



Foundations for Deductive Verification of Continuous Probabilistic Programs

From Lebesgue to Riemann and Back

KEVIN BATZ, RWTH Aachen University, Germany and University College London, United Kingdom

JOOST-PIETER KATOEN, RWTH Aachen University, Germany

FRANCESCA RANDONE, University of Trieste, Italy

TOBIAS WINKLER, RWTH Aachen University, Germany

We lay out novel foundations for the computer-aided verification of guaranteed bounds on expected outcomes of imperative probabilistic programs featuring (i) general *loops*, (ii) *continuous* distributions, and (iii) *conditioning*. To handle loops we rely on user-provided quantitative *invariants*, as is well established. However, in the realm of continuous distributions, *invariant verification* becomes extremely challenging due to the presence of *integrals* in expectation-based program semantics. Our key idea is to soundly *under-* or *over-approximate* these integrals via *Riemann sums*. We show that this approach enables the SMT-based invariant verification for programs with a fairly general control flow structure. On the theoretical side, we prove *convergence* of our Riemann approximations, and establish coRE-completeness of the central verification problems. On the practical side, we show that our approach enables to use existing automated verifiers targeting *discrete* probabilistic programs for the verification of programs involving *continuous sampling*. Towards this end, we implement our approach in the recent quantitative verification infrastructure CAESAR by encoding Riemann sums in its intermediate verification language. We present several promising case studies.

CCS Concepts: • **Theory of computation** → **Probabilistic computation; Logic and verification; Invariants; Pre- and post-conditions; Program verification**; • **Mathematics of computing** → *Probabilistic inference problems*.

Additional Key Words and Phrases: probabilistic programs, deductive program verification, continuous distributions, quantitative loop invariants, weakest preexpectations, approximate integration, SMT solving

ACM Reference Format:

Kevin Batz, Joost-Pieter Katoen, Francesca Randone, and Tobias Winkler. 2025. Foundations for Deductive Verification of Continuous Probabilistic Programs: From Lebesgue to Riemann and Back. *Proc. ACM Program. Lang.* 9, OOPSLA1, Article 95 (April 2025), 28 pages. <https://doi.org/10.1145/3720429>

1 Introduction

Probabilistic programs (PP) are usual programs with the additional ability to

- (1) *branch* their flow of control based on the outcomes of coin flips,
- (2) *sample* numbers from predefined continuous and/or discrete probability distributions, and
- (3) *condition* their current state on observations using dedicated observe-instructions.

Authors' Contact Information: [Kevin Batz](mailto:kevin.batz@cs.rwth-aachen.de), RWTH Aachen University, Aachen, Germany and University College London, London, United Kingdom, kevin.batz@cs.rwth-aachen.de; [Joost-Pieter Katoen](mailto:joost-pieter.katoen@cs.rwth-aachen.de), RWTH Aachen University, Aachen, Germany, katoen@cs.rwth-aachen.de; [Francesca Randone](mailto:francesca.randone@units.it), University of Trieste, Trieste, Italy, francesca.randone@units.it; [Tobias Winkler](mailto:tobias.winkler@cs.rwth-aachen.de), RWTH Aachen University, Aachen, Germany, tobias.winkler@cs.rwth-aachen.de.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2025 Copyright held by the owner/author(s).

ACM 2475-1421/2025/4-ART95

<https://doi.org/10.1145/3720429>

Programs with one or more of these (partly overlapping) abilities have been extensively studied in different contexts and throughout various communities: *Randomized algorithms* typically rely on ability (1) to speed up computations on average [38] or for breaking symmetries in distributed computing [33]. *Sampling* and *conditioning* (abilities (2) and (3)) are extensively employed in connection with Bayes' rule in machine learning, AI, and cognitive science to statistically infer information from observed data [28, 29, 54]. *Continuous distributions* are pivotal for these applications.

Many applications require *expressive* PP with the usual constructs from traditional programs – including unbounded while-loops. The latter pose a major challenge: For instance, reasoning about probabilistic termination is provably harder than reasoning about classical termination [37].

Guaranteed Bounds on Expected Outcomes via Weakest Pre-Expectations. In recent years, there has been considerable interest in analyzing expressive PP in a mathematically rigorous manner – with *hard* rather than statistical guarantees – see e.g., [2, 15, 26, 27, 43] as representative examples. Such stringent analyses are motivated by the fact that more conventional statistical techniques can be far off even for simple programs [11], and may be inappropriate for safety-critical applications.

The *weakest pre-expectation* (wp) *calculus* is a prominent means to specify and establish a broad spectrum of PP properties in a rigorous manner [35, 39, 41]. Generalizing Dijkstra's classical weakest pre-conditions, the central objects are so-called *expectations* f – random variables over a program's state space mapping program states to numbers – which take over the role of predicates from classical program verification. More precisely, given a probabilistic program C , an expectation f , and an initial program state σ , we have

$$\text{wp}\llbracket C \rrbracket (f) (\sigma) = \begin{array}{l} \text{expected value of } f \text{ w.r.t. the distribution of } \textit{final} \text{ states} \\ \text{reached after executing } C \text{ on } \textit{initial} \text{ state } \sigma \\ \text{and aborting all executions violating an observation.} \end{array}$$

Hence, $\text{wp}\llbracket C \rrbracket (f)$ is a map from *initial* program states to *expected final* values of f . The weakest *liberal* pre-expectation $\text{wlp}\llbracket C \rrbracket (f)$ adds to the above quantity the probability of C 's divergence¹. Both wp and wlp can be defined by induction on the structure of C . Notice that neither of these calculi yield *conditional* expected values as one might expect in the presence of conditioning. One can, however, define *conditional weakest pre-expectations* [46, 47] as²

$$\text{cwp}\llbracket C \rrbracket (f) (\sigma) = \frac{\text{wp}\llbracket C \rrbracket (f) (\sigma)}{\text{wlp}\llbracket C \rrbracket (1) (\sigma)} = \frac{\text{expected value of } f \dots}{\text{probability of not violating an observation}},$$

which indeed yields the sought-after conditional expected values. We refer to quantities such as $\text{wp}\llbracket C \rrbracket (f) (\sigma)$, $\text{wlp}\llbracket C \rrbracket (f) (\sigma)$, and $\text{cwp}\llbracket C \rrbracket (f) (\sigma)$ as (*conditional*) *expected outcomes* of probabilistic programs. Reasoning about expected outcomes of loops is typically tackled by means of *quantitative loop invariants*, which are, naturally, often hard to find. However, in the presence of continuous sampling, even *verifying* a *given* candidate loop invariant poses severe challenges as reasoning about the required expected values involves possibly complex *integrals*. The aim of this paper is to lay the foundations for automated techniques tackling these challenges.

Problem Statement. We study *imperative* PP with all three abilities (1), (2), and (3). Our goal is to *semi-automatically verify bounds on (conditional) expected outcomes of programs featuring continuous uniform sampling, unbounded while-loops with user-provided quantitative invariants, and conditioning.*

¹Provided f is upper-bounded by 1. See Section 4.2 for details.

²Notice that this quantity is undefined if $\text{wlp}\llbracket C \rrbracket (1) (\sigma) = 0$ as is natural when conditioning on a probability-0-event.

“Semi-automatic” means that we focus on verification of *user-provided* quantitative loop invariants in the sense of [41]. Unlike fully automatic approaches we do not *synthesize* such invariants automatically, which is a promising direction for future work but outside the scope of this paper.

Approach. Our key idea is simple, yet powerful:

We replace the *integrals* occurring in the definition of the programs’ exact weakest pre-expectations by simpler *sound approximations*, namely lower and upper *Riemann sums*.

This results in a family of approximate *lower* and *upper Riemann wp transformers*, denoted by $\underline{\text{wp}}^N$ and $\overline{\text{wp}}^N$, where N is the number of sub-intervals used to discretize the domain of integration. Crucially, these Riemann expectation transformers ultimately give rise to automated SMT-based techniques for verifying quantitative loop invariants. Let us now consider an introductory example.

Illustrative Example: A Monte Carlo Approximator. The program shown in Figure 1 draws M (x, y) -samples uniformly at random from the unit square. Each time a sample lands in the quarter unit circle, the variable count is incremented. The program hence approximates the transcendental number $\frac{\pi}{4} \approx 0.785$ (the area of the quarter unit circle) in a Monte Carlo manner.

We will first focus on the *blue* fragment C_{inner} of the program from Figure 1. C_{inner} contains two uniform continuous sampling instructions. Suppose we aim to compute the expected value of count after executing C_{inner} . Applying the rules for determining wp’s [53], we obtain the following double Lebesgue integral, where $[\dots]$ denotes the indicator function of the enclosed predicate:

$$\text{wp}[\![C_{\text{inner}}]\!] (\text{count}) = \underbrace{\text{count} + \int_0^1 \int_0^1 [x^2 + y^2 \leq 1] d\lambda(x), d\lambda(y)}_{\text{expected final value of count in terms of its initial value}} = \text{count} + \frac{\pi}{4}. \quad (\dagger)$$

This reflects the fact that, in *expectation*, C_{inner} increments count by $\frac{\pi}{4}$.

Symbolic integration-based tools such as PSI [26] can solve the above integral directly. However, we observe two significant issues: (i) there are integrals that cannot be evaluated symbolically in closed form, and (ii) even if all guards and assignments in the program are polynomial and distributions are restricted to uniform distributions in $[0, 1]$, the resulting integrals may evaluate to transcendental numbers, making automation notoriously difficult.

As outlined, instead of solving integrals exactly, our strategy is to soundly *under-* or *over-approximate* them by lower or upper Riemann sums. In our example, the resulting upper sum

$$\frac{1}{N^2} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} \sup_{\xi \in [\frac{i}{N}, \frac{i+1}{N}]} \sup_{\zeta \in [\frac{j}{N}, \frac{j+1}{N}]} [\xi^2 + \zeta^2 \leq 1],$$

where $N \geq 1$, is a guaranteed upper bound on the double integral in (\dagger) . For instance, with $N = 16$, we can verify that the upper sum is at most 0.85, and hence that $\text{wp}[\![C_{\text{inner}}]\!] (\text{count}) (\sigma) \leq \sigma(\text{count}) + 0.85$ for all initial states σ . Crucially, to prove upper bounds on upper Riemann sums, we do *not* have to evaluate any suprema explicitly because for all $A \subseteq \mathbb{R}$ and $b \in \mathbb{R}$,

$$\sup A \leq b \quad \text{iff} \quad \forall a \in A: a \leq b.$$

In practice, we can thus drop the sup’s in upper Riemann sums (and, dually, the inf’s in lower sums) by introducing \forall -quantifiers, which is highly beneficial in the context of SMT-based automation.

We stress that even though we discretize the integrals’ domains uniformly, our approach is *not* the same as replacing the continuous $\text{unif}_{[0,1]}$ distributions by discrete uniform distributions of N point-masses. Indeed, the latter would neither yield under- nor over-approximations in general, which is, however, essential for invariant verification.

```

i := 1; count := 0;
while (i ≤ M) {
  x ≈ unif[0,1]; y ≈ unif[0,1];
  if(x2 + y2 ≤ 1) {count := count + 1} else {skip};
  i := i + 1 }

```

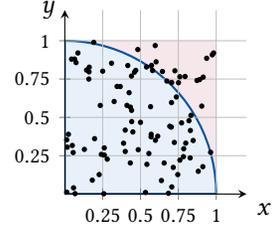


Fig. 1. Left: Monte-Carlo approximator C for π using ability (2). Right: 100 random samples.

Loops and Invariants. Now, consider the entire program C in Figure 1. We have $\text{wp}\llbracket C \rrbracket(\text{count}) = \frac{\pi}{4} \cdot M$ because, intuitively, the percentage of samples landing in the quarter circle equals its area. Our techniques enable to soundly bound $\text{wp}\llbracket C \rrbracket(\text{count})$ in a semi-automated manner for arbitrary initial values of M . This is done, as is standard in deductive verification, using quantitative loop invariants, which we detail in Section 6. Suffices to say here that, given a suitable invariant I , our techniques enable to verify its validity automatically. This in turn allows us to conclude that:

$$\forall \text{ initial states } \sigma: \quad \text{wp}\llbracket C \rrbracket(\text{count})(\sigma) \leq 0.85 \cdot \sigma(M).$$

See Example 6.3 (page 15) and Example 9.1 (page 22) for details.

Contributions. In summary, this paper introduces the novel *lower* and *upper Riemann weakest pre-expectation transformers* and demonstrates their applicability to the semi-automated verification of PP with *continuous uniform sampling* and general while-loops. Put more concretely:

Verification of Invariants We show that externally provided loop invariant candidates can be verified using the Riemann wp transformers in a way that is suitable for SMT-based automation. Such invariants, once verified, imply (one-sided) bounds on wp, wlp, and cwp.

Convergence and Complexity On the more theoretical side, our approximate Riemann wp's are shown to *converge* to the exact wp's under mild assumptions as the discretization becomes finer. As a consequence, we obtain *coRE-completeness*³ results for the problems of verifying upper bounds on wp and lower bounds on wlp for fairly general loopy programs.

Implementation and Case Studies We incorporate our method in the PP verification infrastructure CAESAR [52] by encoding Riemann sums in its *intermediate verification language*. This enables transferring principles from classical SMT-based program verification [44], such as custom first-order theories for reasoning about exponentials, to the verification of PPs involving continuous sampling. We provide various case studies and an empirical evaluation.

Paper Structure. Section 2 provides a bird's eye view on the distinctive features and technicalities of our method. Section 3 introduces basic fixed point theory and Section 4 summarizes the existing definitions of wp, wlp, and cwp. In Section 5, we introduce our Riemann wp transformers and prove them sound. Section 6 is devoted to loop invariants. The convergence results are presented in Section 7. Towards automation, Section 8 introduces a concrete, effective syntax for expectations. Section 9 details the implementation in CAESAR and presents case studies. Additional related work is surveyed in Section 10. We conclude in Section 11. An extended version of this article containing the omitted proofs and additional material is available [10].

³This means that the refutation of such bounds is semi-decidable.

2 A Bird's Eye View

In this section, we explain what makes our method unique by providing a compressed overview of its characteristic features. We also discuss its limitations and point out a subtle technical challenge.

2.1 Highlights and Comparison to Other Approaches

Since we are not the first to address the problem of proving bounds on quantities expressible as $\text{wp}(\text{lp})$'s, we now highlight five distinctive features of our approach – their combination is unique in the literature. A more in-depth review of related work is deferred to Section 10.

Of the following five items, the first two are specific to the way we handle integrals via Riemann sums, whereas the latter three tend to apply to wp -based approaches in general.

An Alternative to Moment-Based Analyses. A major thread of existing work, e.g., [2, 16, 43], circumvents explicit integration altogether. They achieve this by focusing on classes of programs and properties where, simply speaking, the analysis can soundly replace the distributions in the program by their *mean* [1, Appendix A, p. 31] or some higher moment [43]. For example, $\text{wp}[\![x \approx \text{unif}_{[0,1]}\!] \!](xy)$ is equal to⁴ $\text{wp}[\![x := \frac{1}{2}]\!] (xy) = \frac{1}{2}y$, where $\frac{1}{2}$ is the mean of $\text{unif}[0, 1]$. However, such mean-based analyses do not always suffice. Indeed, reconsidering the example from Section 1, the inequality $\text{wp}[\![C_{\text{inner}}]\!] (\text{count}) \sqsubseteq \text{count} + 0.85$, proved correct with our approach, becomes false if we substitute $x \approx \text{unif}_{[0,1]}$ and $y \approx \text{unif}_{[0,1]}$ in C_{inner} by $x := \frac{1}{2}$ and $y := \frac{1}{2}$, respectively. In fact, the so-obtained program C_{inner}' satisfies $\text{wp}[\![C_{\text{inner}}']\!] (\text{count}) = \text{count} + 1$ as the point $(\frac{1}{2}, \frac{1}{2})$ is *certainly inside* the quarter unit circle.

General Conditional Branching. We make relatively mild assumptions regarding the shape of `if`- and `while`-guards in the programs – in practice, general Boolean combinations of the polynomial (in)equalities over all program variables are supported, as in Figure 1. This is different from, e.g., [43] which restricts to finitely-valued variables in guards or [16, 18] which restrict to linear guards. We achieve this high level of generality thanks to the simplicity of the Riemann approximation which treats integration over the resulting indicator functions in a uniform manner.

