{0, 1, \dots\}$ as the discrete time index set, and a general Polish (complete and separable metric) space \mathcal{X} as the state space. Let $\mathcal{U} \subseteq \mathbb{R}^p$ be a compact space of control inputs. We introduce some standard concepts for fully observed controlled Markov processes.

A. CANONICAL SETUP FOR DISCRETE-TIME CONTROLLED MARKOV PROCESSES

The canonical setup for discrete-time controlled processes is provided in [44]. In brief, without loss of generality, we assume that a stochastic process $X := \{X_t\}_{t \in \mathbb{N}}$ and a process of control values $u := \{u_t\}_{t \in \mathbb{N}}$ are defined on some (unknown) probability space where the noise is generated. Given any measurable process u , the probability law of the joint process $(X, u) := \{(X_t, u_t)\}_{t \in \mathbb{N}}$ can be determined on the canonical space $((\mathcal{X} \times \mathcal{U})^\infty, \mathcal{F}, \mathbf{P})$, where

$$\mathcal{F} := \sigma\{(X_t, u_t) \in (\Gamma, \mathcal{C}), (\Gamma, \mathcal{C}) \in \mathcal{B}(\mathcal{X}) \otimes \mathcal{B}(\mathcal{U}), t \in \mathbb{N}\}.$$

We also denote X^u by the controlled process if we emphasize on the state-space marginal of (X, u) .

We consider (X, u) to be obtained from Markov models, whose transition probabilities, unlike control-free systems, have an extra dependence of the current control input, i.e.,

$$\Theta_t^u(x, \Gamma) = \mathbf{P}[X_{t+1} \in \Gamma \mid X_t = x, u_t = u]. \quad (1)$$

Now we suppose that u_t is provided according to some rule at each instant of time $t \in \mathbb{N}$. It is natural to suppose that the selection of a control at time t is based on the history $X_{[0,t]}$ and $u_{[0,t-1]}$, where

$$X_{[0,t]} := \{X_s\}_{s \in [0,t]} \text{ and } u_{[0,t]} := \{u_s\}_{s \in [0,t]}. \quad (2)$$

For each fixed $t > 0$, let $\kappa_t(\cdot \mid \cdot)$ be such that, for any $\mathcal{C} \in \mathcal{B}(\mathcal{U})$,

$$\kappa_t(\mathcal{C} \mid X_{[0,t]}; u_{[0,t-1]}) = \mathbf{P}[u_t \in \mathcal{C} \mid X_{[0,t]}; u_{[0,t-1]}. \quad (3)$$

A control policy is defined as follows.

Definition 2.1: An admissible control policy is the sequence $\kappa = \{\kappa_t, t \in \mathbb{N}\}$, where κ_t is given in the form of (3) for each $t \in \mathbb{N}$.

If u is generated based on a control policy κ , we replace the notation X^u by X^κ .

Assumption 2.2: We assume that u is deterministic *a priori* or generated by deterministic control policies.

This class of policies is considered to be sufficient for achieving optimality with respect to the specifications under consideration in the context of optimal control [34], [44].

B. CONTROLLED MARKOV SYSTEMS

We are interested in controlled Markov processes with discrete labels of states, which is done by assigning abstract labeling functions over a finite set of atomic propositions. Now we consider an abstract family of labelled controlled Markov processes as follows.

Definition 2.3 (Controlled Markov system): A controlled Markov system is a tuple $\mathbb{XU} = (\mathcal{X}, \mathcal{U}, \{\Theta\}, \text{AP}, L)$, where

- $\mathcal{X} = \mathcal{W} \cup \Delta$, where \mathcal{W} is a bounded working space, $\Delta := \mathcal{W}^c$ represents all the out-of-domain states;
- \mathcal{U} is the set of actions;
- $\{\Theta\} := \{[\Theta^u]\}_{u \in \mathcal{U}}$ contains all collections of control-dependent transition probabilities: for every t , given a realization $u \in \mathcal{U}$ of the signal u_t , the transition Θ_t^u is chosen from the collection $[\Theta^u]$ accordingly;
- AP is the finite set of atomic propositions;
- $L : \mathcal{X} \rightarrow 2^{\text{AP}}$ is the (Borel-measurable) labelling function, i.e. for every $A \in \mathcal{B}(2^{\text{AP}})$, $L^{-1}(A) \in \mathcal{F}$.

Note that for every given u and initial condition $X_0 = x_0$ (resp. initial distribution $\nu_0 \in \mathfrak{P}(\mathcal{X})$), we can generate a process $X^u \in \mathbb{XU}^u$, whose probability law is denoted by $\mathbf{P}_X^{x_0, u}$ (resp. $\mathbf{P}_X^{\nu_0, u}$), and \mathbb{XU}^u denotes all the processes that are generated by $\{\Theta\}$ given u . The collection of all the probability laws of such controlled processes is denoted by $\{\mathbf{P}_X^{x_0, u}\}_{X^u \in \mathbb{XU}^u}$ (resp. $\{\mathbf{P}_X^{\nu_0, u}\}_{X^u \in \mathbb{XU}^u}$). We denote by $\{\mathbf{P}_n^{x_0, u}\}_{n=0}^{\infty}$ (resp. $\{\mathbf{P}_n^{\nu_0, u}\}_{n=0}^{\infty}$) a sequence of $\{\mathbf{P}_X^{x_0, u}\}_{X^u \in \mathbb{XU}^u}$ (resp. $\{\mathbf{P}_X^{\nu_0, u}\}_{X^u \in \mathbb{XU}^u}$). We simply use \mathbf{P}_X^u (resp. $\{\mathbf{P}_X^u\}_{X^u \in \mathbb{XU}^u}$) if we do not emphasize the initial condition (resp. distribution).

If u is known to be generated according to some deterministic control policy κ , the previously mentioned notations are changed correspondingly by replacing the superscripts u by κ . If κ is not emphasized in the context, we use the superscripts u to indicate the general controlled quantities.

Definition 2.4 (Clarification of Notation): In the specific context of discrete state space \mathcal{X} , given a controlled Markov process X^u on \mathcal{X}^∞ , we use the notation $(\Omega, \mathcal{F}, \mathbf{P}_X^u)$ for the discrete canonical spaces of some discrete-state controlled process. We would like to still use the notation $(\Omega, \mathcal{F}, \mathbf{P}_X^u)$ if the continuity of \mathcal{X} is not clear or not emphasized.

For a path of controlled state $\varpi := \varpi_0\varpi_1\varpi_2\cdots \in \mathcal{X}^\infty$, define by $L_\varpi := L(\varpi_0)L(\varpi_1)L(\varpi_2)\cdots$ its trace. The space of infinite words is denoted by

$$(2^{\text{AP}})^\omega = \{A_0A_1A_2\cdots : A_i \in 2^{\text{AP}}, i = 0, 1, 2, \dots\}.$$

A linear-time (LT) property is a subset of $(2^{\text{AP}})^\omega$. We are only interested in LT properties Ψ such that $\Psi \in \mathcal{B}((2^{\text{AP}})^\omega)$, i.e., those are Borel-measurable².

To connect with ω -regular specifications, we introduce the semantics of path satisfaction as well as probabilistic satisfaction as follows.

Definition 2.5: Suppose Ψ is a formula of our interest. For a given labelled controlled Markov process X^u from $\mathbb{X}\mathbb{U}^u$ with initial distribution ν_0 , we formulate the canonical space $(\Omega, \mathcal{F}, \mathbf{P}_X^{\nu_0, u})$. For a controlled path $\varpi \in \mathcal{X}^\infty$, we define the path satisfaction as

$$\varpi \models \Psi \iff L_\varpi \models \Psi.$$

We denote by $\{X^u \models \Psi\} := \{\varpi : \varpi \models \Psi\} \in \mathcal{F}$ the events of path satisfaction. Given a specified probability $\rho \in [0, 1]$, we define the probabilistic satisfaction of Ψ as

$$X^u \models \mathbf{P}_{\triangleright \triangleleft \rho}^{\nu_0, u}[\Psi] \iff \mathbf{P}_X^{\nu_0, u}\{X^u \models \Psi\} \triangleright \triangleleft \rho,$$

where $\triangleright \triangleleft \in \{\leq, <, \geq, >\}$.

C. WEAK TOPOLOGY

Since our purpose is to investigate the relation between continuous-state and finite-state controlled Markov systems and then demonstrate probabilistic regularities, it is natural to work on the dual space of the state space, i.e., we consider the set of possibly uncertain measures within the topological space of probability measures.

Consider any separable and complete state space (Polish space) \mathcal{X} . The following concepts on the space of probability measures $\mathfrak{P}(\mathcal{X})$ ³ are frequently used later. Note that, ‘if a space is metrisable, the topology is determined by convergences of sequences, which explains why we sometimes only

define the concept of convergence, without explicitly mention the topology.’ [47]

Definition 2.6 (Weak convergence⁴): A sequence $\{\mu_n\}_{n=0}^\infty \subseteq \mathfrak{P}(\mathcal{X})$ is said to converge weakly to a probability measure μ , denoted by $\mu_n \rightharpoonup \mu$, if

$$\int_{\mathcal{X}} h(x)\mu_n(dx) \rightarrow \int_{\mathcal{X}} h(x)\mu(dx), \quad \forall h \in C_b(\mathcal{X}). \quad (4)$$

We frequently use the following alternative condition [48, Proposition 2.2]:

$$\mu_n(A) \rightarrow \mu(A), \quad \forall A \in \mathcal{B}(\mathcal{X}) \text{ s.t. } \mu(\partial A) = 0. \quad (5)$$

Correspondingly, the weak equivalence of any two measures μ and ν on \mathcal{X} is such that

$$\int_{\mathcal{X}} h(x)\mu(dx) = \int_{\mathcal{X}} h(x)\nu(dx), \quad \forall h \in C_b(\mathcal{X}). \quad (6)$$

We kindly refer readers to [46], [49] and [50, Remark 3] for examples on the weak topology.

Definition 2.7 (Tightness of set of measures): Let \mathcal{X} be any topological state space and $M \subseteq \mathfrak{P}(\mathcal{X})$ be a set of probability measures on \mathcal{X} . We say that M is tight if, for every $\varepsilon > 0$ there exists a compact set $K \subseteq \mathcal{X}$ such that $\mu(K) \geq 1 - \varepsilon$ for every $\mu \in M$.

The following theorem provides an alternative criterion for verifying the compactness of family of measures w.r.t. the corresponding metric space using tightness. Note that, on a compact metric space \mathcal{X} , every family of probability measures is tight.

Theorem 2.8 (Prokhorov): Let \mathcal{X} be a complete separable metric space. A family $\Lambda \subseteq \mathfrak{P}(\mathcal{X})$ is relatively compact if and only if it is tight. Consequently, for each sequence $\{\mu_n\}$ of tight Λ , there exists a $\mu \in \bar{\Lambda}$ and a subsequence $\{\mu_{n_k}\}$ such that $\mu_{n_k} \rightharpoonup \mu$.

D. PROBABILITY METRICS

The space of probability measures on a complete, separable, metric (metrisable) space endowed with the topology of weak convergence is itself a complete, separable, metric (metrisable) space [51]. While not easy to compute, the Prohorov metric can be used to metrize weak topology. We prefer to use Wasserstein metric since it also implies weak convergence and provides more practical meanings in applications. The total variation, on the other hand, implies setwise (conventional) convergence on a continuous base state space \mathcal{X} .

We first recall some basic concepts established in [40] regarding the complete analysis.

