

Sound and Complete Proof Rules for Probabilistic Termination

RUPAK MAJUMDAR, Max Planck Institute for Software Systems (MPI-SWS), Germany

V. R. SATHIYANARAYANA, Max Planck Institute for Software Systems (MPI-SWS), Germany

Termination is a fundamental question in the analysis of probabilistic imperative programs. We consider the *qualitative* and *quantitative* probabilistic termination problems for an imperative programming model with discrete probabilistic choice and demonic bounded nondeterminism. The qualitative question asks if the program terminates almost surely, no matter how nondeterminism is resolved; the quantitative question asks for a bound on the probability of termination. Despite a long and rich literature on the topic, no sound and relatively complete proof systems were known for this problem. We provide the first sound and relatively complete proof rules for proving qualitative and quantitative termination in the assertion language of arithmetic. Our proof rules use supermartingales as estimates of likelihood of the program's evolution—the key insight is to use appropriately defined finite-state sub-instances. Our completeness result shows how to construct a suitable supermartingale from an almost-surely terminating program. We also show that proofs of termination in many existing proof systems can be transformed to proofs in our system, pointing to its applicability in practice. As an application of our proof rule, we show a proof of almost sure termination for the two-dimensional random walker.

ACM Reference Format:

Rupak Majumdar and V. R. Sathiyararayanan. 2024. Sound and Complete Proof Rules for Probabilistic Termination. 1, 1 (May 2024), 26 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Probabilistic programming languages extend the syntax of usual deterministic computation with primitives for random choice. Thus, probabilistic programs express randomized computation and have found applications in many domains where randomization is essential.

We study the *termination problem* for probabilistic programs with discrete probabilistic and non-deterministic choice. Termination is a fundamental property of programs and formal reasoning about (deterministic) program termination goes back to Turing [45]. Its extension to the probabilistic setting can be either qualitative or quantitative. *Qualitative* termination, or Almost-Sure Termination (AST) asks if the program terminates almost-surely, no matter how the nondeterminism is resolved. *Quantitative* termination, on the other hand, relates to finding upper and lower bounds on the probability of termination that hold across all resolutions of nondeterminism.

For finite-state probabilistic programs, both qualitative and quantitative termination problems are well understood: there are sound and complete *algorithmic* procedures for termination that operate by analyzing the underlying finite-state Markov decision processes. Intuitively, every run

Authors' addresses: [Rupak Majumdar](#), Max Planck Institute for Software Systems (MPI-SWS), Paul-Ehrlich-Straße, Building G26, Kaiserslautern, 67663, Germany, rupak@mpi-sws.org; [V. R. Sathiyararayanan](#), Max Planck Institute for Software Systems (MPI-SWS), Paul-Ehrlich-Straße, Building G26, Kaiserslautern, 67663, Germany, sramesh@mpi-sws.org.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM XXXX-XXXX/2024/5-ART

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

of the system eventually arrives in an end component and thus, computing the termination probabilities reduces to computing the reachability probabilities for the appropriate end components [7, 13–15, 28, 46].

The story is different for infinite state spaces. Existing techniques for deducing termination typically take the form of *sound* proof rules over a program logic [9–11, 15, 28, 36, 37]. These rules ask for certificates consisting of a variety of mathematical entities that satisfy locally-checkable properties. None of them, however, are known to be complete; that is, we do not know if certificates can always be found for terminating programs. The search for relatively complete proof rules has been a long-standing open problem. Note that, since the termination problem is undecidable, one can only hope for completeness *relative* to an underlying logic.

In this paper, we describe the first sound and relatively complete proof rules for qualitative and quantitative termination of probabilistic programs. We present our rules in a simple proof system in the style of Floyd [22] that applies naturally to our program model. This proof system uses arithmetic as its assertion language, interpreted over the standard model of rational numbers. Soundness means that if our proof rule applies, then indeed the system satisfies the (qualitative or quantitative) termination criterion. Completeness of our rules is relative to the completeness of a proof system for the underlying assertion language, i.e., arithmetic. Accordingly, we show an effective reduction from the validity of our program logic to the validity of a finite number of assertions in arithmetic. Whenever the original program terminates (qualitatively or quantitatively), one can construct a proof in our program logic in such a way that all relevant certificates can be expressed in the assertion language. This is important: merely knowing that certain semantic certificates exist may not be sufficient for a proof system, e.g., if these certificates are provided non-constructively or require terms that cannot be expressed in the assertion language.

Let us be more precise. We work in an imperative programming model with variables ranging over rationals. Our model fixes a finite set of program locations, and defines a guarded transition relation between the locations representing computational steps. At marked locations, the model contains primitives for probability distributions over available transitions. This allows for the expression of *bounded* nondeterministic and probabilistic choice; we assume the nondeterminism is resolved demonically. We fix the language of arithmetic as our expression and assertion language, and interpret formulas over the standard model of the rational numbers.¹ The semantics of our programming language is given by a Markov decision process on countably many states, where a demonic scheduler resolves the nondeterminism. Since the language has bounded nondeterminism, we note that each state has a finite number of immediate successor states and so, for every scheduler, the number of states reached in a bounded number of steps is finite.

Given a program and a terminal state, the *qualitative termination* question asks if the infimum over all schedulers resolving nondeterminism of the probability of reaching the terminal state is one, that is, if the program almost surely terminates under all possible schedulers. The *quantitative termination* question asks if the probability of reaching the terminal state is bounded above or below by a given probability p .

For the special case of programs without probabilistic choice, sound and relatively complete proof systems for termination are known [1, 26, 35]: they involve finding a *variant* function from (reachable) program states to a well-founded domain that decreases on every step of the program, and which maps the terminal state to a minimal element.

A natural generalization of variant functions is a *ranking supermartingale*: a function from states to reals that reduces in expectation by some amount on each step of the program. Ranking supermartingales of various flavors are the workhorse of existing proof rules for qualitative termination.

¹One can generalize our result to *arithmetical* structures [27], but we stick to the simpler setting for clarity.

Unfortunately, an example from [34] shows that a proof rule based only on a ranking supermartingale is incomplete: there may not exist a ranking supermartingale that decreases in expectation on each step and one may require transfinite ordinals in proofs.

Basic Ingredients. We take a different perspective. Instead of looking for a mapping that represents the distance to a terminal state that goes down in each step, as in previous approaches, we consider modeling the relative likelihood of the program’s evolution instead. We provide two different proof rules. In the first, we certify almost sure termination using a function that is unbounded and non-increasing in expectation on “most” states. In the second, we certify almost sure termination using a family of functions, one for each reachable state, each non-increasing in expectation. Both certificates track the execution’s relative likelihood in subtly different ways, and both require additional side conditions.

We prove an *unrolling lemma* that is a central tool for our completeness results. It states that if the infimum over all schedulers of the probability of reaching a terminal state is at least p , then for every ϵ , there is a finite upper bound k such that the infimum over all schedulers of the probability mass of reaching the terminal state within k steps is at least $p - \epsilon$. In particular, the set of states reachable in k steps defines a finite state space. The unrolling lemma appears as a basic ingredient in characterizing the complexity of almost sure termination [31, 33]. We show that it provides a surprisingly powerful tool in proving completeness of proof systems by “carving out” finite state systems out of infinite-state termination problems.

Our Proof Rules for Qualitative Termination. I. Our first rule asks for a supermartingale [18] V that is non-increasing in expectation on all states except for some set containing the terminal state. In addition, our rule asks for a variant function U that certifies that every reachable state has some finite path to the terminal state. We also require a few compatibility conditions on U from V to let us conclude the almost-sure escape from sets of states within which V is bounded.

We show the rule is sound by partitioning the collection of all runs and strategically employing variant arguments and/or martingale theory within each partition. The completeness of the rule uses the following observation. Fix an enumeration of the reachable states $s_1, s_2, \dots$ of an almost-surely terminating program. Let $\Pr_s[\Diamond \mathbb{1}_{>n}]$ denote the probability that a run starting from state s reaches some state in $\{s_n, s_{n+1}, \dots\}$ in the enumeration. For a fixed s , we first show that $\Pr_s[\Diamond \mathbb{1}_{>n}]$ goes to zero as $n \rightarrow \infty$. Following a diagonal-like construction from the theory of countable Markov chains [39], we next define a sequence $n_1, n_2, \dots$ such that $\Pr_j[\Diamond \mathbb{1}_{s_{n_k}}] \leq \frac{1}{2^k}$ for all $j \leq k$ (this is possible since the limit goes to zero and by the unrolling lemma). This lets us define a supermartingale defined over the states s as

$$\sum_{k \in \mathbb{N}} \Pr_s[\Diamond \mathbb{1}_{>n_k}]$$

This supermartingale satisfies the requirements of our rule. Moreover, we show that this supermartingale can be expressed in arithmetic, granting the rule relative completeness.

II. We provide a second dual proof rule that takes a more local view. Our first rule required certificates inferred from the global behaviour of the program. Our second rule, by contrast, requires proofs of *near termination*, i.e., termination with some non-zero probability, from *every* reachable state. If these proofs together indicate a non-zero lower bound of termination across all states, a zero-one law indicates almost-sure termination. Therefore, our rule asks for a proof of termination with probability at least $1 - \epsilon$, for some $\epsilon > 0$.

How does our rule certify termination from a given state with probability $1 - \epsilon$? It incorporates a proof rule for quantitative termination by Chatterjee et al. [10] that builds finite supermartingales that take on non-trivial values for only finitely many states. This rule employs *stochastic invariants*

[11]: pairs (SI, p) such that the probability with which executions leave the set of states SI is bounded above by p . Our proof rule needs a family of stochastic invariants, one for each reachable state.

The completeness of this rule uses the unrolling lemma in a crucial way. The unrolling lemma implies a finite state space around every state that accumulates a termination probability mass of $1 - \epsilon$. We use this fact to show that the proof rule of Chatterjee et al. [10] is complete for finite-state systems. This completeness in turn induces the finite nature of these supermartingales describing the stochastic invariants around each state.

Quantitative Termination. All our rules for quantitative termination again use the stochastic invariants of Chatterjee et al. [11]. A stochastic invariant easily implies an upper bound rule: if there is a stochastic invariant (SI, p) that avoids terminal states, the probability of termination is upper bounded by p . For lower bounds, our starting point is the proof rule proposed by Chatterjee et al. [10] using stochastic invariants: if there is a stochastic invariant (SI, p) such that runs almost surely terminate within SI or leave SI , then the probability of termination is at least $1 - p$. While they claimed soundness and completeness for their rule, there were two issues. First, their rule was parameterized by a certificate for qualitative termination. In the absence of a relatively complete proof rule for qualitative termination, one could not achieve relative completeness. Second, we show in Section 5.3 an explicit example where their rule cannot apply.

We show a sound and complete rule that is a modification of their rule: we require that for each $n \in \mathbb{N}$, there is a stochastic invariant $(SI_n, p + \frac{1}{n})$ such that all runs almost surely terminate within SI_n or leave SI_n . Sound and relatively complete certificates for almost sure termination are now given using our previous technique for qualitative termination.

In summary, we provide the first sound and relatively complete proof rules for qualitative and quantitative termination, culminating the substantial body of work on probabilistic termination in the last four decades.

Other Related Work. Our proof rules use supermartingales that are closely related to Lyapunov functions in stability of dynamical systems. Lyapunov functions have been used to characterize recurrence and transience in infinite-state Markov chains, going back to the work of Foster [23, 24]. Completeness of Lyapunov functions was shown in general by Mertens et al. [39]. Our proof of soundness and completeness uses insights from Mertens et al. [39], but we have to overcome several technical issues. First, we have demonic nondeterminism and therefore require the unrolling lemma to deal with infimums over all schedulers. Second, we do not have irreducibility. Finally, whereas a Markov chain is either recurrent or transient, we cannot assume that a program that is not almost sure terminating has a strong transience property. Thus, we have to prove these properties ab initio.

In the literature, there exist sound proof rules for AST that use supermartingales. One rule by Huang et al. [30] uses supermartingales that exhibit a lower bound on their variation in each step. A much more closely related work is the excellent AST proof rule by McIver et al. [37]. Their work shows that a proof rule consisting of a supermartingale function that also acts as a distance variant is sound for almost-sure termination. While the former is believed to be incomplete [37], it is not known if the latter rule is complete. Our work indicates that one can achieve completeness by separating the roles of the distance variant and the supermartingale into two functions.

The issue of an appropriate assertion language for proof rules for termination have been mostly elided in the literature, and most rules are presented in an informal language of “sufficiently expressive” mathematical constructs. Important exceptions are the assertion languages of Batz et al. [6] and den Hartog and de Vink [16]. Their work shows that the language of arithmetic extended with suprema and infima of functions over the state space is relatively complete (in the sense of

Cook [12]) for weakest-preexpectation style reasoning of probabilistic programs without nondeterminism. That is, given a function f definable in their language and a program P , they show that their language is expressive enough to represent the weakest pre-expectation of f with respect to P . The need for suprema and infima is motivated by demonstrating simple probabilistic programs whose termination probabilities involve transcendental numbers. We note that arithmetic is sufficient for relative completeness because suprema and infima arising in probabilistic termination can be encoded (through quantifiers). We believe that presenting the rules directly in the language of arithmetic allows greater focus on the nature of the certificates required by the rules. Note that for extensions of the programming model, such as with unbounded nondeterministic choice, arithmetic is no longer sufficient for relative completeness, and this holds already without probabilistic choice in the language [2, 3, 29].

While we focus on almost sure termination, there are related qualitative termination problems: *positive* almost sure termination (PAST) and *bounded* almost sure termination (BAST). These problems strengthen almost sure termination by requiring that the expected time to termination is finite. (Note that a program may be almost surely terminating but the expected run time may be infinite: consider a one-dimensional symmetric random walk where 0 is an absorbing state.) The difference is that PAST allows the expected run time to depend on the scheduler that resolves nondeterminism, and BAST requires a global bound that holds for every scheduler. Sound and complete proof rules for BAST have been studied extensively [4, 8, 20, 25]. More recently, a sound and complete proof rule for PAST was given [34]. Completeness in these papers are semantic, and relative completeness in the sense of Cook was not studied. Our techniques would provide a relative completeness result for BAST. In contrast, Majumdar and Sathiyarayanan [34] show that PAST is Π_1^1 -complete (AST and BAST are arithmetical, in comparison); thus Peano arithmetic would be insufficient as a (relatively complete) assertion language (see [3] for similar issues and appropriate assertion languages for nondeterministic programs with countable nondeterminism).