Moving Backwards: Verification and Parametric Models. We conduct a *backward analysis*: Given a random variable f over final program states (called *post-expectation* in this paper), we approximate the (conditional) weakest pre-expectation $\text{wp}[\![C]\!] (f)$, which is *a function of the initial state*. This is dual to works like [11, 26, 58] that conduct a *forward analysis* to solve a classic Bayesian inference problem, namely to compute a program's posterior density – *a function of the final state* – for a given prior distribution. Let us contrast these two paradigms in greater detail.

Backward analysis, as done in this paper, allows verifying program properties that hold for *all initial states*. For instance, our framework supports questions such as, “*Is the posterior mean of x at most twice the initial value of y ?*”, expressed as $\text{cwp}[C](x) \leq 2y$. The backward approach is thus well-suited for tasks such as *verification* – ensuring a program behaves as intended *for all inputs* – and *parametric analysis* – examining how expected behavior changes when program parameters are altered (parameters can be modeled as uninitialized program variables).

Forward analyses, on the other hand, are typically motivated by probabilistic programming languages (PPLs) explicitly designed for automating Bayesian inference, see e.g., [12, 14, 28, 54]. We remark that there are some fundamental discrepancies between these PPLs and the one studied in this paper. For example, the former usually offer extensive support for general continuous distributions, but do not always support general loops. Moreover, Bayesian inference problems

⁴In general, replacing $x \approx \text{unif}_{[0,1]}$ by $x := \frac{1}{2}$ is sound whenever the post-expectation is linear in x . While our approach is compatible with such optimizations, we shall not discuss them any further.

often require *soft conditioning*, which our approach only encodes as syntactic sugar (see Remark 2). Finally, some applications of Bayesian inference involve loading and processing datasets with thousands of observations. In principle, such data could be hardcoded into our programs, but the result would likely be too large for current SMT-based analysis techniques (as indicated by our experiments in Section 9.3).

In summary, our backward approach is particularly suited for proving properties of complex stochastic processes described as probabilistic programs, especially those involving unbounded loops. However, it is less effective for data-intensive statistical analysis.

A Zoo of Quantities. The majority of existing methods for (semi-)automated exact analysis of PP with loops and continuous distributions focus on one of these tasks: (i) bound the posterior distribution [11, 26, 58], (ii) prove various flavors of (non-)termination, e.g., [2, 15, 17], and (iii) bound assertion violation probabilities [15, 56]. An exception is [57] which tackles *cost analysis*. In contrast, in the wp-based framework, we can express and reason about a remarkably large variety of quantities, including the following (assume that C does not contain conditioning for simplicity):

- Termination probabilities: $\text{wp}\llbracket C \rrbracket (1)$
- The probability of terminating in a predicate⁵ φ : $\text{wp}\llbracket C \rrbracket (\llbracket \varphi \rrbracket)$
- The expected value of variable x on termination: $\text{wp}\llbracket C \rrbracket (x)$
- Higher moments of x after termination: $\text{wp}\llbracket C \rrbracket (x^k)$ for $k \geq 1$
- Expected distance between variables x and y after termination: $\text{wp}\llbracket C \rrbracket (|x - y|)$

Example 2.1. Let us illustrate the above quantities. Consider the following program C :

$$\text{if}(x = 1) \{ \{y := 0\} [1/2] \{y := 2\} \} \text{ else } \{ \{y := 0\} [4/5] \{y := 3\} \}$$

- $\text{wp}\llbracket C \rrbracket (1) = 1$, i.e., C terminates with probability 1 for each initial state.
- $\text{wp}\llbracket C \rrbracket (\llbracket y = 0 \rrbracket) = 1/2 \cdot [x = 1] + 4/5 \cdot [x \neq 1]$, i.e., if initially $x = 1$, then the probability to terminate in $y = 0$ is $1/2$. For all other initial states, this probability is $4/5$.
- $\text{wp}\llbracket C \rrbracket (y) = 1 \cdot [x = 1] + 3/5 \cdot [x \neq 1]$, i.e., the expected final value of y is either 1 or $3/5$, depending on whether $x = 1$ holds initially or not.
- $\text{wp}\llbracket C \rrbracket (y^2) = 2 \cdot [x = 1] + 9/5 \cdot [x \neq 1]$, i.e., the second moment of y on termination is either 2 or $9/5$, depending on whether $x = 1$ holds initially or not.
- $\text{wp}\llbracket C \rrbracket (|x - y|) = [x = 1] \cdot (1/2 \cdot |x| + 1/2 \cdot |x - 2|) + [x \neq 1] \cdot (4/5 \cdot |x| + 1/5 \cdot |x - 3|)$, which is the expected difference between the final values of x and y in terms of their initial values.

The above weakest pre-expectations can be determined by applying the rules in Table 1 (page 10).

Conditioning, Everywhere. Following [47], we allow conditioning in the form of (hard) observe-statements at arbitrary places in the program – even inside loops. All of the above quantities can be generalized sensibly to the conditional setting. We can go further: For instance, the probability to never violate an observe, even if the program does not terminate, is given by $\text{wlp}\llbracket C \rrbracket (1)$.

2.2 Limitations and Assumptions

As mentioned, to analyze loops, we assume *user-provided invariants*. Moreover, we restrict to native support for continuous *uniform* distributions and discrete coin flips. Further, while we support two-sided bounds for loop-free programs, the kind of invariants we study in this paper can only prove *upper* bounds on wp and cwp, and *lower* bounds on wlp. We conjecture that our approach extends to invariant-based methods for the other directions [31] as well, but leave the details for future work. To keep the presentation simple, we only discuss invariant verification for *non-nested loops*. The case of general loop structures, including sequential and/or nested, can be dealt with as

⁵Recall that $\llbracket \varphi \rrbracket$ is the indicator function of the predicate φ .

in [6]. We assume non-negative real-valued program variables and do not support data structures. Soft-conditioning is only supported indirectly. See Remarks 1 to 3 for more details.

2.3 From Lebesgue Integrals to Riemann Sums and Back: A Technical Challenge

The Lebesgue integral is *the* standard integral in probability theory. There is good reason for this: Lebesgue integrals can be defined for a broader class of functions than, say, Riemann integrals, and have favorable mathematical properties like Monotone Convergence Theorems.

Therefore, to define the semantics of PPs, most works (e.g., [23, 53]) have resorted to Lebesgue integrals, and we do not question that this is the way to go. To illustrate the usefulness of the Lebesgue integral as opposed to the Riemann integral⁶, consider the following program D :

```
x := 0; y := 1;
while (y · z ≠ x) { y := y + 1; x := 0; while (y · z ≠ x ∧ x < y) { x := x + 1 } }
```

This program “searches” for non-negative integers x and y such that $z = \frac{x}{y}$ and stops once it finds such x and y . It thus terminates if and only if z is a *rational number* in $[0, 1]$ initially (recall that we allow actual *real-valued* variables). In symbols, we have $\text{wp}[D](1) = [z \in \mathbb{Q} \wedge 0 \leq z \leq 1]$ — this is the well-known *Dirichlet function* restricted to the unit interval. Now, consider the program $z := \text{unif}_{[0,1]}; D$. Its termination probability is zero, the Lebesgue integral of the Dirichlet function over the interval $[0, 1]$. Intuitively, this is because almost all real numbers are irrational, i.e., the rational numbers have Lebesgue measure zero. However, the Dirichlet function is *not Riemann-integrable* — it is “too discontinuous”. As a consequence, one cannot easily define wp — neither in general nor in this specific example — relying on Riemann integrals.

In light of the above example, the following question arises naturally:

How sensible is an analysis of loops based on Riemann sums given the fact that wp 's of loops are not Riemann-integrable in general, not even for programs using only polynomial arithmetic and constant post-expectations?

This question has at least two dimensions.

(1) Regarding *soundness*, we prove that our Riemann wp 's are *always* sound under- or over-approximations, i.e., the following inequalities hold for all $N \geq 1$ in a very general setting:

$$\forall \text{ initial states } \sigma: \quad \underline{\text{wp}}^N[[C]](f)(\sigma) \leq \text{wp}[[C]](f)(\sigma) \leq \overline{\text{wp}}^N[[C]](f)(\sigma)$$

and similarly for wlp . This works because the *lower Riemann integral* — the limit of the lower Riemann sums as the discretization becomes finer — is a lower bound on the Lebesgue integral, and similarly for the upper Riemann integral. This is always true, even if the lower and upper integral are different, i.e., if the function at hand is not Riemann-integrable.

(2) Will the Riemann approximation always converge (in some sense) to the exact wp as the discretization becomes finer? We prove that, under mild assumptions such as polynomial arithmetic in the program, the answer is *yes* for loop-free programs. As one of our theoretical main results for loopy programs, we show that (Theorem 7.6)

$$\sup_{n \geq 1} \underline{\text{wp}}^n[[C^n]](f) = \text{wp}[[C]](f) \quad \text{but} \quad \inf_{n \geq 1} \overline{\text{wp}}^n[[C^n]](f) \stackrel{\text{in general}}{\neq} \text{wp}[[C]](f), \quad (\ddagger)$$

where C^n arises from C by unfolding all loops up to depth n by taking a simultaneous limit of the unfolding depth and the fineness of the lower Riemann sum approximation (the n in $\underline{\text{wp}}^n$). Via equation (\ddagger) we *recover* the exact wp — defined in terms of *Lebesgue* integrals — as a limit of our approximate wp 's based on lower Riemann sums for a general class of probabilistic loops.

⁶The Riemann integral is the limit of both the lower and upper Riemann sums as $N \rightarrow \infty$, provided the two limits coincide.

3 Preliminaries

In this section, we set up our notations and treat the fixed point-theoretic foundations we rely on.

3.1 General Notation

\mathbb{N} is the set of non-negative integers. The set of non-negative *extended reals* is $\overline{\mathbb{R}}_{\geq 0} = \mathbb{R}_{\geq 0} \cup \{\infty\}$. We adopt the following standard conventions: For all $x \in \overline{\mathbb{R}}_{\geq 0}$, we let $x \leq \infty$ and $x + \infty = \infty + x = \infty$. Moreover, we define $\infty \cdot 0 = 0 \cdot \infty = 0$ and $x \cdot \infty = \infty \cdot x = \infty$ for all $x > 0$. Given $a, b \in \mathbb{R}$, we denote by $[a, b]$ the real closed interval with endpoints a and b . The set of truth values is $\mathbb{B} = \{\text{false}, \text{true}\}$. Given a predicate $\varphi: A \rightarrow \mathbb{B}$ over a set A , we define the *Iverson bracket*

$$[\varphi]: A \rightarrow \{0, 1\}, a \mapsto \begin{cases} 0 & \text{if } \varphi(a) = \text{false}, \\ 1 & \text{if } \varphi(a) = \text{true}. \end{cases}$$

For example, assuming that it is understood from context that x is a real number, then $[x \in \mathbb{Q}]$ denotes the *Dirichlet function* that sends every $x \in \mathbb{R}$ to 1 if x is rational, and to 0 otherwise.

Lambda notation $\lambda x.E$, where E is some mathematical expression with free variable x , is used to introduce unnamed functions whose domain and codomain will be clear from the context.

3.2 Fixed Point Theory

We introduce the foundations from fixed point theory [59, Section 5.4] required to sensibly reason about the semantics of loops and to obtain suitable notions of quantitative loop invariants.

Let (L, \preceq) and (M, \preceq) be partial orders. A function $f: L \rightarrow M$ is called *monotonic*, if for all $a \preceq b$, we have $f(a) \preceq f(b)$. An ω -chain in (L, \preceq) is a monotonic function $a: \mathbb{N} \rightarrow L$ (where \mathbb{N} is ordered by the usual \leq relation), i.e., an ω -chain is a non-decreasing sequence of elements $a(0) \preceq a(1) \preceq \dots$ from L . The partial order (L, \preceq) is an ω -complete partial order (ω -cpo, for short), if all ω -chains a in L have a *supremum* (least upper bound) $\sup_{i \in \mathbb{N}} a(i)$ in (L, \preceq) . An ω -cpo (L, \preceq) with an element $\perp \in L$ satisfying $\perp \preceq a$ for all a is called ω -cpo with *bottom*. A function $f: L \rightarrow M$ between ω -cpo's (L, \preceq) and (M, \preceq) is called ω -continuous, if for all ω -chains a in L it holds that $\sup_{i \in \mathbb{N}} f(a(i)) = f(\sup_{i \in \mathbb{N}} a(i))$. Let $f: L \rightarrow L$ be a function where L is an arbitrary set. An element $a \in L$ satisfying $f(a) = a$ is called a *fixed point* of f . The following is often attributed to Kleene:

THEOREM 3.1 (KLEENE'S FIXED POINT THEOREM [59, Theorem 5.11]). *Let (L, \preceq) be an ω -cpo with bottom and let $f: L \rightarrow L$ be ω -continuous. Then f has a least fixed point $\text{lfp } f \in L$ and $\text{lfp } f = \sup_{i \in \mathbb{N}} f^i(\perp)$ where $f^i(\perp)$ denotes the i -fold application of f to \perp .*

A partial order (L, \preceq) is said to be ω -cocomplete (ω -cocpo) if all ω -cochains, i.e., non-increasing sequences $a(0) \succeq a(1) \succeq \dots$ in L have an infimum $\inf_{i \in \mathbb{N}} a(i) \in L$. An ω -cocpo with a greatest element \top is called ω -cocpo with *top*. Note that (L, \preceq) is an ω -cocpo (with top) iff the reversed partial order (L, \succeq) is an ω -cpo (with bottom). Similarly, a function $f: L \rightarrow M$ between ω -cocpos (L, \preceq) and (M, \preceq) is called ω -cocontinuous, if $\inf_{i \in \mathbb{N}} f(a(i)) = f(\inf_{i \in \mathbb{N}} a(i))$ for all non-increasing a . Note that for an ω -cocpo (L, \preceq) with top, Theorem 3.1 reads as follows: If $f: L \rightarrow L$ is ω -cocontinuous, then f has a *greatest fixed point* $\text{gfp } f = \inf_{i \in \mathbb{N}} f^i(\top)$.

A partial order is called ω -bicomplete (ω -bicpo), if it is both an ω -cpo and ω -cocpo. Similarly, a function f between ω -bicpos is called ω -bicontinuous, if it is both ω -continuous and ω -cocontinuous.

A partial order (L, \preceq) is a *complete lattice* if for all $A \subseteq L$ there exists $\sup A \in L$ and $\inf A \in L$. Note that every complete lattice is an ω -bicpo with bottom $\perp = \sup \emptyset$ and top $\top = \inf \emptyset$.

THEOREM 3.2 (KNASTER-TARSKI THEOREM [59, Theorems 5.15 and 5.16]). *Let $f: L \rightarrow L$ be a monotonic function on the complete lattice (L, \preceq) . Then f has a least and greatest fixed point $\text{lfp } f \in L$*

and $\text{gfp } f \in L$. Moreover, for all $a \in L$ it holds that

$$f(a) \preceq a \quad \text{implies} \quad \text{lfp } f \preceq a \quad \text{and} \quad a \preceq f(a) \quad \text{implies} \quad a \preceq \text{gfp } f.$$

The above implications are often referred to as *Park (co)induction* [48].

4 Weakest Pre-Expectations for Probabilistic Programs

In this section we define programming language `pWhile` and its weakest pre-expectation semantics based on *Lebesgue* integrals as defined in [53].

4.1 Program Syntax

For the rest of the paper we fix a finite⁷ set $V = \{x, y, \dots\}$ of program variables. A (*program*) *state* is a variable valuation $\sigma \in \mathbb{R}_{\geq 0}^V$, where $\mathbb{R}_{\geq 0}^V$ is a shorthand for the set of functions $V \rightarrow \mathbb{R}_{\geq 0}$. Notice that our program variables range over the *non-negative* reals. The restriction to non-negative variables is for technical convenience and not essential, see Remark 3.