Definition 2.9 (Wasserstein distance): Let $\mu, \nu \in \mathfrak{P}(\mathcal{X})$ for $(\mathcal{X}, |\cdot|)$, the Wasserstein distance is defined by $\|\mu - \nu\|_W = \inf \mathbf{E}|X - Y|$, where the infimum is taken over all joint distributions of the random variables X and Y with marginals μ and ν respectively.

²By [32] and [45, Proposition 2.3], any ω -regular language of labelled (controlled) Markov processes is measurable. The proof relies on the properties of the canonical space as well as the connection with Büchi automaton.

³ $\mathfrak{P}(\mathcal{X})$ is always metrisable given \mathcal{X} is a Polish space [46].

⁴Technically, this should be weak* convergence as probability measures are linear bounded functionals of bounded continuous functions. We use ‘weak convergence’ due to no conflict of concept in this article.

We frequently use the following duality form of definition⁵,

$$\|\mu - \nu\|_W := \sup \left\{ \left| \int_{\mathcal{X}} h(x) d\mu(x) - \int_{\mathcal{X}} h(x) d\nu(x) \right|, \right. \\ \left. h \in C(\mathcal{X}), \text{Lip}(h) \leq 1 \right\}.$$

The discrete case, $\|\cdot\|_W^d$, is nothing but to change the integral to summation. Let $\mathcal{B}_W = \{\mu \in \mathfrak{P}(\mathcal{X}) : \|\mu - \delta_0\|_W < 1\}$. Given a set $\mathcal{G} \subseteq \mathfrak{P}(\mathcal{X})$, let $\|\mu\|_{\mathcal{G}} = \inf_{\nu \in \mathcal{G}} \|\mu - \nu\|_W$ be the distance from μ to \mathcal{G} , and $\mathcal{G} + r\mathcal{B}_W := \{\mu : \|\mu\|_{\mathcal{G}} < r\}$ be the r -neighborhood of \mathcal{G} .

Note that \mathcal{B}_W is dual to \mathcal{B} . For any $\mu \in \mathcal{B}_W$, the associated random variable X should satisfy $\mathbf{E}|X| \leq 1$, and vice versa.

We also frequently use the following inequalities to bound the Wasserstein distance between two Gaussians.

Proposition 2.10: Let $\mu \sim \mathcal{N}(m_1, \Sigma_1)$ and $\nu \sim \mathcal{N}(m_2, \Sigma_2)$ be two Gaussian measures on \mathbb{R}^n . Then

$$\|m_1 - m_2\| \leq \|\mu - \nu\|_W \\ \leq \left(\|m_1 - m_2\|_2^2 + \|\Sigma_1^{1/2} - \Sigma_2^{1/2}\|_F^2 \right)^{1/2}, \quad (7)$$

where $\|\cdot\|_F$ is the Frobenius norm.

Definition 2.11 (Total variation distance): Given two probability measures μ and ν on $\mathcal{B}(\mathcal{X})$, the total variation distance is defined as

$$\|\mu - \nu\|_{TV} = 2 \sup_{\Gamma \in \mathcal{B}(\mathcal{X})} |\mu(\Gamma) - \nu(\Gamma)|. \quad (8)$$

In particular, if \mathcal{X} is a discrete space, $\|\mu - \nu\|_{TV}^d = \sum_{q \in \mathcal{X}} |\mu(q) - \nu(q)|$.

Remark 2.12: Total variation distance implies a much stronger uniform norm given \mathcal{X} is continuous. However, working on discrete topology of a finite set, we have the following well known connection.

$$\|\mu - \nu\|_W^d = \frac{1}{2} \|\mu - \nu\|_{TV}^d. \quad (9)$$

This equivalence [52, Theorem 4] on the discrete topology implies that abstractions already exist unnecessarily in a functional space with stronger convergence concept.

III. PROBLEM FORMULATION

Our objective is to develop abstractions for specific controlled Markov systems. In this section, we introduce the concrete systems of our interests, and subsequently outline the requirements for a well-posed abstraction. The soundness and completeness analysis of these abstractions will be provided in the remaining sections of this article.

A. THE CONCRETE CONTROLLED MARKOV SYSTEMS

Throughout Sections III to V, we focus on controlled Markov processes determined by the following fully-observed Markov system

$$X_{t+1} = f(X_t, u_t) + b(X_t)\mathbf{w}_t + \vartheta \xi_t, \quad (10)$$

⁵Lip(h) is the Lipschitz constant of h such that $|h(x_2) - h(x_1)| \leq \text{Lip}(h)|x_2 - x_1|$.

the sample state $X_t^u(\varpi) \in \mathcal{X} \subseteq \mathbb{R}^n$ for all $t \in \mathbb{N}$ given a signal process u , the stochastic inputs $\{\mathbf{w}_t\}_{t \in \mathbb{N}}$ are i.i.d. Gaussian random variables with covariance $I_{k \times k}$. Mappings $f : \mathbb{R}^n \times \mathbb{R}^p \rightarrow \mathbb{R}^n$ is locally Lipschitz continuous in both arguments, and $b : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times k}$ is locally Lipschitz continuous. The memoryless perturbation $\xi_t \in \bar{\mathcal{B}}$ are independent random variables with intensity $\vartheta \geq 0$ and unknown distributions. We can translate (10) into the form of a controlled Markov system

$$\mathbb{XU} = (\mathcal{X}, \mathcal{U}, \{\mathcal{T}\}, \text{AP}, L_{\mathbb{XU}}), \quad (11)$$

where $\{\mathcal{T}\} := \{\{\mathcal{T}^u\}\}_{u \in \mathcal{U}}$ is defined in the same way as the $\{\Theta\}$ in Definition 2.3. We use notation \mathcal{T} instead of Θ to indicate the continuity of the transition probability in $x \in \mathcal{X}$.

Remark 3.1: For $\vartheta \neq 0$, (10) defines a family \mathbb{XU} of controlled Markov processes. As to simulate the probability laws at the observation times, the above system can be regarded as a discrete-time numerical scheme of controlled stochastic differential equations (SDEs) driven by Brownian motions, which demonstrates practical meanings in physical sciences and finance. The real noise with bounded supports that are considered in [41] is a special type. Gaussian variables in (10) do not lose any generalities in view of \mathcal{L}^1 properties and are in favor of our formal analysis. In addition, compared to f being mixed-monotone and b being constant in [38], [41], the choice of f and b in this article fits more general dynamics in applications.

For real-world applications, we only care about the behaviors in the bounded working space \mathcal{W} . It is desired to trap the sample paths at the out-of-domain states once Δ they reach Δ . By defining stopping time $\tau = \tau(u) := \inf\{t \in \mathbb{N} : X_t^u \notin \mathcal{W}\}$ for each X^u , it is equivalent to study the probability law of the corresponding stopped process $\{X_{t \wedge \tau}^u\}_{t \in \mathbb{N}}$ for any initial condition (or distribution), which coincides with \mathbf{P}_X^u on \mathcal{W} . In view of the corresponding transitions probability, for each realization of control input u and for all $x \in \mathcal{X} \setminus \mathcal{W}$, the transition probability should satisfy $\mathcal{T}^u(x, \Gamma) = 0$ for all Γ such that $\Gamma \cap \mathcal{W} \neq \emptyset$.

We also would like to point out that the above set-ups are inline with conventional stochastic analysis. It is important to note that for solutions to stochastic differential equations or inclusions, first-exit time of a bounded region is not generally to be triggered within a finite time with probability one, as shown in several studies [53], [54], [55], [56], [57], [58]. For solutions to stochastic difference equations or inclusions of the form (10), it can also be shown that $\mathbf{P}_X^u[\tau < \infty]$ is not necessarily always 1. This can be attributed to factors such as stability or controlled stability [59], [60], [61].

On the other hand, ω -regular formulas themselves are independent of the underlying dynamics. However, in terms of satisfaction of a formula Ψ with ω -regular property, occurrences of events $\{X^u \models \Psi\}$ can be independent of the sigma-algebra $\sigma\{\tau\}$ or not. Even in the latter situations, the probability associated with these events is not trivial (only 0 or 1). We use a control-free 2-dim discrete-time process B (a

discrete-time analogue to a diffusion process) driven by

$$B_{t+1} = B_t + b\mathbf{w}_t$$

as an example, where b is a diagonal constant matrix with entries Δt , and Δt is the time-step size. Note that $\tau < \infty$ almost surely. Given $B_0 = x \in \mathcal{W} := \{x \in \mathbb{R}^2 : \frac{R}{2^k} < |x| < R\}$ a.s. for each k , the probability $p_{x,k}$ of the event Ψ_k : ‘ B reaching $\{x \in \mathbb{R}^2 : |x| = R\}$ before $\{x \in \mathbb{R}^2 : |x| = \frac{R}{2^k}\}$ ’ can be analytically resolved based on a martingale argument. Let $f(x) = e^{-x/\Delta t}$, then one can verify that $f(|B|)$ is a martingale, and hence $f|x| = \mathbf{E}[f|B_{\tau \wedge \tau_k}|]$, where τ_k is the first hitting time of the boundary of $\{x \in \mathbb{R}^2 : |x| = \frac{R}{2^k}\}$. We can conclude that $p_{x,k} = 1 - \mathbf{P}[\tau_k < \tau] = 1 - \frac{1}{2^k}$ for each $k \geq 1$, which is a non-trivial probability.

Remark 3.2: It is worth noting that, in the numerical examples in [37], [38], the pseudo-Gaussian noise with a bounded support is obtained by normalizing real Gaussian distribution on \mathcal{W} by the probability $\mathcal{N}(0, 1)(\mathcal{W})$, which significantly distorts the shape of Gaussian density within $\text{Int}(\mathcal{W})$. The treatment of out-of-domain transitions in this article should preserve the density of \mathbf{w}_t and hence that of X_t for each t on $\text{Int}(\mathcal{W})$. The densities can be recovered to the true densities given the stopping time τ not being triggered.

Definition 3.3 (Clarification of Notations): To avoid any complexity, we use the same notation X^u and \mathbf{P}_X^u to denote the stopped processes and the associated laws.

Assumption 3.4: We assume that $\mathbf{in} \in L(x)$ for any $x \notin \Delta$ and $\mathbf{in} \notin L(\Delta)$. We can also include ‘always (\mathbf{in})’ in the specifications to observe sample paths for ‘inside-domain’ behaviors, which is equivalent to verifying $\{\tau = \infty\}$.

B. ROBUST ABSTRACTIONS

We define a notion of abstraction between continuous-state and finite-state controlled Markov systems via state-level relations and measure-level relations.

Definition 3.5: A (binary) relation γ from A to B is a subset of $A \times B$ satisfying (i) for each $a \in A$, $\gamma(a) := \{b \in B : (a, b) \in \gamma\}$; (ii) for each $b \in B$, $\gamma^{-1}(b) := \{a \in A : (a, b) \in \gamma\}$; (iii) for $A' \subseteq A$, $\gamma(A') = \cup_{a \in A'} \gamma(a)$; (iv) and for $B' \subseteq B$, $\gamma^{-1}(B') = \cup_{b \in B'} \gamma^{-1}(b)$.

Definition 3.6: Given a continuous-state controlled Markov system

$$\mathbb{X}\mathcal{U} = (\mathcal{X}, \mathcal{U}, \{\mathcal{T}\}, \text{AP}, L_{\mathbb{X}\mathcal{U}})$$

with a compact $\mathcal{U} \in \mathbb{R}^p$, and a finite-state Markov system

$$\mathbb{I}\mathbb{A} = (\mathcal{Q}, \text{Act}, \{\Theta\}, \text{AP}, L_{\mathbb{I}\mathbb{A}}),$$

where $\mathcal{Q} = (q_1, \dots, q_N)^T$, $\text{Act} = \{a_1, \dots, a_M\}$, and $\{\Theta\} := \{[\Theta^a]\}_{a \in \text{Act}}$ contains all collections of $n \times n$ stochastic matrices that are also dependent on a .

