



An Assertion-Based Program Logic for Probabilistic Programs

Gilles Barthe, Thomas Espitau, Marco Gaboardi, Benjamin Grégoire, Justin Hsu, Pierre-Yves Strub

► To cite this version:

Gilles Barthe, Thomas Espitau, Marco Gaboardi, Benjamin Grégoire, Justin Hsu, et al.. An Assertion-Based Program Logic for Probabilistic Programs. Lecture Notes in Computer Science, Apr 2018, Thessaloniki, Greece. pp.117-144, 10.1007/978-3-319-89884-1_5 . hal-01959567

HAL Id: hal-01959567

<https://hal.science/hal-01959567>

Submitted on 23 Dec 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

An Assertion-Based Program Logic for Probabilistic Programs^{*}

Gilles Barthe, Thomas Espitau, Marco Gaboardi
Benjamin Grégoire, Justin Hsu, and Pierre-Yves Strub

¹ IMDEA Software Institute, Spain

² Université Paris 6, France

³ University at Buffalo, SUNY, USA

⁴ Inria Sophia Antipolis–Méditerranée, France

⁵ University College London, UK

⁶ École Polytechnique, France

Abstract. Research on deductive verification of probabilistic programs has considered expectation-based logics, where pre- and post-conditions are real-valued functions on states, and assertion-based logics, where pre- and post-conditions are boolean predicates on state distributions. Both approaches have developed over nearly four decades, but they have different standings today. Expectation-based systems have managed to formalize many sophisticated case studies, while assertion-based systems today have more limited expressivity and have targeted simpler examples. We present ELLORA, a sound and relatively complete assertion-based program logic, and demonstrate its expressivity by verifying several classical examples of randomized algorithms using an implementation in the EASYCRYPT proof assistant. ELLORA features new proof rules for loops and adversarial code, and supports richer assertions than existing program logics. We also show that ELLORA allows convenient reasoning about complex probabilistic concepts by developing a new program logic for probabilistic independence and distribution law, and then smoothly embedding it into ELLORA. Our work demonstrates that the assertion-based approach is not fundamentally limited and suggests that some notions are potentially easier to reason about in assertion-based systems.

1 Introduction

The most mature systems for deductive verification of randomized algorithms are *expectation-based* techniques; seminal examples include PDDL [28] and pGCL [34]. These approaches reason about *expectations*, functions E from states to real numbers,⁷ propagating them backwards through a program until they are transformed into a mathematical function of the input. Expectation-based systems are both theoretically elegant [24,16,35,23] and practically useful;

^{*} This is the full version of the paper.

⁷ Treating a program as a function from input states s to output distributions $\mu(s)$, the expected value of E on $\mu(s)$ is an expectation.

implementations have verified numerous randomized algorithms [19,21]. However, properties involving multiple probabilities or expected values can be cumbersome to verify—each expectation must be analyzed separately.

An alternative approach envisioned by Ramshaw [37] is to work with predicates over distributions. A direct comparison with expectation-based techniques is difficult, as the approaches are quite different. In broad strokes, assertion-based systems can verify richer properties in one shot and have specifications that are arguably more intuitive, especially for reasoning about loops, while expectation-based approaches can transform expectations mechanically and can reason about non-determinism. However, the comparison is not very meaningful for an even simpler reason: existing assertion-based systems such as [8,18,38] are not as well developed as their expectation-based counterparts.

Restrictive assertions. Existing probabilistic program logics do not support reasoning about expected values, only probabilities. As a result, many properties about average-case behavior are not even expressible.

Inconvenient reasoning for loops. The Hoare logic rule for deterministic loops does not directly generalize to probabilistic programs. Existing assertion-based systems either forbid loops, or impose complex semantic side conditions to control which assertions can be used as loop invariants. Such side conditions are restrictive and difficult to establish.

No support for external or adversarial code. A strength of expectation-based techniques is reasoning about programs combining probabilities and *non-determinism*. In contrast, Morgan and McIver [30] argue that assertion-based techniques cannot support compositional reasoning for such a combination. For many applications, including cryptography, we would still like to reason about a commonly-encountered special case: programs using external or adversarial code. Many security properties in cryptography boil down to analyzing such programs, but existing program logics do not support adversarial code.

Few concrete implementations. There are by now several independent implementations of expectation-based techniques, capable of verifying interesting probabilistic programs. In contrast, there are only scattered implementations of probabilistic program logics.

These limitations raise two points. Compared to expectation-based approaches:

1. Can assertion-based approaches achieve similar expressivity?
2. Are there situations where assertion-based approaches are more suitable?

In this paper, we give positive evidence for both of these points.⁸ Towards the first point, we give a new assertion-based logic ELLORA for probabilistic programs, overcoming limitations in existing probabilistic program logics. ELLORA supports a rich set of assertions that can express concepts like expected values

⁸ Note that we do not give mathematically precise formulations of these points; as we are interested in the practical verification of probabilistic programs, a purely theoretical answer would not address our concerns.

and probabilistic independence, and novel proof rules for verifying loops and adversarial code. We prove that ELLORA is sound and relatively complete.

Towards the second point, we evaluate ELLORA in two ways. First, we define a new logic for proving probabilistic independence and distribution law properties—which are difficult to capture with expectation-based approaches—and then embed it into ELLORA. This sub-logic is more narrowly focused than ELLORA, but supports more concise reasoning for the target assertions. Our embedding demonstrates that the assertion-based approach can be flexibly integrated with intuitive, special-purpose reasoning principles. To further support this claim, we also provide an embedding of the Union Bound logic, a program logic for reasoning about accuracy bounds [4]. Then, we develop a full-featured implementation of ELLORA in the EASYCRYPT theorem prover and exercise the logic by mechanically verifying a series of complex randomized algorithms. Our results suggest that the assertion-based approach can indeed be practically viable.

Abstract logic. To ease the presentation, we present ELLORA in two stages. First, we consider an abstract version of the logic where assertions are general predicates over distributions, with no compact syntax. Our abstract logic makes two contributions: reasoning for loops, and for adversarial code.

Reasoning about Loops. Proving a property of a probabilistic loop typically requires establishing a loop invariant, but the class of loop invariants that can be soundly used depends on the termination behavior—stronger termination assumptions allows richer loop invariants. We identify three classes of assertions that can be used for reasoning about probabilistic loops, and provide a proof rule for each one:

- arbitrary assertions for *certainly terminating* loops, i.e. loops that terminate in a finite amount of iterations;
- *topologically closed* assertions for *almost surely* terminating loops, i.e. loops terminating with probability 1;
- *downwards closed* assertions for arbitrary loops.

The definition of topologically closed assertion is reminiscent of Ramshaw [37]; the stronger notion of downwards closed assertion appears to be new.

Besides broadening the class of loops that can be analyzed, our rules often enable simpler proofs. For instance, if the loop is certainly terminating, then there is no need to prove semantic side-conditions. Likewise, there is no need to consider the termination behavior of the loop when the invariant is downwards and topologically closed. For example, in many applications in cryptography, the target property is that a “bad” event has low probability: $\Pr[E] \leq k$. In our framework this assertion is downwards and topologically closed, so it can be a loop invariant regardless of the termination behavior.

Reasoning about Adversaries. Existing assertion-based logics cannot reason about probabilistic programs with *adversarial* code. *Adversaries* are special probabilistic procedures consisting of an interface listing the concrete procedures

that an adversary can call (*oracles*), along with restrictions like how many calls an adversary may make. Adversaries are useful in cryptography, where security notions are described using experiments in which adversaries interact with a challenger, and in game theory and mechanism design, where adversaries can represent strategic agents. Adversaries can also model inputs to *online* algorithms.

