

Deductive Proofs of Almost Sure Persistence and Recurrence Properties

Aleksandar Chakarov^(✉), Yuen-Lam Voronin^(✉),
and Sriram Sankaranarayanan^(✉)

University of Colorado Boulder, Boulder, CO 80303, USA
{aleksandar.chakarov,yuen-lam.voronin,
sriram.sankaranarayanan}@colorado.edu

Abstract. Martingale theory yields a powerful set of tools that have recently been used to prove quantitative properties of stochastic systems such as stochastic safety and qualitative properties such as almost sure termination. In this paper, we examine proof techniques for establishing almost sure persistence and recurrence properties of *infinite-state* discrete time stochastic systems. A persistence property $\diamond\Box(P)$ specifies that almost all executions of the stochastic system eventually reach P and stay there forever. Likewise, a recurrence property $\Box\diamond(Q)$ specifies that a target set Q is visited infinitely often by almost all executions of the stochastic system. Our approach extends classic ideas on the use of Lyapunov-like functions to establish qualitative persistence and recurrence properties. Next, we extend known constraint-based invariant synthesis techniques to deduce the necessary supermartingale expressions to partly mechanize such proofs. We illustrate our techniques on a set of interesting examples.

Keywords: Temporal logic · Stochastic systems · Markov processes · Stochastic control · Sum-of-squares programming

1 Introduction

In this paper, we study persistence ($\diamond\Box(\cdot)$) and recurrence ($\Box\diamond(\cdot)$) properties for stochastic systems. Stochastic systems are commonly used to model the effect of noise or random uncertainties on systems. Examples include probabilistic programs [17,20] with random number generating constructs for modeling uncertainties, cyber-physical systems under the influence of external stochastic disturbances, financial process models, and biological models. Given a stochastic system, we attempt to find proofs that for a subset of states T , the behaviors of the system satisfy a persistence property $\diamond\Box(T)$ with probability 1 (almost surely), i.e., almost every behavior of the system eventually enters T , and stays in T , forever. Similarly, we present an approach to prove $\Box\diamond(T)$, i.e., almost every behavior of the system hits T infinitely often. Such persistence properties effectively prove facts about the asymptotic behavior of these processes

which may take many forms, including convergence towards an “equilibrium” region, or asymptotic divergence away from a particular set. Recurrence properties can be used, for instance, to show that the system keeps returning to a set of desirable configurations even if forced to leave under the influence of stochastic disturbances.

Persistence and recurrence properties are also of independent interest to the verification and control theory community [5,21]. In standard model checking approaches, we rely on showing that the system is forced to almost surely reach a strongly connected component [11,13] which is a subset of T , and additionally for proving $\diamond\Box(T)$, that it has no outgoing transition from it. However, Baier et al. have demonstrated that this technique is restricted to finite state stochastic transition systems [5].

In comparison, the technique we propose here can handle infinite state, discrete time, polynomial stochastic systems by automatically deriving supermartingale expressions over the system variables and leveraging properties of these supermartingale expressions. Specifically, our work extends the infinite state probabilistic transition systems used in our earlier work [8], or the probabilistic guarded command language proposed by McIver and Morgan [23] with polynomial guards and updates. However, the ideas we present here can extend to a larger class of Markov models.

In this paper, we introduce two types of proof rule arguments for proving persistence and recurrence that derive directly from classic rules for Markov chains such as the Foster-Lyapunov condition [16,24]: (a) “Geometric” rules involve finding a nonnegative function $V(\mathbf{x})$ over the state variables \mathbf{x} , whose expectation in the next time step is some multiplicative factor $\alpha \in (0, 1)$ of its current value $V(\mathbf{x})$. These are inspired by Lyapunov functions used in control theory to prove exponential stability. (b) “Additive” rules are analogous to ranking function arguments for (nondeterministic) program termination. These conditions were studied for program termination as “supermartingale ranking functions” (SMRFs) in our previous work [8] and proven complete for certain classes of probabilistic systems by Fioriti et al. [15] and more recently by Chatterjee et al. [10]. However, SMRFs are designed to prove almost-sure termination properties of the form $\diamond(T)$. In the current work, we show—rather counterintuitively—that SMRFs cannot in general prove $\diamond\Box(T)$ properties. We provide a suitable technical condition (of bounded increase) under which SMRFs can prove $\diamond\Box(T)$ properties. Next, we also show that both types of proofs are equivalent under some technical conditions: a proof using a geometric proof rule can be transformed into an equivalent proof using an additive rule, and vice-versa. Nevertheless, both forms are useful when searching for certificates of a given form such as a polynomial function over the state-variables.

Finally, we examine the problem of automatically synthesizing the functions $V(\mathbf{x})$ for polynomial, probabilistic transition systems to prove persistence and recurrence properties. Assuming a parameterized “template” form of this function, we derive conditions that must be satisfied over the parameters for proving a target property. We conclude by illustrating our approach on a variety of small, but interesting examples.

1.1 Motivating Example 1: Room Temperature Control

In [3] Abate et al. present a room temperature control problem subject to stochastic temperature perturbations. Suppose that there are two adjacent rooms, whose temperatures change according to the following stochastic difference equation:

$$x'_i := x_i + b_i(x_0 - x_i) + a \cdot \sum_{i \neq j} (x_j - x_i) + c_i(1 - \sigma(x_i)) + \nu_i, \text{ for } i \in \{1, 2\},$$

where $a = 0.0625$, $b_1 = 0.0375$, $b_2 = 0.025$ are respectively the inter-room and external heat convection constants, $x_0 = 6^\circ C$ is the outdoor temperature, $c_1 = 0.65$, $c_2 = 0.6$ are the heat units supplied to the two rooms by the heater, and ν_1, ν_2 are i.i.d. stochastic noise. The behavior of the heater is governed by the controller unit term σ . We focus on the evolution of the room temperatures within the range $[6, 33]^2$.

Abate et al. construct a (nonlinear) sigmoidal controller $\sigma(t)$ that keeps the temperatures within a comfortable range $S : [17, 22] \times [16, 23]$ and focus on bounding the probability that the system leaves S (i.e., *stochastic safety*) within finitely many steps under the influence of Gaussian noise. Figure 1 shows 100 sample executions of the system when the controller is approximated with a degree-7 polynomial:

$$\sigma(t) : 29.2 - 13.42t + 2.55t^2 - 0.26t^3 + 0.015t^4 - 5.13 \times 10^{-4}t^5 + 9.23 \times 10^{-6}t^6 - 6.87 \times 10^{-8}t^7$$

under two different types of random noise: uniform \mathcal{U} on a given range and normal \mathcal{N} .

Controller σ was originally designed to keep the system in S with finite-time (100 min) guarantees in mind. We prove that under mild stochastic disturbances (LEFT) the system satisfies the *almost sure persistence* property $\diamond \square S$, i.e., with probability 1 the system eventually enters S and stays there forever. This is demonstrated by the proof rule PERSIST-GEOM of Sect. 3.1 and the certificate $V(x_1, x_2) : (x_1 - 18.3)^2 + (x_2 - 18.8)^2$. When the level of stochastic disturbance is increased (RIGHT) the almost sure persistence property no longer holds. This is consistent with the results in [3]; however, a weaker, almost sure *recurrence*




Fig. 1. 100 simulations of the two-room controller system, with initial temperatures x_1, x_2 uniformly drawn from $[15, 22]^2$, under two different types of stochastic noise: (LEFT) $\nu_i \sim \mathcal{U}(-0.01, 0.01)$, with the red horizontal lines indicating the intervals $[17.8, 18.7]$ (for room 1) and $[18.4, 19.3]$ (for room 2); (RIGHT) $\nu_i \sim \mathcal{N}(0, 0.25)$, with the red horizontal lines indicating the intervals $[16.9, 19.6]$ (for room 1) and $[17.3, 20.2]$ (for room 2).

property: $\square\Diamond(16.9 \leq x_1 \leq 19.6 \wedge 17.3 \leq x_2 \leq 20.2)$ holds, i.e., with probability 1 the system visits the region infinitely often. This is demonstrated by proof rule REC of Sect. 3.2 using the same certificate function V .