To obtain a well-defined weakest pre-expectation semantics of programs involving continuous sampling, we need to have some measure-theoretic fundamentals in mind, provided in [10]. Suffice it to say here that we consider the standard Borel σ -algebra and Lebesgue measure λ on $\mathbb{R}_{\geq 0}^V$. Lebesgue integrals of a measurable function $f: \mathbb{R} \rightarrow \overline{\mathbb{R}}_{\geq 0}$ over a measurable set $A \subseteq \mathbb{R}$ are denoted by $\int_A f d\lambda$, or $\int_A f(x) d\lambda(x)$. We explicitly allow Lebesgue integrals to evaluate to ∞ .

Now let \mathcal{E} be a set of measurable functions of type $\mathbb{R}_{\geq 0}^V \rightarrow \mathbb{R}_{\geq 0}$ and let \mathcal{G} be a set of measurable functions of type $\mathbb{R}_{\geq 0}^V \rightarrow \mathbb{B}$. Elements of \mathcal{E} and \mathcal{G} are called *arithmetic expressions* and *guards*, resp.

Definition 4.1 (Probabilistic Programs). Programs C in the set `pWhile`(\mathcal{E}, \mathcal{G}) of programs with arithmetic expressions from \mathcal{E} and guards from \mathcal{G} adhere to the following grammar:

$C ::=$	<code>skip</code>	(effectless program)
	<code>diverge</code>	(nonterminating program)
	<code>x := E</code>	(assignment; $x \in V, E \in \mathcal{E}$)
	<code>observe(φ)</code>	(conditioning; $\varphi \in \mathcal{G}$)
	<code>x \approx unif_[0,1]</code>	(sample from real interval $[0, 1]$; $x \in V$)
	<code>if(φ) {C} else {C}</code>	(conditional choice; $\varphi \in \mathcal{G}$)
	<code>{C} [p] {C}</code>	(probabilistic choice; $p \in [0, 1] \cap \mathbb{Q}$)
	<code>C; C</code>	(sequential composition)
	<code>while (φ) {C}</code>	(while loop; $\varphi \in \mathcal{G}$)

If \mathcal{E} and \mathcal{G} are the sets of *all* measurable functions of the corresponding type, we write `pWhile` instead of `pWhile`(\mathcal{E}, \mathcal{G}). A program not containing while-loops is called *loop-free*. \triangle

Let us briefly go over each construct, all of which are standard. `skip` does nothing. `diverge` is a non-terminating program, i.e., behaves like `while (true) {skip}`. `x := E` assigns the value of the arithmetic expression E evaluated in the current state to variable x . `x \approx unif[0,1]` assigns to x a value drawn from the continuous uniform $[0, 1]$ -distribution. `observe(φ)` *conditions* the program execution on the guard φ being true. `{C1} [p] {C2}` executes C_1 with probability p , otherwise C_2 . The assumption $p \in [0, 1] \cap \mathbb{Q}$ is to avoid defining a dedicated syntax for probabilities later on in Section 8. `if(φ) {C1} else {C2}`, `C1; C2`, and `while (φ) {C}` are standard conditional choices, sequential compositions, and while-loops, respectively. See [10] for additional remarks.

⁷Once V is fixed we can only write programs with at most $|V|$ distinct variables. However, as we never make any assumptions about the size of V , our theory applies to programs with arbitrarily many variables. Previous work [53] has considered an infinite V , but this requires defining a measure space on the infinite-dimensional $\mathbb{R}_{\geq 0}^V$, which is somewhat more involved.

Table 1. Inductive definition of weakest (liberal) pre-expectations for post-expectation f [53].

C	$\text{wp}\llbracket C \rrbracket (f)$ where $f \in \mathbb{E}_{\text{meas}}$	$\text{wlp}\llbracket C \rrbracket (f)$ where $f \in \mathbb{E}_{\text{meas}}^{\leq 1}$
skip	f	f
diverge	0	1
$x := E$	$f[x/E]$	$f[x/E]$
$x \approx \text{unif}_{[0,1]}$	$\lambda\sigma. \int_{[0,1]} f(\sigma[x \mapsto \xi]) d\lambda(\xi)$	$\lambda\sigma. \int_{[0,1]} f(\sigma[x \mapsto \xi]) d\lambda(\xi)$
observe(φ)	$[\varphi] \cdot f$	$[\varphi] \cdot f$
if(φ) { C_1 } else { C_2 }	$[\varphi] \cdot \text{wp}\llbracket C_1 \rrbracket (f) + [\neg\varphi] \cdot \text{wp}\llbracket C_2 \rrbracket (f)$	$[\varphi] \cdot \text{wlp}\llbracket C_1 \rrbracket (f) + [\neg\varphi] \cdot \text{wlp}\llbracket C_2 \rrbracket (f)$
{ C_1 } [p] { C_2 }	$p \cdot \text{wp}\llbracket C_1 \rrbracket (f) + (1-p) \cdot \text{wp}\llbracket C_2 \rrbracket (f)$	$p \cdot \text{wlp}\llbracket C_1 \rrbracket (f) + (1-p) \cdot \text{wlp}\llbracket C_2 \rrbracket (f)$
$C_1; C_2$	$\text{wp}\llbracket C_1 \rrbracket (\text{wp}\llbracket C_2 \rrbracket (f))$	$\text{wlp}\llbracket C_1 \rrbracket (\text{wlp}\llbracket C_2 \rrbracket (f))$
while (φ) { B }	$\text{lfp } \lambda Y. [\neg\varphi] \cdot f + [\varphi] \cdot \text{wp}\llbracket B \rrbracket (Y)$	$\text{gfp } \lambda Y. [\neg\varphi] \cdot f + [\varphi] \cdot \text{wlp}\llbracket B \rrbracket (Y)$

Remark 1 (More General Distributions). In theory, the restriction to uniform $[0, 1]$ -distributions does not limit expressiveness: Arbitrary distributions can be simulated by sampling from $[0, 1]$ and applying the inverse *cumulative distribution function* (CDF) of the target distribution [53]. However, to obtain decidability results, we further restrict the syntax of arithmetic expressions and guards to a class of functions expressible in first-order (FO) real arithmetic, see Section 8. Distributions whose inverse CDF belongs to this class include triangular, trapezoidal, U-quadratic, and Kumaraswamy distributions. On the other hand, distributions with transcendental inverse CDF (Gaussian, Laplace, etc.) do not reside in this class. We mention two future directions to address this issue: (i) leverage heuristics implemented in modern SMT solvers to discharge the generated verification conditions, even if they do not belong to a decidable theory, (ii) soundly over/under-approximate the transcendental inverse CDF by FO-expressible algebraic functions.

Remark 2 (Soft Conditioning). Our syntax has native support for *hard conditioning* (observe). Bounded *soft conditioning* (scoring) – multiplying the current execution with a weight in $[0, 1]$ – can be simulated using $x \approx \text{unif}_{[0,1]}$ and observe, see [53, Lemma 5] for details.

4.2 Weakest Pre-Expectation Semantics

We now unify the weakest pre-expectation calculi for continuous probabilistic programs from [53] with the calculi proposed in [47]. The latter calculi take *renormalization* – for conditional expected outcomes – into account. The central objects these calculi operate on are *expectations*:

Definition 4.2 (Expectations). We distinguish between the following sets of functions:

- (1) The set of *expectations* is $\mathbb{E} = \{f \mid f: \mathbb{R}_{\geq 0}^V \rightarrow \overline{\mathbb{R}}_{\geq 0}\}$.
- (2) The set of *1-bounded expectations* is $\mathbb{E}^{\leq 1} = \{f \mid f: \mathbb{R}_{\geq 0}^V \rightarrow [0, 1]\}$.
- (3) The set of *measurable (1-bounded) expectations* $\mathbb{E}_{\text{meas}}^{\leq 1}$ ($\mathbb{E}_{\text{meas}}^{\leq 1}$) is the subset of \mathbb{E} ($\mathbb{E}^{\leq 1}$) containing exactly the Borel-measurable functions.

We equip all of these sets with the partial order \sqsubseteq defined as $f \sqsubseteq g$ iff $\forall \sigma \in \mathbb{R}_{\geq 0}^V: f(\sigma) \leq g(\sigma)$. \triangle

Crucially, we have (see [53, Lemma 2] and [10]):

LEMMA 4.3. (\mathbb{E}, \sqsubseteq) and ($\mathbb{E}^{\leq 1}, \sqsubseteq$) are complete lattices. ($\mathbb{E}_{\text{meas}}, \sqsubseteq$) and ($\mathbb{E}_{\text{meas}}^{\leq 1}, \sqsubseteq$) are ω -bicos with bottom ($0 = \lambda\sigma.0$) and top ($\infty = \lambda\sigma.\infty$ and $1 = \lambda\sigma.1$, respectively).

The arithmetic operations $+$ (addition) and \cdot (multiplication) on \mathbb{E} are defined pointwise, i.e., $\forall \sigma \in \mathbb{R}_{\geq 0}^V: (f+g)(\sigma) = f(\sigma) + g(\sigma)$, and analogously for multiplication. These operations preserve

measurability [49, Theorem 11.18], i.e., \mathbb{E}_{meas} is closed under $+$ and \cdot . Moreover, addition (for both arguments) and multiplication by constants⁸ are ω -bicontinuous functions.

For state $\sigma \in \mathbb{R}_{\geq 0}^V$, program variable $x \in V$, $\xi \in \mathbb{R}_{\geq 0}$, we define the *updated state as*

$$\sigma[x \mapsto \xi] = \lambda y. \begin{cases} \xi & \text{if } y = x \\ \sigma(y) & \text{otherwise} \end{cases}.$$

Further, given an expectation $f \in \mathbb{E}$, a variable $x \in V$, and an arithmetic expression $E: \mathbb{R}_{\geq 0}^V \rightarrow \mathbb{R}_{\geq 0}$ we define the *substitution of x by E in f* as the expectation $f[x/E] = \lambda \sigma. f(\sigma[x \mapsto E(\sigma)])$. If we have syntactic expressions for f and E , then we can obtain a syntactic representation of $f[x/E]$ by substituting all free occurrences of x in f by E in a capture-avoiding manner, see Section 8.

Definition 4.4 (Weakest (Liberal) Pre-expectation Transformers [53]). For all $C \in \text{pWhile}$, the *weakest (liberal) pre-expectation transformers* $\text{wp}\llbracket C \rrbracket: \mathbb{E}_{meas} \rightarrow \mathbb{E}_{meas}$ and $\text{wlp}\llbracket C \rrbracket: \mathbb{E}_{meas}^{\leq 1} \rightarrow \mathbb{E}_{meas}^{\leq 1}$ are defined inductively on the structure of C according to the rules in Table 1. \triangle

Let us briefly explain the inductive definition in Table 1. The effectless program `skip` leaves the post-expectation unchanged. An assignment $x := E$ substitutes the expression E for the variable x in the post-expectation f , while uniform assignment $x \approx \text{unif}_{[0,1]}$ integrates the expectation over all possible values of x in the interval $[0, 1]$, capturing the averaging effect of drawing x uniformly. `observe(φ)` scales the expectation by the Iverson bracket of the guard φ . Proper renormalization is performed in a second step, see Section 4.4. For `if(φ) { C_1 } else { C_2 }` and the probabilistic choice $\{C_1\} [p] \{C_2\}$, the weakest (liberal) pre-expectation is a weighted sum of the wp 's from both branches, either weighted by $[\varphi]$ and $[\neg\varphi]$ in the former case, or by the probabilities p and $1-p$ in the latter. The $\text{w}(l)\text{p}$ of a sequential composition $C_1 ; C_2$ is function composition $\text{w}(l)\text{p}\llbracket C_1 \rrbracket \circ \text{w}(l)\text{p}\llbracket C_2 \rrbracket$. Notably, C_2 is evaluated *before* C_1 – $\text{w}(l)\text{p}$'s are thus computed in a backward manner. Note that wp and wlp only differ in the handling of divergence and loops. The `diverge` command, representing non-termination, sets the wp to 0 and the wlp to 1. The `while (φ) { C }` loop involves computing the *least* fixed point (lfp) for wp and the *greatest* fixed point (gfp) for wlp . The difference between wp and wlp is further explained in Section 4.3.

Remark 3 (On Non-Negativity). We follow the classic line of work on weakest pre-expectations [39, 41] and restrict attention to *non-negative expectations* (see [37] for a discussion of the mixed-sign case). This means that we can only reason about expected values of non-negative random variables measured in a program's final state. As a consequence, in order to reason about the expected final value of a program variable x on termination, we have to assume that x is unsigned. We have opted to ensure this by simply requiring *all* variables to be unsigned real numbers.

Given a loop $C = \text{while } (\varphi) \{B\}$ and $f \in \mathbb{E}_{meas}$, we denote by

$$\text{wp}_C^{\text{p}}\Phi_f: \mathbb{E}_{meas} \rightarrow \mathbb{E}_{meas}, \quad \text{wp}_C^{\text{p}}\Phi_f(g) = [\varphi] \cdot \text{wp}\llbracket B \rrbracket(g) + [\neg\varphi] \cdot f$$

the *wp-characteristic function of C w.r.t. f* (analogously for $f' \in \mathbb{E}_{meas}^{\leq 1}$ and $\text{wlp}_C^{\text{p}}\Phi_{f'}$). Hence, weakest pre-expectations of loops can be denoted more concisely as

$$\text{wp}\llbracket C \rrbracket(f) = \text{lfp}_C^{\text{p}}\Phi_f \quad \text{and} \quad \text{wlp}\llbracket C \rrbracket(f') = \text{gfp}_C^{\text{p}}\Phi_{f'}.$$

Due to the fixed points in Table 1 it is not immediately obvious that $\text{wp}\llbracket C \rrbracket$ and $\text{wlp}\llbracket C \rrbracket$ are well-defined. This will be ensured by Kleene's Theorem 3.1 as shown in the next lemma (parts of which have been proved in [53]).

⁸Formally, for every $g \in \mathbb{E}$ these are the functions $\lambda f. f + g$ and $\lambda f. f \cdot g$.

THEOREM 4.5 (WELL-DEFINEDNESS OF wp AND wlp). *For all programs $C \in \text{pWhile}$, $\text{wp}\llbracket C \rrbracket$ and $\text{wlp}\llbracket C \rrbracket$ are well-defined. In particular, $\text{wp}\llbracket C \rrbracket$ is ω -continuous and $\text{wlp}\llbracket C \rrbracket$ is ω -cocontinuous.*

4.3 Probabilistic Termination: wp vs. wlp

In general, for all $C \in \text{pWhile}$ and $f \in \mathbb{E}_{meas}^{\leq 1}$ we have $\text{wp}\llbracket C \rrbracket(f) \sqsubseteq \text{wlp}\llbracket C \rrbracket(f)$ since the former relies on a least and the latter on a greatest fixed point. More specifically, we have [53, Section 5]

$$\text{wp}\llbracket C \rrbracket(f) + \text{wlp}\llbracket C \rrbracket(0) = \text{wlp}\llbracket C \rrbracket(f). \quad (1)$$

We call C *almost-surely terminating* (AST) if $\text{wlp}\llbracket C \rrbracket(0) = 0$, i.e., if the program does not admit any initial state for which the program's infinite runs that do not violate any observe have positive probability mass. It follows from equation (1) that C is AST iff $\text{wp}\llbracket C \rrbracket(f) = \text{wlp}\llbracket C \rrbracket(f)$.

4.4 Conditional Weakest Pre-Expectations

For a program $C \in \text{pWhile}$, an expectation $f \in \mathbb{E}_{meas}$ and a state σ , following [47], we define:

$$\text{cwp}\llbracket C \rrbracket(f)(\sigma) = \begin{cases} \frac{\text{wp}\llbracket C \rrbracket(f)(\sigma)}{\text{wlp}\llbracket C \rrbracket(1)(\sigma)} & \text{if } \text{wlp}\llbracket C \rrbracket(1)(\sigma) \neq 0 \\ \text{undefined} & \text{else.} \end{cases}$$

The above definition factors out the probability mass of runs violating an observe, i.e., divides by $\text{wlp}\llbracket C \rrbracket(1)(\sigma)$, see Section 4.3. $\text{cwp}\llbracket C \rrbracket(f)(\sigma)$ is thus the expected value of f after termination of C started with initial state σ , *conditioned* on all observe statements in C being successful.