Our results apply to discrete probabilistic choice. While discrete choice and computation captures many randomized algorithms, our proofs of completeness do not apply to programs with, e.g., sampling from continuous probability distributions. The use of real values introduce measure-theoretic complexities in the semantics [44]. These can be overcome, but whether there is a sound and relatively complete proof rule for an appropriate assertion language remains open.

While we focus on the theoretical aspects here, there is a large body of work on *synthesizing* certificates automatically and tools for probabilistic verification [9, 10, 19, 37]. We show a “compilation” of many existing rules into our rules, thus, such tools continue to work in our proof system.

2 PROBABILISTIC PROGRAMS

2.1 Syntax and Semantics

Syntax. We work with *probabilistic control flow graphs*, a program model used by Chatterjee et al. [10] to detail proof rules for quantitative termination. Variables in this model range over the rationals. Assignment statements and loop guards are terms and boolean combinations of atomic formulae expressible in the language of arithmetic. This is standard in program logics [1], and facilitates the use of the language of rational arithmetic with addition, multiplication, and order to make assertions about desirable program properties. For sake of conciseness, we augment this assertion language with additional computable predicates as “syntactic sugar” in our proofs. We interpret assertions over the standard model of rationals.

Definition 2.1 (Control Flow Graphs). A Control Flow Graph (CFG) $\mathcal{G}$ is a tuple $(L, V, l_{init}, x_{init}, \mapsto, G, \text{Pr}, \text{Upd})$, where

- L is a finite set of program locations, partitioned into assignment, nondeterministic, and probabilistic locations L_A , L_N , and L_P , respectively.
- $V = \{x_1, x_2, \dots, x_n\}$ is a finite set of program variables.
- $l_{init} \in L$ is the initial program location, and $\mathbf{x}_{init} \in \mathbb{Q}^V$ is the initial variable valuation.
- $\mapsto \subseteq L \times L$ is a finite set of transitions. If $\tau = (l, l') \in \mapsto$, then l and l' are respectively referred to as the source and target locations of τ .
- G is a function mapping each $\tau \in \mapsto$ to a Boolean expression over $V \cup L$.
- Pr is a function assigning each $(l, l') \in \mapsto$ with $l \in L_P$ a fractional expression p over the variables V representing a rational probability value.
- Upd is a map assigning each $(l, l') \in \mapsto$ with $l \in L_A$ an update pair (j, u) where $j \in \{1, \dots, |V|\}$ is the target variable index and u is an arithmetic expression over the variables V .
- At assignment locations, there is at most one outgoing transition.
- At probabilistic locations l , it must be that $\text{Pr}(l, _)[\mathbf{x}] > 0$ and $\sum G(l, _)[(\mathbf{l}, \mathbf{x})] \times \text{Pr}((l, _))[\mathbf{x}] = 1$ over all transitions $(l, _) \in \mapsto$ for all $\mathbf{x} \in \mathbb{Q}^V$.

We use the boldface notation $\mathbf{x}$ for variable assignments and write $\mathbf{0}$ for the assignment that maps every variable to zero. $\text{Pr}(l, l')[\mathbf{x}]$, $G(l, l')[\mathbf{x}]$, and $\text{Upd}(l, l')[\mathbf{x}]$ refer to the output of the expressions $\text{Pr}(l, l')$, $G(l, l')$, and $\text{Upd}(l, l')$ on the assignment $\mathbf{x}$. Note that the finiteness of L implies that the branching at both nondeterministic and probabilistic locations is bounded. Without loss of generality, we assume simple structural conditions that ensures that every state has a successor. This follows similar assumptions made in prior work [10]. Observe that, while the probabilistic choice is simple, it is sufficient to model probabilistic Turing machines and some quite sophisticated probabilistic phenomena [21]. However, we explicitly forbid unbounded nondeterministic choice or sampling from continuous distributions.

Remark 2.2. While we use CFGs as our formal model of programs, we could have equivalently used probabilistic guarded command language (pGCL). pGCL is the probabilistic extension of the Guarded Command Language of Dijkstra [17], and is a convenient language for specifying probabilistic computation. There is a large body of work [6, 19, 31, 36, 37] that uses pGCL syntax. Our choice of CFGs follows the same choice made by Chatterjee et al. [10] to describe quantitative termination. It is standard to compile pGCL programs into CFGs and vice versa. For readability, we employ the syntax of pGCL in some of our examples.

States, Runs, and Reachable States. Fix a CFG $\mathcal{G} = (L, V, l_{init}, \mathbf{x}_{init}, \mapsto, G, \text{Pr}, \text{Upd})$. A *state* is a tuple $(l, \mathbf{x})$, where $l \in L$ and $\mathbf{x} \in \mathbb{Q}^V$. A state $(l, \mathbf{x})$ is termed *assignment* (resp., *nondeterministic*, or *probabilistic*) if the location l is assignment (resp., nondeterministic, or probabilistic). We will refer to assignment locations l where the updates of all transitions sourced at l don't change the variable values as *deterministic* locations. Accordingly, states $(l, \mathbf{x})$ are termed *deterministic* when l is deterministic. A transition $(l, l') \in \mapsto$ is *enabled* at a state $(l, \mathbf{x})$ if the guard $G(l, l')$ evaluates to true under $(l, \mathbf{x})$. The vector $\mathbf{x}'$ is the *result* of the update pair (j, u) from the state $(l, \mathbf{x})$ if (a) for all $i \neq j$, $\mathbf{x}'[i] = \mathbf{x}[i]$, and (b) $\mathbf{x}'[j] = u(\mathbf{x})$. A state $(l', \mathbf{x}')$ is a *successor* to $(l, \mathbf{x})$ if the transition $(l, l') \in \mapsto$ is enabled at $(l, \mathbf{x})$ and $\mathbf{x}'$ is the result of $\text{Upd}(l, l')$ on $\mathbf{x}$. A *finite path* is a sequence of states $(l_1, \mathbf{x}_1), (l_2, \mathbf{x}_2), \dots, (l_n, \mathbf{x}_n)$ with $(l_{k+1}, \mathbf{x}_{k+1})$ being a successor to $(l_k, \mathbf{x}_k)$. A *run* (or *execution*) of $\mathcal{G}$ is a sequence of states that (a) begins with the initial state $(l_{init}, \mathbf{x}_{init})$, and (b) only induces finite paths as prefixes.

A state $(l', \mathbf{x}')$ is said to be *reachable* from a state $(l, \mathbf{x})$ if there exists a finite path beginning at $(l, \mathbf{x})$ and ending at $(l', \mathbf{x}')$. We write $\text{Reach}(\mathcal{G}, (l, \mathbf{x}))$ for the set of states reachable from $(l, \mathbf{x})$; we simply write $\text{Reach}(\mathcal{G})$ when the initial state is $(l_{init}, \mathbf{x}_{init})$. An CFG is said to be *finite state* if the set of states reachable from its initial state is finite.

Probability Theory. We now introduce some basic probability theoretic notions. The tuple $(X, \mathcal{F}, \mathbb{P})$ is called a *probability space* where X is the sample space, $\mathcal{F}$ is a σ -algebra over X , and $\mathbb{P}$ is the probability measure over $\mathcal{F}$. The sequence $(\mathcal{F}_n)$, where n ranges over $\mathbb{N}$, is a *filtration* of the probability space $(X, \mathcal{F}, \mathbb{P})$ if each $\mathcal{F}_n$ is a sub σ -algebra of $\mathcal{F}$ and $\mathcal{F}_i \subseteq \mathcal{F}_j$ for all $i \leq j$.

A *random variable* is a measurable function from X to $\mathbb{R}^+$, the set of positive real numbers. The *expected value* of the random variable X is denoted by $\mathbb{E}[X]$. A *stochastic process* over the filtered probability space is a sequence of random variables (X_n) such that each X_n is measurable over $\mathcal{F}_n$. A *supermartingale* is a stochastic process (X_n) that does not increase in expectation, i.e., $\mathbb{E}[X_{n+1}] \leq \mathbb{E}[X_n]$ for each $n \in \mathbb{N}$. We will abuse notation slightly and refer to any function over the state space of a CFG that doesn't increase in expectation at each execution step as a *supermartingale function*.

We shall make use of the following version of Doob's Martingale Convergence Theorem.

THEOREM 2.3 ([18]). *If a supermartingale (X_n) is bounded below, then there almost-surely exists a random variable X_∞ such that*

$$\mathbb{P}\left(X_\infty = \lim_{n \rightarrow \infty} X_n\right) = 1 \quad \text{and} \quad \mathbb{E}[X_\infty] \leq \mathbb{E}[X_0]$$

Schedulers and Probabilistic Semantics. Now, we introduce the operational semantics of our programs. These are standard operational semantics that can be found in many other places [5, 10]. A *scheduler* is a mapping from finite paths ending at nondeterministic states to successors from these states. Note that, since the state space for any CFG is countable, we do not require measurability conditions on the scheduler.

The semantics of $\mathcal{G}$ is understood through a probability space over the runs of $\mathcal{G}$. Formally, let $\text{Runs}_{\mathcal{G}}$ be the collection of all executions of $\mathcal{G}$. Further, for a finite path π , denote by $\text{Cyl}_{\mathcal{G}}(\pi)$ the *cylinder set* containing all runs $\rho \in \text{Runs}_{\mathcal{G}}$ such that π is a prefix of ρ . Now, call $\mathcal{F}_{\mathcal{G}}$ the smallest σ -algebra on $\text{Runs}_{\mathcal{G}}$ containing all cylinder sets of finite paths of $\mathcal{G}$.

A scheduler $\mathfrak{s}$ induces a probability space over the collection of all runs of the CFG $\mathcal{G}$. A finite path (or run) π is said to be *consistent* with $\mathfrak{s}$ if for every prefix π' of π ending at a nondeterministic state, the finite path (or run) obtained by appending the successor state $\mathfrak{s}(\pi')$ to π' is a prefix of π . A scheduler is said to induce a finite path (or run) if the path (or run) is consistent with the scheduler. The semantics of the CFG $\mathcal{G}$ under the scheduler $\mathfrak{s}$ is captured by the probability space $(\text{Runs}_{\mathcal{G}}, \mathcal{F}_{\mathcal{G}}, \mathbb{P}_{\mathfrak{s}})$, where for every consistent finite path $\pi = ((l_1, \mathbf{x}_1), (l_2, \mathbf{x}_2), \dots, (l_n, \mathbf{x}_n))$ with probabilistic locations at indices $i_1, i_2, \dots, i_n$,

$$\mathbb{P}_{\mathfrak{s}}(\pi) = \Pr(l_{i_1}, l_{i_1+1})[\mathbf{x}_{i_1}] \times \dots \times \Pr(l_{i_n}, l_{i_n+1})[\mathbf{x}_{i_n}]$$

We analogously define a probability space $(\text{Runs}_{\mathcal{G}(l, \mathbf{x})}, \mathcal{F}_{\mathcal{G}(l, \mathbf{x})}, \mathbb{P}_{\mathfrak{s}})$ to refer to the probability space induced by the scheduler $\mathfrak{s}$ on the CFG obtained from $\mathcal{G}$ by setting the initial state to $(l, \mathbf{x})$.

For a scheduler $\mathfrak{s}$, the *canonical filtration* of $\mathcal{G}$ is the sequence $(\mathcal{F}_n)_{n \in \mathbb{N}}$ such that $\mathcal{F}_n$ is the smallest sub- σ -algebra of $\mathcal{F}_{\mathcal{G}(\sigma)}$ that contains the cylinder sets $\text{Cyl}_{\mathcal{G}(\sigma)}(\pi_{\leq n})$ of all finite paths $\pi_{\leq n}$ of length at most n . Under this filtration, the semantics of $\mathcal{G}$ under $\mathfrak{s}$ can also be viewed as a stochastic process $(X_n^{\mathfrak{s}})_{n \in \mathbb{N}}$ measurable against $(\mathcal{F}_n)_{n \in \mathbb{N}}$ such that $X_n^{\mathfrak{s}}$ takes on an encoding of the state of the execution after n steps. When convenient, we will use this view as well.

2.2 The Termination Problem

Fix an CFG $\mathcal{G}$ and a scheduler $\mathfrak{s}$. Let l_{out} be a distinguished location we refer to as *terminal*. Denote by $\diamond(l_{\text{out}}, \mathbf{0})$ the set of all runs that reach $(l_{\text{out}}, \mathbf{0})$; we call these the *terminating runs*. Observe that $\diamond(l_{\text{out}}, \mathbf{0})$ is measurable. The CFG $\mathcal{G}$ is said to *terminate* with probability p under the scheduler $\mathfrak{s}$ if $\mathbb{P}_{\mathfrak{s}}[\diamond(l_{\text{out}}, \mathbf{0})] = p$.

```

1  x, y := 1, 1
2  while (x ≠ 0 ∨ y ≠ 0):
3      { x := x + 1 ⊕1/2 x := x - 1 } ⊕1/2
4      { y := y + 1 ⊕1/2 y := y - 1 }

```

Prg. 1. The 2D symmetric random walker. The symbol $\oplus$ is a probabilistic choice operator.

Definition 2.4 (Termination Probability). Let $\mathcal{G}$ be an CFG and for a scheduler $\mathfrak{s}$, let $(\text{Runs}_{\mathcal{G}}, \mathcal{F}_{\mathcal{G}}, \mathbb{P}_{\mathfrak{s}})$ be the probability space induced by $\mathfrak{s}$ on the executions of $\mathcal{G}$. The termination probability of $\mathcal{G}$, denoted by $\text{Pr}_{\text{term}}(\mathcal{G})$, is the infimum of $\mathbb{P}_{\mathfrak{s}}[\Diamond(l_{\text{out}}, \mathbf{0})]$ over all schedulers $\mathfrak{s}$.