1.2 Motivating Example 2: Nonlinear Markov Jump System

Figure 2 shows a nonlinear Markov jump system with two modes q_1, q_2 and two state variables $\mathbf{x} : (x, y)$ that evolve according to the mode-dependent difference equations. The system jumps between modes with equal probability.

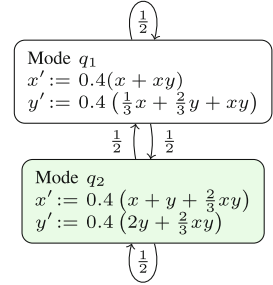


Fig. 2. A nonlinear Markov jump system with 2 modes.

Observe that 0 is an equilibrium and $X : [-0.5, 0.5]^2$ is an invariant of the system, i.e., all sample paths starting in X , stay in the set forever. Figure 3(a) shows the sample paths that start inside X converge towards 0. We establish that the persistence property $\Diamond\square(|x| \leq 0.1 \wedge |y| \leq 0.1)$ holds almost surely over all executions of the system, by synthesizing the nonnegative certificate function $V(\mathbf{x}) : 2.3x^2 + 4.15xy + 3.7y^2$. After one time step, the expected value of V is at most $\frac{1}{2}$ -fraction of its original value, i.e., $(\forall \mathbf{x} \in X) \mathbb{E}(V(\mathbf{x}')) \leq \frac{1}{2}V(\mathbf{x})$. Figure 3(b) plots the function V over the sample paths, showing its convergence. We use the certificate $V(\mathbf{x})$ in PERSIST-GEOM (Sect. 3.1) to establish the required property.

Outside X , the system appears unstable as shown in Fig. 3(c), yet the behaviors approach $x = y$ asymptotically. Using the certificate $\hat{V}(\mathbf{x}) : (x - y)^2$ in SPERSIST-GEOM (Strong Persistence of Sect. 3.1) we can prove that $(\forall \varepsilon > 0) \Diamond\square(|x - y| \leq \varepsilon)$.

Organization. Section 2 presents an infinite state discrete time stochastic system model and formally states the problem of proving persistence and recurrence (Sect. 2.2). Section 3 presents our main contribution in the form of proof rules for persistence and recurrence properties with the soundness results of our analysis

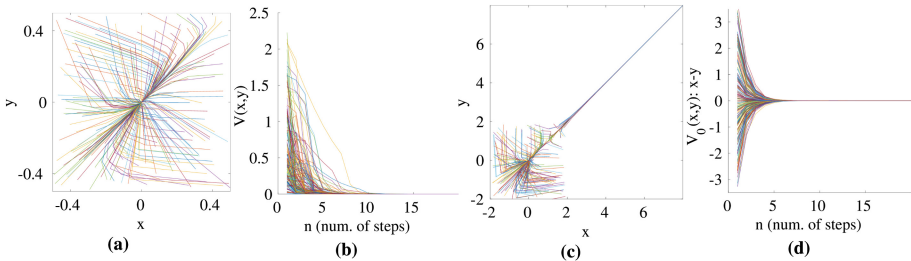


Fig. 3. Sample paths of the Markov jump system described in Fig. 2.

presented in Sect. 4. Section 5 presents the results of our prototype implementation on a set of benchmarks, followed by a summary of relevant related work (Sect. 6) and conclusion (Sect. 7).

2 Preliminaries

In this section, we present the basic computational model of discrete-time stochastic dynamical systems that we study here and introduce the notion of tail invariant properties in the form of persistence and recurrence. We formulate a problem statement for the rest of the paper and finally introduce supermartingales. The proofs of all statements can be found in the extended version of the paper.

2.1 Discrete Time Stochastic Systems

We present a simple yet general model of an infinite-state discrete time stochastic system, and then examine a model of polynomial stochastic transition systems. A purely deterministic system with state-space X is described by a rule of the form $\mathbf{x}' := F(\mathbf{x})$ and $\mathbf{x}_0 \in X_0$, where $\mathbf{x} \in X$ and \mathbf{x}_0 is the initial state belonging to initial set X_0 . Stochastic systems studied here are described by $\mathbf{x}' := F(\mathbf{x}, \mathbf{r})$ and $\mathbf{x}_0 \sim \mathcal{D}_0$, where \mathbf{r} is a vector of random variables and initial state \mathbf{x}_0 is drawn from initial distribution \mathcal{D}_0 .

Definition 1 (DTSS). A discrete-time stochastic system (DTSS) Π is defined as the tuple $\langle \Sigma, \mathcal{R}, \mathcal{F}, \mathcal{D}_0 \rangle$ with the following components:

1. a state space Σ and an associated Borel σ -algebra on it,
2. a probability space $\mathcal{R} : \langle R, \mathcal{F}_R, P \rangle$ (with individual samples denoted by \mathbf{r}),
3. a transition function $\mathcal{F} : \Sigma \times R \rightarrow \Sigma$, wherein $\mathcal{F}(\mathbf{x}, \mathbf{r})$ denotes the next state obtained from a state $\mathbf{x} \in \Sigma$ and random sample $\mathbf{r} \in R$,
4. an initial probability distribution \mathcal{D}_0 over Σ .

Let Ω denote the sample set $\Sigma \times R^\omega$, which consists of tuples $\langle \mathbf{x}_0, \mathbf{r}_0, \mathbf{r}_1, \dots, \mathbf{r}_n, \dots \rangle$. Here $\mathbf{x}_0 \in \Sigma$ denotes the starting state sample, and $\mathbf{r}_0, \mathbf{r}_1, \dots, \mathbf{r}_n, \dots$ denote successive draws of random variables from the sample set R of the probability space \mathcal{R} . Given a discrete-time Markov process, the model maps each $\omega : \langle \mathbf{x}_0, \mathbf{r}_0, \dots \rangle \in \Omega$ to a sample path (or trace) $\pi(\omega)$ as follows $\pi(\omega) : \mathbf{x}_0 \xrightarrow{\mathbf{r}_0} \mathbf{x}_1 \xrightarrow{\mathbf{r}_1} \mathbf{x}_2 \xrightarrow{\mathbf{r}_2} \dots$.

The stochastic process is defined by applying the Kolmogorov extension theorem starting from finite-dimensional distributions $P_{t_1, \dots, t_k}(B)$ where $B \subseteq \Sigma^k$ is a measurable set. The semantics of the stochastic system is, therefore, equivalent to an infinite-state discrete-time Markov chain (see the extended version of the paper for details).

Independence of Samples. The formulation above naturally assumes that the samples of the random variable \mathbf{r}_i are independent of the current state \mathbf{x}_i and from previous samples $\mathbf{r}_0, \dots, \mathbf{r}_{i-1}$. While this is somewhat restrictive, in practice it encompasses nearly all example instances that we are aware of.

No Demonic Nondeterminism. The formulation also precludes any demonic nondeterminism since \mathcal{F} is a function.

In this work, we focus on *polynomial* (stochastic) systems, which are instances of stochastic systems with piecewise polynomial update maps.

Definition 2 (Polynomial Stochastic System). A polynomial stochastic system Π is a tuple $\langle X, \mathcal{R}, \mathcal{T}, \mathcal{D}_0 \rangle$, where (a) the state-space $X \subseteq \mathbb{R}^n$ is a semi-algebraic set (i.e., X is the solution set of finitely many polynomial inequalities), (b) \mathcal{R} is a probability space for the stochastic inputs written collectively as $\mathbf{r} : (\mathbf{r}_c, \mathbf{r}_b)$, wherein \mathbf{r}_c denotes the (possibly multivariate) continuous random variable and \mathbf{r}_b denote the (discrete) random variables that take on finitely many values, (c) $\mathcal{T} : \{\tau_1, \dots, \tau_m\}$ is a finite set of transitions, and (d) \mathcal{D}_0 is an initial state probability distribution over X .