5 Approximate Riemann Weakest Pre-Expectations

We start by recalling lower and upper Riemann sums and integrals.⁹ Let $[a, b] \subseteq \mathbb{R}$ and $N \geq 1$. A *partition* of $[a, b]$ is a tuple of at least $N + 1$ real numbers $P = (x_0, x_1, \dots, x_N)$ such that

$$a = x_0 < x_1 < \dots < x_N = b.$$

The set of all partitions of the interval $[a, b]$ is denoted $\mathfrak{P}[a, b]$. For a bounded $f: [a, b] \rightarrow \mathbb{R}$ and a partition $P = (x_0, \dots, x_N) \in \mathfrak{P}[a, b]$, we define the *lower* and *upper sums* of f w.r.t. P as

$$L_{f,P} = \sum_{i=1}^N (x_i - x_{i-1}) \inf_{\xi \in [x_{i-1}, x_i]} f(\xi) \quad \text{and} \quad U_{f,P} = \sum_{i=1}^N (x_i - x_{i-1}) \sup_{\xi \in [x_{i-1}, x_i]} f(\xi).$$

Note that $U_{f,P}$ and $L_{f,P}$ are well-defined real numbers because f is bounded.

Definition 5.1 (Riemann integral). Let $f: [a, b] \rightarrow \mathbb{R}$ be bounded. The *lower-* and *upper Riemann integrals* of f are defined as follows:

$$\int_a^b f(x) dx = \sup \{L_{f,P} \mid P \in \mathfrak{P}[a, b]\} \quad \text{and} \quad \overline{\int_a^b f(x) dx} = \inf \{U_{f,P} \mid P \in \mathfrak{P}[a, b]\}.$$

If the upper integral equals the lower integral, then the common value is written $\int_a^b f(x) dx$ and called *the Riemann integral* of f . In this case, f is called *Riemann-integrable*. \triangle

⁹The definition of an integral in terms of these sums is in fact commonly attributed to Darboux, not to Riemann. However, Darboux's integral is equivalent to Riemann's which, rather than considering lower and upper sums, relies on evaluating f at sample points within the intervals of a partition of the integration domain. We refer to [13, Chapter 3, Theorems 3.3.1 and 3.3.2] for an in-depth comparison. In this paper, we consistently use Darboux's definitions, but refer to them nonetheless as Riemann sums and integrals, since the latter terminology is more widespread.

Finally, we introduce the following terminology: (i) For $g: D \rightarrow \mathbb{R}$, $D \subseteq \mathbb{R}$, we say that g is Riemann-integrable on an interval $[c, d] \subset D$ if the restriction $g: [c, d] \rightarrow \mathbb{R}$ is Riemann-integrable. (ii) In this paper we often consider functions of the form $g: D^V \rightarrow \mathbb{R}$ where V is a finite set, e.g. the set of program variables. We then say that g is Riemann-integrable on $[c, d] \subseteq D$ w.r.t. some $x \in V$ if for all $\sigma \in D^V$ the function $\lambda\xi.g(\sigma[x \mapsto \xi])$ of type $[c, d] \rightarrow \mathbb{R}$ is Riemann-integrable.

5.1 $\underline{\text{wp}}^N$ and $\overline{\text{wp}}^N$: Lower and Upper Riemann Weakest Pre-Expectations

We are now ready to define our approximate expectation transformers. They arise from the standard transformers defined in Table 1 by replacing the *Lebesgue* integrals in the $\text{unif}_{[0,1]}$ case by a lower or upper Riemann sum. Formally:

Definition 5.2 (Lower and Upper Riemann w(l)p-Transformers). For all $C \in \text{pWhile}$ and integers $N \geq 1$, the expectation transformers $\underline{\text{wp}}^N \llbracket C \rrbracket: \mathbb{E} \rightarrow \mathbb{E}$ and $\overline{\text{wp}}^N \llbracket C \rrbracket: \mathbb{E} \rightarrow \mathbb{E}$ are defined by induction over the structure of C as in Table 1, with the only exception that if C is $x \approx \text{unif}_{[0,1]}$, then for all $f \in \mathbb{E}$ we set¹⁰

$$\begin{aligned} \underline{\text{wp}}^N \llbracket x \approx \text{unif}_{[0,1]} \rrbracket (f) &= \frac{1}{N} \sum_{i=0}^{N-1} \inf_{\xi \in [\frac{i}{N}, \frac{i+1}{N}]} f[x/\xi] \quad \text{and, similarly,} \\ \overline{\text{wp}}^N \llbracket x \approx \text{unif}_{[0,1]} \rrbracket (f) &= \frac{1}{N} \sum_{i=0}^{N-1} \sup_{\xi \in [\frac{i}{N}, \frac{i+1}{N}]} f[x/\xi]. \end{aligned}$$

The lower and upper Riemann weakest *liberal* pre-expectation transformers $\underline{\text{wlp}}^N \llbracket C \rrbracket: \mathbb{E}^{\leq 1} \rightarrow \mathbb{E}^{\leq 1}$ and $\overline{\text{wlp}}^N \llbracket C \rrbracket: \mathbb{E}^{\leq 1} \rightarrow \mathbb{E}^{\leq 1}$ are defined analogously. \triangle

Our Riemann transformers approximate the Lebesgue integral in the definition of wp (and wlp) by a lower or upper Riemann sum. We work with the partitions $0 < \frac{1}{N} < \frac{2}{N} < \dots < 1$ of the unit interval for the sake of concreteness. Note that these partitions are *not* successive refinements of each other (see [10] for definitions). Thus, increasing N by 1 does *not necessarily* yield “better” approximations, i.e., we might have $\underline{\text{wp}}^N \llbracket C \rrbracket (f) \not\subseteq \underline{\text{wp}}^{N+1} \llbracket C \rrbracket (f)$.

Given a loop $C = \text{while } (\varphi) \{B\}$, $N \geq 1$, $f \in \mathbb{E}$, and $\mathcal{T} \in \{\overline{\text{wp}}^N, \underline{\text{wp}}^N\}$, we define the \mathcal{T} -characteristic function of C w.r.t. f as

$$\mathcal{T}_C \Phi_f: \mathbb{E} \rightarrow \mathbb{E}, \quad \mathcal{T}_C \Phi_f(g) = [\varphi] \cdot \mathcal{T} \llbracket B \rrbracket (g) + [\neg\varphi] \cdot f.$$

For $g \in \mathbb{E}^{\leq 1}$ and $\mathcal{T} \in \{\overline{\text{wlp}}^N, \underline{\text{wlp}}^N\}$, $\mathcal{T}_C \Phi_g: \mathbb{E}^{\leq 1} \rightarrow \mathbb{E}^{\leq 1}$ is defined analogously. Note that unlike wp and wlp , the approximate transformers are defined on the full sets of expectations \mathbb{E} and $\mathbb{E}^{\leq 1}$, respectively, not just on the measurable ones. All transformers from Definition 5.2 are well-defined: this follows from the Knaster-Tarski Theorem 3.2 using that \mathbb{E} and $\mathbb{E}^{\leq 1}$ are complete lattices and monotonicity of the transformers, see Lemma 5.3 below.

5.2 Healthiness Properties of $\underline{\text{wp}}^N$ and $\overline{\text{wp}}^N$

LEMMA 5.3 (MONOTONICITY OF THE RIEMANN w(l)p-TRANSFORMERS). For all $C \in \text{pWhile}$ and integers $N \geq 1$, the functions $\underline{\text{wp}}^N \llbracket C \rrbracket$, $\overline{\text{wp}}^N \llbracket C \rrbracket$, $\underline{\text{wlp}}^N \llbracket C \rrbracket$, and $\overline{\text{wlp}}^N \llbracket C \rrbracket$ are monotonic w.r.t. \sqsubseteq .

Notably, the Riemann w(l)p-transformers do not possess the same continuity properties as their Lebesgue counterparts from Section 4.2. For instance, due to the presence of infima in the lower Riemann sum, $\underline{\text{wp}}^N \llbracket C \rrbracket$ is *not* ω -continuous in general (see [10] for a counter-example).

¹⁰Notice that $\inf_{\xi \in [a,b]} f[x/\xi]$ is the same as $\lambda\sigma. \inf_{\xi \in [a,b]} f(\sigma[x \mapsto \xi])$.

LEMMA 5.4 (SOUNDNESS OF THE RIEMANN $w(l)p$ -TRANSFORMERS). *For all programs $C \in p\text{While}$, post-expectations $f \in \mathbb{E}_{meas}$, $g \in \mathbb{E}_{meas}^{\leq 1}$, and integers $N \geq 1$,*

$$\begin{aligned} \overline{wp}^N[[C]](f) \sqsubseteq wp[[C]](f) \sqsubseteq \overline{wp}^N[[C]](f) \quad \text{and} \\ \underline{wlp}^N[[C]](g) \sqsubseteq wlp[[C]](g) \sqsubseteq \overline{wlp}^N[[C]](g) . \end{aligned}$$

For *conditional* weakest pre-expectations (Section 4.4), we immediately get the following:

COROLLARY 5.5. *For all programs $C \in p\text{While}$, post-expectations $f \in \mathbb{E}_{meas}$, and integers $N \geq 1$*

$$\frac{\overline{wp}^N[[C]](f)(\sigma)}{\overline{wlp}^N[[C]](1)(\sigma)} \leq cwp[[C]](f)(\sigma) \leq \frac{\overline{wp}^N[[C]](f)(\sigma)}{\underline{wlp}^N[[C]](1)(\sigma)}$$

for all $\sigma \in \mathbb{R}_{\geq 0}^V$ such that none of $\underline{wlp}^N[[C]](1)(\sigma)$, $wlp[[C]](1)(\sigma)$, and $\overline{wlp}^N[[C]](1)(\sigma)$ is zero.

We stress that Lemma 5.4 and Corollary 5.5 hold even for non-Riemann-integrable post-expectations.

6 Invariant-Based Reasoning for Loops

Deductive probabilistic program verification techniques typically bound expected outcomes of loops by means of *quantitative loop invariants*. Intuitively, a quantitative loop invariant I is an expectation whose pre-expectation w.r.t. *one loop iteration* does not increase (or decrease, depending on whether one wants to establish upper or lower bones, see below). While quantitative loop invariant-based reasoning can simplify the verification of loops significantly, the continuous setting poses challenges: Computing the pre-expectation of I w.r.t. a loop's body requires reasoning about possibly complex Lebesgue integrals — involving, e.g., indicator functions of predicates — arising from the continuous sampling instructions. We now develop invariant-based proof rules for our *Riemann pre-expectations* (Definition 5.2), which yield sound bounds on the *Lebesgue pre-expectations* (Definition 4.4). We will see in Sections 8 and 9 that these proof rules give rise to SMT-based techniques for verifying bounds on Lebesgue pre-expectations of loops in a semi-automated fashion.

Our first insight is that — due to $(\mathbb{E}, \sqsubseteq)$ and $(\mathbb{E}^{\leq 1}, \sqsubseteq)$ being complete lattices and the Riemann expectation transformers being monotonic by Lemma 5.3 — bounds on *Riemann pre-expectations* of loops can be established via Park induction (Theorem 3.2):

LEMMA 6.1. *Let $C = \text{while}(\varphi)\{B\} \in p\text{While}$ and $N \geq 1$. Then:*

(1) *For all $f, I \in \mathbb{E}$, we have*
$$\underbrace{\overline{wp}_C^N \Phi_f(I) \sqsubseteq I}_{I \text{ is } \overline{wp}^N\text{-superinvariant of } C \text{ w.r.t. } f} \quad \text{implies} \quad \overline{wp}^N[[C]](f) \sqsubseteq I .$$

(2) *For all $f, I \in \mathbb{E}^{\leq 1}$, we have*
$$\underbrace{I \sqsubseteq \underline{wlp}_C^N \Phi_f(I)}_{I \text{ is } \underline{wlp}^N\text{-subinvariant of } C \text{ w.r.t. } f} \quad \text{implies} \quad I \sqsubseteq \underline{wlp}^N[[C]](f) .$$

More colloquially stated, *superinvariants* yield *upper* bounds on *upper* Riemann pre-expectations of loops and, dually, *subinvariants* yield *lower* bounds on *lower liberal* Riemann pre-expectations. Notice that establishing the premise of the above proof rules only requires reasoning about the loop's body and avoids explicitly computing Lebesgue integrals.

It thus follows from the soundness of our Riemann expectation transformers (Lemma 5.4) that the above proof rules yield sound bounds on *Lebesgue pre-expectations*.

THEOREM 6.2. *Let $C = \text{while}(\varphi)\{B\} \in p\text{While}$ and $N, N' \geq 1$. We have:*

(1) *If $I \in \mathbb{E}$ is a \overline{wp}^N -superinvariant of C w.r.t. $f \in \mathbb{E}_{meas}$, then $wp[[C]](f) \sqsubseteq I$.*

- (2) If $I \in \mathbb{E}^{\leq 1}$ is a wlp^N -subinvariant of C w.r.t. $f \in \mathbb{E}_{meas}^{\leq 1}$, then $I \sqsubseteq \text{wlp}\llbracket C \rrbracket(f)$.
- (3) If $I \in \mathbb{E}$ is a $\overline{\text{wp}}^N$ -superinvariant of C w.r.t. $f \in \mathbb{E}_{meas}$ and $J \in \mathbb{E}^{\leq 1}$ is a $\text{wlp}^{N'}$ -subinvariant of C w.r.t. 1, then, for all $\sigma \in \mathbb{R}_{\geq 0}^V$, $J(\sigma) > 0$ implies that $\text{cwp}\llbracket C \rrbracket(f)(\sigma)$ is defined and

$$\text{cwp}\llbracket C \rrbracket(f)(\sigma) \leq \frac{I(\sigma)}{J(\sigma)}.$$

It is important to note that in Lemma 6.1 we admit *arbitrary* (1-bounded) post-expectations whereas in Theorem 6.2 we have to restrict to *measurable* (1-bounded) post-expectations in order for the Lebesgue pre-expectations $\text{wp}\llbracket C \rrbracket(f)$ and $\text{wlp}\llbracket C \rrbracket(f)$ to be well-defined. The quantitative loop invariant I , on the other hand, must *not* necessarily be measurable since there is no need to plug I into a Lebesgue expectation transformer.

Example 6.3. Consider the Monte Carlo π -approximator C from Figure 1. The expectation

$$I = \text{count} + [\text{i} \leq M] \cdot (0.85 \cdot ((M \dot{-} \text{i}) + 1))$$

is a $\overline{\text{wp}}^{16}$ -superinvariant of C w.r.t. count (here $\dot{-}$ denote the *monus* operator defined as $t_1 \dot{-} t_2 = \max(t_1 - t_2, 0)$, detailed in Section 8.1). Hence, we get by Theorem 6.2, $\text{wp}\llbracket C \rrbracket(\text{count}) \sqsubseteq I$, i.e., if initially $\text{count} = 0$, $\text{i} = 1$, then $0.85 \cdot M$ upper-bounds the expected final value of count .

7 Convergence of The Riemann Weakest Pre-Expectations

We now address the question if and under which conditions $\text{wp}^N\llbracket C \rrbracket(f)$ and $\overline{\text{wp}}^N\llbracket C \rrbracket(f)$ converge to $\text{wp}\llbracket C \rrbracket(f)$ as N goes to ∞ . We focus on *pointwise* convergence. It is easy to see that in general, pointwise convergence does not hold: consider for example the simple program $C = x := \text{unif}_{[0,1]}$ and the expectation $f = [x \in \mathbb{Q}]$, for which we have $\text{wp}\llbracket C \rrbracket(f) = 0$. However, for every $N \geq 1$ it holds that $\overline{\text{wp}}^N\llbracket C \rrbracket(f) = 1$. Our goal is thus to identify a *Riemann-suitable* subset of our framework, by mildly restricting the class of allowed expectations.

We first recall two fundamental properties of the Riemann integral. The first is a well-known characterization of Riemann integrability, the second reveals that Lebesgue integrals conservatively generalize Riemann integrals. In the following we will use λ to denote the Lebesgue measure. For the formal definition and properties of λ we refer the reader to [10].