We also use the notation $\text{Pr}_{\text{term}}(\mathcal{G}(\sigma))$ to refer to the termination probability of $\mathcal{G}$ if its initial state were changed to σ .

An CFG is said to be *almost surely terminating* (AST) if its termination probability is 1. Our work is on sound and complete proof rules for deciding, for an CFG $\mathcal{G}$, (a) the AST problem, i.e., whether $\mathcal{G}$ is almost surely terminating, (b) the *Lower Bound* problem, i.e., whether $\text{Pr}_{\text{term}}(\mathcal{G})$ exceeds some $p < 1$, (c) the *Upper Bound* problem, i.e., whether $\text{Pr}_{\text{term}}(\mathcal{G})$ is bounded above by some $p > 0$. We remark that, for the lower and upper bound problems, the proof rules we describe are applicable to any number p that is representable in our program logic. This means that p can take on irrational and transcendental values; this is important, as termination probabilities can often take on such values [6, 21]. We will elaborate in Section 2.4.

Note that, while we define termination for a specific state $(l_{\text{out}}, \mathbf{0})$, more general termination conditions can be reduced to this case by a syntactic modification.

Example 2.5 (Symmetric Random Walk). A d -dimensional symmetric random walk has d integer variables $x_1, \dots, x_d$. Initially, all variables are 1. In each step, the program updates the variables to move to a “nearest neighbor” in the d -dimensional lattice $\mathbb{Q}^d$; that is, the program picks uniformly at random one of the variables and an element in $\{-1, +1\}$, and adds the element to the chosen variable. Program 1 shows the code for $d = 2$. It is well known [40] that the symmetric random walk is recurrent in dimension 1 and 2, and transient otherwise. Thus, if we set any element in the lattice, say 0, to be an terminal state, then the program is almost surely terminating in dimension 1 and 2 but not almost surely terminating when $d \geq 3$. We shall refer to the $d = 1$ and $d = 2$ cases as 1DRW and 2DRW, respectively. $\square$

2.3 The Unrolling Lemma

Let $\mathcal{G}$ be an CFG such that $\text{Pr}_{\text{term}}(\mathcal{G}) \geq p$ for some rational $p > 0$. Fix a scheduler $\mathfrak{s}$. Let $(\pi_1, \pi_2, \dots)$ be an ordering of the terminating runs of $\mathcal{G}$ consistent with $\mathfrak{s}$ such that $|\pi_1| \leq |\pi_2| \leq \dots$. For some $\epsilon > 0$, let i_n be the smallest number such that

$$\mathbb{P}_{\mathfrak{s}}(\pi_1) + \dots + \mathbb{P}_{\mathfrak{s}}(\pi_{i_n}) \geq p - \epsilon$$

where $\mathbb{P}_{\mathfrak{s}}$ is the probability measure induced by $\mathfrak{s}$ over the set of all runs of $\mathcal{G}$.

We call $|\pi_n|$ the *required simulation time* of $\mathcal{G}$ under $\mathfrak{s}$ to assimilate a termination probability of $p - \epsilon$. The required simulation time of $\mathfrak{s}$ is simply the length of the longest terminating run that must be accounted for in the termination probability series for it to cross $p - \epsilon$.

Define the simulation time of $\mathcal{G}$ w.r.t. ϵ as the supremum over all schedulers $\mathfrak{s}$ of the required simulation time of $\mathcal{G}$ under $\mathfrak{s}$ and ϵ . The following lemma is at the core of showing that the almost sure termination problem is Π_2^0 -complete [33].

LEMMA 2.6 (UNROLLING LEMMA [33]). *Let $\mathcal{G}$ be an CFG such that $\Pr_{\text{term}}(\mathcal{G}) \geq p$. For any ϵ , the simulation time of $\mathcal{G}$ w.r.t. ϵ is bounded above.*

Lemma 2.6 is a generalization of Lemma B.3 of Majumdar and Sathiyararayanan [33]. It holds for CFGs because the branching at nondeterministic locations is bounded. To prove the unrolling lemma, for each $m \in \mathbb{N}$, we consider unrollings of $\mathcal{G}$ for m steps, running under partial schedules that resolve nondeterministic choices for up to m steps. Partial schedules are naturally ordered into a tree, where a partial schedule $\mathfrak{s}$ is extended by $\mathfrak{s}'$ if $\mathfrak{s}'$ agrees with $\mathfrak{s}$ when restricted to the domain of $\mathfrak{s}$. An infinite path in this tree defines a scheduler. For each scheduler $\mathfrak{s}$, we mark the k -th node in its path if k is the minimum number such that the k -step unrolled program amasses termination probability at least $p - \epsilon$. If two schedulers agree up to k steps, then they both mark the same node. The key observation is that, since the nondeterminism is finite-branching, the scheduler tree is finite-branching. Thus, if we cut off the tree at marked nodes and still have an infinite number of incomparable marked nodes, there must be an infinite path in the tree that is not marked. But this is a contradiction, because this infinite path corresponds to a scheduler that never amasses $p - \epsilon$ probability mass for termination.

Corollary 2.7 follows directly, and is used multiple times in the proofs of the soundness and completeness of our rules.

COROLLARY 2.7. *Let σ be a state of an CFG $\mathcal{G}$. Suppose $\Pr_{\text{term}}(\mathcal{G}(\sigma)) > 0$. Then, varied across schedulers, there is an upper bound on the length of the shortest consistent terminal run from σ .*

PROOF. For a scheduler $\mathfrak{s}$, let $\pi_{\mathfrak{s}}$ be the smallest terminal run of $\mathcal{G}(\sigma)$ consistent with $\mathfrak{s}$. The simulation time to assimilate a termination probability of ϵ for some $0 < \epsilon < \Pr_{\text{term}}(\mathcal{G}(\sigma))$ under scheduler $\mathfrak{s}$ is necessarily at least as large as $|\pi_{\mathfrak{s}}|$. If the collection of lengths $|\pi_{\mathfrak{s}}|$ across schedulers $\mathfrak{s}$ wasn't bounded, then this simulation time is unbounded. This contradicts the unrolling lemma. $\square$

2.4 Assertion Language and Program Logic

Our language of choice for specifying assertions is the language of arithmetic with addition, multiplication, and order interpreted over the domain of rationals. We fix the interpretation model for our assertions as the standard model of rationals. Refer to this interpretation by $\mathbb{I}_{\mathbb{Q}}$.² Let $\text{Th}(\mathbb{Q})$ denote the *theory of rationals*, i.e., the collection of assertions that are true in the standard model of rationals. The evaluation of our assertions is tantamount to their implication by $\text{Th}(\mathbb{Q})$. All proof techniques we present in our work are relative to complete proof systems for $\text{Th}(\mathbb{Q})$.

Assertions are evaluated at program states. Fix a CFG $\mathcal{G}$ with transition relation $\mapsto_{\mathcal{G}}$. A state σ of $\mathcal{G}$ *satisfies* an assertion φ if the interpretation $\mathbb{I}_{\mathbb{Q}}$ augmented with the variable valuation encoded in σ models φ . We denote this by $\sigma \models \varphi$. An assertion φ is *valid* if $\sigma \models \varphi$ for all states σ . Valid assertions are contained in $\text{Th}(\mathbb{Q})$.

We employ a program logic inspired by the seminal work of Floyd [22]. Statements in our logic affix assertions as preconditions and postconditions to transitions in $\mathcal{G}$. For example, the transition $\tau \in \mapsto_{\mathcal{G}}$ could be affixed a precondition φ_{τ} and postcondition ψ_{τ} to yield the sentence $\{\varphi_{\tau}\}\tau\{\psi_{\tau}\}$. The precondition φ_{τ} is evaluated at the program state before taking τ , and the postcondition ψ_{τ} is evaluated at states reached immediately after τ . The sentence $\{\varphi_{\tau}\}\tau\{\psi_{\tau}\}$ is *true* for $\mathcal{G}$ if for every state σ with $\sigma \models \varphi$, if τ is enabled at σ and σ' is a successor of σ through τ , then $\sigma' \models \psi_{\tau}$.

We use the notion of *inductive invariants* in our proof rules. An *inductive invariant* is an assertion with $n + 1$ free variables, the first ranging over L and the others over $\mathbb{Q}$, that is closed under the successor operation. That is, an assertion Inv is an inductive invariant if, whenever $(l, \mathbf{x})$ satisfies

²Instead of fixing $\mathbb{I}_{\mathbb{Q}}$, one can use any *arithmetical structure* [27] to specify and interpret assertions. All our proof rules will remain sound and relatively complete with this change.

Inv, and (l', x') is a successor to (l, x) , then (l', x') satisfies Inv. It follows that if (l_{init}, x_{init}) satisfies Inv, then every reachable state satisfies Inv.

Floyd [22] specified axioms for a proof system over this program logic. *Proof rules* extend this system by enabling the deduction of complicated program properties, such as termination. These rules are composed of *antecedents* and *consequents*. Antecedents are finite collections of statements written in the program logic. Consequents detail properties of the program at which the antecedents are evaluated. *Soundness* of a proof rule implies that if the antecedents are true for a program $\mathcal{G}$, then the consequents hold for $\mathcal{G}$. *Completeness* of a proof rule implies that if the consequents are true for some $\mathcal{G}$, then one can come up with proofs for the antecedents of the rule in the underlying proof system.

The completeness of all proof rules in this work is dependent on the existence of a complete proof system for $\text{Th}(\mathbb{Q})$. Such proof rules are said to be complete *relative to a proof system for* $\text{Th}(\mathbb{Q})$. Relative completeness of this kind is standard in program logics.

To show the relative completeness of our proof rules, we will need to be able to encode *computable relations* in our assertion language. A relation is computable if its characteristic function is decidable. It is known that the theory of arithmetic interpreted over natural numbers can encode all computable relations. Let $\mathbb{I}_{\mathbb{N}}$ refer to the interpretation model of the standard model of naturals. Denote by $\text{Th}(\mathbb{N})$ the collection of all true assertions under $\mathbb{I}_{\mathbb{N}}$. $\text{Th}(\mathbb{N})$ is generally referred to as the theory of natural numbers. Thus, for each computable relation $R(x_1, x_2, \dots, x_n)$, there is an assertion $\varphi_R(x_1, x_2, \dots, x_n)$ that is true in $\text{Th}(\mathbb{N})$. To represent computable relations in $\text{Th}(\mathbb{Q})$, we use a result by Robinson [42].

THEOREM 2.8 (ROBINSON [42]). $\mathbb{N}$ is definable in $\text{Th}(\mathbb{Q})$.

We refer to the assertion that encodes $\mathbb{N}$ by Nat . Therefore, $\text{Nat}(x)$ is true in $\mathbb{I}_{\mathbb{Q}}$ iff $x \in \mathbb{N}$. All computable relations can be encoded in our assertion language through liberal usage of Nat . An important implication is that termination probabilities are expressible in our assertion language.

LEMMA 2.9. *For a CFG $\mathcal{G}$ and a $p \in [0, 1]$ with $\text{Pr}_{\text{term}}(\mathcal{G}) = p$, there is an assertion $\psi(x)$ with one free variable x such that $\text{Th}(\mathbb{Q}) \models \psi(x) \Leftrightarrow x \leq p$.*

PROOF. We know that $\text{Pr}_{\text{term}}(\mathcal{G}) \geq p$ iff $\text{Pr}_{\text{term}}(\mathcal{G}) \geq p - \epsilon$ for all $\epsilon > 0$. The unrolling lemma implies that for all $\epsilon > 0$, there is a $k \in \mathbb{N}$ such that the probability mass of the k -unrolled program is at least $p - \epsilon$. Finite unrollings of $\mathcal{G}$ are, by definition, computable, and checking the probability of termination amassed in this finite unrolling is also computable; see Kaminski et al. [31] for details. This means that a relation $R(\epsilon, k)$ representing this relationship between every rational ϵ and natural k . Such computable relations are representable in $\text{Th}(\mathbb{Q})$ through [Theorem 2.8](#), completing the proof. $\square$

Notice that while the termination probabilities p are real numbers, the lower bounds verified by the assertion ψ in the above lemma are entirely rational. However, by representing the set of rational numbers under p , ψ has effectively captured the Dedekind cut of p . This expressibility shows how irrational lower bounds on termination probabilities can be deduced using our proof techniques.

3 ALMOST-SURE TERMINATION: MARTINGALES

In all rules in this work, we fix a CFG $\mathcal{G} = (L, V, l_{init}, x_{init}, \mapsto, G, \text{Pr}, \text{Upd})$. We also abuse notation slightly and use the inductive invariant Inv as a shorthand for all states of $\mathcal{G}$ satisfying the predicate Inv.

Our proof rules consist of sets and functions over the state spaces that satisfy certain properties. Each of these entities must be representable in our assertion language; therefore, they must be

arithmetical expressions over the program variables and program locations. Instead of specifying each condition in our rules as formal statements in our program logic, we directly describe the properties these entities must satisfy. We do so to emphasize these entities themselves over the formalism surrounding them. It is nevertheless possible to write each of the following proof rules as finite sets of statements in the program logic.

Recall that a proof rule is *sound* if, whenever we can find arithmetical expressions in our assertion language that satisfy the conditions outlined in the premise of a rule, the conclusion of the rule holds. A proof rule is *relatively complete* if, whenever the conclusion holds (e.g., a program $\mathcal{G}$ is AST), we can find certificates in the assertion language that satisfy all the premises.

3.1 McIver and Morgan's Variant Rule

We start with a well-known rule for almost-sure termination from [36]. The rule is sound but complete only for finite-state programs McIver and Morgan [36, Lemma 7.6.1].