Each transition $\tau \in \mathcal{T}$ has two parts: a guard predicate φ_τ and an update function $f_\tau : X \times R \rightarrow X$:

1. The guard $\varphi_\tau(\mathbf{x})$ is a conjunction of polynomial inequalities over \mathbf{x} ;
2. The update function $f_\tau(\mathbf{x}, \mathbf{r}) : X \times R \rightarrow X$ is a piecewise polynomial function of the form:

$$f_\tau(\mathbf{x}, \mathbf{r}) : \begin{cases} g_{\tau,1}(\mathbf{x}, \mathbf{r}_c), & \text{if } \psi_{\tau,1}(\mathbf{r}_b) \\ \vdots \\ g_{\tau,j}(\mathbf{x}, \mathbf{r}_c), & \text{if } \psi_{\tau,j}(\mathbf{r}_b), \end{cases}$$

where $g_{\tau,1}, \dots, g_{\tau,j}$ are multivariate polynomials over \mathbf{x}, \mathbf{r}_c and $\psi_{\tau,1}(\mathbf{r}_b), \dots, \psi_{\tau,j}(\mathbf{r}_b)$ represent mutually exclusive and exhaustive predicates over the random variables with probability $p_{\tau,i} : \text{Prob}(\psi_{\tau,i}(\mathbf{r}_b))$.

We refer to each function $g_{\tau,i}$ as a *fork* of f_τ guarded by $\psi_{\tau,i}$ with corresponding *fork probability* $p_{\tau,i}$, for all $i, 1 \leq i \leq j$.

For a polynomial system to represent a stochastic system over X according to Definition 1, we require that transitions together form a function over the state-space X :

1. The guards are pairwise *mutually exclusive*: $\varphi_{\tau_i} \wedge \varphi_{\tau_j}$ is unsatisfiable for all $i \neq j$.
2. The guards are *mutually exhaustive*: $\bigvee_{j=1}^k \varphi_{\tau_j} \equiv \text{true}$.

With these conditions, it is easy to define an overall piecewise polynomial transition function over \mathcal{F} that casts any polynomial transition system as a stochastic system.

Example 1 (Strange Random Walk). Let $\{Y_i\}$ be a sequence of random variables over \mathbb{R} with Y_0 distributed uniformly over $[0, 1]$. For all $n \geq 0$, define:

$$Y_{n+1} = \begin{cases} Y_n^2, & \text{with probability } \frac{1}{2}, \\ 2Y_n - Y_n^2, & \text{with probability } \frac{1}{2}. \end{cases}$$

The corresponding polynomial stochastic system is $\Pi : \langle \mathbb{R}, \mathcal{R}, \{\tau\}, \mathcal{D}_0 \rangle$, where \mathcal{R} is the probability space for the uniform distribution $\mathcal{U}(0, 1)$, the initial probability distribution \mathcal{D}_0 is $\mathcal{U}(0, 1)$, and transition $\tau : \langle true, f_\tau \rangle$, has update mapping

$$f_\tau(\mathbf{x}, \mathbf{r}) : \begin{cases} g_{\tau,1}(\mathbf{x}) : \mathbf{x}^2, & \text{if } \psi_{\tau,1}(\mathbf{r}_b) : \mathbf{r}_b \leq 1/2, \\ g_{\tau,2}(\mathbf{x}) : 2\mathbf{x} - \mathbf{x}^2, & \text{if } \psi_{\tau,2}(\mathbf{r}_b) : \mathbf{r}_b > 1/2, \end{cases}$$

and defines corresponding fork probabilities $p_1 = p_2 = 1/2$.

Pre-Expectations. Key to the analysis is the notion of *pre-expectation*. The definitions below are inspired by [8, 19, 23] and are related to drift operators of Markov processes [24]. We first formalize the notion of pre-expectations over general stochastic systems and then provide a specialized definition for polynomial stochastic systems.

Consider a stochastic system $\langle \Sigma, \mathcal{R}, \mathcal{F}, \mathcal{D}_0 \rangle$, and a function $h : \Sigma \rightarrow \mathbb{R}$ over the state-space. The *pre-expectation* of h w.r.t to \mathcal{F} yields another function $\hat{h} : \Sigma \rightarrow \mathbb{R}$ such that for any state $\mathbf{x} \in \Sigma$, $\hat{h}(\mathbf{x})$ yields the expected value of $h(\mathbf{x}')$, where the expectation is taken over all states \mathbf{x}' reached in one step from \mathbf{x} . Formally, $\hat{h}(\mathbf{x}) : \mathbb{E}_R(h(\mathcal{F}(\mathbf{x}, \mathbf{r})))$. The pre-expectation can be difficult to compute for a stochastic system, even if $h(\mathbf{x})$ is of a simple form, for example, polynomial.

Now, we translate this definition to polynomial stochastic transition systems. We first define pre-expectations across transitions.

Definition 3 (Pre-Expectation across a Transition). *Given a polynomial stochastic transition system $\langle X, \mathcal{R}, \mathcal{T}, \mathcal{D}_0 \rangle$, a function $h : X \rightarrow \mathbb{R}$ and a transition $\tau \in \mathcal{T}$ with forks $g_{\tau,1}, \dots, g_{\tau,j}$ and corresponding fork probabilities $p_{\tau,1}, \dots, p_{\tau,j}$, the pre-expectation of h across transition τ is a function $\text{pre}\mathbb{E}(h, \tau) : X \rightarrow \mathbb{R}$ defined as follows: $\forall \mathbf{x} \in X$, $\text{pre}\mathbb{E}(h, \tau)(\mathbf{x}) : \mathbb{E}_R[h(f_\tau(\mathbf{x}, \mathbf{r}_c)) | \mathbf{x}] = \sum_{i=1}^j p_{\tau,i} \mathbb{E}_{R_c}[h(g_{\tau,i}(\mathbf{x}, \mathbf{r}_c))]$, with expectation taken over R from which the random choices $\mathbf{r} = (\mathbf{r}_b, \mathbf{r}_c)$ are drawn.*

We now define the pre-expectation transformation over an entire stochastic system.

Definition 4. *The pre-expectation of a function $h : X \rightarrow \mathbb{R}$ w.r.t. a polynomial stochastic system $\Pi : \langle X, \mathcal{R}, \{\tau_1, \dots, \tau_n\}, \mathcal{D}_0 \rangle$ is a function $\text{pre}\mathbb{E}(h, \Pi) : X \rightarrow \mathbb{R}$ defined by*

$$\text{pre}\mathbb{E}(h, \Pi)(\mathbf{x}) : \sum_{i=1}^n \mathbb{1}(\varphi_i(\mathbf{x})) \cdot \text{pre}\mathbb{E}(h, \tau_i)(\mathbf{x}), \text{ for all } \mathbf{x} \in X,$$

where φ_i is the guard of transition τ_i and $\mathbb{1}(\varphi)$ is the indicator function of predicate φ .

Related to the pre-expectation is the notion of a drift operator.

Definition 5 (Drift Operator). *Let Π be a stochastic transition system and h be a function over the state-space. The drift of h w.r.t. Π is the function $\mathcal{D}_\Pi h : \text{pre}\mathbb{E}(h, \Pi) - h$.*

Wherever the system Π is clear from the context, we use $\mathcal{D}h$ to denote the drift operator \mathcal{D}_Π applied to the function h .

Example 2. We return to Example 1 and compute the pre-expectation of $h(x) = x$:

$$\begin{aligned} \text{pre}\mathbb{E}(h, \Pi) &= \mathbb{1}(\text{true}) \cdot \text{pre}\mathbb{E}(h, \tau) = \mathbb{E}_R[p_1 h(g_1(x)) + p_2 h(g_2(x))] \\ &= \mathbb{E}_R[\tfrac{1}{2}(x^2) + \tfrac{1}{2}(2x - x^2)] = x. \end{aligned}$$

It is clear that for any state $x \in X$, the value of h equals the value of $\text{pre}\mathbb{E}(h, \Pi)$, or equivalently, the drift $\mathcal{D}h = 0$. This function is an example of a *martingale expression* and such functions are central to our analysis. We give their definition in Sect. 2.3 and present properties of martingale expressions relevant to our analysis.

Assume that for any polynomial $p(\mathbf{r}_c)$ involving the continuous random variables \mathbf{r}_c the expectation $\mathbb{E}(p(\mathbf{r}_c))$ is finite. Then the drift of a polynomial is a polynomial.