We provide proof rules for reasoning about adversary calls. Our rules are significantly more general than previously considered rules for reasoning about adversaries. For instance, the rule for adversary used by [4] is restricted to adversaries that cannot make oracle calls.

Metatheory. We show soundness and relative completeness of the core abstract logic, with mechanized proofs in the COQ proof assistant.⁹

Concrete logic. While the abstract logic is conceptually clean, it is inconvenient for practical formal verification—the assertions are too general and the rules involve semantic side-conditions. To address these issues, we flesh out a concrete version of ELLORA. Assertions are described by a grammar modeling a two-level assertion language. The first level contains state predicates—deterministic assertions about a single memory—while the second layer contains probabilistic predicates constructed from probabilities and expected values over discrete distributions. While the concrete assertions are theoretically less expressive than their counterparts in the abstract logic, they can already encode common properties and notions from existing proofs, like probabilities, expected values, distribution laws and probabilistic independence. Our assertions can express theorems from probability theory, enabling sophisticated reasoning about probabilistic concepts.

Furthermore, we leverage the concrete syntax to simplify verification.

- We develop an automated procedure for generating pre-conditions of non-looping commands, inspired by expectation-based systems.
- We give syntactic conditions for the closedness and termination properties required for soundness of the loop rules.

Implementation and case studies. We implement ELLORA on top of EASY-CRYPT, a general-purpose proof assistant for reasoning about probabilistic programs, and we mechanically verify a diverse collection of examples including textbook algorithms and a randomized routing procedure. We develop an EASY-CRYPT formalization of probability theory from the ground up, including tools like concentration bounds (e.g., the Chernoff bound), Markov’s inequality, and theorems about probabilistic independence.

Embeddings. We propose a simple program logic for proving *probabilistic independence*. This logic is designed to reason about independence in a lightweight way, as is common in paper proofs. We prove that the logic can be embedded into ELLORA, and is therefore sound. Furthermore, we prove an embedding of the Union Bound logic [4].

⁹ The formalization is available at <https://github.com/strub/xhl>.

2 Mathematical Preliminaries

As is standard, we will model randomized computations using *sub-distributions*.

Definition 1. A sub-distribution over a set A is defined by a mass function $\mu : A \rightarrow [0, 1]$ that gives the probability of the unitary events $a \in A$. This mass function must be s.t. $\sum_{a \in A} \mu(a)$ is well-defined and $|\mu| \triangleq \sum_{a \in A} \mu(a) \leq 1$. In particular, the support $\text{supp}(\mu) \triangleq \{a \in A \mid \mu(a) \neq 0\}$ is discrete.¹⁰ The name “sub-distribution” emphasizes that the total probability may be strictly less than 1. When the weight $|\mu|$ is equal to 1, we call μ a distribution. We let $\mathbf{SDist}(A)$ denote the set of sub-distributions over A . The probability of an event $E(x)$ w.r.t. a sub-distribution μ , written $\Pr_{x \sim \mu}[E(x)]$, is defined as $\sum_{x \in A \mid E(x)} \mu(x)$.

Simple examples of sub-distributions include the *null sub-distribution* $\mathbf{0}$, which maps each element of the underlying space to 0; and the *Dirac distribution* centered on x , written δ_x , which maps x to 1 and all other elements to 0. The following standard construction gives a monadic structure to sub-distributions.

Definition 2. Let $\mu \in \mathbf{SDist}(A)$ and $f : A \rightarrow \mathbf{SDist}(B)$. Then $\mathbb{E}_{a \sim \mu}[f] \in \mathbf{SDist}(B)$ is defined by

$$\mathbb{E}_{a \sim \mu}[f](b) \triangleq \sum_{a \in A} \mu(a) \cdot f(a)(b).$$

We use notation reminiscent of expected values, as the definition is quite similar.

We will need two constructions to model branching statements.

Definition 3. Let $\mu_1, \mu_2 \in \mathbf{SDist}(A)$ such that $|\mu_1| + |\mu_2| \leq 1$. Then $\mu_1 + \mu_2$ is the sub-distribution μ such that $\mu(a) = \mu_1(a) + \mu_2(a)$ for every $a \in A$.

Definition 4. Let $E \subseteq A$ and $\mu \in \mathbf{SDist}(A)$. Then the restriction $\mu|_E$ of μ to E is the sub-distribution such that $\mu|_E(a) = \mu(a)$ if $a \in E$ and 0 otherwise.

Sub-distributions are partially ordered under the pointwise order.

Definition 5. Let $\mu_1, \mu_2 \in \mathbf{SDist}(A)$. We say $\mu_1 \leq \mu_2$ if $\mu_1(a) \leq \mu_2(a)$ for every $a \in A$, and we say $\mu_1 = \mu_2$ if $\mu_1(a) = \mu_2(a)$ for every $a \in A$.

We use the following lemma when reasoning about the semantics of loops.

Lemma 1. If $\mu_1 \leq \mu_2$ and $|\mu_1| = 1$, then $\mu_1 = \mu_2$ and $|\mu_2| = 1$.

Sub-distributions are stable under pointwise-limits.

¹⁰ We work with discrete distributions to keep measure-theoretic technicalities to a minimum, though we do not see obstacles to generalizing to the continuous setting.

Definition 6. A sequence $(\mu_n)_{n \in \mathbb{N}} \in \mathbf{SDist}(A)$ sub-distributions converges if for every $a \in A$, the sequence $(\mu_n(a))_{n \in \mathbb{N}}$ of real numbers converges. The limit sub-distribution is defined as

$$\mu_\infty(a) \triangleq \lim_{n \rightarrow \infty} \mu_n(a)$$

for every $a \in A$. We write $\lim_{n \rightarrow \infty} \mu_n$ for μ_∞ .

Lemma 2. Let $(\mu_n)_{n \in \mathbb{N}}$ be a convergent sequence of sub-distributions. Then for any event $E(x)$, we have:

$$\forall n \in \mathbb{N}. \Pr_{x \sim \mu_\infty} [E(x)] = \lim_{n \rightarrow \infty} \Pr_{x \sim \mu_n} [E(x)].$$

Any bounded increasing real sequence has a limit; the same is true of sub-distributions.

Lemma 3. Let $(\mu_n)_{n \in \mathbb{N}} \in \mathbf{SDist}(A)$ be an increasing sequence of sub-distributions. Then, this sequence converges to μ_∞ and $\mu_n \leq \mu_\infty$ for every $n \in \mathbb{N}$. In particular, for any event E , we have $\Pr_{x \sim \mu_n} [E] \leq \Pr_{x \sim \mu_\infty} [E]$ for every $n \in \mathbb{N}$.

3 Programs and Assertions

Now, we introduce our core programming language and its denotational semantics.