THEOREM 7.1 (RIEMANN-LEBESGUE THEOREM [49, Thm. 11.33]). *Let $f: [a, b] \rightarrow \mathbb{R}$ be bounded. Consider the set $D = \{x \in [a, b] \mid f \text{ is discontinuous in } x\}$. Then f is Riemann-integrable if and only if $\lambda(D) = 0$, i.e., f is continuous almost everywhere.*

THEOREM 7.2 ([4, Theorem 1.7.1]). *Let $f: [a, b] \rightarrow \mathbb{R}$ be bounded and Riemann-integrable. Then $\int_a^b f(x) dx = \int_{[a,b]} f(x) d\lambda(x)$, i.e., the Riemann integral coincides with the Lebesgue integral.*

7.1 Convergence of Approximation for Loop-free Programs

By Theorems 7.1 and 7.2, the Riemann integral coincides with the Lebesgue integral for “sufficiently continuous” functions. There is thus hope to establish convergence if we restrict our setting to almost everywhere continuous expectations. However, there is an additional complication: A function needs to be *bounded* on the integration domain in order to be Riemann-integrable – otherwise the upper sum could be “stuck” at ∞ . We address this issue using the concept of *local boundedness*: A function $f: \mathbb{R}_{\geq 0}^V \rightarrow \mathbb{R} \cup \{\pm\infty\}$ is called *locally bounded* if

$$\text{for all compact } K \subseteq \mathbb{R}_{\geq 0}^V: \quad \sup_{\sigma \in K} |f(\sigma)| < \infty.$$

We recall that a subset K of \mathbb{R}^V is compact iff it is closed and bounded. Many common functions such as polynomials are locally bounded. An example of a function that is *not* locally bounded is $[x > 0] \cdot \frac{1}{x}$. Based on these considerations, we define the following:

Definition 7.3 (Riemann-suitable). A combination $(\mathcal{E}, \mathcal{G}, \mathcal{F})$ of a class \mathcal{E} of locally bounded arithmetic expressions, a class \mathcal{G} of guards, and a class $\mathcal{F} \subseteq \mathbb{E}$ of locally bounded expectations if for all $f \in \mathcal{F}$ and $x \in V$, f is Riemann-integrable on $[0, 1]$ w.r.t. x . Moreover, \mathcal{F} contains the constant expectations $0, 1$ and satisfies the following closure properties:

- (i) \mathcal{F} is closed under taking infima and suprema over closed subintervals of $[0, 1]$, i.e., for all $f \in \mathcal{F}$, $x \in V$, and $[a, b] \subseteq [0, 1]$ we have $\inf_{\xi \in [a, b]} f[x/\xi] \in \mathcal{F}$ and $\sup_{\xi \in [a, b]} f[x/\xi] \in \mathcal{F}$.
- (ii) For all $f \in \mathcal{F}$, $E \in \mathcal{E}$ and $x \in V$, we have $f[x/E] \in \mathcal{F}$.
- (iii) For all $f, g \in \mathcal{F}$ and $p \in [0, 1] \cap \mathbb{Q}$, we have $p \cdot f + (1 - p) \cdot g \in \mathcal{F}$.
- (iv) For all $f, g \in \mathcal{F}$ and $\varphi \in \mathcal{G}$, we have $[\varphi] \cdot f + [\neg\varphi] \cdot g \in \mathcal{F}$. △

The conditions (i) - (iv) in the above definition ensure that \mathcal{F} is closed under $\underline{\text{wp}}^N[[C]]$ and $\overline{\text{wp}}^N[[C]]$ for *loop-free* C (see (2) in Lemma 7.4 below). We remark that it is not immediately clear if any “interesting” Riemann-suitable instantiations exist. Fortunately, the answer is yes, as we show in Section 8. In the Riemann-suitable setting, we can show the following convergence result:

LEMMA 7.4 (CONVERGENCE – LOOP-FREE CASE). *Let $(\mathcal{E}, \mathcal{G}, \mathcal{F})$ be Riemann-suitable. Then for all loop-free $C \in \text{pWhile}(\mathcal{E}, \mathcal{G})$ and post-expectations $f \in \mathcal{F}$ the following holds:*

$$\forall N \geq 1: \quad \underline{\text{wp}}^N[[C]](f) \in \mathcal{F} \quad \text{and} \quad \overline{\text{wp}}^N[[C]](f) \in \mathcal{F}. \quad (2)$$

$$\text{If } f \in \mathbb{E}_{\text{meas}}: \quad \text{wp}[[C]](f) = \sup_{N \geq 1} \underline{\text{wp}}^N[[C]](f) = \inf_{N \geq 1} \overline{\text{wp}}^N[[C]](f). \quad (3)$$

An analogous wlp-version of Lemma 7.4 is stated in [10].

7.2 Convergence of Approximation for Programs with Loops

Recall the *Dirichlet program* D from Section 2. For all N we have $\overline{\text{wp}}^N[[D]](1) = 1$, showing that Lemma 7.4 does not hold for loopy programs in general. The added difficulty stems from the fact that the semantics of loops is itself a limit. Specifically, we can imagine to “unroll” every loop in a program C up to a certain depth $\ell \in \mathbb{N}$, denoted C^ℓ (see [10] for the formal definition), and view the semantics C as the limit of the semantics of C^ℓ when ℓ tends to infinity. Notably, wp and wlp behave differently with respect to this limit: when increasing ℓ the (non-liberal) pre-expectation of f increases, while the liberal pre-expectation decreases. This is formalized in the next lemma:

LEMMA 7.5. *For all $C \in \text{pWhile}$ and post-expectations $f \in \mathbb{E}_{\text{meas}}$ and $g \in \mathbb{E}_{\text{meas}}^{\leq 1}$, $\text{wp}[[C^\ell]](f)$ and $\text{wlp}[[C^\ell]](g)$ are non-decreasing (non-increasing, respectively) sequences in $\ell \in \mathbb{N}$ and it holds that*

$$\sup_{\ell \in \mathbb{N}} \text{wp}[[C^\ell]](f) = \text{wp}[[C]](f) \quad \text{and} \quad \inf_{\ell \in \mathbb{N}} \text{wlp}[[C^\ell]](g) = \text{wlp}[[C]](g).$$

It should now be clear that, when considering the convergence of $\overline{\text{wp}}^N[[C]](f)$ for a program C with loops, we are implicitly considering two limits, the one over N and the one over ℓ . Unfortunately, when the two limits have different monotonic behaviors (i.e., for $\overline{\text{wp}}^N[[C]]$, where the limit in ℓ is the supremum but the limit in N is the infimum, and symmetrically for $\underline{\text{wlp}}^N[[C]]$) convergence cannot be guaranteed. Intuitively, this explains why the next result only holds for $\underline{\text{wp}}^N$ and $\overline{\text{wlp}}^N$.

THEOREM 7.6 (CONVERGENCE – GENERAL CASE). *Let $(\mathcal{E}, \mathcal{G}, \mathcal{F})$ be Riemann-suitable. Then for all $C \in \text{pWhile}(\mathcal{E}, \mathcal{G})$ and post-expectations $f \in \mathcal{F} \cap \mathbb{E}_{\text{meas}}$ and $g \in \mathcal{F} \cap \mathbb{E}_{\text{meas}}^{\leq 1}$ it holds that:*

$$\sup_{n \geq 1} \underline{\text{wp}}^n[[C^n]](f) = \text{wp}[[C]](f) \quad \text{and} \quad \text{wlp}[[C]](g) = \inf_{n \geq 1} \overline{\text{wlp}}^n[[C^n]](g).$$

Note that Theorem 7.6 ensures that when considering a Riemann suitable class, the L.H.S. of the inequality in Corollary 5.5 converges to the exact conditional weakest pre-expectation cwp .

8 Effective Verification of pWhile-Programs

In the previous sections, we have assumed that guards and arithmetical expressions in programs as well as the resulting (pre/post)-expectations are purely mathematical objects. To enable *effective*, i.e., automated, verification we now define, inspired by [8, 52], concrete syntaxes for these objects.

8.1 A Formal Language of Expressions

Let LV be a countably infinite set of logical variables ranged over by x, y, \dots , etc. The sets Terms and Guards of (syntactic) *terms* and *guards* are defined as follows:

$$t ::= q \in \mathbb{Q}_{\geq 0} \mid x \in LV \mid t + t \mid t \dot{-} t \mid t \cdot t \quad \varphi ::= t < t \mid \neg \varphi \mid \varphi \wedge \varphi$$

Note that the other standard comparison relations $\leq, \neq, =, >, \geq$ and Boolean connectives $\vee, \rightarrow, \leftrightarrow$ can be expressed in terms of $<, \neg, \wedge$. We allow further syntactic sugar such as t^2 for $t \cdot t$ and $1 \leq x \leq 2$ for $1 \leq x \wedge x \leq 2$, etc. Additional parentheses are admitted to clarify the order of precedence; to minimize the use of parentheses we assume that \cdot takes precedence over $+$ and $\dot{-}$, and \neg binds stronger than the binary Boolean connectives, as is standard. We let $\text{free}(t)$ and $\text{free}(\varphi)$ be the variables occurring in term t and guard φ .

For every $LV' \supseteq \text{free}(t)$, the semantics $\llbracket t \rrbracket : \mathbb{R}_{\geq 0}^{LV'} \rightarrow \mathbb{R}_{\geq 0}$ of $t \in \text{Terms}$ is standard except that $\dot{-}$ is interpreted as “monus”, i.e., $\llbracket t_1 \dot{-} t_2 \rrbracket = \max(t_1 - t_2, 0)$. The semantics of $\varphi \in \text{Guards}$ can be viewed similarly as a function $\llbracket \varphi \rrbracket : \mathbb{R}_{\geq 0}^{LV'} \rightarrow \mathbb{B}$ for every $LV' \supseteq \text{free}(\varphi)$. For $\sigma \in \mathbb{R}_{\geq 0}^{LV'}$, we write $\sigma \models \varphi$ and $\sigma \not\models \varphi$ to indicate that $\llbracket \varphi \rrbracket(\sigma) = \text{true}$ and $\llbracket \varphi \rrbracket(\sigma) = \text{false}$, respectively.

Note that every $t \in \text{Terms}$ without monus is a non-negative polynomial in the — finitely many — variables $\text{free}(t)$ with rational coefficients, possibly written in (partially) factorized form, whereas terms with monus can be seen as a piecewise defined non-negative polynomial. Similarly, every $\varphi \in \text{Guards}$ is a Boolean combination of (in)equations between such (piecewise) polynomials.

We are now ready to define our expression language similar to [8]. The restriction to infima and suprema over *compact* intervals $[a, b]$ resembles the definition of lower and upper Riemann sums.

Definition 8.1 (The Expression Language Expr). The set of Expr of (syntactic) *expressions* is defined according to the following grammar:

$$f ::= t \mid [\varphi] \cdot f \mid q \cdot f \mid f + f \mid \mathcal{Z}_{[a,b]} x : f \mid \mathcal{L}_{[a,b]} x : f$$

Here, $t \in \text{Terms}$, $\varphi \in \text{Guards}$, $q \in \mathbb{Q}_{\geq 0}$, $a, b \in \mathbb{Q}_{\geq 0}$, $a \leq b$, and $x \in LV$. △

We allow $[\varphi]$ as syntactic sugar for $[\varphi] \cdot 1$ and adopt the usual rules regarding parentheses and orders of precedence. For $f \in \text{Expr}$ we let $\text{free}(f)$ be the variables in f that are not bound by a \mathcal{Z} (supremum) or \mathcal{L} (infimum) “quantifier”. As in standard first-order logic, it is possible for a variable to have both a free and a non-free occurrence in f . We may write $f(x_1, \dots, x_n)$ to indicate that f contains at most the pairwise distinct free variables $x_1, \dots, x_n \in LV$.

Definition 8.2 (Semantics of Expressions). We define the semantics of expressions $\llbracket f \rrbracket : \mathbb{R}_{\geq 0}^{LV'} \rightarrow \mathbb{R}_{\geq 0}$ for every finite $LV' \subseteq LV$ with $\text{free}(f) \subseteq LV'$ inductively as follows. For all $\sigma \in \mathbb{R}_{\geq 0}^{LV'}$:

- $\llbracket t \rrbracket(\sigma)$ is defined in the standard manner, see above.
- $\llbracket [\varphi] \cdot f \rrbracket(\sigma) = \llbracket f \rrbracket(\sigma)$ if $\sigma \models \varphi$; $\llbracket [\varphi] \cdot f \rrbracket(\sigma) = 0$ if $\sigma \not\models \varphi$.
- $\llbracket q \cdot f \rrbracket(\sigma) = q \cdot \llbracket f \rrbracket(\sigma)$.
- $\llbracket f_1 + f_2 \rrbracket(\sigma) = \llbracket f_1 \rrbracket(\sigma) + \llbracket f_2 \rrbracket(\sigma)$
- $\llbracket \mathcal{Z}_{[a,b]} x : f \rrbracket(\sigma) = \sup_{\xi \in [a,b]} \llbracket f \rrbracket(\sigma[x \mapsto \xi])$.

- $\llbracket \mathcal{L}_{[a,b]} x : f \rrbracket (\sigma) = \inf_{\xi \in [a,b]} \llbracket f \rrbracket (\sigma[x \mapsto \xi])$.

In the last two cases, if x is not already in the domain of σ , we tacitly assume that the operation $\sigma[x \mapsto \xi]$ extends the domain of σ by x . \triangle

It is not immediately clear that $\llbracket f \rrbracket$ is well-defined due to the suprema in Definition 8.2 – we have to ensure that those exist in $\mathbb{R}_{\geq 0}$. We show this now.

LEMMA 8.3. *For every $f \in \text{Expr}$ and every finite $LV' \subseteq LV$ with $\text{free}(f) \subseteq LV'$ the semantics $\llbracket f \rrbracket : \mathbb{R}_{\geq 0}^{LV'} \rightarrow \mathbb{R}_{\geq 0}$ is a well-defined locally bounded function (cf. Section 7.1).*

Example 8.4. $f = \mathcal{Z}_{[0,5]} x : [w = x^2] \cdot x$ denotes the function $\llbracket f \rrbracket = \lambda \sigma. [0 \leq \sigma(w) \leq 25] \cdot \sqrt{\sigma(w)}$.

8.2 Properties of Syntactic Expressions

The set FO of first-order formulae¹¹ is defined according to the following grammar:

$$\psi ::= \varphi \mid \exists x : \psi \mid \forall x : \psi \mid \neg \psi \mid \psi \wedge \psi \mid \psi \vee \psi \mid \psi \rightarrow \psi$$

where φ is a Boolean combination of polynomial (in)equalities (similar to the guards defined in Section 8.1 but with standard minus instead of “monus”), and $x \in LV$. As usual, $\text{free}(\psi)$ is the set of free variables in ψ , and we write $\psi(x_1, \dots, x_n)$ to indicate that ψ contains at most the (pairwise distinct) free variables $x_1, \dots, x_n \in LV$. An FO formula ψ is called *quantifier-free* if it does not contain any \exists and \forall quantifiers. Moreover, ψ is said to be in *prenex normal form* (PNF) if $\psi = \mathcal{O}_1 x_1 : \dots \mathcal{O}_n x_n : \psi'$ where $n \geq 0$, $\mathcal{O}_1, \dots, \mathcal{O}_n \in \{\exists, \forall\}$, and ψ' is quantifier-free. Similarly, ψ is in *existential (universal) prenex form* if for all $1 \leq i \leq n$ we have $\mathcal{O}_i = \exists$ ($\mathcal{O}_i = \forall$, respectively).

The semantics of a formula $\psi \in \text{FO}$ is standard and can be viewed as a function $\llbracket \psi \rrbracket : \mathbb{R}^{LV'} \rightarrow \mathbb{B}$ for all $LV' \subseteq LV$ with $\text{free}(\psi) \subseteq LV'$. For $\sigma \in \mathbb{R}^{LV'}$ we write $\sigma \models \psi$ to indicate that σ is a model of ψ , i.e., the *sentence* (formula without free variables) obtained by substituting every free occurrence of $x \in \text{free}(\psi)$ in ψ by $\sigma(x)$ evaluates to true in \mathbb{R} . A formula ψ is called *satisfiable* if it has a model, *unsatisfiable* if it does not have a model, and *valid* if $\neg \psi$ is unsatisfiable.