Lemma 1. *Assume that all cross moments exist for the random variable \mathbf{r}_c . For a polynomial $h(\mathbf{x})$, the pre-expectation $\text{pre}\mathbb{E}(h, \tau)$ across a transition τ is also a polynomial. Moreover, the pre-expectation $\text{pre}\mathbb{E}(h, \Pi)$ is a piecewise polynomial function of the form: $\sum_{j=1}^m \mathbb{1}(\varphi_j) q_j(\mathbf{x})$, where φ_j is a transition guard and $q_j(\mathbf{x})$ is a polynomial.*

2.2 Persistence and Recurrence

Let Π be a polynomial stochastic system with a sample set $\Omega : X \times R^\omega$ and an associated σ -algebra generated by the Borel sets over X and R . Let Pr be the associated measure that maps a measurable subset $S \subseteq \Omega$ to its probability $\text{Pr}(S)$. Let π be a function that maps each sample $\omega \in \Omega$ to the corresponding sample path of the system $\pi(\omega) : \langle \mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_m, \dots \rangle$. Likewise, let π_m map each sample $\omega \in \Omega$ to the state encountered at time m , i.e., $\pi_m(\omega) : \mathbf{x}_m$.

For a predicate φ over the system states, the persistence property $\diamond\Box\varphi$ is a collection of sample paths: $\llbracket \diamond\Box\varphi \rrbracket : \{\omega \in \Omega \mid \exists n \geq 0, \forall m \geq n, \pi_m(\omega) \models \varphi\}$. It is easy to show that this is a measurable set. The probability of the persistence property $\diamond\Box\varphi$ is denoted $\text{Pr}(\diamond\Box\varphi)$. We say the persistence property $\diamond\Box\varphi$ holds almost surely (a.s.) iff $\text{Pr}(\diamond\Box\varphi) = 1$. Such a property is also known as qualitative (probability 1) persistence property and can be stated in PCTL as $\mathbb{P}_{=1}(\diamond\Box\varphi)$.

Similarly, a recurrence property $\Box\diamond\varphi$ is a collection of sample paths: $\llbracket \Box\diamond\varphi \rrbracket : \{\omega \in \Omega \mid \forall n \geq 0, \exists m \geq n, \pi_m(\omega) \models \varphi\}$. We say that the recurrence property $\Box\diamond\varphi$ holds almost surely iff $\text{Pr}(\Box\diamond\varphi) = 1$.

Problem Statement. Let Π be a polynomial stochastic system with state-space X and let $T \subseteq X$ be a measurable set of states. In this paper we are interested in two related problems: (i) Establish that the persistence property $\diamond\Box(T)$ holds a.s.; and/or, (ii) Establish that the recurrence property $\Box\diamond(T)$ holds a.s.

2.3 Supermartingales and Their Properties

Following [32] we recall the notion of supermartingales and some key properties.

Definition 6. A discrete time real-valued stochastic process $\mathcal{M} = \{M_i\}_{i=0}^\infty$ is a supermartingale if $\mathbb{E}(M_{n+1} | M_n = m_n, \dots, M_0 = m_0) \leq m_n$, for all $n \geq 0$ and m_n .

Since our work mostly concerns Markov processes, we will write $\mathbb{E}(M_{n+1} | M_n = m_n)$ to mean $\mathbb{E}(M_{n+1} | M_n = m_n, \dots, M_0 = m_0)$. We now propose the key definitions of additive and multiplicative supermartingales that will be used in our work.

Definition 7 (Additive and Multiplicative Supermartingales). A supermartingale $\mathcal{M} = \{M_i\}_{i=0}^\infty$ is called (ε -)additive iff $\mathbb{E}(M_{n+1} | M_n = m_n) \leq m_n - \varepsilon$, for all $n \geq 0$ and m_n , for some $\varepsilon > 0$. The supermartingale \mathcal{M} is called (α -)multiplicative iff $\mathbb{E}(M_{n+1} | M_n = m_n) \leq \alpha m_n$, for all $n \geq 0$, $m_n \geq 0$, and for some $0 < \alpha < 1$.

First we state the following simple result about α -multiplicative supermartingales.

Lemma 2. Let $\mathcal{M} = \{M_i\}_{i=0}^\infty$ be a nonnegative α -multiplicative supermartingale for some $\alpha \in (0, 1)$. Then $\widehat{\mathcal{M}} = \{\widehat{M}_i : \frac{M_i}{\alpha^i}\}_{i=0}^\infty$ is a nonnegative supermartingale.

Following [8], we relate supermartingales to polynomial stochastic systems.

Definition 8 (Supermartingale Expressions). Let e be an expression over the state variables of a polynomial transition system Π (i.e., e is a real-valued function over the state-space X of Π). The expression e is a supermartingale expression for Π iff

$$(\forall \mathbf{x} \in X) \text{pre}\mathbb{E}(e, \Pi) \leq e(\mathbf{x}), \text{ or equivalently, } (\forall \mathbf{x} \in X) \mathcal{D}e(\mathbf{x}) \leq 0.$$

The expression e is called (c -)additive or alternatively, a supermartingale ranking function (SMRF) iff there exists $c > 0$ such that

$$(\forall \mathbf{x} \in X) \text{pre}\mathbb{E}(e, \Pi) \leq e(\mathbf{x}) - c, \text{ or equivalently, } (\forall \mathbf{x} \in X) \mathcal{D}e(\mathbf{x}) \leq -c.$$

Similarly, e is an (α -)multiplicative supermartingale expression iff

$$(\exists \alpha \in (0, 1)) (\forall \mathbf{x} \in X) \text{pre}\mathbb{E}(e, \Pi) \leq \alpha e(\mathbf{x}).$$

By definition any (α -)multiplicative supermartingale expression of Π induces a(n) (α -)multiplicative supermartingale when evaluated along the sample paths of Π .

3 Proof Rules for Persistence and Recurrence

In this section, we describe the main proof rules for persistence and recurrence properties. All proof rules involve finding a suitable “certificate” in the form of a stochastic analogue of a Lyapunov-like function over the state-space X . The soundness of our approach (presented in Sect. 4) relies on certificate functions behaving as supermartingale expressions over the state variables of the stochastic system.

Let X be the state-space of interest and $T \subseteq X$ be a target set.

3.1 Proof Rules for Persistence

We provide a series of proof rules for proving persistence properties. The relation between these rules is examined in the extended version of the paper.

PERSIST-GEOM: Geometric rule for persistence	
$(p1) (\forall \mathbf{x} \in X) V(\mathbf{x}) \geq 0,$	Positive semidef. of V .
$(p2) (\exists \varepsilon > 0) (\forall \mathbf{x} \in X \setminus T) V(\mathbf{x}) \geq \varepsilon,$	Lower bnd. outside T .
$(p3) (\forall \mathbf{x} \in T) \mathcal{D}V(\mathbf{x}) \leq 0,$	Drift cond. inside T .
$(p4) (\exists \alpha \in (0, 1)) (\forall \mathbf{x} \in X \setminus T) \mathcal{D}V(\mathbf{x}) \leq (\alpha - 1)V(\mathbf{x}),$	Drift cond. outside T .
$\diamond \square(T)$ almost surely.	
PERSIST-ADD: Additive rule for persistence	
$(p1) (\forall \mathbf{x} \in X) V(\mathbf{x}) \geq 0,$	Positive semidef. of V .
$(p2) (\exists \varepsilon > 0) (\forall \mathbf{x} \in X \setminus T) V(\mathbf{x}) \geq \varepsilon,$	Lower bnd. outside T .
$(p3) (\forall \mathbf{x} \in T) \mathcal{D}V(\mathbf{x}) \leq 0,$	Drift condition inside T .
$(p5) (\exists c < 0) (\forall \mathbf{x} \in X \setminus T) \mathcal{D}V(\mathbf{x}) \leq -c,$	Drift condition outside T .
$\diamond \square(T)$ almost surely.	

Both PERSIST-GEOM and PERSIST-ADD state that a polynomial stochastic system Π satisfies $\diamond \square(T)$ almost surely if there exists a nonnegative certificate function V (condition (p1)) whose value outside T is lower bounded by some $\varepsilon > 0$ (condition (p2)). Moreover, the drift conditions ensure that in expected value V in the next step does not increase inside T (condition (p3)), and decreases by some fixed non-zero quantity outside T (an additive constant in (p5), or, a multiplicative factor in (p4)). Intuitively, these conditions together guarantee that V is a supermartingale whose drift condition outside T forces its value to decrease along almost all sample paths and eventually reach a value ε at which point the sample path is “forced” to enter T and persists forever.