Proof Rule 3.1: Variant Rule for AST [36]

If there is

- (1) an inductive invariant Inv containing the initial state $(l_{\text{init}}, \mathbf{x}_{\text{init}})$,
- (2) a function $U : \text{Inv} \rightarrow \mathbb{Z}$,
- (3) bounds Lo and Hi such that for all states $(l, \mathbf{x}) \in \text{Inv}$, $\text{Lo} \leq U(l, \mathbf{x}) < \text{Hi}$, and
- (4) an $\epsilon > 0$,

such that, for each state $(l, \mathbf{x}) \in \text{Inv}$,

- (a) if $(l, \mathbf{x})$ is a terminal state, $U(l, \mathbf{x}) = 0$.
- (b) if $(l, \mathbf{x})$ is an assignment, or nondeterministic state, $U(l', \mathbf{x}') < U(l, \mathbf{x})$ for every successor $(l', \mathbf{x}')$.
- (c) if $(l, \mathbf{x})$ is a probabilistic state, $\sum \text{Pr}(l, l')[\mathbf{x}] > \epsilon$ over all successor states $(l', \mathbf{x}')$ with $U(l', \mathbf{x}') < U(l, \mathbf{x})$.

Then, $\mathcal{G}$ is AST.

LEMMA 3.1 (MCIVER AND MORGAN [36]). *Rule 3.1 is sound for all AST programs. It is relatively complete for finite-state AST CFGs.*

While McIver and Morgan [36] claim completeness and not relative completeness, their proof trivially induces relative completeness. This rule is not complete, however. This is because, if the rule is applicable, the program is guaranteed a terminal run of length at most $\text{Hi} - \text{Lo}$ from any state. But the 1D random walk (outlined in Example 2.5) does not satisfy this property, even though it terminates almost-surely.

Over the years, McIver and Morgan's proof rule has been extended many times [30, 37]. The most significant extension is the proof rule of McIver et al. [37], where they require the function U to additionally be a supermartingale. None of these extensions have managed to be proven complete. Because we do not use ideas from these extensions, we do not present them here.

3.2 Our Rule

We present a martingale-based proof rule for AST that exploits the fact that AST programs, when repeatedly run, are recurrent.

Proof Rule 3.2: Martingale Rule for AST

If there exists

- (1) an inductive invariant Inv containing the initial state,
- (2) a set $A \subset \text{Inv}$ containing the terminal state,
- (3) a supermartingale function $V : \text{Inv} \rightarrow \mathbb{R}$ that assigns 0 to the terminal state and at all states $(l, \mathbf{x}) \in \text{Inv} \setminus A$,
 - (a) $V(l, \mathbf{x}) > 0$,
 - (b) $V(l, \mathbf{x}) > V(l_A, \mathbf{x}_A)$ for each $(l_A, \mathbf{x}_A) \in A$,
 - (c) if $(l, \mathbf{x})$ is an assignment or nondeterministic state, then $V(l, \mathbf{x}) \geq V(l', \mathbf{x}')$ for all possible successor states $(l', \mathbf{x}')$, and
 - (d) if $(l, \mathbf{x})$ is a probabilistic state, then, over all successor states $(l', \mathbf{x}')$, $V(\sigma) \geq \sum \text{Pr}(l, l') V(l', \mathbf{x}')$,
- (4) a variant function $U : \text{Inv} \rightarrow \mathbb{N}$ that
 - (a) assigns 0 to the terminal state,
 - (b) ensures that at nondeterministic and assignment states $(l, \mathbf{x}) \in \text{Inv}$, $U(l, \mathbf{x}) > U(l', \mathbf{x}')$ for all possible successor states $(l', \mathbf{x}')$, and
 - (c) satisfies the following compatibility criteria with the sublevel sets $V_{\leq r} = \{\sigma \in \text{Inv} \mid V(\sigma) \leq r\}$ for each $r \in \mathbb{R}$:
 - (i) the set $\{u \in \mathbb{N} \mid \sigma \in V_{\leq r} \wedge u = U(\sigma)\}$ is bounded, and
 - (ii) there exists an $\epsilon_r > 0$ such that, for all probabilistic states $(l, \mathbf{x}) \in V_{\leq r}$, the sum $\sum \text{Pr}(l, l') [\mathbf{x}] > \epsilon_r$ over all successor states $(l', \mathbf{x}')$ with $U(l', \mathbf{x}') < U(l, \mathbf{x})$.

Under these conditions, $\mathcal{G}$ is AST.

In this rule, U is meant to play the role of the variant function from [Rule 3.1](#). A is meant to form a “ball” around the terminal state; it is useful in applications where the supermartingale properties of V are difficult to establish at all states. If the execution were to be restricted within this ball A , the rule makes it easy to establish almost-sure termination. This is because while V must only be a supermartingale outside of A , the variant U must still decrease within A . Observe that A must be a strict subset of Inv ; this is to enforce an upper bound on the collection of V -values of states in A . It’s easy to see that this rule reduces to [Rule 3.1](#) if the supermartingale V was bounded.

Intuitively, V can be thought of as a measure of *relative likelihood*. The probability that a transition increases V by an amount v reduces as v increases. Unlikely transitions are associated with greater increments to V , and (relatively) unlikely states have greater V values. Separately, the variant U reprises its role from [Rule 3.1](#): it effectively measures the shortest distance to a terminal state.

At a high level, [Rule 3.2](#) works for the following reasons. Suppose V is unbounded and executions begin at some initial state $\sigma_0 \in \text{Inv}$. The supermartingale property of V implies that from σ_0 , the probability of reaching a state σ with $V(\sigma) > V(\sigma_0)$ approaches 0 as $V(\sigma)$ grows to $+\infty$. Now, fix an unlikely state σ with $V(\sigma) \gg V(\sigma_0)$. Let’s now restrict our attention to the executions that remain in states $\gamma \in \text{Inv}$ with $V(\gamma) \leq V(\sigma)$. The compatibility conditions satisfied by the variant U with V at the sublevel set $V_{\leq V(\sigma)}$ implies the almost-sure termination of these executions. The remaining executions must reach some unlikely state γ' with $V(\gamma') \geq V(\sigma)$.

Thus, as the probability of reaching unlikely states γ' reduces the “further away” (from the perspective of V) they are, the probability of terminating approaches 1. Since V is unbounded, the probability of termination is 1.

Remark. We note that our rule is quite similar to the AST rule of McIver et al. [37]. Their rule consisted of a single supermartingale V that, with the help of a few antitone functions, also exhibited the properties of a distance variant. In other words, they combined the duties of the functions V and U into a single function V . It is not known if their rule is complete.

LEMMA 3.2 (SOUNDNESS). *Rule 3.2 is sound.*

PROOF. Let us first dispense of the case where V is bounded. If V is bounded, the compatibility criteria forces a bound on the variant function U . The soundness of Rule 3.1 implies $\mathcal{G} \in \text{AST}$. Therefore, from now on, V is assumed to be unbounded.

Denote the initial state by σ_0 . For each $n \in \mathbb{N}$, define Π_n to be the collection of runs from σ_0 that reach a maximum V value of n . This means that for each state σ encountered in executions in Π_n , $V(\sigma) \leq n$. Define Π_∞ to be the remaining collection of executions beginning at σ_0 that don't have a bound on the V values that they reach. This means that for each execution $\pi \in \Pi_\infty$ and each $n \in \mathbb{N}$, there are states $\sigma \in \pi$ such that $V(\sigma) > n$. We have thus partitioned the collection of executions of the CFG $\mathcal{G}$ to $\Pi_\infty \cup (\bigcup_{i \in \mathbb{N}} \Pi_i)$.

We will now argue that under every scheduler, the probability measure of all non-terminating executions in each Π_n is 0. By definition, all executions in Π_n lie entirely within the sublevel set $V_{\leq n}$. The compatibility of U with $V_{\leq n}$ implies that the variant U is bounded across states in Π_n . Consider an CFG $\mathcal{G}_{\leq n}$ that mirrors $\mathcal{G}$ inside $V_{\leq n}$, but marks states in $\mathcal{G}$ outside $V_{\leq n}$ as terminal. Applying Rule 3.1 using the now bounded variant U allows us to deduce that $\mathcal{G}_{\leq n}$ is almost-surely terminating. Observe now that the collection of non-terminating runs of $\mathcal{G}_{\leq n}$ is precisely the collection of non-terminating runs in Π_n . This immediately gives us what we need.

We now turn our attention to the final collection Π_∞ . Observe that Π_∞ must only contain non-terminal executions. Suppose that, under some scheduler $\mathfrak{s}$, the probability measure of Π_∞ wasn't 0. Let the probability space defining the semantics of $\mathcal{G}$ under $\mathfrak{s}$ (see Section 2.1) be $(\text{Runs}_{\mathcal{G}(\sigma)}, \mathcal{F}_{\mathcal{G}(\sigma)}, \mathbb{P}_{\mathfrak{s}})$, and let its canonical filtration be $\{\mathcal{F}_n\}$. Define a stochastic process $\{X_n^{\mathfrak{s}}\}$ over the aforementioned probability space augmented with the filtration $\{\mathcal{F}_n\}$ that tracks the current state of the execution of the program. Define another stochastic process $\{Y_n^{\mathfrak{s}}\}$ as

$$Y_n^{\mathfrak{s}} \triangleq \begin{cases} V(X_n^{\mathfrak{s}}) & X_n^{\mathfrak{s}} \notin A \\ 0 & \text{otherwise} \end{cases}$$

It's easy to see that $Y_n^{\mathfrak{s}}$ is a non-negative supermartingale. Since $Y_n^{\mathfrak{s}}$ is non-negative, Doob's Martingale Convergence Theorem [18] implies the almost-sure existence of a random variable $Y_\infty^{\mathfrak{s}}$ that the process $\{Y_n^{\mathfrak{s}}\}$ converges to. This means that $\mathbb{E}[Y_\infty^{\mathfrak{s}}] \leq Y_0^{\mathfrak{s}}$.

Under the condition that Π_∞ occurs, $Y_\infty^{\mathfrak{s}} = +\infty$. Since the probability measure of these non-terminal executions isn't 0, we have that $\mathbb{E}[Y_\infty^{\mathfrak{s}}] = +\infty > Y_0^{\mathfrak{s}} = V(X_0^{\mathfrak{s}})$. This raises a contradiction, completing the proof. $\square$

To show completeness, we adapt a technique by Mertens et al. [39] to build the requisite supermartingale V . Suppose $\mathcal{G}$ is AST. Let $\text{Reach}(\mathcal{G})$ be the set of its reachable states. Fix a computable enumeration Enum of $\text{Reach}(\mathcal{G})$ that assigns 0 to its terminal state. Intuitively, Enum is meant to order states in a line so that the probability of reaching a state that's far to the right in this order is small. This is because the AST nature of $\mathcal{G}$ forces executions to "lean left" toward the terminal state. Note that we place no other requirements on Enum ; these intuitions will work no matter how Enum orders the states. A state σ is said to be indexed i if $\text{Enum}(\sigma) = i$. From now on, we will refer to the state indexed i by σ_i .

A crucial part of our construction is the following function $R : (\mathbb{N} \times \mathbb{N}) \rightarrow [0, 1]$. Intuitively, R measures the ability of executions beginning from a state to reach states that are far to the right

of it in the Enum order. Let $\mathcal{G}_i$ be the CFG obtained from $\mathcal{G}$ by switching its initial state to σ_i . Let the semantics of $\mathcal{G}_i$ under a scheduler $\mathfrak{s}$ be the probability space $(\text{Runs}_{\mathcal{G}_i}, \mathcal{F}_{\mathcal{G}_i}, \mathbb{P}_{\mathfrak{s}}^i)$. Define $R(i, n)$ at indices i and n to be

$$R(i, n) \triangleq \inf_{\mathfrak{s}} \mathbb{P}_{\mathfrak{s}}^i (\diamond (\{\sigma_m \in \text{Reach}(\mathcal{G}) \mid m \geq n\})) \quad (1)$$

Where $\diamond(C)$ represents the event of eventually reaching the set C . We will refer to the first argument i as the source index and the second argument n as the minimum target index. Put simply, $R(i, n)$ measures the infimum probability of reaching the target indices $\{n, n+1, \dots\}$ from the source σ_i .

LEMMA 3.3. $R(i, n) \rightarrow 0$ as $n \rightarrow \infty$ at every $i \in \mathbb{N}$, i.e., $\forall i \in \mathbb{N} \cdot \lim_{n \rightarrow \infty} R(i, n) = 0$.

PROOF. Denote by E_n the event that executions beginning from σ_i reach states with index $\geq n$. Clearly, $R(i, n)$ measures the infimum probability of E_n . Denote by E_∞ the event that executions beginning from σ_i increase the maximum observed state index infinitely often. It's easy to see that, for every $n \in \mathbb{N}$, the event E_n contains E_∞ . Also, each execution outside E_∞ must be inside some $E_n \setminus E_{n+1}$, as it must yield a maximum state index contained in it. Additionally, $E_{n+1} \subseteq E_n$ for all n . These three facts imply

$$E_\infty = \bigcap_{i \in \mathbb{N}} E_n = \lim_{n \rightarrow \infty} E_n$$

Suppose that, $\lim_{n \rightarrow \infty} R(i, n) > 0$ for some index i . Since $R(i, n) = \inf_{\mathfrak{s}} \mathbb{P}_{\mathfrak{s}}^i [E_n]$, we have

$$\inf_{\mathfrak{s}} \mathbb{P}_{\mathfrak{s}}^i [E_\infty] = \inf_{\mathfrak{s}} \left(\lim_{n \rightarrow \infty} \mathbb{P}_{\mathfrak{s}}^i [E_n] \right) = \lim_{n \rightarrow \infty} \inf_{\mathfrak{s}} \mathbb{P}_{\mathfrak{s}}^i [E_n] = \lim_{n \rightarrow \infty} R(i, n) > 0$$

As all executions in E_∞ are non-terminating, this contradicts the AST nature of $\mathcal{G}$. $\square$

It turns out that, if we fix the minimum target index n , the function R becomes a supermartingale. Define $V_n(\sigma) = R(\text{Enum}(\sigma), n)$ for every $n \in \mathbb{N}$. It's easy enough to see that V_n is a supermartingale; for assignment / non-deterministic states σ with possible successors σ' , $V_n(\sigma) \geq V_n(\sigma')$, and for probabilistic $\sigma = (l, \mathbf{x})$, $V_n(l, \mathbf{x}) \geq \sum \Pr(l, l') V_n(l', \mathbf{x})$ across all successors $(l', \mathbf{x})$. However, V_n isn't the supermartingale we need, as we may not always be able to construct a compatible U for any V_n . This is because every V_n is bounded above (by 1), whereas U typically isn't bounded above.