Programs. We base our development on PWHILE, a strongly-typed imperative language with deterministic assignments, probabilistic assignments, conditionals, loops, and an **abort** statement which halts the computation with no result. Probabilistic assignments $x \xleftarrow{s} g$ assign a value sampled from a distribution g to a program variable x . The syntax of statements is defined by the grammar:

$$\begin{aligned} s ::= & \text{skip} \mid \text{abort} \mid x \leftarrow e \mid x \xleftarrow{s} g \mid s; s \\ & \mid \text{if } e \text{ then } s \text{ else } s \mid \text{while } e \text{ do } s \mid x \leftarrow \mathcal{I}(e) \mid x \leftarrow \mathcal{A}(e) \end{aligned}$$

where x , e , and g range over typed variables in \mathcal{X} , expressions in \mathcal{E} and distribution expressions in \mathcal{D} respectively. The set \mathcal{E} of well-typed expressions is defined inductively from \mathcal{X} and a set \mathcal{F} of function symbols, while the set \mathcal{D} of well-typed distribution expressions is defined by combining a set of distribution symbols \mathcal{S} with expressions in \mathcal{E} . Programs may call a set \mathcal{I} of internal procedures as well as a set \mathcal{A} of external procedures. We assume that we have code for internal procedures but not for external procedures—we only know indirect information, like which internal procedures they may call. Borrowing a convention from cryptography, we call internal procedures *oracles* and external procedures *adversaries*.

$$\begin{aligned}
\llbracket \text{skip} \rrbracket_m &= \delta_m \\
\llbracket \text{abort} \rrbracket_m &= \mathbf{0} \\
\llbracket x \leftarrow e \rrbracket_m &= \delta_{m[x:=\llbracket e \rrbracket_m]} \\
\llbracket x \leftarrow^{\$} g \rrbracket_m &= \mathbb{E}_{v \sim \llbracket g \rrbracket_m} [\delta_{m[x:=v]}] \\
\llbracket s_1; s_2 \rrbracket_m &= \mathbb{E}_{m' \sim \llbracket s_1 \rrbracket_m} [\llbracket s_2 \rrbracket_{m'}] \\
\llbracket \text{if } e \text{ then } s_1 \text{ else } s_2 \rrbracket_m &= \text{if } \llbracket e \rrbracket_m \text{ then } \llbracket s_1 \rrbracket_m \text{ else } \llbracket s_2 \rrbracket_m \\
\llbracket \text{while } e \text{ do } s \rrbracket_m &= \lim_{n \rightarrow \infty} \llbracket (\text{if } e \text{ then } s)^n; \text{if } e \text{ then abort} \rrbracket_m \\
\llbracket x \leftarrow \mathcal{I}(e) \rrbracket_m &= \llbracket f_{\text{arg}} \leftarrow e; f_{\text{body}}; x \leftarrow f_{\text{res}} \rrbracket_m \\
\llbracket x \leftarrow \mathcal{A}(e) \rrbracket_m &= \llbracket a_{\text{arg}} \leftarrow e; a_{\text{body}}; x \leftarrow a_{\text{res}} \rrbracket_m
\end{aligned}$$

$$\llbracket s \rrbracket_{\mu} = \mathbb{E}_{m \sim \mu} [\llbracket s \rrbracket_m]$$

Fig. 1. Denotational semantics of programs

Semantics. The denotational semantics of programs is adapted from the seminal work of [27] and interprets programs as sub-distribution transformers. We view states as type-preserving mappings from variables to values; we write **State** for the set of states and **SDist(State)** for the set of probabilistic states. For each procedure name $f \in \mathcal{I} \cup \mathcal{A}$, we assume a set $\mathcal{X}_f^{\$} \subseteq \mathcal{X}$ of *local variables* s.t. $\mathcal{X}_f^{\$}$ are pairwise disjoint. The other variables $\mathcal{X} \setminus \bigcup_f \mathcal{X}_f^{\$}$ are *global variables*.

To define the interpretation of expressions and distribution expressions, we let $\llbracket e \rrbracket_m$ denote the interpretation of expression e with respect to state m , and $\llbracket e \rrbracket_{\mu}$ denote the interpretation of expression e with respect to an initial sub-distribution μ over states defined by the clause $\llbracket e \rrbracket_{\mu} \triangleq \mathbb{E}_{m \sim \mu} [\llbracket e \rrbracket_m]$. Likewise, we define the semantics of commands in two stages: first interpreted in a single input memory, then interpreted in an input sub-distribution over memories.

Definition 7. *The semantics of commands are given in Fig. 1.*

- The semantics $\llbracket s \rrbracket_m$ of a statement s in initial state m is a sub-distribution over states.
- The (lifted) semantics $\llbracket s \rrbracket_{\mu}$ of a statement s in initial sub-distribution μ over states is a sub-distribution over states.

We briefly comment on loops. The semantics of a loop **while** e **do** c is defined as the limit of its lower approximations, where the n -th *lower approximation* of $\llbracket \text{while } e \text{ do } c \rrbracket_{\mu}$ is $\llbracket (\text{if } e \text{ then } s)^n; \text{if } e \text{ then abort} \rrbracket_{\mu}$, where **if** e **then** s is shorthand for **if** e **then** s **else** **skip** and c^n is the n -fold composition $c; \dots; c$. Since the sequence is increasing, the limit is well-defined by Lemma 3. In contrast, the n -th *approximation* of $\llbracket \text{while } e \text{ do } c \rrbracket_{\mu}$ defined by $\llbracket (\text{if } e \text{ then } s)^n \rrbracket_{\mu}$ may not converge, since they are not necessarily increasing. However, in the special case

where the output distribution has weight 1, the n -th lower approximations and the n -th approximations have the same limit.

Lemma 4. *If the sub-distribution $\llbracket \mathbf{while } e \mathbf{ do } c \rrbracket_\mu$ has weight 1, then the limit of $\llbracket (\mathbf{if } e \mathbf{ then } s)^n \rrbracket_\mu$ is defined and*

$$\lim_{n \rightarrow \infty} \llbracket (\mathbf{if } e \mathbf{ then } s)^n; \mathbf{if } e \mathbf{ then abort} \rrbracket_\mu = \lim_{n \rightarrow \infty} \llbracket (\mathbf{if } e \mathbf{ then } s)^n \rrbracket_\mu.$$

This follows by Lemma 1, since lower approximations are below approximations so the limit of their weights (and the weight of their limit) is 1. It will be useful to identify programs that terminate with probability 1.

Definition 8 (Lossless). *A statement s is lossless if for every sub-distribution μ , $\llbracket s \rrbracket_\mu = |\mu|$, where $|\mu|$ is the total probability of μ . Programs that are not lossless are called lossy.*

Informally, a program is lossless if all probabilistic assignments sample from full distributions rather than sub-distributions, there are no **abort** instructions, and the program is almost surely terminating, i.e. infinite traces have probability zero. Note that if we restrict the language to sample from full distributions, then losslessness coincides with almost sure termination.

Another important class of loops are loops with a uniform upper bound on the number of iterations. Formally, we say that a loop **while** e **do** s is *certainly terminating* if there exists k such that for every sub-distribution μ , we have $|\llbracket \mathbf{while } e \mathbf{ do } s \rrbracket_\mu| = |\llbracket (\mathbf{if } e \mathbf{ then } s)^k \rrbracket_\mu|$. Note that certain termination of a loop does not entail losslessness—the output distribution of the loop may not have weight 1, for instance, if the loop samples from a sub-distribution or if the loop aborts with positive probability.

Semantics of Procedure Calls and Adversaries. The semantics of internal procedure calls is straightforward. Associated to each procedure name $f \in \mathcal{I}$, we assume a designated input variable $f_{\mathbf{arg}} \in \mathcal{X}_f^\Sigma$, a piece of code $f_{\mathbf{body}}$ that executes the function call, and a result expression $f_{\mathbf{res}}$. A function call $x \leftarrow \mathcal{I}(e)$ is then equivalent to $f_{\mathbf{arg}} \leftarrow e; f_{\mathbf{body}}; x \leftarrow f_{\mathbf{res}}$. Procedures are subject to well-formedness criteria: procedures should only use local variables in their scope and after initializing them, and should not perform recursive calls.