Applications. We present an application of each rule and defer soundness to Sect. 4.

Example 3. Consider a stochastic system Π with a single variable x over \mathbb{R} , and a single transition: $x' := 0.1(1 + w)x$, where w is a standard Gaussian random

variable. We show that the almost sure persistence property $\diamond\Box(T : |x| \leq 0.1)$ holds.

Consider the function $m(x) : x^2$, which is nonnegative on X and $m(x) \geq 0.01$ for all $x \in X \setminus T$ (conditions (p1), (p2)). Moreover, for all $x \in X$, $\text{preE}(m, \Pi) = \mathbb{E}_w(m(x_{n+1})|x_n) = 0.02x_n^2$, so $\mathcal{D}m(x) \leq -0.98m(x)$. Hence, $m(x)$ defines a 0.02-multiplicative supermartingale expression (conditions (p3), (p4)).

Applying PERSIST-GEOM, we conclude that $\diamond\Box(-0.1 \leq x \leq 0.1)$ holds a.s.

Example 4. For the polynomial stochastic system of Example 1 over the state-space $X = [0, 1]$, we establish the almost sure persistence property $\diamond\Box(x \leq 0.05 \vee x \geq 0.95)$.

Consider the certificate function $V(x) = x(1 - x)$. For all $x \in X$, $V(x) \geq 0$, and for all $x \in X \setminus T = (0.05, 0.95)$, $V(x) \geq 0.0475$ (conditions (p1), (p2)). Next, note that $\text{preE}(V, \Pi) = x(1 - x)(1 - x + x^2)$, and $\mathcal{D}V(x) = x(1 - x)(x^2 - x)$. It is easy to check that for all $x \in (0.05, 0.95)$, $\mathcal{D}V(x) \leq -0.00225625$, and for all $x \in [0, 0.05] \cup [0.95, 1]$, $\mathcal{D}V(x) \leq 0$ (conditions (p3), (p5)).

Applying PERSIST-ADD, we conclude that $\diamond\Box(x \leq 0.05 \vee x \geq 0.95)$ holds a.s.

Note 1. In both examples, the certificates $m(x)$ and $V(x)$ can be used in both PERSIST-GEOM and PERSIST-ADD: $\mathcal{D}m(x) \leq -0.98x^2 \leq -0.0098$, for all $x \in X \setminus T$, and $\mathcal{D}V(x) \leq x(1 - x)(x^2 - x) \leq -0.0475V(x)$. We expand on this point next.

Strong Persistence. Rules PERSIST-GEOM and PERSIST-ADD present sufficient conditions under which certificates V prove that $\diamond\Box(T)$ holds almost surely. Unfortunately, the difference in constraints inside and outside T may force V to be a high degree polynomial (or a piecewise polynomial function). To simplify constraints and make the search for certificates for almost sure persistence properties tractable we propose a stronger version of proof rules for persistence of the form: $(\forall \varepsilon > 0) \diamond\Box(V(\mathbf{x}) \leq \varepsilon)$.

SPERSIST-GEOM: Geometric rule for strong persistence

$$\frac{\begin{array}{l} (p1) (\forall \mathbf{x} \in X) V(\mathbf{x}) \geq 0, \quad \text{Positive semidef. of } V. \\ (p6) (\exists \alpha \in (0, 1)) (\forall \mathbf{x} \in X) \mathcal{D}V(\mathbf{x}) \leq (\alpha - 1)V(\mathbf{x}), \text{ Drift condition.} \end{array}}{(\forall \varepsilon > 0) \diamond\Box(V(\mathbf{x}) \leq \varepsilon) \text{ almost surely.}}$$

Similarly, we provide an additive version of strong persistence rule. We say a function $V(\mathbf{x})$ has *bounded increase* over Π iff there is a constant $C > 0$ such that, for every possible next state \mathbf{x}' reached from \mathbf{x} (i.e., $\mathbf{x} \xrightarrow{\Pi} \mathbf{x}'$), $|V(\mathbf{x}') - V(\mathbf{x})| \leq C$.

SPERSIST-ADD: Additive rule for strong persistence

$$\frac{\begin{array}{l} (p7) (\exists c < 0) (\forall \mathbf{x} \in X) \mathcal{D}V(\mathbf{x}) \leq -c, \quad \text{Drift condition.} \\ (p8) (\forall \mathbf{x} \in X) V(\mathbf{x}) \text{ has bounded increase, (See above, and Def. 9 on page 273).} \end{array}}{(\forall K) \diamond\Box(V(\mathbf{x}) \leq K) \text{ almost surely.}}$$

SPERIST-GEOM presents a stronger, yet simpler to state and encode, version of the drift requirement dictated by PERSIST-GEOM (viz. inside the region T). Similarly, SPERSIST-ADD does not insist on V being positive definite but only V decreasing in expectation by $-c$ everywhere. The benefit of the stronger formulations of the persistence rule is that each level set of the Lyapunov-like certificate V acts as a *tail invariant*: a set S that almost all traces of the stochastic system reach and asymptotically confine to.

Relations Between Proof Rules. In Note 1 we allude to the fact that certificates for rule PERSIST-GEOM can equivalently be used for rule PERSIST-ADD to prove persistence properties of polynomial stochastic systems. We state the main result of importance here and defer all proofs and relationship between proof rules to the extended version.

Theorem 1. *Let $m(\mathbf{x})$ be an ε -additive supermartingale expression that has bounded increase in Π . Then there exist positive constants $\lambda > 1$ and $\alpha < 1$ such that $\lambda^{m(\mathbf{x})}$ is an α -multiplicative supermartingale expression. Moreover, let $\kappa \in \mathbb{R}$ be such that $\{\mathbf{x} \in X \mid m(\mathbf{x}) \leq \kappa\}$ is nonempty. Then the system Π satisfies the tail invariance property $m(\mathbf{x}) \leq \kappa$ almost surely.*

Under some technical conditions, it is possible to prove the converse of Theorem 1. This shows that any positive α -multiplicative supermartingale expression used to prove a tail invariant property has an equivalent additive supermartingale (more precisely, SMRF) formulation and vice versa.

Incompleteness. We demonstrate that although sound, our approach is incomplete. The existence of a nonnegative α -multiplicative supermartingale expression or a SMRF of bounded increase is sufficient but not a necessary condition for the system to almost surely satisfy a tail invariant property. Example 8 in Sect. 4.2 demonstrates this result.

3.2 Proof Rule for Recurrence

We now focus on proof rules for proving the almost sure recurrence property: $\square\lozenge(T)$, i.e., T is visited infinitely often by almost all sample paths. The proof rule is almost identical to a related rule that establishes “positive recurrence” in Markov chains [24].

REC: Rule for Recurrence

- | | |
|---------------------------------------------------------------------------------------------------------|-------------------------------|
| $(r1) (\forall \mathbf{x} \in X) V(\mathbf{x}) \geq 0,$ | Positive semidef. of V . |
| $(r2) (\exists \varepsilon > 0) (\forall \mathbf{x} \in X \setminus T) V(\mathbf{x}) \geq \varepsilon,$ | Lower bnd. outside T . |
| $(r3) (\exists H) (\forall \mathbf{x} \in T) \mathcal{D}V(\mathbf{x}) \leq H,$ | Drift condition inside T . |
| $(r4) (\exists c > 0) (\forall \mathbf{x} \in X \setminus T) \mathcal{D}V(\mathbf{x}) \leq -c,$ | Drift condition outside T . |

$\square\lozenge(T)$ almost surely

4 Soundness of Proof Rules

4.1 Supermartingales as Certificates of Geometric Persistence Rules

We state two theorems that formally relate supermartingales and persistence properties. The first establishes the convergence of nonnegative α -multiplicative supermartingales.

Theorem 2. *Let $\mathcal{M} = \{M_i\}_{i=0}^\infty$ be nonnegative α -multiplicative supermartingale for some $\alpha \in (0, 1)$. Then \mathcal{M} converges almost surely (samplewise) to 0.*

This result can be applied directly to stochastic transition systems. Let $m(\mathbf{x})$ be a nonnegative α -multiplicative supermartingale expression for a transition system Π for some $\alpha \in (0, 1)$. For a sample path $\{\mathbf{x}_i\}_{i=0}^\infty$ of Π , we say that $m(\mathbf{x}_i)$ *upcrosses* a level $\kappa > 0$ iff $m(\mathbf{x}_i) \leq \kappa$ and $m(\mathbf{x}_{i+1}) > \kappa$. If $m(\mathbf{x})$ converges almost surely to 0 on all sample paths, then the number of upcrossings of $m(\mathbf{x})$ on any sample path is a.s. finite.