We say that $\mathbb{I}\mathbb{A}$ abstracts $\mathbb{X}\mathcal{U}$, and write $\mathbb{X}\mathcal{U} \leq_{\Sigma_\alpha} \mathbb{I}\mathbb{A}$, if there exist

- 1) a state-level relation $\alpha \subseteq \mathcal{X} \times \mathcal{Q}$ from $\mathbb{X}\mathcal{U}$ to $\mathbb{I}\mathbb{A}$ such that, for all $x \in \mathcal{X}$, there exists $q \in \mathcal{Q}$ such that $(x, q) \in \alpha$ ($\alpha(x) \neq \emptyset$) and $L_{\mathbb{I}\mathbb{A}}(q) = L_{\mathbb{X}\mathcal{U}}(x)$;
- 2) a measure-level relation $\Sigma_\alpha \subseteq \mathfrak{P}(\mathcal{X}) \times \mathfrak{P}(\mathcal{Q})$ from $\mathbb{X}\mathcal{U}$ to $\mathbb{I}\mathbb{A}$ such that, for all $i \in \{1, 2, \dots, N\}$ and $a \in$

Act , there exists $u \in \mathcal{U}$ such that for any $\mathcal{T}^a \in \llbracket \mathcal{T}^a \rrbracket$ and all $x \in \alpha^{-1}(q_i)$, there exists $\Theta^a \in \llbracket \Theta^a \rrbracket$ satisfying $(\mathcal{T}^a(x, \cdot), \Theta^a) \in \Sigma_\alpha$ and $\mathcal{T}^a(x, \alpha^{-1}(q_j)) = \Theta^a_{ij}$ for all $j \in \{1, 2, \dots, n\}$.

The converse abstraction is defined in a similar way.

Remark 3.7: Heuristically, we stand from the side of the original system and require an abstraction to

- contain states with the same labels as states of the original system;
- include transitional measures with the same measuring results on all the discrete states given any starting point of the original system that can be mapped to an abstract state.

Given a rectangular partition and the existence of an abstraction, one immediate consequence is that the transition matrices are able to recover all possible transition probabilities (of the original system) from a grid to another.

Assumption 3.8: Without loss of generality, we assume that the labelling function is amenable to a rectangular partition⁶. In other words, a state-level abstraction can be obtained from a rectangular partition.

IV. SOUNDNESS OF ROBUST BMDP ABSTRACTIONS

BMDPs are quasi-controlled Markov systems on a discrete state space with upper/under approximations ($\hat{\Theta}^u/\check{\Theta}^u$) of the real transition matrices.

Definition 4.1: A BMDP is a tuple $\mathcal{I}\mathcal{A} = (\mathcal{Q}, \text{Act}, \{\check{\Theta}\}, \{\hat{\Theta}\}, \text{AP}, L_{\mathcal{I}\mathcal{A}})$, where

- \mathcal{Q} is an $(N + 1)$ -dimensional state-space for any N , which is obtained by a finite state-space partition containing $\{\Delta\}$, i.e., $\mathcal{Q} = (q_1, q_2, \dots, q_N, q_{N+1} := \Delta)^T$;
- Act is a finite-dimensional actions;
- AP and $L_{\mathcal{I}\mathcal{A}}$ are the same as in Definition 2.3;
- $\{\check{\Theta}\} := \{\check{\Theta}^a\}_{a \in \text{Act}}$ is a family of $N \times N$ matrix such that $\check{\Theta}^a_{ij}$ is the lower bound of transition probability from the state number i to j for each $i, j \in \{1, 2, \dots, N\}$ and action $a \in \text{Act}$;
- $\{\hat{\Theta}\} := \{\hat{\Theta}^a\}_{a \in \text{Act}}$ is a family of $N \times N$ matrix such that $\hat{\Theta}^a_{ij}$ is the upper bound of transition probability from the state number i to j for each $i, j \in \{1, 2, \dots, N\}$ and action $a \in \text{Act}$.

By adding constraints

$$\llbracket \Theta^a \rrbracket = \{\Theta^a : \text{stochastic matrices with } \check{\Theta}^a \leq \Theta^a \leq \hat{\Theta}^a \text{ component wisely}\}, \quad (12)$$

we are able to transfer an $\mathcal{I}\mathcal{A}$ into a controlled Markov system $\mathbb{I}\mathbb{A}$ as in Definition 3.6, whose $\llbracket \Theta^a \rrbracket$'s are well defined sets of stochastic matrices for each $a \in \text{Act}$. We call the induced $\mathbb{I}\mathbb{A}$, which is verified to satisfy Definition 3.6, the abstraction generated by the BMDP $\mathcal{I}\mathcal{A}$, or simply the BMDP abstraction.

Remark 4.2: To make $\mathbb{I}\mathbb{A}$ an abstraction for (11), we can discretize both \mathcal{X} and \mathcal{U} , such that each node $a \in \text{Act}$ represents a grid of $u \in \mathcal{U}$. We then need the approximation to be such that $\check{\Theta}^a_{ij} \leq \int_{\alpha^{-1}(q_j)} \mathcal{T}^a(x, dy) \leq \hat{\Theta}^a_{ij}$ for a $u \in a$,

⁶See e.g. [41, Definition 1].

for all $x \in \alpha^{-1}(q_i)$ and $i, j = 1, \dots, N$, as well as $\Theta_{N+1} = (0, 0, \dots, 1)$.

For any realization of a sequence of actions $a := \{a_i\}_{i \in \mathbb{N}}$, the controlled Markov system \mathbb{IA}^a is reduced to a family of perturbed Markov chains generated by the uncertain choice of $\{\Theta\}$ for each t . The n -step transition are derived based on $\llbracket \Theta^{a_i} \rrbracket$:

$$\begin{aligned} \llbracket \Theta^{(2)} \rrbracket &:= \{\Theta_0^{a_0} \Theta_1^{a_1} : \Theta_i^{a_i} \in \llbracket \Theta^{a_i} \rrbracket, i = 0, 1\}, \\ &\vdots \\ \llbracket \Theta^{(n)} \rrbracket &:= \{\Theta_0^{a_0} \Theta_1^{a_1} \dots \Theta_n^{a_n} : \Theta_i^{a_i} \in \llbracket \Theta^{a_i} \rrbracket, \\ &\quad i = 0, 1, \dots, n\}. \end{aligned}$$

The weak compactness and convexity of the probability laws of \mathbb{IA}^a are proved in [40, Section 3.2]. We also kindly refer readers to the arXiv version [50, Section 3.1] for more details on the weak topology properties.

Taking the advantages of the above properties, we now show the soundness of BMDP abstractions.

Definition 4.3: Given a state-level relation α and a measure-level relation Σ_α from \mathbb{XU} to \mathbb{IA} . Let ϕ and κ be some control policies of \mathbb{XU} and \mathbb{IA} , respectively. Recall notations in (2). We call ϕ a Σ_α -implementation of κ if, for each $t \in \mathbb{N}$,

$$u_t = \phi_t(X_{[0,t]}, u_{[0,t-1]}), \quad X \in \mathbb{XU}^\phi$$

is chosen according to

$$a_t = \kappa_t(I_{[0,t]}, a_{[0,t-1]}), \quad I \in \mathbb{IA}^\kappa$$

in a way that, for any realization u and a of u_t and a_t , for any $\mathcal{T}^a \in \llbracket \mathcal{T}^a \rrbracket$ and all $x \in \alpha^{-1}(q_i)$, there exists $\Theta^a \in \llbracket \Theta^a \rrbracket$ satisfying $(\mathcal{T}^a(x, \cdot), \Theta_i^a) \in \Sigma_\alpha$ and $\mathcal{T}^a(x, \alpha^{-1}(q_j)) = \Theta_{ij}^a$ for all $j \in \{1, 2, \dots, n\}$.

We can define the converse implementation from \mathbb{IA} to \mathbb{XU} based on a converse measure-level relation (from \mathbb{IA} to \mathbb{XU}) in a similar way.

Remark 4.4: Heuristically, a control policy κ is generated in the finite-state finite-action abstraction model within \mathbb{IA} to ensure a probabilistic satisfaction of some specification. The selection of the control policy ϕ is subjected to κ and hence \mathbb{IA}^κ according to the abstraction relation, such that (2) of Definition 3.6 is guaranteed.

Proposition 4.5: Let \mathbb{IA} be a controlled Markov system that is derived from a BMDP with any initial distribution μ_0 . Then for any ω -regular specification Ψ , given any admissible deterministic control policy κ , the set

$$S^{\mu_0, \kappa} = \{\mathcal{P}_I^{\mu_0, \kappa} (I^\kappa \models \Psi)\}_{I^\kappa \in \mathbb{IA}^\kappa}$$

is a compact interval.

Proof: The proof is similar to [40, Theorem 2]. We only show the sketch. Let a be the control input process generated by κ such that $a_t = \kappa_t(I_{[0,t]}, a_{[0,t-1]})$ for each t . Note that $a \in \mathcal{B}(Act^\infty)$ and $a_t \in \mathcal{B}(Act)$, where the set of actions Act admits a discrete topology. The weak compactness of

the probability law $\{\mathcal{P}_I^{\mu_0, \kappa}\}_{I^\kappa \in \mathbb{IA}^\kappa}$ follows exactly the same reasoning as in [40, Proposition 1]. The convexity of every finite-dimensional distribution of I^a can be obtained in similar way as in [40, Theorem 2] based on the transition procedure, i.e., for any $q_0, q_{n_1}, \dots, q_{n_t} \in \mathcal{Q}$,

$$\begin{aligned} &\mathcal{P}_I^{q_0, \kappa} [I_0 = q_0, \dots, I_t = q_{n_t}, I_{t+1} = q_{n_{t+1}}] \\ &\in \{\Theta_{n_{t+1}, n_t}^{a_t} \Theta_{n_t, n_{t-1}}^{a_{t-1}} \dots \Theta_{n_1, 0}^{a_0} \delta_{q_0} : \Theta^{a_i} \in \llbracket \Theta^{a_i} \rrbracket, \\ &\quad i \in \{0, \dots, t\}, \text{ and } a_t = \kappa(I_{[0,t]} = q_{[0,t]}, a_{[0,t-1]})\}. \end{aligned}$$

By a standard monotone class argument, the convexity for any Borel measurable set $A \in \mathcal{F}$ measured in the set of laws $\mathcal{P}_I^{q_0, \kappa}$ are guaranteed, which implies the convexity of $S^{q_0, \kappa}$, and hence that of $S^{\mu_0, \kappa}$. ■

The soundness regularity is provided as follows.