External procedure calls, also known as adversary calls, are a bit more involved. Each name $a \in \mathcal{A}$ is parametrized by a set $a_{\mathbf{ocl}} \subseteq \mathcal{I}$ of internal procedures which the adversary may call, a designated input variable $a_{\mathbf{arg}} \in \mathcal{X}_a^\Sigma$, a (unspecified) piece of code $a_{\mathbf{body}}$ that executes the function call, and a result expression $a_{\mathbf{res}}$. We assume that adversarial code can only access its local variables in \mathcal{X}_a^Σ and can only make calls to procedures in $a_{\mathbf{ocl}}$. It is possible to impose more restrictions on adversaries—say, that they are lossless—but for simplicity we do not impose additional assumptions on adversaries here.

4 Proof System

In this section we introduce a program logic for proving properties of probabilistic programs. The logic is abstract—assertions are arbitrary predicates on sub-distributions—but the meta-theoretic properties are clearest in this setting. In the following section, we will give a concrete version suitable for practical use.

Assertions and Closedness Conditions. We use predicates on state distribution.

Definition 9 (Assertions). *The set Assn of assertions is defined as $\mathcal{P}(\mathbf{SDist}(\mathbf{State}))$. We write $\eta(\mu)$ for $\mu \in \eta$.*

Usual set operations are lifted to assertions using their logical counterparts, e.g., $\eta \wedge \eta' \triangleq \eta \cap \eta'$ and $\neg\eta \triangleq \bar{\eta}$. Our program logic uses a few additional constructions. Given a predicate ϕ over states, we define

$$\Box\phi(\mu) \triangleq \forall m. m \in \text{supp}(\mu) \implies \phi(m)$$

where $\text{supp}(\mu)$ is the set of all states with non-zero probability under μ . Intuitively, ϕ holds deterministically on all states that we may sample from the distribution. To reason about branching commands, given two assertions η_1 and η_2 , we let

$$(\eta_1 \oplus \eta_2)(\mu) \triangleq \exists \mu_1, \mu_2. \mu = \mu_1 + \mu_2 \wedge \eta_1(\mu_1) \wedge \eta_2(\mu_2).$$

This assertion means that the sub-distribution is the sum of two sub-distributions such that η_1 holds on the first piece and η_2 holds on the second piece.

Given an assertion η and an event $E \subseteq \mathbf{State}$, we let $\eta|_E(\mu) \triangleq \eta(\mu|_E)$. This assertion holds exactly when η is true on the portion of the sub-distribution satisfying E . Finally, given an assertion η and a function F from $\mathbf{SDist}(\mathbf{State})$ to $\mathbf{SDist}(\mathbf{State})$, we define $\eta[F] \triangleq \lambda\mu. \eta(F(\mu))$. Intuitively, $\eta[F]$ is true in a sub-distribution μ exactly when η holds on $F(\mu)$.

Now, we can define the closedness properties of assertions. These properties will be critical to our rules for **while** loops.

Definition 10 (Closedness properties). *A family of assertions $(\eta_n)_{n \in \mathbb{N}^\infty}$ is:*

- *u-closed if for every increasing sequence of sub-distributions $(\mu_n)_{n \in \mathbb{N}}$ such that $\eta_n(\mu_n)$ for all $n \in \mathbb{N}$ then $\eta_\infty(\lim_{n \rightarrow \infty} \mu_n)$;*
- *t-closed if for every converging sequence of sub-distributions $(\mu_n)_{n \in \mathbb{N}}$ such that $\eta_n(\mu_n)$ for all $n \in \mathbb{N}$ then $\eta_\infty(\lim_{n \rightarrow \infty} \mu_n)$;*
- *d-closed if it is t-closed and downward closed, that is for every sub-distributions $\mu \leq \mu'$, $\eta_\infty(\mu')$ implies $\eta_\infty(\mu)$.*

When $(\eta_n)_n$ is constant and equal to η , we say that η is u-/t-/d-closed.

Note that *t*-closedness implies *u*-closedness, but the converse does not hold. Moreover, *u*-closed, *t*-closed and *d*-closed assertions are closed under arbitrary

intersections and finite unions, or in logical terms under finite boolean combinations, universal quantification over arbitrary sets and existential quantification over finite sets.

Finally, we introduce the necessary machinery for the frame rule. The set $\text{mod}(s)$ of *modified* variables of a statement s consists of all the variables on the left of a deterministic or probabilistic assignment. In this setting, we say that an assertion η is *separated* from a set of variables X , written $\text{separated}(\eta, X)$, if $\eta(\mu_1) \iff \eta(\mu_2)$ for any distributions μ_1, μ_2 s.t. $|\mu_1| = |\mu_2|$ and $\mu_1|_{\bar{X}} = \mu_2|_{\bar{X}}$ where for a set of variables X , the restricted sub-distribution $\mu|_X$ is

$$\mu|_X : m \in \mathbf{State}|_X \mapsto \Pr_{m' \sim \mu} [m = m'|_X]$$

where $\mathbf{State}|_X$ and $m|_X$ restrict \mathbf{State} and m to the variables in X .

Intuitively, an assertion is separated from a set of variables X if every two sub-distributions that agree on the variables outside X either both satisfy the assertion, or both refute the assertion.

Judgments and Proof Rules. Judgments are of the form $\{\eta\} s \{\eta'\}$, where the assertions η and η' are drawn from \mathbf{Assn} .

Definition 11. A judgment $\{\eta\} s \{\eta'\}$ is *valid*, written $\models \{\eta\} s \{\eta'\}$, if $\eta'(\llbracket s \rrbracket_\mu)$ for every interpretation of adversarial procedures and every probabilistic state μ such that $\eta(\mu)$.

Figure 2 describes the structural and basic rules of the proof system. Validity of judgments is preserved under standard structural rules, like the rule of consequence [CONSEQ]. As usual, the rule of consequence allows to weaken the post-condition and to strengthen the pre-condition; in our system, this rule serves as the interface between the program logic and mathematical theorems from probability theory. The [EXISTS] rule is helpful to deal with existentially quantified pre-conditions.

The rules for **skip**, assignments, random samplings and sequences are all straightforward. The rule for **abort** requires $\Box \perp$ to hold after execution; this assertion uniquely characterizes the resulting null sub-distribution. The rules for assignments and random samplings are semantical.

The rule [COND] for conditionals requires that the post-condition must be of the form $\eta_1 \oplus \eta_2$; this reflects the semantics of conditionals, which splits the initial probabilistic state depending on the guard, runs both branches, and recombines the resulting two probabilistic states.

The next two rules ([SPLIT] and [FRAME]) are useful for local reasoning. The [SPLIT] rule reflects the additivity of the semantics and combines the pre- and post-conditions using the \oplus operator. The [FRAME] rule asserts that lossless statements preserve assertions that are not influenced by modified variables.

The rule [CALL] for internal procedures is as expected, replacing the procedure call f with its definition.