To construct an unbounded supermartingale, one could consider the sum $\sum V_n$ varied across all $n \in \mathbb{N}$. However, this sum could be ∞ for certain states. To combat this, we carefully choose an infinite subset of $\mathbb{N}$ to form the domain for $\sum V_n$. Consider the sequence $(n_j)_{j \in \mathbb{N}}$ such that n_j is the smallest number so that $R(i, n_j) \leq 2^{-j}$ for all $i \leq j$. Each element in this sequence is certain to exist due to the monotonically non-increasing nature of $R(i, n)$ for fixed i and the limit result of Lemma 3.3. Furthermore, restricting the domain of $\sum V_n$ to elements in $(n_j)_{j \in \mathbb{N}}$ will mean that no state is assigned ∞ by the sum. This is because for each σ , the values of $V_{n_j}(\sigma) = R(\text{Enum}(\sigma), n_j)$ will certainly repeatedly halve after $j \geq \text{Enum}(\sigma)$. Further note that the supermartingale nature of the V_n implies that this sum is also a supermartingale. We thus have our required supermartingale

$$V(\sigma) = \sum_{j \in \mathbb{N}} V_{n_j}(\sigma) = \sum_{j \in \mathbb{N}} R(\text{Enum}(\sigma), n_j) \quad (2)$$

LEMMA 3.4 (COMPLETENESS). *Rule 3.2 is relatively complete.*

PROOF. Take an AST CFG $\mathcal{G}$, and set Inv and A to $\text{Reach}(\mathcal{G})$ and the singleton set containing the terminal state respectively. We first describe our choice for the variant function U . Since $\mathcal{G}$ is AST, for every $\sigma \in \text{Reach}(\mathcal{G})$, every scheduler must induce a finite path to a terminal state. Corollary 2.7 implies an upper bound on the length of the shortest terminal run from every $\sigma \in \text{Inv}$. Set U to map each $\sigma \in \text{Inv}$ to this upper bound.

If U is bounded, setting $V(\sigma) = 1$ for all $\sigma \in \text{Inv}$ suffices. Otherwise, set V to the supermartingale function defined in Eq. (2). It is easy to observe that for every r , the sublevel set $V_{\leq r} = \{\sigma \mid V(\sigma) \leq r\}$ is finite. This implies that U is bounded within every sublevel set, and is hence compatible with this V . This completes the construction of the certificates required by the proof rule.

We now argue that the invariant Inv , the set A , the supermartingale V and variant U can each be represented in our assertion language of arithmetic interpreted over the rationals. We do this by encoding them first encoding them in the theory of natural numbers, and then using the relation Nat from Theorem 2.8 to insert them into our assertion language. Recall that all computable relations can be encoded in $\text{Th}(\mathbb{Q})$. We present techniques with which one can augment computable relations with first-order quantifiers to represent these entities. By doing so, we demonstrate that these sets are *arithmetical*; see the works of Kozen [32] and Rogers Jr. [43] for detailed accounts on arithmetical sets.

A can trivially be represented in $\text{Th}(\mathbb{Q})$. For Inv , consider the relation I that contains tuples of the form (k, σ_1, σ_2) where $k \in \mathbb{N}$ and σ_1 and σ_2 are states of $\mathcal{G}$. Require $(k, \sigma_1, \sigma_2) \in I$ iff there is a finite path of length $\leq k$ from σ_1 to σ_2 . Clearly, I is a computable relation and is thus representable in $\text{Th}(\mathbb{Q})$. $\text{Inv}(\sigma)$ can be represented from I as $\exists k \cdot I(k, \sigma_0, \sigma)$ where σ_0 is the initial state of $\mathcal{G}$. Similarly, the output of $U(\sigma)$ at every σ can be represented using I as $U(\sigma) = k \iff I(\sigma, \sigma_{\perp}, k) \wedge (\forall n < k \cdot \neg I(\sigma, \sigma_{\perp}, n))$, where $\sigma_{\perp}$ is the terminal state. If U were bounded, representing V is trivial; we focus our attention on representing V when U isn't bounded.

Representations of R (defined in Eq. (1) and used to derive V) and V are complicated slightly because they can output real numbers. Instead of capturing the precise values of these functions, we capture their Dedekind cuts instead. In other words, we show that the collections of rational numbers $\leq V(\sigma)$, $\geq V(\sigma)$, $\leq \varphi_i(n)$ and $\geq \varphi_i(n)$ are each representable for each σ , i , and n .

The unrolling lemma implies that if the probability of termination is p , then for all $n \in \mathbb{N}$, assimilating a termination probability mass of at least $p - 1/n$ requires finitely many steps. It is simple to generalize this to observe that assimilating a probability mass of at least $R(i, n) - 1/n$ for the event $\diamond (\{\sigma_m \in \text{Inv} \mid m \geq n\})$ when σ_i is the initial state also requires finitely many steps. Furthermore, computing the probability of the occurrence of this event within k steps is computable for every natural number k . These two facts indicate that lower bounds on $R(i, n)$ can be represented in $\text{Th}(\mathbb{Q})$. Upper bounds on $R(i, n)$ can be represented by simply negating this lower bound representation.

Using these, enable the representation of each member of the sequence $(n_j)_{j \in \mathbb{N}}$ that forms the domain of the sum that defines V . This enables representations of lower bounds on $V(\sigma)$, which in turn enables representations of upper bounds on $V(\sigma)$. Thus, the Dedekind cut of $V(\sigma)$ is representable in $\text{Th}(\mathbb{Q})$. This completes the proof. $\square$

Example 3.5 (Random Walks). For the 1DRW example from Example 2.5, we take Inv to be all program states, A to be the set containing the single terminal state $\{x_1 := 0\}$, and we set $V(x) = U(x) := |x|$. It's trivial to observe that all conditions required in Rule 3.2 are met, and therefore, the 1DRW is AST.

Let us now consider the 2-D Random Walker (2DRW). The AST nature of this program has been notoriously hard to prove using prior rules. The principal enabler of our rule on the 2DRW is its set A that forms a circle around the origin.

Begin by setting Inv to the set of all states. Declare the distance variant U to be the Manhattan distance $|x| + |y|$ of any state (x, y) from the origin. Clearly, U has a $\geq 1/4$ probability of reducing in a single step from every state. Now, define V as

$$V(x, y) \triangleq \sqrt{\ln \left(1 + \sqrt{x^2 + y^2} \right)}$$

It is difficult to prove that V is a supermartingale for all non-terminal states. However, using Taylor series expansions, one can show that V is a supermartingale for “sufficiently large” values of $x^2 + y^2$. Menshikov et al. [38] (see also Popov [41, Section 2.3]) showed precisely this; their proof showed that the error terms in the Taylor series expansion of $V(x, y)$ cease to matter when $x^2 + y^2$ grow large. They prove that there must exist a number k such that the error terms in the Taylor expansion do not affect the non-negativity conditions at states where $x^2 + y^2 > k$. We now declare A to be the set of states (x, y) where $x^2 + y^2 \leq k$. Thus, V satisfies the supermartingale conditions of Rule 3.2 outside of A . Note that we don’t need to precisely decipher the value of k for the soundness of the proof rule to work; we just need A to be smaller than the invariant Inv . The finiteness of A satisfies this criterion. Furthermore, since the sublevel set $V_{\leq r}$ is finite, V is compatible with U . Therefore, the 2DRW is AST. $\square$

4 ALMOST-SURE TERMINATION: STOCHASTIC INVARIANTS

Next, we give a different proof rule that takes a dual view. Instead of a single unbounded supermartingale, we consider several bounded supermartingales that each focus on different finite parts of the program’s state space. This focus means that each of these supermartingales takes on non-trivial values (i.e., between 0 and 1) at only finitely many states. They are thus “local” supermartingales, i.e., they are only meaningful at particular parts of the state space.

The key to our rule is the observation that a program is almost surely terminating if, for every $\epsilon > 0$, we can show that it terminates with probability at least ϵ . We characterize this “ ϵ -wiggle room” using the stochastic invariants of Chatterjee et al. [11].

Definition 4.1 (Stochastic Invariants [11]). Let $\mathcal{G} = (L, V, l_{\text{init}}, \mathbf{x}_{\text{init}}, \mapsto, G, \text{Pr}, \text{Upd})$ be a CFG. Suppose Ψ is a subset of states and let p be a probability value. The tuple (Ψ, p) is a *stochastic invariant* (SI) if, under any scheduler $\mathfrak{s}$, the probability mass of the collection of runs beginning from $(l_{\text{init}}, \mathbf{x}_{\text{init}})$ leaving Ψ is bounded above by p , i.e.,

$$\sup_{\mathfrak{s}} \mathbb{P}_{\mathfrak{s}} [\rho \in \text{Runs}_{\mathcal{G}} \mid \exists n \in \mathbb{N} \cdot \rho[n] \notin \Psi] \leq p$$

Intuitively, stochastic invariants generalize the standard notion of invariants to the probabilistic setting. Given a stochastic invariant (Ψ, p) , the program execution is expected to hold Ψ (i.e., remain inside Ψ) with probability $\geq 1 - p$. As with invariants, the collection of states in stochastic invariants is typically captured by a predicate written in the assertion language of the program logic. In this work however, we do not characterize stochastic invariants directly; we instead use stochastic invariant indicators.

Definition 4.2 (Stochastic Invariant Indicator [10]). Let $\mathcal{G}$ be the CFG $(L, V, l_{\text{init}}, \mathbf{x}_{\text{init}}, \mapsto, G, \text{Pr}, \text{Upd})$. A tuple (SI, p) is a *stochastic invariant indicator* (SI-indicator) if p is a probability value and $\text{SI} : L \times \mathbb{Z}^V \rightarrow \mathbb{R}$ is a partial function such that $\text{SI}(l_{\text{init}}, \mathbf{x}_{\text{init}}) \leq p$, and for all states $(l, \mathbf{x})$ reachable from $(l_{\text{init}}, \mathbf{x}_{\text{init}})$,

- (1) $\text{SI}(l, \mathbf{x}) \geq 0$.
- (2) if $(l, \mathbf{x})$ is an assignment or nondeterministic state, then $\text{SI}(l, \mathbf{x}) \geq \text{SI}(l', \mathbf{x}')$ for every successor $(l', \mathbf{x}')$.
- (3) if $(l, \mathbf{x})$ is a probabilistic state, then $\text{SI}(l, \mathbf{x}) \geq \sum \text{Pr}((l, l'))[\mathbf{x}] \times \text{SI}(l', \mathbf{x}')$ over all possible successor states $(l', \mathbf{x}')$.

Observe that functions SI in the SI-indicators are supermartingale functions. These SI are typically most interesting at states σ where $\text{SI}(\sigma) < 1$; in fact, the collection of states σ with this property corresponds to an underlying stochastic invariant with the same probability value as the SI-indicator. This was formally proven by Chatterjee et al. [10].

LEMMA 4.3 ([10]). Let $\mathcal{G}$ be an CFG. For each stochastic invariant (Ψ, p) of $\mathcal{G}$, there exists a stochastic invariant indicator (SI, p) of $\mathcal{G}$ such that $\Psi \supseteq \{\gamma \in \Sigma_{\mathcal{G}} \mid \text{SI}(\gamma) < 1\}$. Furthermore, for each stochastic invariant indicator (SI, p) of $\mathcal{G}$, there is a stochastic invariant (Ψ, p) such that $\Psi = \{\gamma \in \Sigma_{\mathcal{G}} \mid \text{SI}(\gamma) < 1\}$.

The SI-indicator SI corresponding to the stochastic invariant Ψ maps each state σ the probability with which runs beginning from σ exit Ψ . Thus, the SI-indicator tracks the probability of violating the stochastic invariant. Observe that $\text{SI}(\sigma) \geq 1$ for all states $\sigma \notin \Psi$.

We will use SI-indicators in our proof rules. Note that we cannot use a single stochastic invariant: this is too weak to ensure soundness. Instead, we use a family of stochastic invariants, one for each reachable state. Unlike stochastic invariants, representing SI-indicators is more complicated. Because SI-indicators map states to reals, we cannot always write them directly in our assertion language. At all AST programs, we will nevertheless show that one can always find expressions in our assertion language to effectively represent the SI-indicators our proof rule needs.

AST is a property that holds when the initial state is changed to any reachable state. If $\mathcal{G}$ is almost surely terminating, then $\mathcal{G}$ is almost surely terminating from any state σ reachable from the initial state $\sigma_0 = (l_{\text{init}}, \mathbf{x}_{\text{init}})$. Furthermore, if from every reachable state σ , $\mathcal{G}$ is known to terminate with some minimum probability $\epsilon > 0$, then the program is AST. We exploit these facts in our rule.