Theorem 4.6: Let \mathbb{XU} as in (11) be a controlled Markov system driven by (10). Suppose that there exist a state-level relation α , a measure-level relation Σ_α , and a BMDP abstraction \mathbb{IA} such that $\mathbb{XU} \preceq_{\Sigma_\alpha} \mathbb{IA}$. Let Ψ be an ω -regular specification. Suppose the initial distribution ν_0 of \mathbb{XU} is such that $\nu_0(\alpha^{-1}(q_0)) = 1$. Then, given an admissible deterministic control policy κ , there exists a Σ_α -implementation policy ϕ of κ such that

$$\mathbf{P}_X^{\nu_0, \phi} (X^\phi \models \Psi) \in \{\mathcal{P}_I^{q_0, \kappa} (I^\kappa \models \Psi)\}_{I^\kappa \in \mathbb{IA}^\kappa}, \quad X^\phi \in \mathbb{XU}^\phi.$$

Proof: We denote by μ_t and ν_t , respectively, the marginal probability measures on $\mathcal{B}(\mathcal{Q})$ and $\mathcal{B}(\mathcal{X})$ for $t \in \mathbb{N}$. We also use the shorthand notation $\mu_t^a(\cdot) := \mu_t(\cdot | a_{t-1} = a)$ and $\nu_t^u(\cdot) := \nu_t(\cdot | u_{t-1} = u)$ to indicate the conditional probabilities. We consider $\nu_0 = \delta_{x_0}$ a.s. for simplicity. Note that, at $t = 1$, by the definition of BMDP abstraction and Remark 4.2, there exists a $u_0 \in \mathcal{U}$ such that,

$$\begin{aligned} \check{\Theta}_{ij}^{a_0} &\leq \nu_1^{u_0}(\alpha^{-1}(q_j)) = \int_{\alpha^{-1}(q_j)} \delta_{x_0} \mathcal{T}^{u_0}(x_0, dy) \\ &\leq \hat{\Theta}_{ij}^{a_0}, \quad \forall x_0 \in q_0 \text{ and } \forall j \in \{1, 2, \dots, N+1\}, \end{aligned}$$

where $a_0 = \kappa_0(q_0)$, and u_0 is selected accordingly such that the above relation is satisfied. We can easily check that

$$\mu_1^{a_0} = (\nu_1^{u_0}(\alpha^{-1}(q_1)), \dots, \nu_1^{u_0}(\alpha^{-1}(q_{N+1})))^T$$

is a proper marginal probability measure of \mathbb{IA} at $t = 1$. In particular, $\mu_1^{a_0}(q_j) = \nu_1^{u_0}(\alpha^{-1}(q_j))$ for each $j \in \{1, 2, \dots, N+1\}$.

Similarly, at $t = 2$, we have

$$\begin{aligned} \check{\Theta}_{ij}^{a_{1i}} \mu_1^{u_0}(q_i) &\leq \nu_2^{u_{1i}}(\alpha^{-1}(q_j)) \\ &= \int_{\alpha^{-1}(q_j)} \int_{\alpha^{-1}(q_i)} \nu_1^{u_0}(dx) \mathcal{T}^{u_{1i}}(x, dy) \\ &\leq \hat{\Theta}_{ij}^{a_{1i}} \mu_1^{u_0}(q_i), \quad \forall i, j \in \{1, 2, \dots, N+1\}, \end{aligned}$$

where $a_{1i} = \kappa_1(q_i)$ for each $i \in \{1, 2, \dots, N+1\}$, and u_{1i} is selected accordingly such that the above relation is satisfied. Then, $\mu_2^{a_{1i}} = (\nu_2^{u_{1i}}(\alpha^{-1}(q_1)), \dots, \nu_2^{u_{1i}}(\alpha^{-1}(q_{N+1})))^T$ is again a proper marginal probability measure of \mathbb{IA} at $t = 2$ for each

$i \in \{1, 2, \dots, N+1\}$. In addition, there also exists a $\mathcal{P}^{q_0, \kappa}$ such that its one-dimensional marginals up to $t = 2$ admit μ_1 and μ_2 , and satisfies

$$\begin{aligned} & \mathcal{P}^{q_0, \kappa} [I_0 = q_0, \mathbf{a}_0 = a_0, I_1 = q_i, \mathbf{a}_1 = a_{1i}, I_2 = q_j] \\ &= \mu_0(q_0) \mu_1^{a_0}(q_i) \mu_2^{a_{1i}}(q_j) \\ &= \delta_{q_0}(q_0) \int_{\alpha^{-1}(q_i)} v_1^{u_0}(dx) \int_{\alpha^{-1}(q_j)} v_2^{u_{1i}}(dy) \\ &= \mathbf{P}_X^{q_0, u} [X_0 = x_0, u_0 = u_0, X_1 \in \alpha^{-1}(q_i), u_1 = u_{1i}, \\ & \quad X_2 \in \alpha^{-1}(q_j)] \end{aligned}$$

for all $i, j \in \{1, 2, \dots, N+1\}$. We then propagate the process inductively according to the above machinery by

- 1) selecting $u_t := u_{tj}$ at each time according to the realization $a_t := a_{tj} = \kappa_t(q_j)$;
- 2) selecting \mathcal{T}^{a_t} and $\Theta^{a_t} \in \llbracket \Theta^{a_t} \rrbracket$ at each time via the connection as the above.

We can verify that, by the above selection procedure, there exists $\mathcal{P}^{q_0, \kappa}$ such that

$$\begin{aligned} & \mathcal{P}^{q_0, \kappa} [I_0 = q_0, \mathbf{a}_0 = a_0, I_1 = q_i, \mathbf{a}_1 = a_{1i}, \dots] \\ &= \mathbf{P}_X^{x_0, u} [X_0 = x_0, u_0 = u_0, X_1 \in \alpha^{-1}(q_i), u_2 = u_2, \dots] \end{aligned}$$

holds for any finite-dimensional distribution. By Kolmogorov extension theorem, there exists a unique probability law $\mathcal{P}_I^{q_0, \kappa}$ for (I, \mathbf{a}) or $I^\kappa \in \mathbb{I}\mathbb{A}^\kappa$ such that it has the same measuring results on any \mathcal{F} -measurable sets (recall that $\mathcal{F} = \mathcal{B}(\mathcal{Q}^\infty)$) as the probability law $\mathbf{P}_X^{x_0, u}$ of the generated process (X, u) or X^u .

The Σ_α -implementation $\phi = \{\phi_t\}_{t \in \mathbb{N}}$ exists and is given as $\phi_t(\cdot \mid X_{[0,t]}, u_{[0,t-1]}) = \mathbf{P}_X^{x_0, u} [u_t = (\cdot) \mid X_{[0,t]}, u_{[0,t-1]}]$ after averaging out along \mathcal{X}^∞ . ■

Based on Theorem 4.6, we can immediately show whether a control strategy exists based on the BMDP abstraction such that the controlled process satisfy the probabilistic specification.

Corollary 4.7: Let $\mathbb{X}\mathbb{U}$, its BMDP abstraction $\mathbb{I}\mathbb{A}$, an ω -regular formula Ψ , and a constant $\rho \in [0, 1]$ be given. Suppose there exists a control policy κ such that $I^\kappa \models \mathcal{P}_{\triangleright \triangleleft \rho}^{q_0} [\Psi]$ for all $I^\kappa \in \mathbb{I}\mathbb{A}^\kappa$, then there exists a policy ϕ such that $X^\phi \models \mathbf{P}_{\triangleright \triangleleft \rho}^{v_0, \phi} [\Psi]$ for all $X^\phi \in \mathbb{X}\mathbb{U}^\phi$ with $v_0(\alpha^{-1}(q_0)) = 1$.

Remark 4.8: The purpose of abstraction-based formal methods is in general different from constructing numerical solutions for SDEs. The numerical analysis for SDEs is to determine how good the approximation is and in what sense it is close to the exact solution [62].

The stochastic driving forces in discrete-time numerical simulations are given with discrete distributions *in a priori*. Consequently, there is a unique solution in the discrete canonical space driven by this discrete noise. The numerical simulation provides a much smaller set of measurable sample paths, i.e. a natural filtration $\mathcal{F}^{\mathbf{w}^d}$ subjected to the discrete version of noise \mathbf{w}^d rather than \mathcal{F} (recall Definition 2.4). The missing transitions or measurable sample paths from \mathcal{F} cannot be recovered given a fixed discretized noise at a time.

In comparison with the numerical solutions, the stochastic abstractions in this article do not use the spatially discretized noise as the driving force. Instead, we directly work on generating a relation based on the state-space discretization such that the transition kernel of the original system is ‘included’ in the discrete family of transition matrices in the sense of Theorem 4.6. Even though a refinement of grid size can lead to a convergence for both numerical simulations and stochastic abstractions (see [40, Proposition 3] for details), they converge from different ‘directions’. In other words, the family of the discrete probability laws from an abstraction reduces to a singleton whilst the missing transitions in a numerical simulation become empty as the size of the grids converges to 0.

V. ROBUST COMPLETENESS OF BMDP ABSTRACTIONS

In this section, we propose the concept of robustly complete abstractions of discrete-time nonlinear stochastic systems of the form (10) and provide computational procedures for constructing sound and robustly complete abstractions for this class of controlled stochastic systems under mild conditions.

Note that, in view of the soundness analysis given in Theorem 4.6, the formulation of a formal inclusion of transition probabilities, and consequently, the inclusion of the ‘reachable set’ of marginal probability measures, can be achieved as long as BMDP abstractions can be constructed. This guarantees that the real satisfaction probability is preserved as in Corollary 4.7, however, creates a deviation from the original concrete system. In order to prevent a sound BMDP abstraction from generating significant deviation, we perform a completeness analysis in this section. The purpose is to investigate whether such deviations can be reflected in some continuous-state systems with arbitrarily small perturbations when compared to the original concrete systems. By conducting this type of comparison, we aim to release the abstraction from the ‘virtual’ systems and establish a perception that even small deviations would not lead to substantial issues for the original system.

To do this, we work on the space of probability measure metricized by the Wasserstein metric (recall Section II-D) to quantify this extra perturbation.

A. CONSTRUCTION OF ROBUSTLY COMPLETE BMDP ABSTRACTIONS

We consider two continuous-state systems with parameters $\vartheta_2 > \vartheta_1 \geq 0$. The first system, denoted by $\mathbb{X}\mathbb{U}_1$, is given by

$$X_{t+1} = f(X_t, u_t) + b(X_t) \mathbf{w}_t + \vartheta_1 \xi_t^{(1)}, \quad \xi_t^{(1)} \in \bar{\mathbb{B}}, \quad (13)$$

and the second system, denoted by $\mathbb{X}\mathbb{U}_2$, is driven by

$$X_{t+1} = f(X_t, u_t) + b(X_t) \mathbf{w}_t + \vartheta_2 \xi_t^{(2)}, \quad \xi_t^{(2)} \in \bar{\mathbb{B}}. \quad (14)$$

We construct a sound and robustly complete BMDP abstraction $\mathbb{I}\mathbb{A}$ for $\mathbb{X}\mathbb{U}_1$ in a similar way as in [40], i.e., we build a state-level relation α and a measure-level Σ_α such that

$$\mathbb{X}\mathbb{U}_1 \preceq_{\Sigma_\alpha} \mathbb{I}\mathbb{A}, \quad \mathbb{I}\mathbb{A} \preceq_{\Sigma_\alpha^{-1}} \mathbb{X}\mathbb{U}_2.$$

We define the set of transition probabilities of $\mathbb{X}U_i$, for each fixed $u \in \mathcal{U}$, from any box $[x] \subseteq \mathbb{R}^n$ as

$$\mathbb{T}_i^u([x]) = \{\mathcal{T}^u(x, \cdot) : \mathcal{T}^u \in [\mathcal{T}^u]_i, x \in [x], i = 1, 2.\}$$

The following lemma is to straightforward based on [40, Lemma 3].

Lemma 5.1: Fix any $\vartheta_1 \geq 0$, any box $[x] \subseteq \mathbb{R}^n$, and $u \in \mathcal{U}$. For all $k > 0$, there exists a finitely terminated algorithm to compute an over-approximation of the set of (Gaussian) transition probabilities from $[x]$, such that

$$\mathbb{T}_1^u([x]) \subseteq \widehat{\mathbb{T}}_1^u([x]) \subseteq \mathbb{T}_1^u([x]) + k\bar{B}_W,$$

where $\widehat{\mathbb{T}}_1^u([x])$ is the computed over-approximation set of Gaussian measures.