Figure 3 presents the rules for loops. We consider four rules specialized to the termination behavior. The [WHILE] rule is the most general rule, as it deals with

$$\begin{array}{c}
\frac{\eta_0 \Rightarrow \eta_1 \quad \{\eta_1\} s \{\eta_2\} \quad \eta_2 \Rightarrow \eta_3}{\{\eta_0\} s \{\eta_3\}} [\text{CONSEQ}] \quad \frac{\forall x : T. \{\eta\} s \{\eta'\}}{\{\exists x : T. \eta\} s \{\eta'\}} [\text{EXISTS}] \\
\\
\frac{}{\{\eta\} \mathbf{abort} \{\Box \perp\}} [\text{ABORT}] \quad \frac{\eta' \triangleq \eta[\llbracket x \leftarrow e \rrbracket]}{\{\eta'\} x \leftarrow e \{\eta\}} [\text{ASSGN}] \quad \frac{}{\{\eta\} \mathbf{skip} \{\eta\}} [\text{SKIP}] \\
\\
\frac{\eta' \triangleq \eta[\llbracket x \stackrel{\$}{\leftarrow} g \rrbracket]}{\{\eta'\} x \stackrel{\$}{\leftarrow} g \{\eta\}} [\text{SAMPLE}] \quad \frac{\{\eta_0\} s_1 \{\eta_1\} \quad \{\eta_1\} s_2 \{\eta_2\}}{\{\eta_0\} s_1; s_2 \{\eta_2\}} [\text{SEQ}] \\
\\
\frac{\{\eta_1 \wedge \Box e\} s_1 \{\eta'_1\} \quad \{\eta_2 \wedge \Box \neg e\} s_2 \{\eta'_2\}}{\{(\eta_1 \wedge \Box e) \oplus (\eta_2 \wedge \Box \neg e)\} \mathbf{if} e \mathbf{then} s_1 \mathbf{else} s_2 \{\eta'_1 \oplus \eta'_2\}} [\text{COND}] \\
\\
\frac{\{\eta_1\} s \{\eta'_1\} \quad \{\eta_2\} s \{\eta'_2\}}{\{\eta_1 \oplus \eta_2\} s \{\eta'_1 \oplus \eta'_2\}} [\text{SPLIT}] \\
\\
\frac{\text{separated}(\eta, \text{mod}(s)) \quad s \text{ is lossless}}{\{\eta\} s \{\eta\}} [\text{FRAME}] \\
\\
\frac{\{\eta\} f_{\mathbf{arg}} \leftarrow e; f_{\mathbf{body}} \{\eta'[\llbracket x \leftarrow f_{\mathbf{res}} \rrbracket]\}}{\{\eta\} x \leftarrow f(e) \{\eta'\}} [\text{CALL}]
\end{array}$$

Fig. 2. Structural and basic rules

arbitrary loops. For simplicity, we explain the rule in the special case where the family of assertions is constant, i.e. we have $\eta_n = \eta$ and $\eta'_n = \eta'$. Informally, the η is the loop invariant and η' is an auxiliary assertion used to prove the invariant. We require that η is u -closed, since the semantics of a loop is defined as the limit of its lower approximations. Moreover, the first premise ensures that starting from η , one guarded iteration of the loop establishes η' ; the second premise ensures that restricting to $\neg e$ a probabilistic state μ' satisfying η' yields a probabilistic state μ satisfying η . It is possible to give an alternative formulation where the second premise is substituted by the logical constraint $\eta'_{|\neg e} \implies \eta$. As usual, the post-condition of the loop is the conjunction of the invariant with the negation of the guard (more precisely in our setting, that the guard has probability 0).

The [WHILE-AST] rule deals with lossless loops. For simplicity, we explain the rule in the special case where the family of assertions is constant, i.e. we have $\eta_n = \eta$. In this case, we know that lower approximations and approximations have the same limit, so we can directly prove an invariant that holds after one guarded iteration of the loop. On the other hand, we must now require that the η satisfies the stronger property of t -closedness.

The [WHILE-D] rule handles arbitrary loops with a d -closed invariant; intuitively, restricting a sub-distribution that satisfies a downwards closed assertion η yields a sub-distribution which also satisfies η .

The [WHILE-CT] rule deals with certainly terminating loops. In this case, there is no requirement on the assertions.

We briefly compare the rules from a verification perspective. If the assertion is d -closed, then the rule [WHILE-D] is easier to use, since there is no need to prove any termination requirement. Alternatively, if we can prove certain termination of the loop, then the rule [WHILE-CT] is the best to use since it does not impose any condition on assertions. When the loop is lossless, there is no need to introduce an auxiliary assertion η' , which simplifies the proof goal. Note however that it might still be beneficial to use the [WHILE] rule, even for lossless loops, because of the weaker requirement that the invariant is u -closed rather than t -closed.

Finally, Fig. 4 gives the adversary rule for general adversaries. It is highly similar to the general rule [WHILE-D] for loops since the adversary may make an arbitrary sequence of calls to the oracles in a_{ocl} and may not be lossless. Intuitively, η plays the role of the invariant: it must be d -closed and it must be preserved by every oracle call with arbitrary arguments. If this holds, then η is also preserved by the adversary call. Some framing conditions are required, similar to the ones of the [FRAME] rule: the invariant must not be influenced by the state writable by the external procedures.

It is possible to give other variants of the adversary rule with more general invariants by restricting the adversary, e.g., requiring losslessness or bounding the number of calls the external procedure can make to oracles, leading to rules akin to the almost surely terminating and certainly terminating loop rules, respectively.

$$\begin{array}{c}
\frac{\text{uclosed}((\eta'_n)_{n \in \mathbb{N}^\infty}) \quad \forall n. \{ \eta_n \} \text{ if } e \text{ then } s \{ \eta_{n+1} \} \quad \forall n. \{ \eta_n \} \text{ if } e \text{ then abort } \{ \eta'_n \}}{\{ \eta_0 \} \text{ while } e \text{ do } s \{ \eta'_\infty \wedge \Box \neg e \}} \text{ [WHILE]} \\
\\
\frac{\text{tclosed}((\eta_n)_{n \in \mathbb{N}^\infty}) \quad \forall n. \{ \eta_n \} \text{ if } e \text{ then } s \{ \eta_{n+1} \} \quad \forall \mu. \eta_0(\mu) \implies |\llbracket (\text{while } e \text{ do } s) \rrbracket_\mu| = 1}{\{ \eta_0 \} \text{ while } e \text{ do } s \{ \eta_\infty \wedge \Box \neg e \}} \text{ [WHILE-AST]} \\
\\
\frac{\text{dclosed}((\eta_n)_{n \in \mathbb{N}^\infty}) \quad \forall n. \{ \eta_n \} \text{ if } e \text{ then } s \{ \eta_{n+1} \}}{\{ \eta_0 \} \text{ while } e \text{ do } s \{ \eta_\infty \wedge \Box \neg e \}} \text{ [WHILE-D]} \\
\\
\frac{\forall \mu. \eta_0(\mu) \implies \llbracket (\text{if } e \text{ then } s)^k \rrbracket_\mu = \llbracket (\text{while } e \text{ do } s) \rrbracket_\mu \quad \forall n. \{ \eta_n \} \text{ if } e \text{ then } s \{ \eta_{n+1} \}}{\{ \eta_0 \} \text{ while } e \text{ do } s \{ \eta_k \wedge \Box \neg e \}} \text{ [WHILE-CT]}
\end{array}$$

Fig. 3. Rules for loops

$$\frac{\forall n \in \mathbb{N}^\infty. \text{separated}(\eta_n, \{x, \mathfrak{s}\}) \quad \text{dclosed}((\eta_n)_{n \in \mathbb{N}^\infty}) \quad \forall f \in a_{\text{ocl}}, x \in \mathcal{X}_a^\mathfrak{g}, e \in \mathcal{E}, n \in \mathbb{N}. \{ \eta_n \} x \leftarrow f(e) \{ \eta_{n+1} \}}{\{ \eta_0 \} x \leftarrow a(e) \{ \eta_\infty \}} \text{ [ADV]}$$

Fig. 4. Rules for adversaries

Soundness and Relative Completeness. Our proof system is sound and relatively complete with respect to the semantics; these proofs have also been formalized in the COQ proof assistant.