Lemma 3. *Let $m(\mathbf{x})$ be a nonnegative α -multiplicative supermartingale expression for a polynomial stochastic system Π . Then for all $\kappa > 0$, the number of κ -upcrossings of $m(\mathbf{x})$ is almost surely finite, i.e.,*

$$\Pr(\{\omega \in \Omega \mid \{i \mid m(\pi_i(\omega)) \leq \kappa \wedge m(\pi_{i+1}(\omega)) > \kappa\} \text{ is finite}\}) = 1.$$

Proof. The result follows directly from the almost sure convergence of $m(\mathbf{x})$ to zero on sample paths of Π . □

This means that for any threshold $\kappa > 0$, $m(\mathbf{x}) \leq \kappa$ is a tail invariant property: i.e., $\diamond\Box(\varphi : m(\mathbf{x}) \leq \kappa)$ holds almost surely.

Theorem 3 (Soundness of PERSIST-GEOM). *A polynomial stochastic system Π satisfies the almost sure persistence property $\diamond\Box(T)$ if there exists a function V that satisfies conditions (p1)-(p4) of PERSIST-GEOM.*

Necessity. The two conditions on the multiplicative supermartingale $m(\mathbf{x})$ are $\alpha \in (0, 1)$ and $m(\mathbf{x})$ nonnegative. We show their necessity through the following example.

Example 5. Consider a stochastic transition system with a single variable x defined over the state space $[0, \infty)$ with two transitions τ_1 and τ_2 . Transition τ_1 has a guard $x \geq 1$ and does not alter the value of x . Transition τ_2 has a guard $x < 1$ and chooses between $x' := 2x$ or $x' := \frac{x}{2}$ with equal probabilities. That is:

$$x' := \begin{cases} x & \text{if } x \geq 1, \text{ and} \\ 2x & \text{if } x < 1, \text{ with prob. } \frac{1}{2} \\ \frac{x}{2} & \text{if } x < 1, \text{ with prob. } \frac{1}{2} \end{cases}$$

The initial value of x is exponentially distributed over $[0, \infty)$. Note that $m(x) : x$ is a nonnegative α -multiplicative supermartingale only when $\alpha = 1$. Clearly, m does not converge almost surely to zero.

Consider another transition system involving $x \in \mathbb{R}$ having two forks: $x' := x$ with probability $\frac{2}{3}$, and $x' := -x$ with probability $\frac{1}{3}$. Clearly, $m(x) : x$ is a $\frac{1}{3}$ -multiplicative supermartingale. However, m is not nonnegative over the state-space \mathbb{R} , so m does not prove any persistence property. Indeed, $\diamond\Box(x \leq 1)$ is not a tail invariant for the system.

4.2 Supermartingales as Certificates of Additive Persistence Rules

Recall that an additive supermartingale expressions $m(\mathbf{x})$ of Π satisfies the condition

$$(\forall \mathbf{x} \in X) \text{pre}\mathbb{E}(m(\mathbf{x}), \Pi) \leq m(\mathbf{x}) - \varepsilon,$$

for some constant $\varepsilon > 0$. (See Definition 7.) Given an additive supermartingale expression m , let $M_\kappa : \{\mathbf{x} \in X \mid m(\mathbf{x}) \leq \kappa\}$. For any κ where $M_\kappa \neq \emptyset$, we can prove $\diamond(M_\kappa)$ holds a.s. [8,10,15]. Yet in general, the property $\diamond\Box(M_\kappa)$ does not hold a.s.

Example 6. (MOONWALK) A MOONWALK system consists of a random walk over the state-space $X : \mathbb{Z}_{\leq 0}$ of the nonpositive integers:

$$x_{n+1} := \begin{cases} x_n - 1 & \text{with prob. } p(x_n), \\ 0 & \text{with prob. } 1 - p(x_n), \end{cases}$$

wherein $p(x) : \frac{x-0.5}{x-1} = 1 - \frac{0.5}{1-x}$, for $x < 0$, and $p(0) = 1$. In other words, the random walk either chooses to decrease x by 1 with probability $p(x)$ or jumps to 0 with probability $1 - p(x)$. The initial state follows a negative Poisson distribution.

The function $m(x) : x$ is an additive supermartingale expression for the MOONWALK system: for $x_n < 0$, $\mathbb{E}(m(x_{n+1})|x_n) = x_n - 0.5$, and $\mathbb{E}(m(x_{n+1})|0) = x_n - 1$. Yet the sublevel sets of the function m cannot be used for establishing persistence properties, because of the following result.

Lemma 4. *For any $\eta < 0$, the probability that a sample path of MOONWALK satisfies $\diamond\Box(x \leq \eta)$ is 0.*

Using an additive supermartingale expression m to prove tail invariance properties of the form $\diamond\Box(m(\mathbf{x}) \leq \kappa)$ requires additional assumptions on the expression m . One such condition is the bounded increase (which was assumed in Theorem 1, for establishing the soundness of proving persistence properties via additive supermartingales).

Definition 9 (Bounded Increase Expression). *An expression $m(\mathbf{x})$ has bounded increase for a stochastic transitions system Π iff there exists $M > 0$ so that for all possible states $\mathbf{x} \in X$ and all possible next states \mathbf{x}' reachable from \mathbf{x} , $|m(\mathbf{x}') - m(\mathbf{x})| \leq M$.*

We give an example of bounded increase expressions, which do not have to be bounded functions: whether a particular expression $m(\mathbf{x})$ has bounded increase on a system depends as much on the system itself as on the growth of m .

Example 7. Consider a stochastic system over \mathbb{R} , in which $x_{n+1} := x_n - 1 + w_n$, with w_n a uniform random variable over $[-1, 1]$. Then the function $m(x) : x$ has bounded increase property.

If each w_n is a Gaussian random variable, then $m(x)$ does not satisfy the bounded increase property. Restricting the set of support of distribution w_n to a compact set (by truncation), however, allows x to satisfy the bounded increase property again.

Returning to the MOONWALK system in Example 6, the additive supermartingale $m(x) : x$, whose sublevel sets do not prove any tail invariance property (due to Lemma 4), does not satisfy the bounded increase property since it is possible to move from $x = -j$, for any $j > 0$, to $x = 0$ with nonzero probability.

We close the section by demonstrating the incompleteness: an additive supermartingale does not *always* need the bounded increase property for a tail invariant property to be established.

Example 8 (Incompleteness). Consider the MOONWALK system with modified probability $p(x) : 1 - \frac{0.5}{(x-1)^2}$, for $x < 0$, and $p(0) = 1$. The probability of the event $\{x > \kappa\}$ is $\sum_{j=-\kappa}^{\infty} \frac{0.5}{(j+1)^2}$, which converges. By the Borel-Cantelli Lemma [14, 2.3.1], $\Pr(x > \kappa \text{ i.o.}) = 0$ holds, i.e., the tail invariant $\diamond\Box(x < \kappa)$ holds; however, the system does not have the bounded increase property.

5 Implementation and Evaluation

Given a polynomial stochastic system and a semi-algebraic target set, the problem of finding “certificates” V that prove persistence or recurrence properties is in general intractable. In practice, we impose several restrictions on the proof rules so that their solutions are tractable, based on sum-of-squares (SOS) optimization techniques (see e.g. [4, 6] and the references therein). For illustration, we only focus on SPERSIST-GEOM and PERSIST-GEOM; the formulations for the other proof rules are similar.