Lemma 5.1 is to construct an over-approximation $\widehat{\mathbb{T}}_1^u([x])$ of the set of Gaussian transition probabilities from the original concrete system $\mathbb{X}U_1$, such that any Gaussian measure within $\widehat{\mathbb{T}}_1^u([x])$ will not perturb the mean more than any arbitrarily small k . We skip the proof due to the similarity to [40, Lemma 3] and [7, Lemma 1]. The main step is to find inclusion functions for f and b , as well as a mesh of $[x]$ with appropriate size. The over-approximation of the ‘reachable’ mean and covariance can be obtained by union the regions generated by inclusion functions acting on the mesh.

Note that, recalling Definition 3.3, we are actually working on the quantification for the stopped processes. The introduce a modification that does not affect the law of the stopped processes, i.e., we use a weighted point mass to represent the measures at the boundary, and the mean value should remain the same.

Definition 5.2: For $i = 1, 2$, we introduce the modified transition probabilities for $\mathbb{X}U_i = (\mathcal{X}, \mathcal{U}, \{\mathcal{T}_i\}, \text{AP}, L_{\mathbb{X}U})$. For any $u \in \mathcal{U}$, for all $\mathcal{T}_i^u \in [\mathcal{T}_i^u]$, let

$$\tilde{\mathcal{T}}_i^u(x, \Gamma) = \begin{cases} \mathcal{T}_i^u(x, \Gamma), & \forall \Gamma \subseteq \mathcal{W}, \forall x \in \mathcal{W}, \\ \mathcal{T}_i^u(x, \mathcal{W}^c), & \Gamma = \partial\mathcal{W}, \forall x \in \mathcal{W}, \\ 1, & \Gamma = \partial\mathcal{W}, x \in \partial\mathcal{W}. \end{cases} \quad (15)$$

Correspondingly, let $[\tilde{\mathcal{T}}^u]$ denote the collection. Likewise, we also use $(\cdot)^u$ to denote the induced quantities of any other types w.r.t. such a modification.

We are now ready to show the existence of a robustly complete abstraction given (13) and (14).

Theorem 5.3: For any $0 \leq \vartheta_1 < \vartheta_2$, we consider $\mathbb{X}U_i = (\mathcal{X}, \mathcal{U}, \{\tilde{\mathcal{T}}_i\}, \text{AP}, L_{\mathbb{X}U})$, $i = 1, 2$, that are driven by (13) and (14), respectively. Then, under Assumption 3.8, there exists a rectangular partition \mathcal{Q} (state-level relation $\alpha \subseteq \mathcal{X} \times \mathcal{Q}$), a measure-level relation Σ_α and a finite-state abstraction system $\mathbb{I}\mathbb{A} = (\mathcal{Q}, \text{Act}, \{\Theta\}, \text{AP}, L_{\mathbb{I}\mathbb{A}})$ such that

$$\mathbb{X}U_1 \leq_{\Sigma_\alpha} \mathbb{I}\mathbb{A}, \quad \mathbb{I}\mathbb{A} \leq_{\Sigma_\alpha^{-1}} \mathbb{X}U_2. \quad (16)$$

Proof: We construct a finite-state BMDP abstraction in a similar way as in [40, Theorem 4]. Aside from the additional dependence on the control inputs, we also provide tighter estimations on sets of probability measures. By Assumption 3.8, we use uniform rectangular partition \mathcal{Q} on \mathcal{W} . We

then let the state-level relation be $\alpha = \{(x, q) : q = \eta \lfloor \frac{x}{\eta} \rfloor\} \cup \{(\Delta, \Delta)\}$, and $\text{Act} = \{a : \varrho \lfloor \frac{u}{\varrho} \rfloor\}$, where $\lfloor \cdot \rfloor$ is the floor function. The parameters η, ϱ are to be chosen later. Denote the number of discrete nodes by $N + 1$.

We construct the measure-level abstraction as follows. We repeat the procedure with updated notations for the control systems. For any fixed $u = a \in \text{Act}$, for any $\tilde{\mathcal{T}}^u \in [\mathcal{T}^u]_1$ and $q \in \mathcal{Q}$,

- 1) for all $\tilde{v}^u \sim \tilde{\mathcal{N}}(m, s^2) \in \tilde{\mathbb{T}}_1^u(\alpha^{-1}(q), \cdot)$, store $\{(m_l, s_l) = (\eta \lfloor \frac{m}{\eta} \rfloor, \eta^2 \lfloor \frac{s^2}{\eta^2} \rfloor)\}_l$;
- 2) for each l , define $\tilde{v}_l^{u, \text{ref}} \sim \tilde{\mathcal{N}}(m_l, s_l)$ (implicitly, we need to compute $\tilde{v}_l^{u, \text{ref}}(\alpha^{-1}(\Delta))$); compute $\tilde{v}_l^{u, \text{ref}}(\alpha^{-1}(q_j))$ for each $q_j \in \mathcal{Q} \setminus \Delta$;
- 3) for each l , define $\mu_l^{u, \text{ref}} = [\tilde{v}_l^{u, \text{ref}}(\alpha^{-1}(q_1)), \dots, \tilde{v}_l^{u, \text{ref}}(\alpha^{-1}(q_N)), \tilde{v}_l^{u, \text{ref}}(\alpha^{-1}(\Delta))]$;
- 4) compute $\mathbf{ws} := (\sqrt{2n} + 2)\eta$ and $\mathbf{tv} := 2 \cdot \mathbf{ws}$;
- 5) construct $[\mu^u] = \bigcup_l \{\mu : \|\mu - \mu_l^{u, \text{ref}}\|_{\text{TV}} \leq \mathbf{tv}(\eta), \mu(\Delta) + \sum_j \mu(q_j) = 1\}$;
- 6) let $\Sigma_\alpha := \{\{\tilde{v}^u, \mu^u\}, \mu^u \in [\mu^u]\}$ be a relation between $\tilde{v}^u \in \tilde{\mathbb{T}}^u(\alpha^{-1}(q))$ and the generated $[\mu^u]$.

Repeat the above step for all q and then for all $u = a \in \text{Act}$, the relation Σ_α is obtained. We denote $\mathfrak{G}_i^u := \tilde{\mathbb{T}}_1^u(\alpha^{-1}(q_i), \cdot)$ and $\hat{\mathfrak{G}}_i^u := \widehat{\tilde{\mathbb{T}}_1^u}(\alpha^{-1}(q_i), \cdot)$.

Step 1: For each $u = a \in \text{Act}$, for $i \leq N$, let $[\Theta_i^u] = \Sigma_\alpha(\hat{\mathfrak{G}}_i^u)$ and the transition collection be $[\Theta^u]$. It can be shown that the finite-state BMDP $\mathbb{I}\mathbb{A}$ abstracts $\mathbb{X}U_1$ based on Definition 3.6: for each a , there exists $u \in \mathcal{U}$ (where we set it to be a), such that for any $\tilde{v}^u \in \mathfrak{G}_i^u$ and hence in $\hat{\mathfrak{G}}_i^u$, there exists a discrete measures in $\Theta_i^u \in \Sigma_\alpha(\hat{\mathfrak{G}}_i^u)$ such that for all q_j we have $\tilde{v}^u(\alpha^{-1}(q_j)) = \Theta_{ij}^u$.

The proof is done by the exact same way as the proof of [40, Claim 1, Theorem 4] for each fixed control input. We summarize the methodology as follows:

- a) To not miss any possible transition of $\mathbb{X}U_1$ from each $x \in \alpha^{-1}(q_i)$, we work on the over-approximation set $\hat{\mathfrak{G}}_i^u$ of Gaussian measures. It can be easily verified that, within the abstractions, we also have $\Sigma_\alpha(\hat{\mathfrak{G}}_i^u) \supseteq \Sigma_\alpha(\mathfrak{G}_i^u)$. Now we verify that Σ_α given in 6) is indeed a valid relation that creates an abstraction.
- b) For any modified Gaussian $\tilde{v}^u \in \hat{\mathfrak{G}}_i^u$, there exists a $\tilde{v}^{u, \text{ref}}$ such that the distance is bounded: $\|\tilde{v} - \tilde{v}^{u, \text{ref}}\|_W \leq \|\tilde{v} - \tilde{v}^{u, \text{ref}}\|_W \leq \sqrt{2n}\eta$. This is estimated by Proposition 2.10.
- c) Reflecting on the space of discrete measures (a row of an abstraction matrix), we have the following inflation

$$\begin{aligned} & \|\mu - \mu^{u, \text{ref}}\|_W^d \\ & \leq \|\mu - \tilde{v}\|_W + \|\tilde{v} - \tilde{v}^{u, \text{ref}}\|_W + \|\tilde{v}^{u, \text{ref}} - \mu^{u, \text{ref}}\|_W \\ & \leq \mathbf{ws}, \end{aligned} \quad (17)$$

where the first and third term above is to connect a discretized measure with a continuous measure. Note that for any continuous measure \mathbf{v} and its discretized version \mathbf{m} ,

$$\begin{aligned} & \|\mathbf{m} - \mathbf{v}\|_{\mathbb{W}} \\ &= \sup_{h \in C(\mathcal{X}), \text{Lip}(h) \leq 1} \left| \int_{\mathcal{X}} h(x) d\mathbf{m}(x) - \int_{\mathcal{X}} h(x) d\mathbf{v}(x) \right| \\ &\leq \sup_{h \in C(\mathcal{X}), \text{Lip}(h) \leq 1} \sum_{j=1}^n \int_{\alpha^{-1}(q_j)} |h(x) - h(q_j)| d\mathbf{v}(x) \quad (18) \\ &\leq \eta \sum_{j=1}^n \int_{\alpha^{-1}(q_j)} d\mathbf{v}(x) \leq \eta. \end{aligned}$$

By 5) an 6) and Definition 3.6, we have stored such μ centered at the reference measure w.r.t. the total variation distance, and this collection has sufficient amount of transition matrices as a valid abstraction by definition.

Step 2: Now we choose of η and ϱ such that the constructed BMDP abstraction can be abstracted by \mathbb{XU}_2 via the converse relation Σ_{α}^{-1} . Note that $a \in u + \varrho \overline{\mathbb{B}}$ for any $u \in \mathcal{U}$. We need to choose η , ϱ and k sufficiently small such that

$$2\eta + 1/2 \cdot \mathbf{tv}(\eta) + L\rho + k \leq \vartheta_2 - \vartheta_1, \quad (19)$$

where L is the Lipschitz constant of f . Then, we have

$$\begin{aligned} \Sigma_{\alpha}^{-1}(\Sigma_{\alpha}(\widehat{\mathfrak{G}}_i^u)) &\subseteq \widehat{\mathfrak{G}}_i^u + (2\eta + 1/2 \cdot \mathbf{tv}(\eta)) \cdot \overline{\mathbb{B}}_{\mathbb{W}} + L\varrho \overline{\mathbb{B}} \\ &\subseteq \mathfrak{G}_i^u + (2\eta + 1/2 \cdot \mathbf{tv}(\eta) + L\varrho + k) \cdot \overline{\mathbb{B}}_{\mathbb{W}} \quad (20) \end{aligned}$$

for each i . Note that all the ‘ref’ information is recorded, and, particularly, for any $\mu^u \in \Sigma_{\alpha}(\mathfrak{G}_i^u)$ there exists a $\mu^{u,\text{ref}}$ within a total variation radius $\mathbf{tv}(\eta)$.