Theorem 1 (Soundness). *Every judgment $\{\eta\} s \{\eta'\}$ provable using the rules of our logic is valid.*

Completeness of the logic follows from the next lemma, whose proof makes an essential use of the [WHILE] rule. In the sequel, we use $\mathbf{1}_\mu$ to denote the characteristic function of a probabilistic state μ , an assertion stating that the current state is equal to μ .

Lemma 5. *For every probabilistic state μ , the following judgment is provable using the rule of the logic:*

$$\{\mathbf{1}_\mu\} s \{\mathbf{1}_{\llbracket s \rrbracket_\mu}\}.$$

Proof. By induction on the structure of s .

- $s = \mathbf{abort}$, $s = \mathbf{skip}$, $x \leftarrow e$ and $s = x \stackrel{\$}{\leftarrow} g$ are trivial;
- $s = s_1; s_2$, we have to prove

$$\{\mathbf{1}_\mu\} s_1; s_2 \{\mathbf{1}_{\llbracket s_2 \rrbracket_{\llbracket s_1 \rrbracket_\mu}}\}.$$

We apply the [SEQ] rule with $\eta_1 = \mathbf{1}_{\llbracket s_1 \rrbracket_\mu}$ premises can be directly proved using the induction hypothesis;

- $s = \mathbf{if } e \mathbf{ then } s_1 \mathbf{ else } s_2$, we have to prove

$$\{\mathbf{1}_\mu\} \mathbf{if } e \mathbf{ then } s_1 \mathbf{ else } s_2 \{(\mathbf{1}_{\llbracket s_1 \rrbracket_{\mu|e}} \oplus \mathbf{1}_{\llbracket s_2 \rrbracket_{\mu|\neg e}})\}.$$

We apply the [CONSEQ] rule to be able to apply the the [COND] rule with $\eta_1 = \mathbf{1}_{\llbracket s_1 \rrbracket_{\mu|e}}$ and $\eta_2 = \mathbf{1}_{\llbracket s_2 \rrbracket_{\mu|\neg e}}$. Both premises can be proved by an application of the [CONSEQ] rule followed by the application of the induction hypothesis.

- $s = \mathbf{while } e \mathbf{ do } s$, we have to prove

$$\{\mathbf{1}_\mu\} \mathbf{while } e \mathbf{ do } s \{\mathbf{1}_{\lim_{n \rightarrow \infty} \llbracket (\mathbf{if } e \mathbf{ then } s)^n; \mathbf{if } e \mathbf{ then abort} \rrbracket_\mu}\}.$$

We first apply the [WHILE] rule with $\eta'_n = \mathbf{1}_{\llbracket (\mathbf{if } e \mathbf{ then } s)^n \rrbracket_\mu}$ and

$$\eta_n = \mathbf{1}_{\llbracket (\mathbf{if } e \mathbf{ then } s)^n; \mathbf{if } e \mathbf{ then abort} \rrbracket_\mu}.$$

For the first premise we apply the same process as for the conditional case: we apply the [CONSEQ] and [COND] rules and we conclude using the induction hypothesis (and the [SKIP] rule). For the second premise we follow the same process but we conclude using the [ABORT] rule instead of the induction hypothesis. Finally we conclude since $\mathbf{uclosed}((\eta_n)_{n \in \mathbb{N}^\infty})$. \square

The abstract logic is also relatively complete. This property will be less important for our purposes, but it serves as a basic sanity check.

Theorem 2 (Relative completeness). *Every valid judgment is derivable.*

Proof. Consider a valid judgment $\{\eta\} s \{\eta'\}$. Let μ be a probabilistic state such that $\eta(\mu)$. By the above proposition, $\{\mathbf{1}_\mu\} s \{\mathbf{1}_{\llbracket s \rrbracket_\mu}\}$. Using the validity of the judgment and [CONSEQ], we have $\{\mathbf{1}_\mu \wedge \eta(\mu)\} s \{\eta'\}$. Using the [EXISTS] and [CONSEQ] rules, we conclude $\{\eta\} s \{\eta'\}$ as required. \square

The side-conditions in the loop rules (e.g., `uclosed`/`tclosed`/`dclosed` and the weight conditions) are difficult to prove, since they are semantic properties. Next, we present a concrete version of the logic with give easy-to-check, syntactic sufficient conditions.

5 A Concrete Program Logic

To give a more practical version of the logic, we begin by setting a concrete syntax for assertions

Assertions. We use a two-level assertion language, presented in Fig. 5. A *probabilistic assertion* η is a formula built from comparison of probabilistic expressions, using first-order quantifiers and connectives, and the special connective \oplus . A *probabilistic expression* p can be a logical variable v , an operator applied to probabilistic expressions $o(\mathbf{p})$ (constants are 0-ary operators), or the expectation $\mathbb{E}[\tilde{e}]$ of a state expression \tilde{e} . A *state expression* \tilde{e} is either a program variable x , the characteristic function $\mathbf{1}_\phi$ of a state assertion ϕ , an operator applied to state expressions $o(\tilde{e})$, or the expectation $\mathbb{E}_{v \sim g}[\tilde{e}]$ of state expression \tilde{e} in a given distribution g . Finally, a *state assertion* ϕ is a first-order formula over program variables. Note that the set of operators is left unspecified but we assume that all the expressions in \mathcal{E} and \mathcal{D} can be encoded by operators.

$$\begin{aligned}
\tilde{e} &::= x \mid v \mid \mathbf{1}_\phi \mid \mathbb{E}_{v \sim g}[\tilde{e}] \mid o(\tilde{e}) & (\text{S-expr.}) \\
\phi &::= \tilde{e} \bowtie \tilde{e} \mid FO(\phi) & (\text{S-assn.}) \\
p &::= v \mid o(\mathbf{p}) \mid \mathbb{E}[\tilde{e}] & (\text{P-expr.}) \\
\eta &::= p \bowtie p \mid \eta \oplus \eta \mid FO(\eta) & (\text{P-assn.}) \\
\bowtie &\in \{=, <, \leq\} \quad o \in Ops & (\text{Ops.})
\end{aligned}$$

Fig. 5. Assertion syntax

The interpretation of the concrete syntax is as expected. The interpretation of probabilistic assertions is relative to a valuation ρ which maps logical variables to values, and is an element of **Assn**. The definition of the interpretation is straightforward; the only interesting case is $\llbracket \mathbb{E}[\tilde{e}] \rrbracket_\mu^\rho$ which is defined by $\mathbb{E}_{m \sim \mu}[\llbracket \tilde{e} \rrbracket_m^\rho]$, where $\llbracket \tilde{e} \rrbracket_m^\rho$ is the interpretation of the state expression \tilde{e} in the memory m and valuation ρ . The interpretation of state expressions is a mapping from memories to values, which can be lifted to a mapping from distributions over memories to distributions over values. The definition of the interpretation is straightforward; the most interesting case is for expectation $\llbracket \mathbb{E}_{v \sim g}[\tilde{e}] \rrbracket_m^\rho \triangleq \mathbb{E}_{w \sim \llbracket g \rrbracket_m^\rho}[\llbracket \tilde{e} \rrbracket_m^{\rho[v:=w]}]$. We present the full interpretations in the supplemental materials.