The decision problem of checking whether a given *quantifier-free* $\psi \in \text{FO}$ is satisfiable is called QF_NRA¹². It is known that QF_NRA \in PSPACE. Checking the satisfiability of general FO formulae is decidable as well, but has higher complexity.

LEMMA 8.5 (Expr to FO). *For every $f(x_1, \dots, x_n) \in \text{Expr}$ there exists an FO formula $\psi_f(x_1, \dots, x_n, y)$ encoding $\llbracket f \rrbracket$ in the sense that for all $\sigma \in \mathbb{R}_{\geq 0}^{\text{free}(f)}$ and $r \in \mathbb{R}_{\geq 0}$, we have that $\sigma, r \models \psi_f$ iff $\llbracket f \rrbracket (\sigma) = r$. For quantifier-free f , we can construct ψ_f in existential prenex form in linear time.*

FO-expressible functions in the sense of Lemma 8.5 have been extensively studied [22]. Let $\text{FO}_{\mathbb{R}}$ be defined like FO with the only difference that *arbitrary real numbers* are allowed as coefficients in the polynomials. A set $A \subseteq \mathbb{R}^n$, $n \geq 1$, is called *semi-algebraic* if there exists an $\text{FO}_{\mathbb{R}}$ formula $\psi(x_1, \dots, x_n)$ such that $A = \{\sigma \in \mathbb{R}^n \mid \sigma \models \psi\}$. A function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is called *semi-algebraic* if its graph $\{(\sigma, f(\sigma)) \mid \sigma \in \mathbb{R}^n\}$ is a semi-algebraic set.

LEMMA 8.6 (e.g., [22, Exercise 2.22]). *Every semi-algebraic function $f : \mathbb{R} \rightarrow \mathbb{R}$ has at most finitely many discontinuities. In particular, f is almost everywhere continuous and hence Riemann-integrable on every interval $[a, b] \subset \mathbb{R}$ where f is bounded.*

By combining local boundedness from Lemma 8.3, the translation to FO from Lemma 8.5, and the fact that semi-algebraic functions are Riemann-integrable (Lemma 8.6), we obtain:

¹¹In this paper, *first-order formula* always refers to first-order logic over the signature of polynomial arithmetic with comparison; we interpret such formulae over the fixed structure \mathbb{R} .

¹²Quantifier-free non-linear real arithmetic, a term coined by the SMT-community [5].

THEOREM 8.7. *For all $f \in \text{Expr}$, intervals $[a, b] \subset \mathbb{R}_{\geq 0}$, finite $LV' \subseteq LV$ with $\text{free}(f) \subseteq LV'$, and variables $x \in LV'$, the function $\llbracket f \rrbracket : \mathbb{R}_{\geq 0}^{LV'} \rightarrow \mathbb{R}_{\geq 0}$ is Riemann-integrable on $[a, b]$ w.r.t. x .*

In analogy to standard FO, we say that an expression $f \in \text{Expr}$ is in *prenex normal form* (PNF) if

$$f = \mathcal{Q}_{[a_1, b_1]}^1 x_1 : \dots \mathcal{Q}_{[a_n, b_n]}^n x_n : f'$$

where $n \geq 0$, $\mathcal{Q}^1, \dots, \mathcal{Q}^n \in \{\mathcal{L}, \mathcal{E}\}$, and f' is quantifier-free.

LEMMA 8.8 (PRENEX NORMAL FORM [8, Lemma 8.1]). *Every $f \in \text{Expr}$ can be transformed into an equivalent $f' \in \text{Expr}$ in PNF with the same free variables. Moreover, if f is \mathcal{L} -free (\mathcal{E} -free), then f' is \mathcal{L} -free (\mathcal{E} -free, respectively) as well. This transformation can be done in linear time.*

Next, we formalize one of our key insights: In order to check a “*quantitative entailment*” $\llbracket f \rrbracket \sqsubseteq \llbracket g \rrbracket$ such that suprema occur only in f and infima only in g , it is *not* actually necessary to evaluate them exactly (which amounts to solving optimization problems). Instead, one can, loosely speaking, replace sup and inf by \forall -quantifiers. This idea is based on the following elementary observation: For arbitrary sets $A, B \subseteq \mathbb{R}$, we have $\sup A \leq \inf B$ iff $\forall a \in A, b \in B: a \leq b$.

LEMMA 8.9 (CHECKING QUANTITATIVE ENTAILMENTS, cf. [52, Section 5.2]). *Let*

$$f = \mathcal{E}_{[a_1, b_1]} x_1 : \dots \mathcal{E}_{[a_n, b_n]} x_n : f' \quad \text{and} \quad g = \mathcal{L}_{[c_1, d_1]} z_1 : \dots \mathcal{L}_{[c_m, d_m]} z_m : g'$$

be \mathcal{L} -free (\mathcal{E} -free, respectively) expressions in PNF. Further, let $\psi_{f'}$ ($\text{free}(f')$, y_f) and $\psi_{g'}$ ($\text{free}(g')$, y_g) be the FO formulae in existential prenex form encoding f' and g' from Lemma 8.5. Then:

$$\begin{aligned} & \llbracket f \rrbracket \sqsubseteq \llbracket g \rrbracket \\ \text{iff} & \left(\bigwedge_{\substack{x \in \text{free}(f) \\ \cup \text{free}(g)}} x \geq 0 \wedge \bigwedge_{i=1}^n a_i \leq x_i \leq b_i \wedge \bigwedge_{i=1}^m c_i \leq z_i \leq d_i \wedge \psi_{f'} \wedge \psi_{g'} \right) \rightarrow y_f \leq y_g \quad \text{is valid} \quad (4) \end{aligned}$$

As a consequence, the decision problem “ $\llbracket f \rrbracket \sqsubseteq \llbracket g \rrbracket$?” for f and g is linear-time reducible to QF_NRA.

Example 8.10. Reconsider the expression $f = \mathcal{E}_{[0,5]} x : [w = x^2] \cdot x$ from Example 8.4. Suppose we wish to check whether $\llbracket f \rrbracket$ is upper-bounded by the constant function $\llbracket g \rrbracket$, $g = 4$. Since f and g have the form required by Lemma 8.9 we can achieve this by considering the FO formula

$$(w \geq 0 \wedge 0 \leq x \leq 5 \wedge (w = x^2 \rightarrow y = x) \wedge (w \neq x^2 \rightarrow y = 0) \wedge y' = 4) \rightarrow y \leq y'$$

which is *not* valid as witnessed by the assignment $\{x \mapsto 5, w \mapsto 25, y \mapsto 5, y' \mapsto 4\}$. Hence $\llbracket f \rrbracket \not\sqsubseteq 4$.

8.3 Decidability and Complexity of pWhile Verification

We now assume (w.l.o.g.) that the set of logical variables LV used in syntactic expressions contains our fixed set V of program variables.

Definition 8.11 (Representable and Syntactic Expectations). We define the following terminology:

- An expectation $f \in \mathbb{E}$ is called *representable* if $f = \llbracket f \rrbracket$ for some $f \in \text{Expr}$ with $\text{free}(f) \subseteq V$. The set of all representable (1-bounded) expectations is denoted $\mathbb{E}_{\text{Expr}}^{\leq 1}$ ($\mathbb{E}_{\text{Expr}}^{\leq 1}$, respectively).
- A *syntactic expectation* is an expression $f \in \text{Expr}$ whose free variables are program variables, i.e., $\text{free}(f) \subseteq V$. A *1-bounded syntactic expectation* $\text{Expr}^{\leq 1}$ is $f \in \text{Expr}$ such that $\llbracket f \rrbracket \in \mathbb{E}^{\leq 1}$. The set of all 1-bounded syntactic expectations is denoted $\text{Expr}^{\leq 1}$. \triangle

Membership in $\text{Expr}^{\leq 1}$ is decidable by encoding $f(x_1, \dots, x_n)$ as the FO formula $\psi_f(x_1, \dots, x_n, y)$ from Lemma 8.5 and checking validity of $(x_1 \geq 0 \wedge \dots \wedge x_n \geq 0 \wedge \psi_f(x_1, \dots, x_n, y)) \rightarrow y \leq 1$.

Our next theorem implies that if we restrict our framework to representable expectations and only allow the Terms and Guards defined in Section 8.1 as the arithmetic and Boolean, respectively, expressions in our programs, then we obtain the convergence guarantees from Section 7.

THEOREM 8.12 (RIEMANN-SUITABILITY OF Expr). *The combination $(\mathbb{E}_{\text{Expr}}, \text{Terms}, \text{Guards})$ with Terms and Guards as defined in Section 8.1 is Riemann-suitable (see Definition 7.3).*

8.3.1 Verifying Loop-free Programs. Since Expr is effectively closed under weakest (liberal) Riemann pre-expectation (formalized in [10]), it follows that we can effectively prove upper and lower bounds on weakest (liberal) pre-expectations of loop-free programs:

THEOREM 8.13 (BOUNDS ON LOOP-FREE $w(l)p$). *For all loop-free $C \in \text{pWhile}(\text{Terms}, \text{Guards})$ and integers $N \geq 1$ the following decision problems can be effectively translated¹³ to QF_NRA: Given ...*

- (1) ... an \mathcal{O} -free $f \in \text{Expr}$ and an \mathcal{L} -free $g \in \text{Expr}$, is $\llbracket g \rrbracket \sqsubseteq \underline{wp}^N \llbracket C \rrbracket (\llbracket f \rrbracket)$?
If yes, then $\llbracket g \rrbracket \sqsubseteq wp \llbracket C \rrbracket (\llbracket f \rrbracket)$, provided $\llbracket f \rrbracket \in \mathbb{E}_{\text{meas}}$.
- (2) ... an \mathcal{L} -free $f \in \text{Expr}$ and an \mathcal{O} -free $g \in \text{Expr}$, is $\overline{wp}^N \llbracket C \rrbracket (\llbracket f \rrbracket) \sqsubseteq \llbracket g \rrbracket$?
If yes, then $wp \llbracket C \rrbracket (\llbracket f \rrbracket) \sqsubseteq \llbracket g \rrbracket$, provided $\llbracket f \rrbracket \in \mathbb{E}_{\text{meas}}$.
- (3) ... an \mathcal{O} -free $f \in \text{Expr}^{\leq 1}$ and an \mathcal{L} -free $g \in \text{Expr}^{\leq 1}$, is $\llbracket g \rrbracket \sqsubseteq \underline{wlp}^N \llbracket C \rrbracket (\llbracket f \rrbracket)$?
If yes, then $\llbracket g \rrbracket \sqsubseteq wlp \llbracket C \rrbracket (\llbracket f \rrbracket)$, provided $\llbracket f \rrbracket \in \mathbb{E}_{\text{meas}}^{\leq 1}$.
- (4) ... an \mathcal{L} -free $f \in \text{Expr}^{\leq 1}$ and an \mathcal{O} -free $g \in \text{Expr}^{\leq 1}$, is $\overline{wlp}^N \llbracket C \rrbracket (\llbracket f \rrbracket) \sqsubseteq \llbracket g \rrbracket$?
If yes, then $wlp \llbracket C \rrbracket (\llbracket f \rrbracket) \sqsubseteq \llbracket g \rrbracket$, provided $\llbracket f \rrbracket \in \mathbb{E}_{\text{meas}}^{\leq 1}$.

Moreover, for fixed C and f , the size of the syntactic expectation representing $\mathcal{T} \llbracket C \rrbracket (\llbracket f \rrbracket)$ for $\mathcal{T} \in \{\underline{wp}^N, \overline{wp}^N, \underline{wlp}^N, \overline{wlp}^N\}$ is in $O(N^k)$, where k is the number of uni f-statements in C .

Using Theorem 8.13 and Corollary 5.5, it is straightforward to derive bounds on cwp as well.

8.3.2 Verifying Loops. We now state how to apply the proof rules from Section 6 effectively:

THEOREM 8.14 (VERIFICATION OF INVARIANTS). *Let $C = \text{while}(\varphi) \{B\} \in \text{pWhile}(\text{Terms}, \text{Guards})$ such that B is loop-free. Further, let $I, f \in \text{Expr}$ and $J, g \in \text{Expr}^{\leq 1}$ all be quantifier-free, and let $N \geq 1$ be an integer. Then we can ...*

- (1) ... decide if $\llbracket I \rrbracket$ is a \overline{wp}^N -superinvariant of C w.r.t. $\llbracket f \rrbracket$ by a reduction to QF_NRA.
If the superinvariant property holds, then $wp \llbracket C \rrbracket (\llbracket f \rrbracket) \sqsubseteq \llbracket I \rrbracket$, provided $\llbracket f \rrbracket \in \mathbb{E}_{\text{meas}}$
- (2) ... decide if $\llbracket J \rrbracket$ is a \underline{wlp}^N -subinvariant of C w.r.t. $\llbracket g \rrbracket$ by a reduction to QF_NRA.
If the subinvariant property holds, then $\llbracket J \rrbracket \sqsubseteq wlp \llbracket C \rrbracket (\llbracket g \rrbracket)$, provided $\llbracket g \rrbracket \in \mathbb{E}_{\text{meas}}^{\leq 1}$.

THEOREM 8.15 (COMPLEXITY OF VERIFICATION PROBLEMS). *The following decision problems are coRE-complete¹⁴: Given an arbitrary (not necessarily loop-free) $C \in \text{pWhile}(\text{Terms}, \text{Guards})$ and ...*

- (1) ... $f, g \in \text{Expr}$ such that $\llbracket f \rrbracket \in \mathbb{E}_{\text{meas}}$, does it hold that $wp \llbracket C \rrbracket (\llbracket f \rrbracket) \sqsubseteq \llbracket g \rrbracket$?
- (2) ... $f, g \in \text{Expr}^{\leq 1}$ such that $\llbracket f \rrbracket \in \mathbb{E}_{\text{meas}}^{\leq 1}$, does it hold that $\llbracket g \rrbracket \sqsubseteq wlp \llbracket C \rrbracket (\llbracket f \rrbracket)$?

The proof of membership in coRE heavily relies on the convergence results established in Theorem 7.6. The coRE-hardness proofs are standard [36] and follow from coRE-hardness of the non-halting problem. See [10] for details.

¹³In polynomial time if we assume C to be of constant size and N given in unary. In general, an expression encoding $\underline{wp}^N \llbracket C \rrbracket (\llbracket f \rrbracket)$ for a given f can be exponential in the size of C – consider n sequential if-else constructs. However, our reduction is of practical interest as QF_NRA is supported by many SMT solvers.

¹⁴I.e., Π_1^0 -complete in the arithmetical hierarchy.

9 Implementation and Case Studies

We have automated our techniques using a modern deductive verifier for probabilistic programs. We first describe how to integrate our techniques within that verifier in Section 9.1. We then describe our case studies, i.e., programs and specifications we have verified, in Section 9.2. Finally, we evaluate our approach empirically on these case studies in Section 9.3.

9.1 Implementation

CAESAR¹⁵ [52] is a recent expectation-based automated deductive verifier targeting *discrete* and possibly *nondeterministic*¹⁶ probabilistic programs. As input, CAESAR takes programs written in an intermediate verification language called HeyVL — a quantitative analogue of BOOGIE [40]. HeyVL provides a programmatic means to describe verification conditions such as the validity of quantitative loop invariants. These verification conditions are offloaded to an SMT solver, which enables the semi-automated verification of said discrete probabilistic programs. CAESAR does *not* support continuous sampling instructions. We will, however, now demonstrate that with our approach based on Riemann sums, CAESAR *can* readily be used to verify bounds on expected outcomes of continuous probabilistic programs in a semi-automated fashion.

Our key insight is that our Riemann expectation transformers can be expressed as expectation transformers denoted by *discrete* probabilistic programs featuring angelic (resp. demonic) *nondeterministic choices* for $\mathbb{R}_{\geq 0}$ -valued intervals. The latter features *are* supported by CAESAR: HeyVL supports the discrete fragment of pWhile and, amongst others, the two statements

- (1) $x := [E_1, E_2]$, which, on program state σ , assigns a nondeterministically chosen value from the $\mathbb{R}_{\geq 0}$ -valued interval $[E_1(\sigma), E_2(\sigma)]$ to the variable x , and
- (2) $x \approx \text{discrete_unif}(N)$, which, given a (constant) natural number $N \geq 1$, samples a value from the set $\{0, \dots, N-1\}$ uniformly at random and assigns the result to the variable x .