The inclusions in (20) are to conversely find all possible corresponding measure $\tilde{\nu}^u$ that matches μ^u by their probabilities on discrete nodes. All such $\tilde{\nu}^u$ should satisfy,

$$\begin{aligned} & \|\tilde{\nu}^u - \mu^{u,\text{ref}}\|_{\mathbb{W}} \\ &\leq \|\tilde{\nu} - \mu\|_{\mathbb{W}} + \|\mu - \mu^{u,\text{ref}}\|_{\mathbb{W}}^d + \|\mu^{u,\text{ref}} - \tilde{\nu}^{u,\text{ref}}\|_{\mathbb{W}} \\ &\leq 2\eta + 1/2 \cdot \mathbf{tv}(\eta), \quad (21) \end{aligned}$$

where the bounds for the first and third terms are obtained in the same way as (18). The second term is improved compared to [40] based on the connection (9).

By the construction, we can verify that for each $u \in \mathcal{U}$, there exists an $a \in \text{Act}$ (guaranteed by the finite covering relation $a \in u + \varrho \overline{\mathbb{B}}$) such that the choice in (20) makes $\Sigma_{\alpha}^{-1}(\Sigma_{\alpha}(\widehat{\mathfrak{G}}_i^u)) \subseteq \overline{\mathbb{T}}_2^u(\alpha^{-1}(q_i))$, which completes the proof. \blacksquare

Remark 5.4: As noted in [40], the key point of the construction in Theorem 5.3 is to record the ‘ref’ points and corresponding radii, which form finite coverings of the compact space of measures. We use ‘finite-state’ instead of ‘finite’ abstraction because we do not further discretize the dual space

of the solution processes, which is the space of probability measures.

Remark 5.5: As shown in Step 1.b), we estimate the Wasserstein distance with the reference measure by the second moment difference of the associated random variables. In this view, we can also replace the additional uncertainty $\xi^{(1)}$, which is a sequence of point-mass perturbations, by a sequence of \mathcal{L}^1 independent noise with known, bounded second moment. In this case, we can still eventually obtain a robust complete abstraction by a similar methodology.

Theorem 5.3 shows that, given any $0 \leq \vartheta_1 < \vartheta_2$, there exists a sufficiently and necessarily refined uniform discretization of \mathcal{X} , as well as a measure-level relation such that a robustly complete abstraction \mathbb{IA} can be constructed. We can algorithmically synthesize a control strategy for \mathbb{XU}_1 by generating \mathbb{IA} and then solving a discrete synthesis problem for \mathbb{IA} with some probabilistic specification. In view of Corollary 4.7, if a control strategy κ exists to fulfill the probabilistic specification for \mathbb{IA}^{κ} , then there exists a policy ϕ to guarantee the satisfaction of \mathbb{XU}_1^{ϕ} . On the other hand, if there is no policies to realize a specification for \mathbb{IA} , then the system \mathbb{XU}_2 is also not realizable w.r.t. the same specification. The latter is implied by the following corollary.

Corollary 5.6: Given a specification formula Ψ , let $S_2^{\nu_0, \phi} = \{\mathbf{P}_X^{\nu_0, \phi}(X \models \Psi)\}_{X \in \mathbb{XU}_2}$ be the set of the satisfaction probability of Ψ under a control policy ϕ for the system \mathbb{XU}_2 . Then, for each control policy ϕ of \mathbb{XU}_2 , there exists a policy κ for \mathbb{IA} such that $S_{\mathbb{IA}}^{q_0, \kappa} \subseteq S_2^{\nu_0, \phi}$ for any initial conditions satisfying $\nu_0(\alpha^{-1}(q_0)) = 1$, where $S_{\mathbb{IA}}^{q_0, \kappa} = \{\mathcal{P}_I^{q_0, \kappa}(I \models \Psi)\}_{I \in \mathbb{IA}}$. Both $S_{\mathbb{IA}}^{q_0, \kappa}$ and $S_2^{\nu_0, \phi}$ are compact.

Proof: The inclusion and compactness (for each policy) is done in a similar way as Theorem 4.6 by the inductive construction of probability laws. \blacksquare

VI. A DISCUSSION ON STOCHASTIC CONTROL SYSTEMS WITH NOISY OBSERVATION

Continuing our objective of developing abstractions, in this section, we discuss the case when the observations of the sample paths are corrupted by noise.

Apart from the nonlinear filtering, the philosophy of constructing sound and robustly complete abstractions for such systems maintain the same. We hence do not reiterate the procedure in this section but rather deliver a discussion on the mathematical complexity of the potential abstractions. Before we proceed, we briefly introduce the notion from nonlinear filtering.

A. NONLINEAR FILTERING FOR DISCRETE-TIME SYSTEMS

Consider the discrete-time signal and observation of the following form

$$X_{t+1} = f(X_t, u_t) + b(X_t)w_t, \quad (22a)$$

$$Y_t = h(X_t) + \beta_t, \quad (22b)$$

where Y is a \mathcal{Y} -valued observation via a continuous Borel measurable function h and i.i.d. Gaussian process $\beta := \{\beta_t\}_{t \in \mathbb{N}}$ with proper dimensions. We also set \mathbf{w} and β to be mutually independent.

Similar to (2), for any fixed $t > 0$, we define the short hand notation for the history of observation $Y_{[0,t]} := \{Y_s\}_{s \in [0,t]}$. Unlike the system without corrupted observations, it is natural to suppose that the selection of a control at time t is based on $Y_{[0,t]}$ and $\mathbf{u}_{[0,t-1]}$. An admissible control policy $\kappa = \{\kappa_t\}_{t \in \mathbb{N}}$ in this case is such that, for each fixed $t > 0$, we have, for any $\mathcal{C} \in \mathcal{B}(\mathcal{U})$,

$$\kappa_t(\mathcal{C} \mid Y_{[0,t]}; \mathbf{u}_{[0,t-1]}) = \mathbf{P}[\mathbf{u}_t \in \mathcal{C} \mid Y_{[0,t]}; \mathbf{u}_{[0,t-1]}]. \quad (23)$$

A deterministic admissible policy κ is such that $\mathbf{u}_t = \kappa_t(Y_{[0,t]}; \mathbf{u}_{[0,t-1]})$.

Let $H(Y_t \in A \mid X_t = x_t)$, $A \in \mathcal{B}(\mathcal{Y})$, be the observation channel, which is the transition kernel generated by (22b). Given any initial distribution μ_0 of X , the probability law $\mathbf{P}^{\mu_0, \kappa}$ of $(X, Y, \mathbf{u}) := \{X_t, Y_t, \mathbf{u}_t\}_{t \in \mathbb{N}}$ can be uniquely determined based on the the transition kernel, observation channel, and the control policy.

Given a policy κ (we set it to be deterministic without loss of generality), the estimation of X_t given $Y_{[0,t]}$ that minimizes the mean square error loss is given as

$$\Pi_t(\Gamma) := \mathbf{P}^{\mu_0, \kappa}[X_t \in \Gamma \mid Y_{[0,t]}, \mathbf{u}_{[0,t-1]}], \quad \Gamma \in \mathcal{B}(\mathcal{X}).$$

We call this random measure $\Pi_t \in \mathfrak{P}(\mathcal{X})$ for each t the optimal filter. Using Bayes rule, we have

$$\begin{aligned} \Pi_t(\Gamma) &= \mathbf{P}^{\mu_0, \kappa}[X_t \in \Gamma \mid Y_{[0,t]}, \mathbf{u}_{[0,t-1]}] \\ &= \frac{\int_{\mathcal{X}} H(Y_t \mid X_t = x_t) \Theta_{t-1}^{\mathbf{u}_{t-1}}(x_{t-1}, \Gamma) \cdot \Pi_{t-1}(dx_{t-1})}{\int_{\mathcal{X}} \int_{\mathcal{X}} H(Y_t \mid X_t = x_t) \Theta_{t-1}^{\mathbf{u}_{t-1}}(x_{t-1}, dx_t) \cdot \Pi_{t-1}(dx_{t-1})} \\ &=: F(\Pi_{t-1}, Y_{t-1}, \mathbf{u}_{t-1})(\Gamma), \end{aligned} \quad (24)$$

where \mathbf{u}_{t-1} is determined by $\kappa_{t-1}(Y_{[0,t-1]}, \mathbf{u}_{[0,t-2]})$. It can also be shown that the process $(\Pi, \mathbf{u}) := \{\Pi_t, \mathbf{u}_t\}_{t \in \mathbb{N}}$ is a controlled Markov process [21] with transition probability

$$\begin{aligned} &\mathbf{P}[\Pi_{t+1} \in D \mid \Pi_t = \pi_t, \mathbf{u}_t = \mathbf{u}_t] \\ &= \int_{\mathcal{Y}} \mathbb{1}_{\{F(\pi_t, y_t, \mathbf{u}_t) \in D\}} \cdot \mathbf{n}(dy_t), \quad D \in \mathcal{B}(\mathfrak{P}(\mathcal{X})). \end{aligned} \quad (25)$$

We also use $\Pi^{\mathbf{u}}$ to emphasize the marginal behavior of the process (Π, \mathbf{u}) . Given the observations and the adaptively generated control signal, the optimal estimation of the conditional probability of satisfying any ω -regular formula Ψ is given by

$$\mathbf{P}_{\Pi}^{\mu_0, \mathbf{u}}[X \models \Psi] := \mathbf{P}^{\mu_0, \mathbf{u}}[X \models \Psi \mid Y] = \int_{\mathcal{X}^{\infty}} \mathbb{1}_{\{X \models \Psi\}} \Pi^{\mathbf{u}}(dx). \quad (26)$$

Note that it is difficult to obtain the full knowledge of Y , our goal is to generate control policies such that the optimal estimation $\mathbf{P}^{\mu_0, \mathbf{u}}[X \models \Psi \mid Y]$ possesses certain confidence of satisfying the probabilistic requirement given any realization

of observation. The above derivation converts the problem into a fully observed controlled Markov process (Π, \mathbf{u}) via an enlargement of the state space, where control policies and even optimal control policies can be synthesized accordingly for the (hypothetically) fully observed Π [21]. The policy fulfilling the goal mentioned above is thereby decidable.

The construction of the optimal filter process (or the function F in (24)) can be decomposed into a two-step recursion based on the transition relation in (24).

Prediction (Prior): At time t , $\Pi_{t-1}(dx)$ is feed into the r.h.s. of the prior knowledge of the dynamics for X , i.e., (22a). The prediction of X_t based on $Y_{[0,t-1]}$ as well as the \mathbf{u} determined at t is such that

$$\hat{\Pi}_t(dx) = \int_{\mathcal{X}} \Theta_t^{\mathbf{u}}(\bar{x}, dx) \Pi_{t-1}(d\bar{x}). \quad (27)$$

Filtering (Posterior): This step is to assimilate the observation at the instant t , which is given as

$$\Pi_t(dx) = \mathbf{n}(Y_t) H(Y_t \mid X_t = x) \hat{\Pi}_t(dx), \quad (28)$$

where $\mathbf{n}(Y_t) = \int_{\mathcal{X}} H(Y_t \mid X_t = x) \hat{\Pi}_t(dx)$ is the normalizer.

For numerical approximation, we simulate and propagate the optimal filter process using matrix approximations of each step's transition kernel, whereas for formal abstractions, we need to find the 'inclusion' of the transitions for each step as usual.