Many standard concepts from probability theory have a natural representation in our syntax. For example:

- the probability that ϕ holds in some probabilistic state is represented by the probabilistic expression $\Pr[\phi] \triangleq \mathbb{E}[\mathbf{1}_\phi]$;
- probabilistic independence of state expressions $\tilde{e}_1, \dots, \tilde{e}_n$ is modeled by the probabilistic assertion $\#\{\tilde{e}_1, \dots, \tilde{e}_n\}$, defined by the clause¹¹

$$\forall v_1 \dots v_n, \Pr[\top]^{n-1} \Pr[\bigwedge_{i=1 \dots n} \tilde{e}_i = v_i] = \prod_{i=1 \dots n} \Pr[\tilde{e}_i = v_i];$$

- the fact that a distribution is proper is modeled by the probabilistic assertion $\mathcal{L} \triangleq \Pr[\top] = 1$;
- a state expression \tilde{e} distributed according to a law g is modeled by the probabilistic assertion

$$\tilde{e} \sim g \triangleq \forall w, \Pr[\tilde{e} = w] = \mathbb{E}[\mathbb{E}_{v \sim g}[\mathbf{1}_{v=w}]].$$

The inner expectation computes the probability that v drawn from g is equal to a fixed w ; the outer expectation weights the inner probability by the probability of each value of w .

We can easily define \Box operator from the previous section in our new syntax: $\Box\phi \triangleq \Pr[\neg\phi] = 0$.

Syntactic Proof Rules. Now that we have a concrete syntax for assertions, we can give syntactic versions of many of the existing proof rules. Such proof rules are often easier to use since they avoid reasoning about the semantics of commands and assertions. We tackle the non-looping rules first, beginning with the following syntactic rules for assignment and sampling:

$$\frac{}{\{\eta[x := e]\} \ x \leftarrow e \ \{\eta\}} \text{[ASSGN]} \qquad \frac{}{\{\mathcal{P}_x^g(\eta)\} \ x \stackrel{*}{\leftarrow} g \ \{\eta\}} \text{[SAMPLE]}$$

The rule for assignment is the usual rule from Hoare logic, replacing the program variable x by its corresponding expression e in the pre-condition. The replacement $\eta[x := e]$ is done recursively on the probabilistic assertion η ; for instance for expectations, it is defined by $\mathbb{E}[\tilde{e}][x := e] \triangleq \mathbb{E}[\tilde{e}[x := e]]$, where $\tilde{e}[x := e]$ is the syntactic substitution.

The rule for sampling uses probabilistic substitution operator $\mathcal{P}_x^g(\eta)$, which replaces all occurrences of x in η by a new integration variable t and records that t is drawn from g ; the operator is defined in Fig. 6.

Next, we turn to the loop rule. The side-conditions from Fig. 3 are purely semantic, while in practice it is more convenient to use a sufficient condition in the Hoare logic. We give sufficient conditions for ensuring certain and almost-sure termination in Fig. 7; \tilde{e} is an integer-valued expression. The first side-condition $\mathcal{C}_{\text{CTerm}}$ shows certain termination given a strictly decreasing *variant* \tilde{e}

¹¹ The term $\Pr[\top]^{n-1}$ is necessary since we work with sub-distributions.

$$\begin{aligned}
\mathcal{P}_x^g(v) &\triangleq v \\
\mathcal{P}_x^g(\mathbb{E}[\tilde{e}]) &\triangleq \mathbb{E}[\mathbb{E}_{t \sim g}[\tilde{e}[x := t]]] \\
\mathcal{P}_x^g(o(\boldsymbol{\eta})) &\triangleq o(\mathcal{P}_x^g(\eta_1), \dots, \mathcal{P}_x^g(\eta_n)) \\
\mathcal{P}_x^g(\eta_1 \bowtie \eta_2) &\triangleq \mathcal{P}_x^g(\eta_1) \bowtie \mathcal{P}_x^g(\eta_2)
\end{aligned}$$

for $o \in \mathbf{Ops}, \bowtie \in \{\wedge, \vee, \Rightarrow\}$.

Fig. 6. Syntactic op. \mathcal{P} (main cases)

$$\begin{aligned}
\mathcal{C}_{\text{CTerm}} &\triangleq \{\mathcal{L} \wedge \Box(\tilde{e} = k \wedge 0 < k \wedge b)\} \mid s \{\mathcal{L} \wedge \Box(\tilde{e} < k)\} \\
&\models \eta \Rightarrow (\exists \dot{y}. \Box \tilde{e} \leq \dot{y}) \wedge \Box(\tilde{e} = 0 \Rightarrow \neg b) \\
\mathcal{C}_{\text{ASTerm}} &\triangleq \{\mathcal{L} \wedge \Box(\tilde{e} = k \wedge 0 < k \leq K \wedge b)\} \mid s \{\mathcal{L} \wedge \Box(0 \leq \tilde{e} \leq K) \wedge \Pr[\tilde{e} < k] \geq \epsilon\} \\
&\models \eta \Rightarrow \Box(0 \leq \tilde{e} \leq K \wedge \tilde{e} = 0 \Rightarrow \neg b) \\
&\models \text{tclosed}(\eta)
\end{aligned}$$

Fig. 7. Side-conditions for loop rules

that is bounded below, similar to how a decreasing variant shows termination for deterministic programs. The second side-condition $\mathcal{C}_{\text{ASTerm}}$ shows almost-sure termination given a probabilistic variant \tilde{e} , which must be bounded both above and below. While \tilde{e} may increase with some probability, it must decrease with strictly positive probability. This condition was previously considered by [17] for probabilistic transition systems and also used in expectation-based approaches [33,20]. Our framework can also support more refined conditions (e.g., based on super-martingales [9,31]), but the condition $\mathcal{C}_{\text{ASTerm}}$ already suffices for most randomized algorithms.

While t -closedness is a semantic condition (cf. Definition 10), there are simple syntactic conditions to guarantee it. For instance, assertions that carry a non-strict comparison $\bowtie \in \{\leq, \geq, =\}$ between two bounded probabilistic expressions are t -closed; the assertion stating probabilistic independence of a set of expressions is t -closed.

Precondition Calculus. With a concrete syntax for assertions, we are also able to incorporate syntactic reasoning principles. One classic tool is Morgan and McIver’s *greatest pre-expectation*, which we take as inspiration for a pre-condition calculus for the loop-free fragment of ELLORA. Given an assertion η and a loop-free statement s , we mechanically construct an assertion η^* that is the pre-condition of s that implies η as a post-condition. The basic idea is to replace each expectation expression p inside η by an expression p^* that has the same denotation before running s as p after running s . This process yields an assertion η^* that, interpreted before running s , is logically equivalent to η interpreted after running s .