Notice that the latter statement is syntactic sugar for pWhile as it can be simulated by binary probabilistic choices. The statement $x := [E_1, E_2]$, on the other hand, is not syntactic sugar as nondeterminism is not supported in pWhile.

Now, given a (discrete but possibly nondeterministic) program C , CAESAR employs a transformer $\text{vc}\llbracket C \rrbracket : \mathbb{E} \rightarrow \mathbb{E}$, which is defined¹⁷ as in Table 1 and where for the additional statements, we have

$$\begin{aligned} \text{vc}\llbracket x := [E_1, E_2] \rrbracket (f) &= \lambda \sigma. \sup_{\xi \in [E_1(\sigma), E_2(\sigma)]} f(\sigma[x \mapsto \xi]) \\ \text{vc}\llbracket x \approx \text{discrete_unif}(N) \rrbracket (f) &= \frac{1}{N} \cdot \sum_{i=0}^{N-1} f[x/i]. \end{aligned}$$

Notice that $\text{vc}\llbracket x := [E_1, E_2] \rrbracket (f)$ resolves the nondeterministic choice of x *angelically* by returning the *supremum* of all values obtained from evaluating f at σ where x is set to some value from $[E_1(\sigma), E_2(\sigma)]$. The demonic counterpart is obtained from replacing sup by inf . Now let j be a fresh program variable and observe that for all $N \geq 1$ and all $f \in \mathbb{E}$, we have

$$\underbrace{\overline{\text{wp}}^N \llbracket x \approx \text{unif}_{[0,1]} \rrbracket (f)}_{\text{defined in this paper}} = \underbrace{\text{vc}\llbracket j \approx \text{discrete_unif}(N) ; x := [\frac{j}{N}, \frac{j+1}{N}] \rrbracket (f)}_{\text{expressible in HeyVL and thus supported by CAESAR}}.$$

¹⁵<https://www.caesarverifier.org/>

¹⁶Here we refer to pure nondeterminism which is to be resolved angelically or demonically.

¹⁷We do not need to require f to be measurable since C is does not contain continuous sampling. Moreover, as is standard for deductive verifiers, loops need to be annotated with quantitative loop invariants; see Example 9.1 on page 22.

```

while (i ≤ M) {
  y := unif[0,1]; x := x + y;
  i := i + 1 }

```

```

while (i ≤ M) {
  y := unif[0,1]; observe(y ≤ 1/2); x := x + y;
  i := i + 1 }

```

Fig. 2. Generating the standard Irwin-Hall distribution (left) and a variant with conditioning (right).

Hence, the upper Riemann pre-expectation of $x := \text{unif}_{[0,1]}$ w.r.t. f corresponds to the maximal — under all possible resolutions of the nondeterminism — expected final value of f obtained from (i) sampling one of the N intervals occurring in the Riemann sums uniformly at random and (ii) assigning to x a nondeterministically chosen value from that sampled interval. Lower Riemann pre-expectations can be expressed analogously by resolving the nondeterminism *demonically*.

Example 9.1. Reconsider the Monte Carlo π -approximator C and the superinvariant I of C w.r.t. count from Example 6.3. We can provide CAESAR with the following annotated loop:

```

@invariant (count + [i ≤ M] · (0.85 · ((M ÷ i) + 1)))
while (i ≤ M) {
  j1 := discrete_unif(16); x := [ $\frac{j_1}{16}, \frac{j_1+1}{16}$ ];
  j2 := discrete_unif(16); y := [ $\frac{j_2}{16}, \frac{j_2+1}{16}$ ];
  if(x2 + y2 ≤ 1) {count := count + 1} else {skip}; i := i + 1 }

```

We can then instruct CAESAR to check whether I is a superinvariant of this loop w.r.t. count. Since the loop body encodes appropriate upper Riemann pre-expectations, CAESAR offloads the quantitative entailments corresponding to Theorem 8.14 to an SMT solver to check them automatically. For lower Riemann pre-expectations and subinvariants (for wlp), CAESAR can be used analogously.

9.2 Case Studies

In what follows, we describe the programs and specifications we have verified using CAESAR with the SMT solver Z3 [24] as back-end. All case studies (including Example 9.1) have been verified within 12 seconds on an Apple M2. The precise inputs to CAESAR are provided in [10]. Both the verified bounds and the required partition sizes N were determined manually by guessing the respective weakest pre-expectations and increasing N until CAESAR reports success.

9.2.1 The Irwin-Hall Distribution. The Irwin-Hall distribution (parameterized in M) [34] is the sum of M independent and identically distributed random variables, each of which is distributed uniformly over $[0, 1]$. The loop C depicted in Figure 2 (left) models this family of distributions: On each iteration, C samples a value for y uniformly at random and adds the result to the variable x . Hence, if initially $i = 1$, $x = 0$, and $M \in \mathbb{N}$, then the final distribution of x indeed corresponds to the Irwin-Hall distribution with parameter M .

We now aim to upper-bound the expected final value of x , i.e., the expectation of the Irwin-Hall distribution, for *arbitrary* values of M . Towards this end, we employ our encoding from Section 9.1 and use CAESAR to automatically verify that the expectation $I = x + [i \leq M] \cdot (1.1 \cdot \frac{(M-i)+1}{2})$ is a $\overline{\text{wp}}^{10}$ -superinvariant of C w.r.t. x . Hence, by Theorem 6.2, we get $\text{wp}\llbracket C \rrbracket(x) \sqsubseteq I$ and thus

$$\text{for initial } \sigma \text{ with } \sigma(i) = 1, \sigma(x) = 0, \sigma(M) \in \mathbb{N} \text{ with } \sigma(M) \geq 1: \text{wp}\llbracket C \rrbracket(x)(\sigma) \leq 1.1 \cdot \frac{\sigma(M)}{2}.$$

Now consider a variant of the Irwin-Hall distribution where we condition each of the M random variables to take a value in $[0, 1/2]$. This situation is modeled by the loop C depicted in Figure 2

```

while (x > 0) {
  y := unif[0,1]; y := (b ÷ a) · y + a;
  if(y ≤ (a + b)/2) {diverge} else {skip};
  x := x ÷ 1 }

while (h ≤ t) {
  {x := unif[0,1]; x := 10 · x; h := h + x} [1/2] {skip};
  t := t + 1;
  count := count + 1 }

```

Fig. 3. A probably diverging loop (left) and a race between tortoise and hare (right) [16].

(right), where we employ conditioning inside of the loop. We aim to upper-bound $\text{cwp}\llbracket C \rrbracket(x)$, i.e., the *conditional expected final value* of x . For that, we use CAESAR to verify that the following expectation is a $\overline{\text{wp}}^{19}$ -superinvariant of C w.r.t. count (resp. $\underline{\text{wlp}}^2$ -subinvariant of C w.r.t. 1):

$$I = x + [i \leq M] \cdot \left(1.5 \cdot \frac{(M - i) + 1}{8}\right) \quad \text{and} \quad J = [i \leq M] \cdot 0.5^{(M-i)+1} + [i > M] \cdot 1.$$

Hence, we get by Theorem 6.2 that $\text{cwp}\llbracket C \rrbracket(x)(\sigma)$ is well-defined for all $\sigma \in \mathbb{R}_{\geq 0}^V$ and

$$\text{for initial } \sigma \text{ with } \sigma(i) = 1, \sigma(x) = 0, \sigma(M) \in \mathbb{N} \text{ with } \sigma(M) \geq 1: \text{cwp}\llbracket C \rrbracket(x)(\sigma) \leq \frac{1.5 \cdot \sigma(M)}{8 \cdot 0.5^{\sigma(M)}}.$$

Notice that, even though exponential functions like 0.5^M are not supported by our decidable language of syntactic expectations from Section 8, we can use CAESAR to reason about it by means of standard user-defined domain declaration for exponentials (see [52, Section 5.1] for details). Decidability of the corresponding entailments involving such domains is, however, not guaranteed.

9.2.2 Reasoning about Diverging Programs. Weakest *liberal* pre-expectations enable us to lower-bound divergence probabilities of programs involving continuous sampling. Consider the loop C depicted in Figure 3 (left). In each iteration, we sample uniformly from the interval $[a, b]$ (notice that a, b are uninitialized variables and hence correspond to parameters) and assign the result to y . Whenever we sample a value in $[0, \frac{a+b}{2}]$, the program diverges. Using CAESAR, we verify that

$$I = [a \leq b] \cdot (1 \div 0.5^x)$$

is a $\underline{\text{wlp}}^2$ -subinvariant of C w.r.t. 0. Hence, we have $I \sqsubseteq \underline{\text{wlp}}\llbracket C \rrbracket(0)$ by Theorem 6.2, i.e., whenever $a \leq b$, then C diverges on initial state σ with probability at least $1 - 0.5^{\sigma(x)}$.

9.2.3 Race between Tortoise and Hare. The loop C depicted in Figure 3 (right) is adapted from [16] and models a race between a tortoise (t) and a hare (h). As long as the hare did not overtake the tortoise, the hare flips a fair coin to decide whether to move or not. If the hare decides to move, it samples a distance uniformly at random from $[0, 10]$. The tortoise always moves exactly one step.

We now upper-bound the expected final value of variable count . Towards this end, we use CAESAR to verify that the expectation $I = \text{count} + [h \leq t] \cdot 3.012 \cdot ((t - h) + 2)$ is a $\overline{\text{wp}}^{16}$ -superinvariant of C w.r.t. count . Hence, we get $\text{wp}\llbracket C \rrbracket(\text{count}) \sqsubseteq I$ by Theorem 6.2, i.e., if C is executed on an initial state σ with $\sigma(\text{count}) = 0$ and $\sigma(h) \leq \sigma(t)$, then the expected final value of count is at most $3.012 \cdot (\sigma(t) - \sigma(h) + 2)$.

9.3 Experimental Evaluation

In this section, we empirically evaluate the scalability of our approach based on the programs from the previous sections, i.e., our case studies and the Monte-Carlo approximator from Figure 1. For each of these programs and post-expectations, we verify the tightest upper bound (up to a precision of 3 decimal places) on the respective weakest pre-expectation for increasing values of N (except for diverging where $N = 2$ already yields the precise weakest pre-expectation).

Table 2. Experimental results. Time in seconds. TO = 180 s, MO = 8 GB.

Program	N	Time	AST	Post	Bound
MonteCarlo (Fig. 1)	2	0.03	552	count	$1.000 \cdot M$
	4	0.05	1634		$0.938 \cdot M$
	8	0.38	5448		$0.875 \cdot M$
	16	10.05	19676		$0.837 \cdot M$
	32	TO	74526		–
IrwinHall (Fig. 2 (left))	2	0.02	187	x	$0.750 \cdot M$
	4	0.02	309		$0.625 \cdot M$
	8	0.03	553		$0.565 \cdot M$
	10	0.08	716		$0.55 \cdot M$
	16	2.15	1041		$0.532 \cdot M$
32	TO	2017	–		
diverging (Fig. 3 (left))	2	0.22	162	0	$[a \leq b] \cdot (1 \dot{-} 0.5^x)$
TortoiseHare (Fig. 3 (right))	2	0.02	2	count	∞
	4	.03	400		$8 \cdot (t \dot{-} h + 2)$
	8	.015	680		$3.38 \cdot (t \dot{-} h + 2)$
	16	100.09	1240		$3.012 \cdot (t \dot{-} h + 2)$
	32	TO	2160		–

Our empirical results are depicted in Table 2. Column **Program** depicts the respective program, **N** denotes the partition size for the Riemann pre-expectation, **Time** denotes the verification time in seconds (i.e., verification condition generation and time for SMT solving) required by CAESAR, **|AST|** denotes the size of the formula offloaded to the SMT solver Z3 (i.e., the number of nodes in the formula’s abstract syntax tree), **Post** denotes the post-expectation, and **Bound** depicts the upper (lower) bound on the (liberal) weakest pre-expectation. For the diverging benchmark, we lower-bound a liberal weakest pre-expectation. All other benchmarks upper-bound a non-liberal weakest pre-expectation. The bounds were inferred manually in a binary search-like fashion.

With our encoding of Riemann pre-expectations described in Section 9.1, CAESAR allows verifying increasingly sharper bounds when increasing the partition size N . The size of the generated formula for the verification conditions mostly increases moderately when increasing N . This is to be expected since the formula size grows polynomially in N for a fixed program and a fixed post-expectation (see Section 8.3.1). The time required for verification, however, can increase drastically when doubling the value of N (see entries of Table 2 where $N = 32$).

10 Related Work

Integral Approximation. Our approach relies on under- and over-approximation of Lebesgue integrals via Riemann sums, achieved by discretizing the domain of the continuous uniform distributions in the program. Recently, [25] proposed an efficient *bit blasting* discretization method for continuous mixed-gamma distributions (which subsume the uniform distributions) for programs with bounded (for) loops. In [11], integral approximation is used to derive guaranteed bounds on a program’s posterior distribution. Following up on this, [58] extended the computation of such bounds to a class of programs with soft conditioning (*score-at-end* programs). Previously, [51] and [3] used integral approximations to bound the satisfaction probabilities of properties like safety and fairness.

However, these approaches are not able to deal with unbounded loops in full generality. AQUA [32] uses a different discretization method, where the posterior of a given program is approximated by quantizing the state space. However, this approximation does not necessarily provide an upper or lower bound of the true distribution, and unbounded loops are not supported. Finally, [55] propose a method to derive upper and lower bounds on central moments of cost accumulators.

Loop Invariant Analysis. The main purpose of the current work is to enable automatic verification of inductive invariants for continuous probabilistic programs. Reasoning about probabilistic loops via quantitative invariants was first introduced by McIver and Morgan [41] for probabilistic programs with discrete probabilistic choice and extended by various authors to different settings [30, 37, 53]. A generalization of this framework to programs featuring continuous distribution was proposed in [15] and uses super-martingale ranking functions as the probabilistic analogue of classic ranking functions. The super-martingale approach has been specialized to qualitative and quantitative termination problems for various classes of programs such as polynomial programs [17], affine programs admitting a linear ranking super-martingale [18], and non-deterministic probabilistic programs [2, 21]. Besides termination, this approach has been successfully applied also to cost analysis [57] and equivalence refutation [20]. While building on the martingale approach, [19] overcomes the necessity to compute ranking martingales, by introducing the notion of stochastic invariants indicators, that can be used to prove termination in a sound and relatively complete way. A different generalization was proposed in [16], where expectation invariants are proposed to introduce a generalized fixed-point framework for programs with continuous distributions and unbounded loops. However, this analysis rules out invariants expressed as Iverson brackets and is restricted to piecewise linear assertion guards and updates. Similarly, [56], restricts to affine programs and invariants to find exponential bounds on assertion violation probabilities.

Exact Inference. Exact inference on probabilistic programs involving continuous distributions is especially hard because of the challenge of computing densities via integrals. Exact symbolic inference engines like [26, 27, 45, 50] only support deterministically bounded loops. Even within this constrained context, the occurrence of non-simplified integrals in the output can significantly impede the potential to perform appropriate analysis or further computations. A different, yet still exact approach is pursued in [43], where the focus is on the computation of moments as functions of the loop iteration. In order to compute the moments exactly, the syntax of the programs is restricted quite severely, though parametric models are supported. Under the same syntactic restrictions, it has been shown in [42] that termination probabilities can be computed automatically.

11 Conclusion

We tackled the semi-automatic verification of user-provided quantitative invariants for probabilistic programs involving continuous distributions, conditioning, and potentially unbounded loops. Our approach is to approximate the Lebesgue integrals appearing in a program's weakest pre-expectation semantics by lower and upper Riemann sums. This simple idea allows us to prove sound bounds on expected outcomes by verifying invariants automatically with SMT-solvers. We demonstrated that this method can be readily integrated in the CAESAR verification infrastructure [52].