B. A BRIEF DISCUSSION ON STOCHASTIC ABSTRACTIONS FOR CONTROL SYSTEMS WITH NOISY OBSERVATIONS

Motivated by generating optimal control policies using the knowledge of filter process (Π, \mathbf{u}) , the stochastic abstractions for partially observed processes can be reduced to obtain a sound and robustly complete abstraction for the process (Π, \mathbf{u}) . To convey the idea, we simply consider the following two systems with noisy observations

$$X_{t+1} = f(X_t, \mathbf{u}_t) + b(X_t) \mathbf{w}_t + \vartheta_1 \xi_t^{(1)}, \quad (29a)$$

$$Y_t = X_t + \beta_t + \zeta_1 \zeta_t^{(1)}, \quad (29b)$$

and

$$X_{t+1} = f(X_t, \mathbf{u}_t) + b(X_t) \mathbf{w}_t + \vartheta_2 \xi_t^{(2)}, \quad (30a)$$

$$Y_t = X_t + \beta_t + \zeta_2 \zeta_t^{(2)}, \quad (30b)$$

where $\zeta_t^{(i)} \in \bar{\mathcal{B}}$ are i.i.d. for each t and each $i \in \{1, 2\}$, the intensities satisfy $0 \leq \zeta_1 < \zeta_2$. The rest of the notations are as previously mentioned.

We convert the filter processes that are generated by (29) and (30) into the expression of controlled Markov systems

$$\mathbb{F}\mathbb{U}_i = (\mathcal{X}, \mathcal{Y}, \mathcal{U}, \{T\}_i, \llbracket H \rrbracket_i, \text{AP}, L_{\mathbb{F}\mathbb{U}}), \quad i = 1, 2, \quad (31)$$

where the additional \mathcal{Y} and its collection of observation channel $\llbracket H \rrbracket_i$ are needed in the filtering step for generating the controlled filter process $\Pi^{\mathbf{u}}$. The other notions are the same as previously mentioned.

To find an abstraction for $\mathbb{F}\mathbb{U}_1$, we need a state-level relation or discretization α as usual. Then, we need both $\{T\}_1$ and $\llbracket H \rrbracket_1$ to be abstracted via some measure-level relation, so that the transition probability of Π is abstracted by a set of discrete transition probabilities given the same set of discrete observations in the sense of (2) of Definition 3.6. Now we denote the BMDP abstraction for $\mathbb{F}\mathbb{U}_1$ as

$$\mathbb{I}\mathbb{A} = (\mathcal{Q}, \mathcal{Y}_{\mathcal{Q}}, Act, \{\Theta\}, \llbracket H_{\mathcal{Q}} \rrbracket, AP, L_{\mathbb{I}\mathbb{A}}), \quad (32)$$

where $\mathcal{Y}_{\mathcal{Q}}$ is the discretized observation states that are obtained by the state-level relation α , and $\llbracket H_{\mathcal{Q}} \rrbracket$ is the collection of the discrete observation channels that are obtained based on some measure-level relation Σ_{α} . The intuition of $\mathbb{I}\mathbb{A}$ is that we need ‘more’ transitions in the abstraction for the prior knowledge of the dynamics that are related via the measurability of labelled nodes, as well as ‘more’ transitions for the filtering step to obtain enough observations for decision making.

The soundness of $\mathbb{I}\mathbb{A}$ for the controlled filter process system $\mathbb{F}\mathbb{U}_1$ is in the following sense: given any initial distribution, for each κ (based on the discrete observation $Y_{\mathcal{Q}}$) for $\mathbb{I}\mathbb{A}$, there exists a control policy ϕ such that for each $\Pi^{\phi} \in \mathbb{F}\mathbb{U}_1^{\phi}$,

- there exists a $\Pi^{d,\kappa} \in \mathbb{I}\mathbb{A}^{\kappa}$ whose observation process $Y_{\mathcal{Q}}$ has the same probability with Y of Π^{ϕ} on each discrete node $q \in \mathcal{Q}$, and $\Pi^{d,\kappa}$ has the same evaluation on all the discrete measurable sets $A \in \mathcal{F}$ with Π^{ϕ} ;
- the discrete probability law $\mathbb{P}^{d,\kappa}$ for $\Pi^{d,\kappa} \in \mathbb{I}\mathbb{A}^{\kappa}$ forms a convex and weakly compact set;
- the optimal estimation satisfies, for a given $p \in [0, 1]$,

$$\left. \begin{aligned} & \int_{\Pi^{\phi} \in \mathcal{B}(\mathfrak{P}(\mathcal{X}))} \mathbb{1}_{\{\mathbb{P}_{\Pi}^{\nu_0, \phi} [X^{\phi} \models \Psi] \triangleright \triangleleft p\}} \mathbb{P}^{\phi}(d\Pi^{\phi}) \\ & \in \left\{ \int_{\Pi^{d,\kappa} \in \mathcal{B}(\mathfrak{P}(\mathcal{Q}))} \mathbb{1}_{\{\mathbb{P}_{\Pi}^{\mu_0, \kappa} [X^{\phi} \models \Psi] \triangleright \triangleleft p\}} \mathbb{P}^{d,\kappa}(d\Pi^{d,\kappa}) \right\}. \end{aligned} \right\} \quad (33)$$

A proper task is to find a control policy such that the optimal estimation of the probabilistic specification of $X \models \Psi$ has a confidence at least $q \in [0, 1]$, i.e., $\mathbb{P}^{\phi}(\mathbb{P}_{\Pi}^{\mu_0, \phi} [X \models \Psi \mid Y] \triangleright \triangleleft p) \geq q$. Then we can search control policies κ in $\mathbb{I}\mathbb{A}$ for all the filter process Π^d , such that strategy can make the lower bound of

$$\left\{ \int_{\Pi^{d,\kappa} \in \mathcal{B}(\mathfrak{P}(\mathcal{Q}))} \mathbb{1}_{\{\mathbb{P}_{\Pi}^{\mu_0, \kappa} [X^{\phi} \models \Psi] \triangleright \triangleleft p\}} \mathbb{P}^{d,\kappa}(d\Pi^{d,\kappa}) \right\}_{\Pi^{d,\kappa} \in \mathbb{I}\mathbb{A}^{\kappa}}$$

greater than or equal to q .

The robust completeness can be verified in a similar way as Section V, except now we need to decompose the procedure to guarantee the robust completeness for both prediction and filtering steps. The discretization need to rely on the value of $\vartheta_2 - \vartheta_1$ and $\varsigma_2 - \varsigma_1$.

Note that the counterpart of Fokker-Planck equations for evaluating the probability law of the optimal filter in systems with noisy observations is the famous Zakai’s stochastic partial differential equation⁷. The approximation of such a

solution already suffers from the curse of dimensionality. Using formal abstractions to enlarge the partially observed processes into the filter processes with full observations, based on which control policies can be determined and utilized back to the partially observed cases, seems tedious and impractical. Besides the theoretical formal guarantee of a confidence of a satisfaction probability (i.e., a probabilistic requirement of the probabilistic specification), the abstraction essentially solves the continuous probability law of a continuous conditional expectation (or a random measure) upon some process with discrete labels using discrete inclusions. We hence do not recommend readers to complicate the problem.

VII. CONCLUSION

In this article, we investigated the mathematical properties of formal abstractions for discrete-time controlled nonlinear stochastic systems. We discussed the motivation of constructing sound and complete formal stochastic abstractions and the philosophy in comparison to numerical approximations in Section IV. A brief discussion on the extension of stochastic abstractions for controlled stochastic systems with noisy observation was provided in Section VI. The construction of such abstractions can be analogous to solving a discretized version of Zakai’s equation via formal inclusions, which suffers from a curse of excessive dimensionality.

We provide an appropriate mathematical language to discuss the soundness and approximate completeness of abstractions for stochastic systems using BMDP. We show that abstractions with extra uncertainties are not straightforward extensions of their non-stochastic counterparts [7], [11], and view this as the most significant contribution of our work.

For future work, it would be interesting to design algorithms to construct robustly complete BMDP abstractions for more general robust stochastic systems with \mathcal{L}^1 perturbations based on the weak topology. The size of state discretization can be refined given more specific assumptions on system dynamics and linear-time objectives. It is also of a theoretical interest to construct robustly complete abstractions for continuous-time stochastic system and demonstrate the controllability given mild conditions. Even though we aimed to provide a theoretical foundation of BMDP abstractions for continuous-state stochastic systems, we hope the results can shed some light on designing more powerful robust control synthesis algorithms.

REFERENCES

- [1] C. Baier and J.-P. Katoen, *Principles of Model Checking*. Cambridge, MA, USA: MIT Press, 2008.
- [2] C. Belta, B. Yordanov, and E. A. Gol, *Formal Methods for Discrete-Time Dynamical Systems*, vol. 89. Berlin, Germany: Springer, 2017.
- [3] Y. Li and J. Liu, *Formal Methods for Control of Nonlinear Systems*. London, U.K.: Chapman and Hall/CRC, 2023.
- [4] J. Liu and N. Ozay, “Finite abstractions with robustness margins for temporal logic-based control synthesis,” *Nonlinear Anal.: Hybrid Syst.*, vol. 22, pp. 1–15, 2016.

⁷We omit the content here and kindly refer readers to [63] for details.