Proof Rule 4.1: SI-Indicators for AST

If, for a fixed $0 < p < 1$, there exists

- (1) an inductive invariant Inv containing the initial state,
- (2) a mapping SI from each $\sigma \in \text{Inv}$ to SI-indicator functions $(\text{SI}_{\sigma}, p) : \text{Inv} \rightarrow \mathbb{R}$ such that $\text{SI}_{\sigma}(l, \mathbf{x}) \leq p$ and, for all $(l, \mathbf{x}) \in \text{Inv}$,
 - (a) $\text{SI}_{\sigma}(l, \mathbf{x}) \geq 0$.
 - (b) if $(l, \mathbf{x})$ is an assignment, or nondeterministic state, then $\text{SI}_{\sigma}(l, \mathbf{x}) \geq \text{SI}_{\sigma}(l', \mathbf{x}')$ for every successor $(l', \mathbf{x}')$.
 - (c) if $(l, \mathbf{x})$ is a probabilistic state, then $\text{SI}_{\sigma}(l, \mathbf{x}) \geq \sum \text{Pr}((l, l'))[\mathbf{x}] \times \text{SI}_{\sigma}(l', \mathbf{x}')$ over all possible successor states $(l', \mathbf{x}')$.
- (3) a mapping $\mathcal{E}$ mapping states $\sigma \in \text{Inv}$ to values $\epsilon_{\sigma} \in (0, 1]$,
- (4) a mapping $\mathcal{H}$ mapping states $\sigma \in \text{Inv}$ to values $H_{\sigma} \in \mathbb{N}$,
- (5) a mapping $\mathcal{U}$ from each $\sigma \in \text{Inv}$ to variants $U_{\sigma} : \text{Inv} \rightarrow \mathbb{N}$ that is bounded above by $\mathcal{H}(\sigma)$, maps all states $\{\gamma \mid \text{SI}_{\sigma}(\gamma) \geq 1 \vee \gamma \text{ is terminal}\}$ to 0 and, for other states $(l, \mathbf{x}) \in \text{Inv}$,
 - (a) if $(l, \mathbf{x})$ is an assignment, or nondeterministic state, $U_{\sigma}(l', \mathbf{x}') < U_{\sigma}(l, \mathbf{x})$ for every successor $(l', \mathbf{x}')$.
 - (b) if $(l, \mathbf{x})$ is a probabilistic state, $\sum \text{Pr}(l, l')[\mathbf{x}] > \mathcal{E}(\sigma)$ over all successor states $(l', \mathbf{x}')$ with $U_{\sigma}(l', \mathbf{x}') < U_{\sigma}(l, \mathbf{x})$.

Then, $\mathcal{G}$ is AST.

Intuitively, our rule requires SI-indicator functions SI_{σ} at each $\sigma \in \text{Inv}$ that hold for executions beginning at σ with probability $\geq 1 - p$. Each of these imply stochastic invariants (Ψ_{σ}, p) centered around the state σ . The functions $\mathcal{U}$, $\mathcal{H}$, and $\mathcal{E}$ combine together to form variant functions U_{σ} of the McIver-Morgan kind at each $\sigma \in \text{Inv}$. These U_{σ} further imply that a terminal state is contained within each Ψ_{σ} , and induces paths within each Ψ_{σ} to this terminal state. Feeding U_{σ} , ϵ_{σ} , and H_{σ} into McIver and Morgan's variant [Rule 3.1](#) gives us a proof for the fact that, were the execution to be restricted to Ψ_{σ} , the probability of termination from σ is 1. Therefore, the probability of

termination from each σ is $\geq 1 - p$. Applying the zero-one law of probabilistic processes [36, Lemma 2.6.1] completes the proof of soundness of this rule.

Notice that we don't mandate any locality conditions on the SI-indicators in this rule. This is because they aren't necessary to infer the soundness of the rule. However, we show in our completeness proof that one can always find "local" SI-indicators that only take on values < 1 at finitely many states for AST programs. This is because these SI-indicators are built from appropriate finite stochastic invariants, the existence of which is a consequence of the unrolling lemma.

LEMMA 4.4 (COMPLETENESS). *Rule 4.1 is relatively complete.*

PROOF (SKETCH). Let $\mathcal{G}$ be an AST CFG. Set $\text{Inv} = \text{Reach}(\mathcal{G})$, the set of reachable states of $\mathcal{G}$. Fix a $0 < p < 1$ and a $\sigma \in \text{Reach}(\mathcal{G})$. Since $\mathcal{G}$ is AST, $\text{Pr}_{\text{term}}(\mathcal{G}(\sigma)) = 1$. The unrolling lemma indicates a $k \in \mathbb{N}$ such that the required simulation time to amass a termination probability of $1 - p$ in $\mathcal{G}(\sigma)$ is bounded above by k . Let Σ_k^σ be the collection of all states reachable from σ by a finite path of length at most k . Observe that runs of $\mathcal{G}(\sigma)$ (a) terminate inside Σ_k^σ with a probability of at least $1 - p$, and (b) almost-surely terminate either inside Σ_k^σ or outside Σ_k^σ . (Σ_k^σ, p) is thus the required stochastic invariant for σ . Arguments by Chatterjee et al. [10, Theorem 1] indicate the existence of an SI-Indicator (SI_σ, p) for the stochastic invariant (Σ_k^σ, p) . Furthermore, the finiteness of Σ_k^σ combined with the almost-sure property of either termination inside or exit from Σ_k^σ enables Corollary 2.7, from which it is trivial to extract a suitable variant function U_σ bounded above by some H_σ exhibiting a minimum probability of decrease of $\epsilon_\sigma > 0$. This can be done for all $\sigma \in \text{Inv}$.

Let us now show how we can represent these entities in our assertion language. As in the proof of the completeness of the martingale Rule 3.2, consider the relation I such that $(k, \sigma_1, \sigma_2) \in I$ iff there is a finite path of length $\leq k$ from σ_1 to σ_2 in $\mathcal{G}$. Clearly, I is a computable relation and is thus representable in $\text{Th}(\mathbb{Q})$. Inv and each U_σ can easily be represented in $\text{Th}(\mathbb{Q})$ using this relation I in exactly the same way as with Rule 3.2.

To represent each SI_σ , we note that arguments from Chatterjee et al. [10] show that the output of SI_σ at a state γ is precisely the probability of leaving the stochastic invariant (Σ_k^σ, p) . Encoding lower bounds on this probability in $\text{Th}(\mathbb{Q})$ uses the unrolling lemma, and was essentially shown by Majumdar and Sathiyararayana [33]. The probability with which executions escape Σ_k^σ within m steps lower bounds the probability of leaving Σ_k^σ . The former probability is computable, and the unrolling lemma ensures that in spite of non-determinism, augmenting m with a universal quantifier produces the precise lower bounds over the latter. Hence, each $\text{SI}_\sigma(\gamma)$ is representable in $\text{Th}(\mathbb{Q})$. This completes our proof of relative completeness. $\square$

5 QUANTITATIVE TERMINATION

We now extend our rules to reason about lower and upper bounds on the probability of termination. This notion of termination is referred to in the literature as *quantitative termination*. This is in contrast to *qualitative termination*, the nomenclature employed for AST. As mentioned earlier, the proof rules we specify in this section can be used to show irrational bounds on the termination probability; this is a simple consequence of the fact that all possible termination probabilities can be expressed in our assertion language.

Our upper bound rule is immediate from an observation of the nature of stochastic invariants.

Proof Rule 5.1: Upper Bounds Rule

If there exists

- (1) an inductive invariant Inv containing $(l_{\text{init}}, \mathbf{x}_{\text{init}})$,

- (2) a function $SI : Inv \rightarrow \mathbb{R}$ such that $SI(l_{init}, \mathbf{x}_{init}) \leq p$ and for all $(l, \mathbf{x}) \in Inv$,
- (a) $SI(l, \mathbf{x}) \geq 0$;
 - (b) if $(l, \mathbf{x})$ is an assignment, or nondeterministic state, then $SI(l, \mathbf{x}) \geq SI(l', \mathbf{x}')$ for every successor $(l', \mathbf{x}')$;
 - (c) if $(l, \mathbf{x})$ is a probabilistic state, then $SI(l, \mathbf{x}) \geq \sum \Pr((l, l'))[\mathbf{x}] \times SI(l', \mathbf{x}')$ over all possible successor states $(l', \mathbf{x}')$.
 - (d) if $(l, \mathbf{x})$ is a terminal state, then $SI(l, \mathbf{x}) \geq 1$.

Then, $\Pr_{\text{term}}(\mathcal{G}) \leq p$.

This rule asks for an SI-indicator (and therefore a stochastic invariant) that excludes the terminal state. Therefore, the probability of termination is the probability of escaping the SI-indicator, which is included in the property of the indicator. Notice that, because the SI-indicator is a supermartingale function, the mere existence of a bounded supermartingale that assigns to the terminal state a value ≥ 1 is sufficient to extract an upper bound.

LEMMA 5.1. *Rule 5.1 is sound and relatively complete.*

PROOF. Lemma 4.3 shows that the pair $(\{\gamma \in Inv \mid SI(\gamma) < 1\}, p)$ is a stochastic invariant of $\mathcal{G}$. Since $SI(\gamma) \geq 1$ at the terminal state, this stochastic invariant doesn't contain the terminal state. Soundness of the rule trivially follows from the fact that, in order to terminate, a run must leave this invariant and this probability is bounded above by p .

For completeness, set $Inv = \text{Reach}(\mathcal{G})$ and let $\sigma_{\perp} \in Inv$ be the terminal state. Set $\Psi = Inv \setminus \{\sigma_{bot}\}$, and observe that the collection of runs leaving Ψ is identical to the collection of terminal runs. Therefore, (Ψ, p) must be a stochastic invariant. Lemma 4.3 indicates the existence of the SI-Invariant SI from Ψ . SI immediately satisfies the conditions of the rule.

To represent SI in our assertion language, note again that $SI(\gamma)$ is precisely the probability of termination from γ . The unrolling lemma indicates that this probability is lower bounded by the probability of termination within m steps from γ . The latter probability is computable, and is hence representable in $\text{Th}(\mathbb{Q})$. Prepending an appropriate universal quantifier for m allows us to form lower bounds for SI in $\text{Th}(\mathbb{Q})$. Our upper bound rule is thus relatively complete. $\square$

5.1 Towards a Lower Bound

For a lower bound on the probability of termination, we start with the following rule from Chatterjee et al. [10].

Proof Rule 5.2: SI-indicators for Lower Bounds

If there exists

- (1) an inductive invariant Inv containing $(l_{init}, \mathbf{x}_{init})$,
- (2) a stochastic invariant indicator $SI : Inv \rightarrow \mathbb{R}$ such that $SI(l_{init}, \mathbf{x}_{init}) \geq p$, and for all $(l, \mathbf{x}) \in Inv$,
 - (a) $SI(l, \mathbf{x}) \geq 0$.
 - (b) if $(l, \mathbf{x})$ is an assignment, or nondeterministic state, then $SI(l, \mathbf{x}) \geq SI(l', \mathbf{x}')$ for every successor $(l', \mathbf{x}')$.
 - (c) if $(l, \mathbf{x})$ is a probabilistic state, then $SI(l, \mathbf{x}) \geq \sum \Pr((l, l'))[\mathbf{x}] \times SI(l', \mathbf{x}')$ over all possible successor states $(l', \mathbf{x}')$.
- (3) an $\epsilon > 0$,




Fig. 1. If shortest runs of Σ_{good} were too long. This is a representation of the collection of executions beginning at a good state $\sigma \in \Sigma_{good}$, according to the partitioning system suggested in the proof of Lemma 5.2. The black nodes are the terminal states; they all lead to the single terminal state. The blue states are identical to each other; the same holds for the blue green states. The pathological scheduler s' always takes the red back edges, rendering no terminal runs from σ .

- (4) a variant $U : \text{Inv} \rightarrow \mathbb{N}$ that is bounded above by some H , maps all states $\{\gamma \in \text{Inv} \mid \text{SI}(\gamma) \geq 1 \vee \gamma \text{ is terminal}\}$ to 0, and for other states $(l, \mathbf{x}) \in \text{Inv}$,
 - (a) if $(l, \mathbf{x})$ is an assignment, or nondeterministic state, $U(l', \mathbf{x}') < U(l, \mathbf{x})$ for every successor $(l', \mathbf{x}')$.
 - (b) if $(l, \mathbf{x})$ is a probabilistic state, $\sum \text{Pr}(l, l')[\mathbf{x}] > \epsilon$ over all successor states $(l', \mathbf{x}')$ with $U(l', \mathbf{x}') < U(l, \mathbf{x})$.

Then, $\text{Pr}_{\text{term}}(\mathcal{G}) \geq 1 - p$.

This rule demands a SI-indicator (SI, p) and a bounded variant U such that SI induces a stochastic invariant (Ψ, p) that contains the terminal state. Intuitively, this rule works by splitting the state space into terminating and possibly non-terminating segments. The invariant Ψ represents the terminating section of the state space. The application of McIver & Morgan's Rule 3.1 with the variant U allows us to deduce that, were the execution be restricted to Ψ , the program almost-surely terminates. Observe that the variant U effectively considers all states outside Ψ to be terminal. Therefore, $\mathcal{G}$ almost-surely either escapes Ψ or terminates within Ψ . Non-termination is thus subsumed by the event of escaping the invariant, the probability of which is p . Hence, the probability of termination is $\geq 1 - p$.

Observe that, like our AST Rule 3.2, this rule requires a supermartingale and a variant function that work in tandem. However, unlike Rule 3.2, the supermartingale and the variant are entirely bounded.

This soundness argument was formally shown by Chatterjee et al. [10]. In their original presentation, they do not specify an exact technique for determining the almost-sure property of either termination within or escape from the induced stochastic invariant Ψ . We will explain our choice of the bounded variant Rule 3.1 of McIver and Morgan [36] in a moment. They additionally claimed the completeness of this rule, assuming the usage of a complete rule for almost-sure termination. If their completeness argument were true, it would indicate that all probabilistic programs induce state spaces that can neatly be partitioned into terminating and non-terminating sections. In Section 5.3, we show that this isn't the case using a counterexample where this split isn't possible.

Nevertheless, this rule is complete for finite-state programs. This finite-state completeness pairs well with the finite-state completeness of the bounded variant Rule 3.1 we use to certify the almost-sure property contained in the rule.

LEMMA 5.2. *Rule 5.2 is complete for finite state CFGs.*

PROOF. In this proof, we argue about the stochastic invariants directly. The SI-indicator and variant functions can be derived from them using prior techniques [10, 36].