There are many opportunities for future work. A natural direction is to integrate our approach to invariant verification with techniques for *fully automatic invariant generation* such as [7]. While we can already verify invariants for non-trivial benchmarks, there is some room for improvement on the engineering side. This includes more sophisticated, multi-dimensional, and adaptive partitions of the integration domain, and a combination with direct integration methods, where applicable. We also plan to extend the front end of CAESAR with dedicated support for continuous distributions.

Due to its simplicity, we are confident that our approach is amenable to generalizations such as non-determinism and strategy synthesis as in [6], and programs with pointers and data structures.

Acknowledgments

The authors thank Philipp Schroer for his support with CAESAR, Mirco Tribastone for his help with the initial conceptualization of the paper, and the anonymous referees for their detailed comments. This work was partially funded by the DFG GRK 2236 UnRAVeL, the EU's Horizon 2020 programme under the Marie Skłodowska-Curie grant No. 101008233 (MISSION), the ERC AdG No. 787914 (FRAPPANT), and the PNRR project iNEST (Interconnected Nord-Est Innovation Ecosystem) funded by the European Union Next-GenerationEU (Piano Nazionale di Ripresa e Resilienza (PNRR) – Missione 4 Componente 2, Investimento 1.5 – D.D. 1058 23/06/2022, ECS_00000043).

Data Availability Statement

A software artifact comprising the CAESAR verifier, the encodings of the case studies from Section 9, and instructions for running the experiments is available [9].

References

- [1] Sheshansh Agrawal, Krishnendu Chatterjee, and Petr Novotný. 2017. Lexicographic Ranking Supermartingales: An Efficient Approach to Termination of Probabilistic Programs. *CoRR* abs/1709.04037 (2017). arXiv:1709.04037 <http://arxiv.org/abs/1709.04037>
- [2] Sheshansh Agrawal, Krishnendu Chatterjee, and Petr Novotný. 2018. Lexicographic ranking supermartingales: an efficient approach to termination of probabilistic programs. *Proc. ACM Program. Lang.* 2, POPL (2018), 34:1–34:32. doi:10.1145/3158122
- [3] Aws Albarghouthi, Loris D'Antoni, Samuel Drews, and Aditya V. Nori. 2017. FairSquare: probabilistic verification of program fairness. *Proc. ACM Program. Lang.* 1, OOPSLA (2017), 80:1–80:30. doi:10.1145/3133904
- [4] Robert B Ash and Catherine A Doléans-Dade. 2000. *Probability and measure theory*. Academic press.
- [5] Clark Barrett, Pascal Fontaine, and Cesare Tinelli. 2016. The Satisfiability Modulo Theories Library (SMT-LIB). www.SMT-LIB.org.
- [6] Kevin Batz, Tom Jannik Biskup, Joost-Pieter Katoen, and Tobias Winkler. 2024. Programmatic Strategy Synthesis: Resolving Nondeterminism in Probabilistic Programs. *Proc. ACM Program. Lang.* 8, POPL (2024), 2792–2820. doi:10.1145/3632935
- [7] Kevin Batz, Mingshuai Chen, Sebastian Junges, Benjamin Lucien Kaminski, Joost-Pieter Katoen, and Christoph Matheja. 2023. Probabilistic Program Verification via Inductive Synthesis of Inductive Invariants. In *TACAS (2) (Lecture Notes in Computer Science, Vol. 13994)*. Springer, 410–429. doi:10.1007/978-3-031-30820-8_25
- [8] 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. doi:10.1145/3434320
- [9] Kevin Batz, Joost-Pieter Katoen, Francesca Randone, and Tobias Winkler. 2025. *Artifact for Paper Foundations for Deductive Verification of Continuous Probabilistic Programs*. doi:10.5281/zenodo.15175355
- [10] Kevin Batz, Joost-Pieter Katoen, Francesca Randone, and Tobias Winkler. 2025. Foundations for Deductive Verification of Continuous Probabilistic Programs: From Riemann to Lebesgue and Back. *CoRR* abs/2502.19388 (2025). <https://arxiv.org/abs/2502.19388>
- [11] Raven Beutner, C.-H. Luke Ong, and Fabian Zaiser. 2022. Guaranteed bounds for posterior inference in universal probabilistic programming. In *PLDI*. ACM, 536–551. doi:10.1145/3519939.3523721
- [12] Eli Bingham, Jonathan P. Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul A. Szerlip, Paul Horsfall, and Noah D. Goodman. 2019. Pyro: Deep Universal Probabilistic Programming. *J. Mach. Learn. Res.* 20 (2019), 28:1–28:6. <https://jmlr.org/papers/v20/18-403.html>
- [13] Frank E Burk. 2007. *A garden of integrals*. Vol. 31. American Mathematical Soc. doi:10.7135/UPO9781614442097
- [14] Bob Carpenter, Andrew Gelman, Matthew D. Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. 2017. Stan: A Probabilistic Programming Language. *Journal of Statistical Software* 76, 1 (2017), 1–32. doi:10.18637/jss.v076.i01
- [15] Aleksandar Chakarov and Sriram Sankaranarayanan. 2013. Probabilistic Program Analysis with Martingales. In *CAV (Lecture Notes in Computer Science, Vol. 8044)*. Springer, 511–526. doi:10.1007/978-3-642-39799-8_34
- [16] Aleksandar Chakarov and Sriram Sankaranarayanan. 2014. Expectation Invariants for Probabilistic Program Loops as Fixed Points. In *SAS (Lecture Notes in Computer Science, Vol. 8723)*. Springer, 85–100. doi:10.1007/978-3-319-10936-7_6
- [17] Krishnendu Chatterjee, Hongfei Fu, and Amir Kafshdar Goharshady. 2016. Termination Analysis of Probabilistic Programs Through Positivstellensatz's. In *CAV (1) (Lecture Notes in Computer Science, Vol. 9779)*. Springer, 3–22.

- doi:10.1007/978-3-319-41528-4_1
- [18] Krishnendu Chatterjee, Hongfei Fu, Petr Novotný, and Rouzbeh Hasheminezhad. 2016. Algorithmic analysis of qualitative and quantitative termination problems for affine probabilistic programs. In *POPL*. ACM, 327–342. doi:10.1145/2837614.2837639
- [19] Krishnendu Chatterjee, Amir Kafshdar Goharshady, Tobias Meggendorfer, and Dorde Zikelic. 2022. Sound and Complete Certificates for Quantitative Termination Analysis of Probabilistic Programs. In *CAV (1) (Lecture Notes in Computer Science, Vol. 13371)*. Springer, 55–78. doi:10.1007/978-3-031-13185-1_4
- [20] Krishnendu Chatterjee, Ehsan Kafshdar Goharshady, Petr Novotný, and Dorde Zikelic. 2024. Equivalence and Similarity Refutation for Probabilistic Programs. *Proc. ACM Program. Lang.* 8, PLDI (2024), 2098–2122. doi:10.1145/3656462
- [21] Krishnendu Chatterjee, Petr Novotný, and Dorde Zikelic. 2017. Stochastic invariants for probabilistic termination. In *POPL*. ACM, 145–160. doi:10.1145/3009837.3009873
- [22] Michel Coste. 2000. An introduction to semialgebraic geometry.
- [23] Fredrik Dahlqvist, Alexandra Silva, and Dexter Kozen. 2020. Semantics of probabilistic programming: A gentle introduction. *Foundations of Probabilistic Programming* (2020), 1–42.
- [24] Leonardo Mendonça de Moura and Nikolaj S. Bjørner. 2008. Z3: An Efficient SMT Solver. In *TACAS (Lecture Notes in Computer Science, Vol. 4963)*. Springer, 337–340. doi:10.1007/978-3-540-78800-3_24
- [25] Poorva Garg, Steven Holtzen, Guy Van den Broeck, and Todd D. Millstein. 2024. Bit Blasting Probabilistic Programs. *Proc. ACM Program. Lang.* 8, PLDI (2024), 865–888. doi:10.1145/3656412
- [26] Timon Gehr, Sasa Misailovic, and Martin T. Vechev. 2016. PSI: Exact Symbolic Inference for Probabilistic Programs. In *CAV (1) (Lecture Notes in Computer Science, Vol. 9779)*. Springer, 62–83. doi:10.1007/978-3-319-41528-4_4
- [27] Timon Gehr, Samuel Steffen, and Martin T. Vechev. 2020. λ PSI: exact inference for higher-order probabilistic programs. In *PLDI*. ACM, 883–897. doi:10.1145/3385412.3386006
- [28] Noah D Goodman and Andreas Stuhlmüller. 2014. The Design and Implementation of Probabilistic Programming Languages. <http://dippl.org>. Accessed: 2024-10-15.
- [29] Andrew D. Gordon, Thomas A. Henzinger, Aditya V. Nori, and Sriram K. Rajamani. 2014. Probabilistic programming. In *FOSE*. ACM, 167–181. doi:10.1145/2593882.2593900
- [30] Friedrich Gretz, Joost-Pieter Katoen, and Annabelle McIver. 2014. Operational versus weakest pre-expectation semantics for the probabilistic guarded command language. *Perform. Evaluation* 73 (2014), 110–132. doi:10.1016/J.PEVA.2013.11.004
- [31] Marcel Hark, Benjamin Lucien Kaminski, Jürgen Giesl, and Joost-Pieter Katoen. 2020. Aiming low is harder: induction for lower bounds in probabilistic program verification. *Proc. ACM Program. Lang.* 4, POPL (2020), 37:1–37:28. doi:10.1145/3371105
- [32] Zixin Huang, Saikat Dutta, and Sasa Misailovic. 2022. Automated quantized inference for probabilistic programs with AQUA. *Innov. Syst. Softw. Eng.* 18, 3 (2022), 369–384. doi:10.1007/S11334-021-00433-3
- [33] Alon Itai and Michael Rodeh. 1990. Symmetry breaking in distributed networks. *Inf. Comput.* 88, 1 (1990), 60–87. doi:10.1016/0890-5401(90)90004-2
- [34] N.L. Johnson, S. Kotz, and N. Balakrishnan. 1995. *Continuous Univariate Distributions, Volume 2*. Wiley.
- [35] Benjamin Lucien Kaminski. 2019. *Advanced weakest precondition calculi for probabilistic programs*. Ph. D. Dissertation. RWTH Aachen University, Germany. doi:10.18154/RWTH-2019-01829
- [36] Benjamin Lucien Kaminski and Joost-Pieter Katoen. 2015. On the Hardness of Almost-Sure Termination. In *Mathematical Foundations of Computer Science 2015 - 40th International Symposium, MFCS 2015, Milan, Italy, August 24-28, 2015, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 9234)*, Giuseppe F. Italiano, Giovanni Pighizzini, and Donald Sannella (Eds.). Springer, 307–318. doi:10.1007/978-3-662-48057-1_24
- [37] Benjamin Lucien Kaminski and Joost-Pieter Katoen. 2017. A weakest pre-expectation semantics for mixed-sign expectations. In *32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20-23, 2017*. IEEE Computer Society, 1–12. doi:10.1109/LICS.2017.8005153
- [38] Richard M. Karp. 1991. An introduction to randomized algorithms. *Discret. Appl. Math.* 34, 1-3 (1991), 165–201.
- [39] Dexter Kozen. 1983. A Probabilistic PDL. In *STOC*. ACM, 291–297. doi:10.1145/800061.808758
- [40] K. Rustan M. Leino. 2008. *This Is Boogie 2*.
- [41] Annabelle McIver and Carroll Morgan. 2005. *Abstraction, Refinement and Proof for Probabilistic Systems*. Springer. doi:10.1007/B138392
- [42] Marcel Moosbrugger, Ezio Bartocci, Joost-Pieter Katoen, and Laura Kovács. 2022. The probabilistic termination tool amber. *Formal Methods Syst. Des.* 61, 1 (2022), 90–109. doi:10.1007/S10703-023-00424-Z
- [43] Marcel Moosbrugger, Miroslav Stankovic, Ezio Bartocci, and Laura Kovács. 2022. This is the moment for probabilistic loops. *Proc. ACM Program. Lang.* 6, OOPSLA2 (2022), 1497–1525. doi:10.1145/3563341
- [44] Peter Müller, Malte Schwerhoff, and Alexander J. Summers. 2016. Viper: A Verification Infrastructure for Permission-Based Reasoning. In *VMCAI (Lecture Notes in Computer Science, Vol. 9583)*. Springer, 41–62. doi:10.1007/978-3-662-49122-5_2

- [45] Praveen Narayanan, Jacques Carette, Wren Romano, Chung-chieh Shan, and Robert Zinkov. 2016. Probabilistic Inference by Program Transformation in Hakaru (System Description). In *FLOPS (Lecture Notes in Computer Science, Vol. 9613)*. Springer, 62–79. doi:10.1007/978-3-319-29604-3_5
- [46] Aditya V. Nori, Chung-Kil Hur, Sriram K. Rajamani, and Selva Samuel. 2014. R2: An Efficient MCMC Sampler for Probabilistic Programs. In *AAAI*. AAAI Press, 2476–2482. doi:10.1609/AAAI.V28I1.9060
- [47] Federico Olmedo, Friedrich Gretz, Nils Jansen, Benjamin Lucien Kaminski, Joost-Pieter Katoen, and Annabelle McIver. 2018. Conditioning in Probabilistic Programming. *ACM Trans. Program. Lang. Syst.* 40, 1 (2018), 4:1–4:50. doi:10.1145/3156018
- [48] David Park. 1969. Fixpoint induction and proofs of program properties. *Machine intelligence* 5 (1969).
- [49] Walter Rudin. 1953. *Principles of mathematical analysis*.
- [50] Feras A. Saad, Martin C. Rinard, and Vikash K. Mansinghka. 2021. SPPL: probabilistic programming with fast exact symbolic inference. In *PLDI*. ACM, 804–819. doi:10.1145/3453483.3454078
- [51] Sriram Sankaranarayanan, Aleksandar Chakarov, and Sumit Gulwani. 2013. Static analysis for probabilistic programs: inferring whole program properties from finitely many paths. In *PLDI*. ACM, 447–458. doi:10.1145/2491956.2462179
- [52] Philipp Schröder, Kevin Batz, Benjamin Lucien Kaminski, Joost-Pieter Katoen, and Christoph Matheja. 2023. A Deductive Verification Infrastructure for Probabilistic Programs. *Proc. ACM Program. Lang.* 7, OOPSLA2 (2023), 2052–2082. doi:10.1145/3622870
- [53] Marcin Szymczak and Joost-Pieter Katoen. 2019. Weakest Preexpectation Semantics for Bayesian Inference - Conditioning, Continuous Distributions and Divergence. In *SETSS (Lecture Notes in Computer Science, Vol. 12154)*. Springer, 44–121. doi:10.1007/978-3-030-55089-9_3
- [54] Jan-Willem van de Meent, Brooks Paige, Hongseok Yang, and Frank Wood. 2018. An Introduction to Probabilistic Programming. *CoRR* abs/1809.10756 (2018). arXiv:1809.10756 <http://arxiv.org/abs/1809.10756>
- [55] Di Wang, Jan Hoffmann, and Thomas W. Reps. 2021. Central moment analysis for cost accumulators in probabilistic programs. In *PLDI*. ACM, 559–573. doi:10.1145/3453483.3454062
- [56] Jinyi Wang, Yican Sun, Hongfei Fu, Krishnendu Chatterjee, and Amir Kafshdar Goharshady. 2021. Quantitative analysis of assertion violations in probabilistic programs. In *PLDI*. ACM, 1171–1186. doi:10.1145/3453483.3454102
- [57] Peixin Wang, Hongfei Fu, Amir Kafshdar Goharshady, Krishnendu Chatterjee, Xudong Qin, and Wenjun Shi. 2019. Cost analysis of nondeterministic probabilistic programs. In *PLDI*. ACM, 204–220. doi:10.1145/3314221.3314581
- [58] Peixin Wang, Tengshun Yang, Hongfei Fu, Guanyan Li, and C-H Luke Ong. 2024. Static Posterior Inference of Bayesian Probabilistic Programming via Polynomial Solving. *Proceedings of the ACM on Programming Languages* 8, PLDI (2024), 1361–1386. doi:10.1145/3656432
- [59] Glynn Winskel. 1993. *The formal semantics of programming languages - an introduction*. MIT Press.

Received 2024-10-16; accepted 2025-02-18