Recall that for proving strong persistence properties via the geometric rule SPERSIST-GEOM, we need to find a function V such that conditions (p1), (p6) hold. We impose the following restrictions to make the feasibility problem tractable. First, we require that V is a polynomial of degree at most some integer d . This means that $\mathcal{D}V$ is also a polynomial, which can be expressed in terms of the coefficients of V and the moments of the random variable \mathbf{r}_c . Second, we replace the nonnegativity constraints by the more restrictive *sum-of-squares* (SOS) constraints, i.e., we require that both V and $-\mathcal{D}V$ be sums of squares of some unknown polynomial functions. We also require that V is positive definite, which is a common regularity condition assumed in semidefinite optimization and allows us to find an $\alpha \in (0, 1)$ such that the condition (p6) in SPERSIST-GEOM

holds. Under these two restrictions, the generally intractable feasibility problem from SPERSIST-GEOM is equivalent to a linear semidefinite feasibility problem: a polynomial being a sum of squares (of polynomial functions) is equivalent to its vector of coefficients being the image of an unknown positive semidefinite matrix under a predetermined linear transformation. (For more details on SOS relaxation techniques for solving polynomial feasibility/optimization problems, see e.g. [4, 6].)

For proving persistence properties with respect to a nonempty set T via the geometric rule PERSIST-GEOM, we need to find a function V such that conditions (p1)-(p4) hold. Again, we require that V is a polynomial of degree at most d , and that $T = \{\mathbf{x} \mid g_1(\mathbf{x}) \geq 0 \wedge \cdots \wedge g_\ell(\mathbf{x}) \geq 0\}$ for some polynomials g_1, \dots, g_ℓ . Then we replace those constraints pertaining to the elements in T , by truncated *quadratic module* membership. For instance, we replace the condition (p3) by the tractable constraint:

$$DV = s_0 + s_1 g_1 + \cdots + s_\ell g_\ell, \quad s_0, s_1, \dots, s_\ell \text{ SOS of degree at most some integer } \tilde{d}.$$

(The tractability is due to the fact that the polynomials s_i being SOS can be phrased as semidefinite feasibility constraints.) Similar treatment can be applied on those constraints pertaining to the elements in $X \setminus T$, which is also a semialgebraic set.

Many standard semidefinite optimization solvers¹ and SOS optimization front-ends (such as SOSTOOLS [25]) are available for solving the SOS optimization problems outlined above. Below we present some simple examples on the use of SPERSIST-GEOM rules for proving persistence. In each example, an α -multiplicative supermartingale expression is obtained using SDPT3-4.0 [31] on MATLAB R2014b, taking less than 10 seconds on a Linux machine with Intel(R) Core(TM) i7-4650U CPU @ 1.70GHz.

Example 9 (Rimless wheel model [7, 22, 27]). A rimless wheel with 8 equally spaced inelastic spokes of length L rolls down a hill with stochastic slope angle γ . Let ω_n be the angular velocity at the n -th impact (which occurs when the stance leg is vertical). In [7, 27], the dynamics of the rimless wheel is described as:

$$x_{n+1} := \cos^2 \theta \left(x_n + \frac{2g}{L} (1 - \cos(\frac{\theta}{2} + \gamma)) \right) - \frac{2g}{L} (1 - \cos(\frac{\theta}{2} - \gamma)),$$

where $x_n = \omega_n^2$, g is the gravitational constant, $\theta = 45^\circ$ is the angle between two consecutive spokes and $\gamma \sim \mathcal{N}(8, 1)$ (in degrees). We approximate the functions $\xi \mapsto \cos(\frac{\theta}{2} \pm \xi)$ over the interval [5, 11] by degree 2 polynomials, and find that the angular velocity in the approximated stochastic system goes to 0 almost surely when $L = 2g$: the function $V(x) : 0.00085x^3 + x^4$ satisfies the conditions (p1), (p6) with $X : [0, \infty)$ and $\alpha = 0.95$: V and $-DV$ are nonnegative on X and $DV(x) \leq -0.05V(x)$ for all $x \geq 0$. Hence V is a α -multiplicative supermartingale for this system over X , and $\diamond \square(V(x) \leq \varepsilon)$ holds a.s. for any $\varepsilon > 0$. In other words, despite the randomness in the slope of the terrain, the rolling rimless wheel (with very long spokes) would eventually become stationary almost surely.

¹ See e.g. the list in <http://plato.asu.edu/sub/nlores.html#semidef>.

Example 10 (Room temperature control [3]). In the two-room temperature control example from Sect. 1.1, we are interested in the evolution of the room temperatures within the range $X = [6, 33]^2$. Consider the nonnegative function $V(x_1, x_2) : (x_1 - 18.3)^2 + (x_2 - 18.8)^2$. When the noise follows the uniform distribution $\mathcal{U}(-0.01, 0.01)$, $\mathcal{D}V$ is nonpositive on X , and $V(x_1, x_2) \geq 0.09$ and $\mathcal{D}V(x_1, x_2) \leq -0.01V(x_1, x_2)$ for all $x \in X \setminus [17.8, 18.7] \times [18.4, 19.3]$. Hence conditions (p1)-(p4) hold, implying the persistence property $\diamond\Box(17.8 \leq x_1 \leq 18.7 \wedge 18.4 \leq x_2 \leq 19.3)$.

In the case of Gaussian noise $\mathcal{N}(0, 0.25)$, $\mathcal{D}V(x_1, x_2) \leq 0.25$ for all $(x_1, x_2) \in [16.9, 19.6] \times [17.3, 20.2]$, and $V(x_1, x_2) \geq 0.8$ and $\mathcal{D}V(x_1, x_2) \leq -6 \times 10^{-5}$ for all $(x_1, x_2) \in X \setminus [16.9, 19.6] \times [17.3, 20.2]$. Hence conditions (r1)-(r4) hold, implying the recurrence property $\Box\diamond(16.9 \leq x_1 \leq 19.6 \wedge 17.3 \leq x_2 \leq 20.2)$.

We list some additional examples in which a system is proved to satisfying a persistence or recurrence property via some of the proof rules from Sect. 3. The details of these examples can be found in the extended version of the paper.

ADDITIONAL STOCHASTIC SYSTEMS	NOISE u_j (i.i.d.)	SUPERMARTINGALE $V(x, y)$
$x' := x + \frac{1}{2}y + u_1,$ $y' := \frac{1}{2}x + y - u_2$ (over $X = \mathbb{R}^2$)	$\mathcal{N}(-1, 1)$	$\max(x - y, 0)$ (for proving recurrence)
$x' := 0.5(x + y) + 0.4u_1\sqrt{x^2 + y^2},$ $y' := 0.5(x - y) + 0.4u_2\sqrt{x^2 + y^2},$ (over $X = \mathbb{R}^2$)	$\mathcal{N}(0, 1)$	$x^2 + y^2$ (0.82-multi.)
$x' := 0.75y^4 + 0.1u_1,$ $y' := 0.75x^4 + 0.1u_2,$ (over $X = \{(x, y) \mid x^2 + y^2 \leq 1\}$)	$\mathcal{U}(-1, 1)$	$0.78x^2 + 1.23xy + 0.78y^2$ (0.75-multi.)
$x' := 0.1(y(3x^2 + 2y^2 - 0.5) + u_1\sqrt{x^2 + y^2}),$ $y' := 0.1(y(2x^2 + 4xy + 3y^2 - 0.5) + u_2\sqrt{x^2 + y^2}),$ (over $X = \{(x, y) \mid x^2 + y^2 \leq 1\}$)	$\mathcal{U}(-\sqrt{3}, \sqrt{3})$	$1.55x^2 + 2.36xy + 1.34y^2$ (0.5-multi.)

6 Related Work

Martingale analysis has been used to prove almost sure termination [8, 10, 15], derive inductive invariant expressions [9] in probabilistic programs, and prove stochastic reachability and safety [26, 27] in the context of stochastic hybrid systems. Our paper extends this set of properties to include tail invariant, or qualitative persistence, properties.

Qualitative persistence properties are expressible in PCTL [18] and have been studied by the model-checking community [5, 21] in the context of finite-state Markov processes. The approach leverages the fact that the system ends up a.s. in a *bottom strongly connected component* (BSCC [11–13], a strongly connected component with no outgoing edges), then uses a graph algorithm to efficiently check that all states in the BSCC satisfy φ_{INV} . Baier et al. [5] have shown that

while such results suffice for the analysis of finite-state Markov chains, they, however, do not extend to infinite-state models.