Let $\mathcal{G}$ be a finite state CFG. Partition $\text{Reach}(\mathcal{G})$, the set of states reachable from the initial state of $\mathcal{G}$, into (a) the singleton containing the terminal state $\Sigma_{\perp}$, (b) the *bad* states Σ_{bad} with the property that no finite paths ending on these states can be extended to a terminal run, (c) the *good* states Σ_{good} such that if a scheduler f induces a finite path ending at a good state σ , f induces at least one terminating run passing through σ , and (d) the remaining *neutral* states $\Sigma_{neutral}$, with the property that each neutral state σ is associated with a pathological scheduler $\mathfrak{s}_{\sigma}$ that induces runs that, if they pass through σ , do not terminate. Notice that, as long as $p < 1$ (the case where $p = 1$ is trivial, so we skip it), the initial state of $\mathcal{G}$ is in Σ_{good} . We show that $(\Sigma_{good} \cup \Sigma_{\perp}, p)$ is the required stochastic invariant.

We begin by showing that runs that remain inside $\Sigma_{good} \cup \Sigma_{\perp}$ almost-surely terminate. Fix a state $\sigma \in \Sigma_{good}$. Map to every scheduler $\mathfrak{s}$ that induces runs passing through σ the shortest terminating consistent finite path $\pi_{\mathfrak{s}}$ beginning from σ . Suppose, for some scheduler $\mathfrak{s}$, $|\pi_{\mathfrak{s}}| > |\text{Reach}(\mathcal{G})|$. Then, $\pi_{\mathfrak{s}}$ must visit some $\sigma' \in \text{Reach}(\mathcal{G})$ twice. This indicates a loop from $\sigma' \rightarrow \sigma'$ in $\mathcal{G}$, and there must thus exist a scheduler $\mathfrak{s}'$ that extends $\pi_{\mathfrak{s}}$ by repeating this loop infinitely often. Since $\pi_{\mathfrak{s}}$ is the smallest terminating finite path beginning from σ , the same holds for all terminating finite paths consistent with $\mathfrak{s}$ beginning from σ . Thus, all terminating runs consistent with $\mathfrak{s}$ that pass through σ must contain a loop. There must exist a scheduler $\mathfrak{s}'$ which exploits these loops and yields no terminating runs passing through σ . Fig. 1 depicts the operation of $\mathfrak{s}'$. Observe that the existence of $\mathfrak{s}'$ contradicts $\sigma \in \Sigma_{good}$.

Therefore, from every state $\sigma \in \Sigma_{good}$, there is a terminating run of length $\leq |\text{Reach}(\mathcal{G})|$ no matter which scheduler is used. Thus, the probability of leaving Σ_{good} from σ is bounded below by $q^{|\text{Reach}(\mathcal{G})|}$, where q is the smallest transition probability of $\mathcal{G}$ (note that q only exists because $\mathcal{G}$ is finite state). Further, Σ_{good} only contains non-terminal states. This enables the zero-one law of probabilistic processes [36, Lemma 2.6.1], allowing us to deduce the almost-certain escape from Σ_{good} to $\Sigma_{\perp} \cup \Sigma_{bad} \cup \Sigma_{neutral}$. Hence, under all schedulers, the probability of either terminating inside $\Sigma_{good} \cup \Sigma_{\perp}$ or entering $\Sigma_{bad} \cup \Sigma_{neutral}$ is 1.

We now show that $(\Sigma_{good} \cup \Sigma_{bot}, p)$ is a stochastic invariant. It's easy to see that if a run ever enters Σ_{bad} , it never terminates. We know that if an execution enters some $\sigma \in \Sigma_{neutral}$ under a pathological scheduler $\mathfrak{s}_{\sigma}$, it never terminates. Let σ_1 and σ_2 be neutral states and $\mathfrak{s}_1$ and $\mathfrak{s}_2$ be their corresponding pathological schedulers. Notice that $\mathfrak{s}_1$ may induce terminating executions that pass through σ_2 . One can build a scheduler $\mathfrak{s}_3$ that mimics $\mathfrak{s}_1$ until the execution reaches σ_2 , and once it does, mimics $\mathfrak{s}_2$. Thus, $\mathfrak{s}_3$ would produce the pathological behaviour of both $\mathfrak{s}_1$ and $\mathfrak{s}_2$. In this way, we compose the pathological behavior of all neutral states to produce a scheduler $\mathfrak{f}$ that induces runs that, if they enter a neutral state, never terminate. Under $\mathfrak{f}$, leaving $\Sigma_{good} \cup \Sigma_{\perp}$ is equivalent to non-termination, and all terminating runs are made up of good states until their final states.

Take a scheduler $\mathfrak{s}$ that induces terminating runs that pass through neutral states. Compose the scheduler $\mathfrak{s}'$ that mimics $\mathfrak{s}$ until the execution enters a neutral state and mimics $\mathfrak{f}$ from then on. Let $T_{\mathfrak{s}}$ and $T_{\mathfrak{s}'}$ be the collection of terminating runs consistent with $\mathfrak{s}$ and $\mathfrak{s}'$ respectively. Each run in $T_{\mathfrak{s}'}$ is made up of good and/or terminal states, and is therefore consistent with $\mathfrak{s}$. Hence, $T_{\mathfrak{s}} \supset T_{\mathfrak{s}'}$ and, because $\mathfrak{s}$ and $\mathfrak{s}'$ agree on $T_{\mathfrak{s}'}$, we have $\mathbb{P}_{\mathfrak{s}}(T_{\mathfrak{s}}) > \mathbb{P}_{\mathfrak{s}}(T_{\mathfrak{s}'})$, where $\mathbb{P}_{\mathfrak{s}}$ is the probability measures in the semantics of $\mathcal{G}$ induced by $\mathfrak{s}$.

Leaving $\Sigma_{good} \cup \Sigma_{\perp}$ is equivalent to entering $\Sigma_{neutral} \cup \Sigma_{\perp}$. Observe that the probability of never leaving $\Sigma_{good} \cup \Sigma_{\perp}$ under $\mathfrak{s}$ is the same as the probability of never leaving $\Sigma_{good} \cup \Sigma_{\perp}$ under $\mathfrak{s}'$, as $\mathfrak{s}$ and $\mathfrak{s}'$ agree until then. Furthermore, the probability measure of never leaving $\Sigma_{good} \cup \Sigma_{\perp}$ is just




Fig. 2. Counterexample to the SI-rule for lower bounds. With an initial state of $(l_1, (1, 2, 1/4))$, the termination probability of this program is $1/2$. However, there is no SI that shows this.

$\Pr_{\mathfrak{s}}(T_{\mathfrak{s}'}) = \Pr_{\mathfrak{s}'}(T_{\mathfrak{s}'})$. Add $\Pr_{\text{term}}(\mathcal{G}) \geq 1 - p \implies p_{\mathfrak{s}'}(T_{\mathfrak{s}'}) \geq 1 - p$, and we get that, under $\mathfrak{s}$, the probability of leaving $\Sigma_{\text{good}} \cup \Sigma_{\perp}$ is upper bounded by p . Since this is true for any $\mathfrak{s}$, the lemma is proved. $\square$

We do not show the relative completeness of this rule; nevertheless, this is easy to show using techniques discussed in prior rules.

5.2 Our Rule

We now show a sound and complete rule for lower bounds that fixes the prior [Rule 5.2](#). This rule is implicitly contained in the details of the erroneous proof of [\[10\]](#). It is principally similar to [Rule 4.1](#), in that it identifies finite sub-instances where prior rules can apply. It then combines the proofs of these sub-instances to deduce the desired lower bound.

Proof Rule 5.3: Lower Bounds Rule

If for all $n \in \mathbb{N}$, there are functions SI_n and U_n that enable the application of [Rule 5.2](#) to deduce $\Pr_{\text{term}}(\mathcal{G}) \geq 1 - (p + \frac{1}{n})$, then $\Pr_{\text{term}}(\mathcal{G}) \geq 1 - p$.

Soundness of this rule follows trivially from the soundness of the prior [Rule 5.2](#). The completeness of this rule is derived from the unrolling lemma; to reach a termination probability of p , the program must be able to amass a termination probability of $p - \frac{1}{n}$ within a finite subspace. [Lemma 5.2](#) shows that finiteness can always be captured by the prior [Rule 5.2](#).

LEMMA 5.3. *Rule 5.3 is relatively complete.*

PROOF. Let $\mathcal{G}$ be CFG such that $\Pr_{\text{term}}(\mathcal{G}) \geq 1 - p$ for some $p > 0$. Fix some $n \in \mathbb{N}$. Let k_n be the upper bound over the required simulation times across all schedulers to amass a termination probability of $p - 1/n$. Denote by Σ_n the set of states σ such that there is a finite path of length at most k_n beginning at $(l_{\text{init}}, \mathbf{x}_{\text{init}})$ and ending at σ . Clearly, Σ_n must be finite and, for a fixed scheduler $\mathfrak{s}$, the probability measure of the collection of terminating runs made up of states in Σ_n consistent with $\mathfrak{s}$ must be $\geq p - 1/n$. Additionally, observe that runs of $\mathcal{G}$ either terminate inside Σ_n or leave Σ_n . By the soundness of [Rule 5.2](#), the termination probability of $\mathcal{G}_{\varphi_n}$ must be $\geq p - 1/n$.

Each stochastic invariant $(\Sigma_n, p + 1/n)$ can be transformed into SI-Indicators $(SI_n, p + 1/n)$ using prior techniques [\[10\]](#). Representing each SI_n and U_n in $\text{Th}(\mathbb{Q})$ can be done using techniques described in the proof of [Lemma 4.4](#). Hence, this rule is relatively complete. $\square$

5.3 Counterexample to Completeness for Rule 5.2

[\[10\]](#) claimed that their [Rule 5.2](#) is complete for all programs. As promised, we now demonstrate a counterexample to their claim of completeness.

Rule 5.2 for lower bounds can be applied onto any CFG that induces a set of states Ψ with the property that executions remain within Ψ with exactly the probability of termination. As mentioned previously, not all programs are so well behaved. Consider the program $\mathcal{K}$ defined in Fig. 2.

The initial location of $\mathcal{K}$ is l_1 , and the values of the variables (x_1, x_2, x_3) are $(1, 2, 1/4)$. l_1 is a probabilistic location, l_3 is an assignment location, and l_2 is a terminal location. It isn't difficult to prove that the probability of termination of $\mathcal{K}$ is $1/2$; we leave the details to the diligent reader. The SI-rule for lower bounds requires a stochastic invariant $(\Psi_{\mathcal{K}}, 1/2)$ such that executions almost-surely either terminate or exit $\Psi_{\mathcal{K}}$.

LEMMA 5.4. *There is no stochastic invariant $(\Psi_{\mathcal{K}}, 1/2)$ of $\mathcal{K}$ such that runs almost-surely either terminate or leave $\Psi_{\mathcal{K}}$.*

PROOF. Suppose there does exist a stochastic invariant $(\Psi_{\mathcal{K}}, 1/2)$ that satisfies these properties. Therefore, the probability measure of the union of the collection of runs Leave_{Ψ} that leave $\Psi_{\mathcal{K}}$ and the runs Term_{Ψ} that terminate inside $\Psi_{\mathcal{K}}$ is 1. However, because $\Psi_{\mathcal{K}}$ is a stochastic invariant, the probability measure of Leave_{Ψ} is bounded above by $1/2$. This means that the measure of Term_{Ψ} is bounded below by $1/2$. But, the termination probability of $\mathcal{K}$ is $1/2$. Consequently, the measure of Term_{Ψ} must be exactly $1/2$. This means Term_{Ψ} contains all terminating runs of $\mathcal{K}$.

It is easy to see that from any state $(l, \mathbf{x})$ reachable from the initial state, there is a finite path of length at most 2 that leads it to a terminal state. Therefore, all reachable states $(l, \mathbf{x})$ are a part of some terminating run; meaning that the set of states that make up the runs in Term_{Ψ} must be the set of reachable states. This is only possible when $\Psi_{\mathcal{K}}$ is the set of reachable states. This means no runs leave $\Psi_{\mathcal{K}}$, and therefore, the measure of Term_{Ψ} is 1. This contradicts the fact that the termination probability of $\mathcal{K}$ is $1/2$. $\square$

A note on syntax. The CFG of Chatterjee et al. [10] over which the claim of completeness of Rule 5.2 was made do not feature fractional expressions guiding probabilistic branching. Nevertheless, they can be simulated with small programs that only use the basic coin flip [21]. Therefore, Fig. 2 is a valid counterexample to their claim.

6 TRAVELING BETWEEN PROOF SYSTEMS

A new proof rule, ultimately, is interesting only if one can actually prove the termination of many programs. In order to show that our proof rules, in addition to their theoretical properties, are also applicable in a variety of situations, we demonstrate that proofs in many existing proof systems can be compiled into our proof rules.

From McIver and Morgan [36]. The variant functions from Rule 3.1 immediately form the variant functions required in Rule 3.2. Take A to simply be the singleton containing the terminal state, and set V to 0 at the terminal state and 1 everywhere else. This gives all we need to apply Rule 3.2.

From McIver et al. [37]. The AST proof rule proposed by McIver et al. [37] has been applied onto a variety of programs, and has been shown to be theoretically applicable over the 2D random walker. Applications of their rule effectively requires the construction of a distance variant that is also a supermartingale. We note that their variants can be reused in Rule 3.2 with little alterations as both the supermartingale and variant functions. This means proofs in their rule can be easily translated to proofs that use Rule 3.2.

From Rule 3.2 to Rule 4.1. Hidden in the proof of the soundness of Rule 3.2 are the stochastic invariants that form the basis of Rule 4.1. Fix a p , and take the set $\Psi_{\sigma} = \{\gamma \in \text{Inv} \mid V(\gamma) \leq v_{\uparrow}\}$ for a sufficiently high value of $v_{\uparrow}$ to yield an upper bound of p on the probability of exiting Ψ_{σ} . Then, expand Ψ_{σ} with the states necessary to keep all shortest consistent terminal runs across

schedulers from states in Ψ_σ entirely within Ψ_σ . In spite of these extensions, the value of V will be entirely bounded when restricted to states in Ψ_σ . It is then trivial to build the indicator functions from each stochastic invariant (Ψ_σ, p) and the variant functions from U , completing a translation from Rule 3.2 to Rule 4.1. Note that using this technique, one can translate proofs from McIver et al. [37] and McIver and Morgan [36] to Rule 4.1 as well.