- [5] P. Nilsson, N. Ozay, and J. Liu, "Augmented finite transition systems as abstractions for control synthesis," *Discrete Event Dyn. Syst.*, vol. 27, no. 2, pp. 301–340, 2017.
- [6] N. Ozay, J. Liu, P. Prabhakar, and R. M. Murray, "Computing augmented finite transition systems to synthesize switching protocols for polynomial switched systems," in *Proc. IEEE Amer. Control Conf.*, 2013, pp. 6237–6244.
- [7] J. Liu, "Robust abstractions for control synthesis: Completeness via robustness for linear-time properties," in *Proc. 20th Int. Conf. Hybrid Syst. Comput. Control*, 2017, pp. 101–110.
- [8] Y. Li and J. Liu, "Robustly complete synthesis of memoryless controllers for nonlinear systems with reach-and-stay specifications," *IEEE Trans. Autom. Control*, vol. 66, no. 3, pp. 1199–1206, Mar. 2021.
- [9] M. Kloetzer and C. Belta, "A fully automated framework for control of linear systems from temporal logic specifications," *IEEE Trans. Autom. Control*, vol. 53, no. 1, pp. 287–297, Feb. 2008.
- [10] P. Tabuada and G. J. Pappas, "Linear time logic control of discrete-time linear systems," *IEEE Trans. Autom. Control*, vol. 51, no. 12, pp. 1862–1877, Dec. 2006.
- [11] J. Liu, "Closing the gap between discrete abstractions and continuous control: Completeness via robustness and controllability," in *Proc. 19th Int. Conf. Formal Model. Anal. Timed Syst.*, 2021, pp. 67–83.
- [12] M. Kwiatkowska, G. Norman, and D. Parker, "PRISM: Probabilistic symbolic model checker," in *Proc. Comput. Perform. Eval. Modelling Techn. Tools: 12th Int. Conf.*, 2002, pp. 200–204.
- [13] D. Parker, "Verification of probabilistic real-time systems," in *Proc. Real-time Syst. Summer Sch.*, 2013.
- [14] K. Chatterjee, M. Chmelik, and M. Tracol, "What is decidable about partially observable Markov decision processes with ω -regular objectives," *J. Comput. Syst. Sci.*, vol. 82, no. 5, pp. 878–911, 2016.
- [15] C. Hensel, S. Junges, J.-P. Katoen, T. Quatmann, and M. Volk, "The probabilistic model checker storm," *Int. J. Softw. Tools Technol. Transfer*, vol. 24, pp. 589–610, 2022.
- [16] M. Kwiatkowska, G. Norman, and D. Parker, "Probabilistic model checking and autonomy," *Annu. Rev. Control Robot. Auton. Syst.*, vol. 5, pp. 385–410, 2022.
- [17] X. C. Ding, S. L. Smith, C. Belta, and D. Rus, "MDP optimal control under temporal logic constraints," in *Proc. IEEE 50th Conf. Decis. Control Eur. Control Conf.*, 2011, pp. 532–538.
- [18] X. C. D. Ding, S. L. Smith, C. Belta, and D. Rus, "LTL control in uncertain environments with probabilistic satisfaction guarantees," *IFAC Proc. Volumes*, vol. 44, no. 1, pp. 3515–3520, 2011.
- [19] B. Lacerda, D. Parker, and N. Hawes, "Optimal and dynamic planning for Markov decision processes with co-safe LTL specifications," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2014, pp. 1511–1516.
- [20] A. M. Wells, M. Lahijanian, L. E. Kaviraki, and M. Y. Vardi, "LTLf synthesis on probabilistic systems," 2020, *arXiv:2009.10883*.
- [21] N. Saldi, S. Yüksel, and T. Linder, "Finite model approximations for partially observed Markov decision processes with discounted cost," 2017, *arXiv:1710.07009*.
- [22] G. Norman, D. Parker, and X. Zou, "Verification and control of partially observable probabilistic systems," *Real-Time Syst.*, vol. 53, pp. 354–402, 2017.
- [23] F. Ramponi, D. Chatterjee, S. Summers, and J. Lygeros, "On the connections between PCTL and dynamic programming," in *Proc. 13th ACM Int. Conf. Hybrid Syst.: Comput. Control*, 2010, pp. 253–262.
- [24] S. E. Z. Soudjani and A. Abate, "Adaptive gridding for abstraction and verification of stochastic hybrid systems," in *Proc. IEEE 8th Int. Conf. Quantitative Eval. Syst.*, 2011, pp. 59–68.
- [25] S. Summers and J. Lygeros, "Verification of discrete time stochastic hybrid systems: A stochastic reach-avoid decision problem," *Automatica*, vol. 46, no. 12, pp. 1951–1961, 2010.
- [26] A. Abate, M. Prandini, J. Lygeros, and S. Sastry, "Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems," *Automatica*, vol. 44, no. 11, pp. 2724–2734, 2008.
- [27] A. Abate, J.-P. Katoen, and A. Mereacre, "Quantitative automata model checking of autonomous stochastic hybrid systems," in *Proc. Hybrid Syst.: Comput. Control*, 2011, pp. 83–92.
- [28] I. Tkachev and A. Abate, "On infinite-horizon probabilistic properties and stochastic bisimulation functions," in *Proc. IEEE 50th Conf. Decis. Control Eur. Control Conf.*, 2011, pp. 526–531.
- [29] I. Tkachev and A. Abate, "Regularization of bellman equations for infinite-horizon probabilistic properties," in *Proc. 15th ACM Int. Conf. Hybrid Syst.: Comput. Control*, 2012, pp. 227–236.
- [30] I. Tkachev and A. Abate, "Formula-free finite abstractions for linear temporal verification of stochastic hybrid systems," in *Proc. 16th Int. Conf. Hybrid Syst.: Comput. Control*, 2013, pp. 283–292.
- [31] I. Tkachev and A. Abate, "Characterization and computation of infinite-horizon specifications over Markov processes," *Theor. Comput. Sci.*, vol. 515, pp. 1–18, 2014.
- [32] I. Tkachev, A. Mereacre, J.-P. Katoen, and A. Abate, "Quantitative model-checking of controlled discrete-time Markov processes," *Inf. Comput.*, vol. 253, pp. 1–35, 2017.
- [33] R. Givan, S. Leach, and T. Dean, "Bounded-parameter Markov decision processes," *Artif. Intell.*, vol. 122, no. 1/2, pp. 71–109, 2000.
- [34] M. Lahijanian, S. B. Andersson, and C. Belta, "Formal verification and synthesis for discrete-time stochastic systems," *IEEE Trans. Autom. Control*, vol. 60, no. 8, pp. 2031–2045, Aug. 2015.
- [35] D. Wu and X. Koutsoukos, "Reachability analysis of uncertain systems using bounded-parameter Markov decision processes," *Artif. Intell.*, vol. 172, no. 8/9, pp. 945–954, 2008.
- [36] N. Cauchi, L. Laurenti, M. Lahijanian, A. Abate, M. Kwiatkowska, and L. Cardelli, "Efficiency through uncertainty: Scalable formal synthesis for stochastic hybrid systems," in *Proc. 22nd ACM Int. Conf. Hybrid Syst.: Comput. Control*, 2019, pp. 240–251.
- [37] M. D. H. Dutreix, "Verification and synthesis for stochastic systems with temporal logic specifications," Ph.D. dissertation, Georgia Inst. Technol., Atlanta, GA, USA, 2020.
- [38] M. Dutreix, J. Huh, and S. Coogan, "Abstraction-based synthesis for stochastic systems with omega-regular objectives," *Nonlinear Anal.: Hybrid Syst.*, vol. 45, 2022, Art. no. 101204.
- [39] G. Delimpaltadakis, L. Laurenti, and M. Mazo, "Abstracting the sampling behaviour of stochastic linear periodic event-triggered control systems," in *Proc. IEEE 60th Conf. Decis. Control*, 2021, pp. 1287–1294.
- [40] Y. Meng and J. Liu, "Robustly complete finite-state abstractions for verification of stochastic systems," in *Proc. Formal Model. Anal. Timed Syst.: 20th Int. Conf.*, 2022, pp. 80–97.
- [41] M. Dutreix and S. Coogan, "Specification-guided verification and abstraction refinement of mixed monotone stochastic systems," *IEEE Trans. Autom. Control*, vol. 66, no. 7, pp. 2975–2990, Jul. 2021.
- [42] A. Abate, A. D'Innocenzo, M. D. Di Benedetto, and S. S. Sastry, "Markov set-chains as abstractions of stochastic hybrid systems," in *Proc. Int. Workshop Hybrid Syst.: Comput. Control*, 2008, pp. 1–15.
- [43] L. Laurenti, M. Lahijanian, A. Abate, L. Cardelli, and M. Kwiatkowska, "Formal and efficient synthesis for continuous-time linear stochastic hybrid processes," *IEEE Trans. Autom. Control*, vol. 66, no. 1, pp. 17–32, Jan. 2020.
- [44] I. I. Gihman and A. V. Skorohod, *Controlled Stochastic Processes*. Berlin, Germany: Springer, 2012.
- [45] M. Y. Vardi, "Automatic verification of probabilistic concurrent finite state programs," in *Proc. IEEE 26th Annu. Symp. Found. Comput. Sci.*, 1985, pp. 327–338.
- [46] L. C. G. Rogers and D. Williams, *Diffusions, Markov Processes and Martingales, Volume 1: Foundations*, vol. 1, Cambridge university press, 2000.
- [47] M. Hairer and X.-M. Li, *Markov Processes*, lecture notes, Imperial College London, Sep. 12, 2022.
- [48] G. Da Prato and J. Zabczyk, *Stochastic Equations in Infinite Dimensions*. New York, NY, USA: Cambridge Univ. Press, 2014.
- [49] G. Sagar and D. Ravi, "Compactness of any countable product of compact metric spaces in product topology without using Tychonoff's theorem," 2021, *arXiv:2111.02904*.
- [50] Y. Meng and J. Liu, "Robustly complete finite-state abstractions for verification of stochastic systems," in *Proc. Int. Conf. Formal Model. Anal. Timed Syst.*, 2022, pp. 80–97.
- [51] P. Billingsley, *Convergence of Probability Measures*. New York, NY, USA: Wiley, 2013.
- [52] A. L. Gibbs and F. E. Su, "On choosing and bounding probability metrics," *Int. Stat. Rev.*, vol. 70, no. 3, pp. 419–435, 2002.
- [53] M. T. Barlow and E. A. Perkins, "Brownian motion on the sierpinski gasket," *Probability Theory Related Fields*, vol. 79, no. 4, pp. 543–623, 1988.
- [54] P. H. Baxendale, "A stochastic Hopf bifurcation," *Probability Theory Related Fields*, vol. 99, pp. 581–616, 1994.
- [55] D. Blömker, M. Hairer, and G. Pavliotis, "Multiscale analysis for stochastic partial differential equations with quadratic nonlinearities," *Nonlinearity*, vol. 20, no. 7, 2007, Art. no. 1721.

- [56] D. Blömker and H. Fu, “The impact of multiplicative noise in SPDEs close to bifurcation via amplitude equations,” *Nonlinearity*, vol. 33, no. 8, 2020, Art. no. 3905.
- [57] Y. Meng and J. Liu, “Sufficient conditions for robust probabilistic reach-avoid-stay specifications using stochastic Lyapunov-Barrier functions,” in *Proc. IEEE Amer. Control Conf.*, 2022, pp. 2283–2288.
- [58] Y. Meng, “Bifurcation and robust control of instabilities in the presence of uncertainties,” Ph.D. dissertation, Univ. Waterloo, Waterloo, ON, Canada, 2022.
- [59] A. Subbaraman, *Robust Stability Theory for Stochastic Dynamical Systems*. Univ. California, Santa Barbara, Santa Barbara, CA, USA, 2015.
- [60] A. R. Teel, J. P. Hespanha, and A. Subbaraman, “A converse Lyapunov theorem and robustness for asymptotic stability in probability,” *IEEE Trans. Autom. Control*, vol. 59, no. 9, pp. 2426–2441, Sep. 2014.
- [61] Y. Li, W. Zhang, and X. Liu, “Stability of nonlinear stochastic discrete-time systems,” *J. Appl. Math.*, vol. 2013, pp. 1–8, 2013.
- [62] P. E. Kloeden, E. Platen, P. E. Kloeden, and E. Platen, *Stochastic Differential Equations*. Berlin, Germany: Springer, 1992.
- [63] A. Budhiraja, L. Chen, and C. Lee, “A survey of numerical methods for nonlinear filtering problems,” *Physica D: Nonlinear Phenomena*, vol. 230, no. 1/2, pp. 27–36, 2007.



JUN LIU (Senior Member, IEEE) received the B.S. degree in applied mathematics from Shanghai Jiao Tong University, Shanghai, China, in 2002, the M.S. degree in mathematics from Peking University, Beijing, China, in 2005, and the Ph.D. degree in applied mathematics from the University of Waterloo, Waterloo, ON, Canada, in 2011. Between 2011 and 2012, he held an NSERC Postdoctoral Fellowship with Control and Dynamical Systems, Caltech. Between 2012 and 2015, he was a Lecturer of control and systems engineering with The

University of Sheffield, Sheffield, U.K. In 2015, he joined the Faculty of Mathematics, University of Waterloo, where he is currently an Associate Professor of applied mathematics and directs the Hybrid Systems Laboratory. His main research focuses on the theory and applications of hybrid systems and control, including rigorous computational methods for control design with applications in cyber-physical systems and robotics. He was the recipient of a Marie-Curie Career Integration Grant from the European Commission in 2013, a Canada Research Chair from the Government of Canada in 2017, an Early Researcher Award from the Ontario Ministry of Research, Innovation and Science in 2018, and an Early Career Award from the Canadian Applied and Industrial Mathematics Society and Pacific Institute for the Mathematical Sciences (CAIMS/PIMS) in 2020. His best paper awards include the Zhang Si-Ying Outstanding Youth Paper Award (2010, 2015) and the Nonlinear Analysis: Hybrid Systems Paper Prize (2017). He is a Member of SIAM and a Lifetime Member of CAIMS.



YIMING MENG received the bachelor’s degree in chemical engineering from Tianjin University, Tianjin, China, in 2013, the M.Sc. degree in process system engineering from Chemical Engineering Department, Imperial College London, London, U.K., in 2017, and the Ph.D. degree in applied mathematics from the University of Waterloo, Waterloo, ON, Canada, in 2022. He is currently a Postdoctoral Fellow with the Department of Applied Mathematics, University of Waterloo. His research focuses on stochastic dy-

namical systems and control, including formal methods and data-based decision-making problems with applications in finance, robotics, and other physical sciences.