Abate [1, 2] and Tkachev et al. [29] present approaches to reducing the verification problem of infinite-state Markov processes to that over finite-state Markov chains. Specifically, [1, 2] presents a framework for proving probabilistic bisimulation between the original infinite-state system and its discretized (approximate) finite-state version. Guarantees on the quality of the results are proved using supermartingale bisimulation functions. Tkachev et al. [29] present a framework for analyzing infinite-horizon reach-avoid properties ($\diamond\varphi$ and $\square\varphi$) for Markov processes. They use locally-excessive (i.e., supermartingale) value functions to identify a subset of the state-space where discretization guarantees a precise approximate solution. In [28, 30] Tkachev et al. tackle *quantitative* reachability, invariance and reach-while-avoid properties operating directly over the infinite-state model. [28] provides a characterization of the statespace based on harmonic functions defining absorbing or stochastically attractive sets. Unfortunately, the sufficient conditions for certificates in [28] define problems that in general have no analytical or computation solution. Our paper can be seen as a set of practical sufficient conditions that yield efficiently computable problems (SOS, SDP) for qualitative repeated reachability.

The problem of stability and identifying the limiting behavior of a stochastic system has been well-studied in the theory of Markov chains [24]. Similar “Foster-Lyapunov” [16] drift criteria have been derived to argue recurrence and transience for sets of states. Unfortunately, most results rely on the topological properties of the infinite-state Markov chains that may be difficult to check automatically.

7 Conclusion

We presented an analysis framework capable of proving that the limiting behavior of an infinite-state discrete-time polynomial stochastic system eventually settles almost surely within a region S of the statespace (i.e., $\diamond\square(S)$). Our analysis employs constraint-based invariant generation techniques to efficiently infer polynomial functions over the states of the system: nonnegative α -multiplicative supermartingale expressions and additive supermartingale expressions of bounded increase. We established that both types of functions constitute certificates verifying tail invariant properties and we demonstrated their equivalence. Finally, we highlighted the individual strengths of each of the two types but also the incompleteness of the general approach through the means of numerous simple, yet intricate examples.

Acknowledgements. We thank the anonymous reviewers for their comments. This work was supported by the US National Science Foundation (NSF) under award numbers 1527075 and 1320069. All opinions expressed are those of the authors and not necessarily of the US NSF.

References

1. Abate, A.: A contractivity approach for probabilistic bisimulations of diffusion processes. In: Proceedings of CDC, pp. 2230–2235 (2009)
2. Abate, A.: Probabilistic bisimulations of switching and resetting diffusions. In: Proceedings of CDC, pp. 5918–5923 (2010)
3. Abate, A., Katoen, J.-P., Lygeros, J., Prandini, M.: Approximate model checking of stochastic hybrid systems. *Eur. J. Control* **16**(6), 624–641 (2010)
4. Anjos, M.F., Lasserre, J.B.: Handbook on Semidefinite, Conic and Polynomial Optimization. International Series in Operations Research & Management Science, vol. 166. Springer, New York (2012)
5. Baier, C., Katoen, J.-P., et al.: Principles of Model Checking, vol. 26202649. MIT press, Cambridge (2008)
6. Blekherman, G., Parrilo, P.A., Thomas, R.R.: Semidefinite optimization and convex algebraic geometry. In: MOS-SIAM Series on Optimization. Society for Industrial and Applied Mathematics (SIAM), vol. 13. Mathematical Optimization Society, Philadelphia, PA (2013)
7. Byl, K., Tedrake, R.: Metastable walking machines. *Int. J. Robot. Res.* **28**(8), 1040–1064 (2009)
8. Chakarov, A., Sankaranarayanan, S.: Probabilistic program analysis with martingales. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 511–526. Springer, Heidelberg (2013)
9. Chakarov, A., Sankaranarayanan, S.: Expectation invariants for probabilistic program loops as fixed points. In: Müller-Olm, M., Seidl, H. (eds.) Static Analysis. LNCS, vol. 8723, pp. 85–100. Springer, Heidelberg (2014)
10. Chatterjee, K., Fu, H., Hasheminezhad, R., Novotny, P.: Algorithmic analysis of qualitative and quantitative termination problems for affine probabilistic programs. In: Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL , St. Petersburg, Florida, United States, 20–22 January 2016
11. Courcoubetis, C., Yannakakis, M.: The complexity of probabilistic verification. *J. ACM (JACM)* **42**(4), 857–907 (1995)
12. De Alfaro, L.: Formal verification of probabilistic systems. PhD thesis, Stanford University (1997)
13. De Alfaro, L.: How to specify and verify the long-run average behaviour of probabilistic systems. In: Proceedings of LICS, pp. 454–465. IEEE (1998)
14. Durrett, R.: Probability: theory and examples. Cambridge University Press, Cambridge (2010)
15. Fioriti, L.M.F., Hartmanns, A., Hermann, H.: Probabilistic termination: Soundness, completeness, and compositionality. In: Proceedings of POPL, pp. 489–501. ACM (2015)
16. Foster, F.G.: On the stochastic matrices associated with certain queuing processes. *Ann. Math. Stat.* **24**(3), 355–360 (1953)
17. Gordon, A.D., Henzinger, T.A., Nori, A.V., Rajamani, S.K.: Probabilistic programming. In: Proceedings of FOSE 2014, pp. 167–181 (2014)
18. Hansson, H., Jonsson, B.: A logic for reasoning about time and reliability. *Formal Aspects Comput.* **6**(5), 512–535 (1994)
19. Katoen, J.-P., McIver, A.K., Meinicke, L.A., Morgan, C.C.: Linear-invariant generation for probabilistic programs: In: Cousot, R., Martel, M. (eds.) SAS 2010. LNCS, vol. 6337, pp. 390–406. Springer, Heidelberg (2010)

20. Kozen, D.: Semantics of probabilistic programs. *J. Comput. Syst. Sci.* **22**(3), 328–350 (1981)
21. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: verification of probabilistic real-time systems. In: Gopalakrishnan, G., Qadeer, S. (eds.) *CAV 2011*. LNCS, vol. 6806, pp. 585–591. Springer, Heidelberg (2011)
22. McGeer, T.: Passive dynamic walking. *Int. J. Robot. Res.* **9**(2), 62–82 (1990)
23. McIver, A., Morgan, C.: *Abstraction, Refinement And Proof For Probabilistic Systems (Monographs in Computer Science)*. SpringerVerlag, New York (2004)
24. Meyn, S.P., Tweedie, R.L.: *Markov Chains and Stochastic Stability*. Cambridge University Press, Cambridge (2009)
25. Papachristodoulou, A., Anderson, J., Valmorbida, G., Prajna, S., Seiler, P., Parrilo, P.A.: SOSTOOLS: Sum of squares optimization toolbox for MATLAB (2013). <http://arxiv.org/abs/1310.4716>
26. Prajna, S., Jadbabaie, A., Pappas, G.J.: Stochastic safety verification using barrier certificates. In: 43rd IEEE Conference on Decision and Control, CDC, vol. 1, pp. 929–934. IEEE (2004)
27. Steinhardt, J., Tedrake, R.: Finite-time regional verification of stochastic non-linear systems. *Int. J. Robot. Res.* **31**(7), 901–923 (2012)
28. Tkachev, I., Abate, A.: Stability and attractivity of absorbing sets for discrete-time Markov processes. In: IEEE 51st Annual Conference on Decision and Control (CDC), pp. 7652–7657. IEEE (2012)
29. Tkachev, I., Abate, A.: Characterization and computation of infinite-horizon specifications over markov processes. *Theoret. Comput. Sci.* **515**, 1–18 (2014)
30. Tkachev, I., Mereacre, A., Katoen, J.-P., Abate, A.: Quantitative model-checking of controlled discrete-time Markov processes (2014). arXiv preprint [arXiv:1407.5449](https://arxiv.org/abs/1407.5449)
31. Kim-Chuan, T., Todd, M.J., Tütüncü, R.H.: On the implementation and usage of SDPT3—a matlab software package for semidefinite-quadratic-linear programming, version 4.0. In: Anjos, M.F., Lasserre, J.B. (eds.) *Handbook on Semidefinite, Conic and Polynomial Optimization*. International Series in Operations Research and Management Science, vol. 166, pp. 715–754. Springer, New York (2012)
32. Williams, D.: *Probability with Martingales*. Cambridge University Press, Cambridge (1991)