Using Guard Strengthening [19]. Feng et al. [19] have demonstrated Guard Strengthening as a technique for proving lower bounds on, amongst others, the termination of deterministic probabilistic programs. In principle, their technique can be translated to apply over CFGs as follows: strengthen each transition guard by a suitable predicate φ and add self loops at all locations guarded by $\neg\varphi$. We observe that guards can succinctly overapproximate the set of states forming finite-state stochastic invariants: simply take the highest and lowest values each variable can take while remaining in the invariant, and form a predicate that limits each variable within these bounds. Thus, guards serve as convenience mechanisms for the application of Rule 5.3. Additionally, observe that when proving lower bounds on termination probabilities, relevant finite state stochastic invariants Ψ_σ for each reachable state σ are guaranteed to exist. Therefore, each stochastic invariant in Rule 5.3 can be more succinctly represented using guards.

Using Stochastic Invariants [10]. Separately, Chatterjee et al. [10] have shown the applicability of Rule 5.2 to demonstrate lower bounds on the termination probabilities for a variety of programs, and have also presented template-based synthesis techniques for achieving limited completeness. These proofs are also valid for Rule 5.3, by setting the same stochastic invariant for each n .

7 CONCLUSION

We have presented the first sound and relatively complete proof rules for qualitative and quantitative termination of probabilistic programs with bounded probabilistic and nondeterministic choice. Our proof rules combine the familiar ingredients of supermartingales and variant functions in novel ways to reach completeness. We have demonstrated relative completeness of our rules in the assertion language of arithmetic. Our rules are able to accommodate existing proof techniques in the literature with minimal effort, thus demonstrating their applicability.

REFERENCES

- [1] Krzysztof R. Apt. 1981. Ten Years of Hoare's Logic: A Survey - Part 1. *ACM Trans. Program. Lang. Syst.* 3, 4 (1981), 431–483. <https://doi.org/10.1145/357146.357150>
- [2] Krzysztof R. Apt and Dexter Kozen. 1986. Limits for Automatic Verification of Finite-State Concurrent Systems. *Inf. Process. Lett.* 22, 6 (1986), 307–309. [https://doi.org/10.1016/0020-0190\(86\)90071-2](https://doi.org/10.1016/0020-0190(86)90071-2)
- [3] Krzysztof R. Apt and Gordon D. Plotkin. 1986. Countable nondeterminism and random assignment. *J. ACM* 33, 4 (1986), 724–767. <https://doi.org/10.1145/6490.6494>
- [4] Martin Avanzini, Ugo Dal Lago, and Akihisa Yamada. 2020. On probabilistic term rewriting. *Sci. Comput. Program.* 185 (2020). <https://doi.org/10.1016/j.scico.2019.102338>
- [5] Christel Baier and Joost-Pieter Katoen. 2008. *Principles of model checking*. MIT Press.
- [6] Kevin Batz, Benjamin Lucien Kaminski, Joost-Pieter Katoen, and Christoph Matheja. 2021. Relatively complete verification of probabilistic programs: an expressive language for expectation-based reasoning. *Proc. ACM Program. Lang.* 5, POPL (2021), 1–30. <https://doi.org/10.1145/3434320>
- [7] Andrea Bianco and Luca de Alfaro. 1995. Model Checking of Probabilistic and Nondeterministic Systems. In *Foundations of Software Technology and Theoretical Computer Science, 15th Conference, Bangalore, India, December 18-20, 1995, Proceedings (Lecture Notes in Computer Science, Vol. 1026)*, P. S. Thiagarajan (Ed.). Springer, 499–513. https://doi.org/10.1007/3-540-60692-0_70
- [8] Olivier Bournez and Florent Garnier. 2005. Proving Positive Almost-Sure Termination. In *Term Rewriting and Applications, 16th International Conference, RTA 2005, Nara, Japan, April 19-21, 2005, Proceedings (Lecture Notes in Computer Science, Vol. 3467)*, Jürgen Giesl (Ed.). Springer, 323–337. https://doi.org/10.1007/978-3-540-32033-3_24
- [9] Aleksandar Chakarov and Sriram Sankaranarayanan. 2013. Probabilistic Program Analysis with Martingales. In *Computer Aided Verification - 25th International Conference, CAV 2013, Saint Petersburg, Russia, July 13-19, 2013. Proceedings (Lecture Notes in Computer Science, Vol. 8044)*, Natasha Sharygina and Helmut Veith (Eds.). Springer, 511–526. https://doi.org/10.1007/978-3-642-39799-8_34
- [10] Krishnendu Chatterjee, Amir Kafshdar Goharshady, Tobias Meggendorfer, and Dorde Zikelic. 2022. Sound and Complete Certificates for Quantitative Termination Analysis of Probabilistic Programs. In *Computer Aided Verification - 34th International Conference, CAV 2022, Haifa, Israel, August 7-10, 2022, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 13371)*, Sharon Shoham and Yakir Vizel (Eds.). Springer, 55–78. https://doi.org/10.1007/978-3-031-13185-1_4
- [11] Krishnendu Chatterjee, Petr Novotný, and Dorde Zikelic. 2017. Stochastic invariants for probabilistic termination. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*, Giuseppe Castagna and Andrew D. Gordon (Eds.). ACM, 145–160. <https://doi.org/10.1145/3009837.3009873>
- [12] Stephen A. Cook. 1978. Soundness and Completeness of an Axiom System for Program Verification. *SIAM J. Comput.* 7, 1 (1978), 70–90. <https://doi.org/10.1137/0207005>
- [13] Costas Courcoubetis and Mihalis Yannakakis. 1995. The Complexity of Probabilistic Verification. *J. ACM* 42, 4 (1995), 857–907. <https://doi.org/10.1145/210332.210339>
- [14] Luca de Alfaro and Thomas A. Henzinger. 2000. Concurrent Omega-Regular Games. In *15th Annual IEEE Symposium on Logic in Computer Science, Santa Barbara, California, USA, June 26-29, 2000*. IEEE Computer Society, 141–154. <https://doi.org/10.1109/LICS.2000.855763>
- [15] Luca de Alfaro, Thomas A. Henzinger, and Orna Kupferman. 2007. Concurrent reachability games. *Theor. Comput. Sci.* 386, 3 (2007), 188–217. <https://doi.org/10.1016/J.TCS.2007.07.008>
- [16] Jerry den Hartog and Erik P. de Vink. 2002. Verifying Probabilistic Programs Using a Hoare Like Logic. *Int. J. Found. Comput. Sci.* 13, 3 (2002), 315–340. <https://doi.org/10.1142/S012905410200114X>
- [17] Edsger W. Dijkstra. 1976. *A Discipline of Programming*. Prentice-Hall. <https://www.worldcat.org/oclc/01958445>
- [18] J. L. Doob. 1953. *Stochastic processes*. John Wiley & Sons, New York. viii+654 pages. MR 15,445b. Zbl 0053.26802..
- [19] Shenghua Feng, Mingshuai Chen, Han Su, Benjamin Lucien Kaminski, Joost-Pieter Katoen, and Naijun Zhan. 2023. Lower Bounds for Possibly Divergent Probabilistic Programs. *Proc. ACM Program. Lang.* 7, OOPSLA1 (2023), 696–726. <https://doi.org/10.1145/3586051>
- [20] Luis María Ferrer Fioriti and Holger Hermanns. 2015. Probabilistic Termination: Soundness, Completeness, and Compositionality. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015*, Sriram K. Rajamani and David Walker (Eds.). ACM, 489–501. <https://doi.org/10.1145/2676726.2677001>
- [21] Philippe Flajolet, Maryse Pelletier, and Michèle Soria. 2011. On Buffon Machines and Numbers. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, Dana Randall (Ed.). SIAM, 172–183. <https://doi.org/10.1137/1.9781611973082.15>

- [22] Robert W. Floyd. 1993. *Assigning Meanings to Programs*. Springer Netherlands, Dordrecht, 65–81. https://doi.org/10.1007/978-94-011-1793-7_4
- [23] F.G. Foster. 1951. Markov chains with an enumerable number of states and a class of cascade processes. *Math. Proc. Cambridge Philos. Soc.* 47 (1951), 77–85.
- [24] F.G. Foster. 1953. On the stochastic matrices associated with certain queuing processes. *Ann. Math. Statistics* 24 (1953), 355–360.
- [25] Hongfei Fu and Krishnendu Chatterjee. 2019. Termination of Nondeterministic Probabilistic Programs. In *Verification, Model Checking, and Abstract Interpretation - 20th International Conference, VMCAI 2019, Cascais, Portugal, January 13-15, 2019, Proceedings (Lecture Notes in Computer Science, Vol. 11388)*, Constantin Enea and Ruzica Piskac (Eds.). Springer, 468–490. https://doi.org/10.1007/978-3-030-11245-5_22
- [26] David Harel. 1980. Proving the Correctness of Regular Deterministic Programs: A Unifying Survey Using Dynamic Logic. *Theor. Comput. Sci.* 12 (1980), 61–81. [https://doi.org/10.1016/0304-3975\(80\)90005-5](https://doi.org/10.1016/0304-3975(80)90005-5)
- [27] David Harel, Dexter Kozen, and Jerzy Tiuryn. 2000. *Dynamic Logic*. MIT Press.
- [28] Sergiu Hart, Micha Sharir, and Amir Pnueli. 1983. Termination of Probabilistic Concurrent Program. *ACM Trans. Program. Lang. Syst.* 5, 3 (1983), 356–380. <https://doi.org/10.1145/2166.357214>
- [29] Peter Hitchcock and David Michael Ritchie Park. 1972. Induction Rules and Termination Proofs. In *Automata, Languages and Programming, Colloquium, Paris, France, July 3-7, 1972*, Maurice Nivat (Ed.). North-Holland, Amsterdam, 225–251.
- [30] Mingzhang Huang, Hongfei Fu, and Krishnendu Chatterjee. 2018. New Approaches for Almost-Sure Termination of Probabilistic Programs. In *Programming Languages and Systems - 16th Asian Symposium, APLAS 2018, Wellington, New Zealand, December 2-6, 2018, Proceedings (Lecture Notes in Computer Science, Vol. 11275)*, Sukyoung Ryu (Ed.). Springer, 181–201. https://doi.org/10.1007/978-3-030-02768-1_11
- [31] Benjamin Lucien Kaminski, Joost-Pieter Katoen, and Christoph Matheja. 2019. On the hardness of analyzing probabilistic programs. *Acta Informatica* 56, 3 (2019), 255–285. <https://doi.org/10.1007/s00236-018-0321-1>
- [32] Dexter Kozen. 2006. *Theory of Computation*. Springer. <https://doi.org/10.1007/1-84628-477-5>
- [33] Rupak Majumdar and V. R. Sathiyararayana. 2023. Positive Almost-Sure Termination - Complexity and Proof Rules. *CoRR* abs/2310.16145 (2023). <https://doi.org/10.48550/ARXIV.2310.16145> arXiv:2310.16145
- [34] Rupak Majumdar and V. R. Sathiyararayana. 2024. Positive Almost-Sure Termination: Complexity and Proof Rules. *Proc. ACM Program. Lang.* 8, POPL (2024), 1089–1117. <https://doi.org/10.1145/3632879>
- [35] Zohar Manna and Amir Pnueli. 1974. Axiomatic Approach to Total Correctness of Programs. *Acta Informatica* 3 (1974), 243–263. <https://doi.org/10.1007/BF00288637>
- [36] Annabelle McIver and Carroll Morgan. 2005. *Abstraction, Refinement and Proof for Probabilistic Systems*. Springer. <https://doi.org/10.1007/B138392>
- [37] Annabelle McIver, Carroll Morgan, Benjamin Lucien Kaminski, and Joost-Pieter Katoen. 2018. A new proof rule for almost-sure termination. *Proc. ACM Program. Lang.* 2, POPL (2018), 33:1–33:28. <https://doi.org/10.1145/3158121>
- [38] Mikhail Menshikov, Serguei Popov, and Andrew Wade. 2017. *Non-homogeneous random walks: Lyapunov function methods for near critical stochastic systems*. Cambridge University Press.
- [39] Jean-François Mertens, Ester Samuel-Cahn, and Shmuel Zamir. 1978. Necessary and Sufficient Conditions for Recurrence and Transience of Markov Chains, in Terms of Inequalities. *Journal of Applied Probability* 15, 4 (1978), 848–851. <http://www.jstor.org/stable/3213440>
- [40] George Pólya. 1921. Über eine aufgabe betreffend die irrfahrt im strassennetz. *Math. Ann.* 84 (1921), 149–160.
- [41] Serguei Popov. 2021. *Two-Dimensional Random Walk: From Path Counting to Random Interlacements*. Cambridge University Press. <https://doi.org/10.1017/9781108680134>
- [42] Julia Robinson. 1949. Definability and Decision Problems in Arithmetic. *J. Symb. Log.* 14, 2 (1949), 98–114. <https://doi.org/10.2307/2266510>
- [43] Hartley Rogers Jr. 1987. *Theory of recursive functions and effective computability (Reprint from 1967)*. MIT Press. <https://mitpress.mit.edu/9780262680523/theory-of-recursive-functions-and-effective-computability/>
- [44] Toru Takisaka, Yuichiro Oyabu, Natsuki Urabe, and Ichiro Hasuo. 2021. Ranking and Repulsing Supermartingales for Reachability in Randomized Programs. *ACM Trans. Program. Lang. Syst.* 43, 2 (2021), 5:1–5:46. <https://doi.org/10.1145/3450967>
- [45] Alan M. Turing. 1937. On computable numbers, with an application to the Entscheidungsproblem. *Proc. London Math. Soc.* s2-42, 1 (1937), 230–265. <https://doi.org/10.1112/PLMS/S2-42.1.230>
- [46] Moshe Y. Vardi. 1985. Automatic Verification of Probabilistic Concurrent Finite-State Programs. In *26th Annual Symposium on Foundations of Computer Science, Portland, Oregon, USA, 21-23 October 1985*. IEEE Computer Society, 327–338. <https://doi.org/10.1109/SFCS.1985.12>