The computation rules for pre-conditions are defined in Fig. 8. For a probability assertion η , its pre-condition $\text{pc}(s, \eta)$ corresponds to η where the expectation ex-

pressions of the form $\mathbb{E}[\tilde{e}]$ are replaced by their corresponding *pre-term*, $\text{pe}(s, \mathbb{E}[\tilde{e}])$. Pre-terms correspond loosely to Morgan and McIver’s *pre-expectations*—we will make this correspondence more precise in the next section. The main interesting cases for computing pre-terms are for random sampling and conditionals. For random sampling the result is $\mathcal{P}_x^g(\mathbb{E}[\tilde{e}])$, which corresponds to the [SAMPLE] rule. For conditionals, the expectation expression is split into a part where e is true and a part where e is not true. We restrict the expectation to a part satisfying e with the operator $\mathbb{E}[\tilde{e}]|_e \triangleq \mathbb{E}[\tilde{e} \cdot \mathbf{1}_e]$. This corresponds to the expected value of \tilde{e} on the portion of the distribution where e is true. Then, we can build the pre-condition calculus into ELLORA.

$$\begin{aligned}
\text{pe}(s_1; s_2, \mathbb{E}[\tilde{e}]) &\triangleq \text{pe}(s_1, \text{pe}(s_2, \mathbb{E}[\tilde{e}])) \\
\text{pe}(x \leftarrow e, \mathbb{E}[\tilde{e}]) &\triangleq \mathbb{E}[\tilde{e}][x := e] \\
\text{pe}(x \stackrel{\$}{\leftarrow} g, \mathbb{E}[\tilde{e}]) &\triangleq \mathcal{P}_x^g(\mathbb{E}[\tilde{e}]) \\
\text{pe}(\text{if } e \text{ then } s_1 \text{ else } s_2, \mathbb{E}[\tilde{e}]) &\triangleq \text{pe}(s_1, \mathbb{E}[\tilde{e}]|_e) + \text{pe}(s_2, \mathbb{E}[\tilde{e}]|_{\neg e}) \\
\text{pc}(s, p_1 \bowtie p_2) &\triangleq \text{pe}(s, p_1) \bowtie \text{pe}(s, p_2)
\end{aligned}$$

Fig. 8. Precondition calculus (selected)

Theorem 1. *Let s be a non-looping command. Then, the following rule is derivable in the concrete version of ELLORA:*

$$\frac{}{\{pc(s, \eta)\} s \{\eta\}} \text{ [PC]}$$

6 Case Studies: Embedding Lightweight Logics

While ELLORA is suitable for general-purpose reasoning about probabilistic programs, in practice humans typically use more special-purpose proof techniques—often targeting just a single, specific kind of property, like probabilistic independence—when proving probabilistic assertions. When these techniques apply, they can be a convenient and powerful tool.

To capture this intuitive style of reasoning, researchers have considered lightweight program logics where the assertions and proof rules are tailored to a specific proof technique. We demonstrate how to integrate these tools in an assertion-based logic by introducing and embedding a new logic for reasoning about independence and distribution laws, useful properties when analyzing randomized algorithms. We crucially rely on the rich assertions in ELLORA—it is not clear how to extend expectation-based approaches to support similar, lightweight reasoning. Then, we show to embed the union bound logic [4] for proving accuracy bounds.

6.1 Law and Independence Logic

We begin by describing the law and independence logic IL, a proof system with intuitive rules that are easy to apply and amenable to automation. For simplicity, we only consider programs which sample from the binomial distribution, and have deterministic control flow—for lack of space, we also omit procedure calls.

Definition 12 (Assertions). IL *assertions have the grammar:*

$$\xi := \text{det}(e) \mid \# E \mid e \sim B(e, p) \mid \top \mid \perp \mid \xi \wedge \xi$$

where $e \in \mathcal{E}$, $E \subseteq \mathcal{E}$, and $p \in [0, 1]$.

The assertion $\text{det}(e)$ states that e is deterministic in the current distribution, i.e., there is at most one element in the support of its interpretation. The assertion $\# E$ states that the expressions in E are independent, as formalized in the previous section. The assertion $e \sim B(m, p)$ states that e is distributed according to a binomial distribution with parameter m (where m can be an expression) and constant probability p , i.e. the probability that $e = k$ is equal to the probability that exactly k independent coin flips return heads using a biased coin that returns heads with probability p .

Assertions can be seen as an instance of a logical abstract domain, where the order between assertions is given by implication based on a small number of axioms. Examples of such axioms include independence of singletons, irreflexivity of independence, anti-monotonicity of independence, an axiom for the sum of binomial distributions, and rules for deterministic expressions:

$$\begin{aligned} \#\{x\} \quad \quad \quad \#\{x, x\} &\iff \text{det}(x) & \quad \quad \quad \#(E \cup E') &\implies \# E \\ e \sim B(m, p) \wedge e' \sim B(m', p) \wedge \#\{e, e'\} &\implies e + e' \sim B(m + m', p) \\ \bigwedge_{1 \leq i \leq n} \text{det}(e_i) &\implies \text{det}(f(e_1, \dots, e_n)) \end{aligned}$$

Definition 13. *Judgments of the logic are of the form $\{\xi\} \text{ s } \{\xi'\}$, where ξ and ξ' are IL-assertions. A judgment is valid if it is derivable from the rules of Fig. 9; structural rules and rule for sequential composition are similar to those from § 4 and omitted.*

The rule [IL-ASSGN] for deterministic assignments is as in § 4. The rule [IL-SAMPLE] for random assignments yields as post-condition that the variable x and a set of expressions E are independent assuming that E is independent before the sampling, and moreover that x follows the law of the distribution that it is sampled from. The rule [IL-COND] for conditionals requires that the guard is deterministic, and that each of the branches satisfies the specification; if the guard

is not deterministic, there are simple examples where the rule is not sound.¹² The rule [IL-WHILE] for loops requires that the loop is certainly terminating with a deterministic guard. Note that the requirement of certain termination could be avoided by restricting the structural rules such that a statement s has deterministic control flow whenever $\{\xi\} s \{\xi'\}$ is derivable.

We now turn to the embedding. The embedding of IL assertions into general assertions is immediate, except for $\text{det}(e)$ which is translated as $\Box e \vee \Box \neg e$. We let $\bar{\xi}$ denote the translation of ξ .

Theorem 2 (Embedding and soundness of IL logic). *If $\{\xi\} s \{\xi'\}$ is derivable in the IL logic, then $\{\bar{\xi}\} s \{\bar{\xi}'\}$ is derivable in (the syntactic variant of) ELLORA. As a consequence, every derivable judgment $\{\xi\} s \{\xi'\}$ is valid.*

Proof sketch. By induction on the derivation. The interesting cases are conditionals and loops. For conditionals, the soundness follows from the soundness of the rule:

$$\frac{\{\eta\} s_1 \{\eta'\} \quad \{\eta\} s_2 \{\eta'\} \quad \Box e \vee \Box \neg e}{\{\eta\} \text{ if } e \text{ then } s_1 \text{ else } s_2 \{\eta'\}}$$

To prove the soundness of this rule, we proceed by case analysis on $\Box e \vee \Box \neg e$. We treat the case $\Box e$; the other case is similar. In this case, η is equivalent to $\eta_1 \wedge \Box e \oplus \eta_2 \wedge \Box \neg e$, where $\eta_1 = \eta$ and $\eta_2 = \perp$. Let $\eta'_1 = \eta'$ and $\eta'_2 = \Box \perp$; again, $\eta'_1 \oplus \eta'_2$ is logically equivalent to η' . The soundness of the rule thus follows from the soundness of the [COND] and [CONSEQ] rules. For loops, there exists a natural number n such that **while** b **do** s is semantically equivalent to **(if** b **then** s) ^{n} . By assumption $\{\xi\} s \{\xi\}$ holds, and thus by induction hypothesis $\{\bar{\xi}\} s \{\bar{\xi}\}$. We also have $\xi \implies \text{det}(b)$, and hence $\{\bar{\xi}\} \text{ if } b \text{ then } s \{\bar{\xi}\}$. We conclude by [SEQ]. \square

To illustrate our system IL, consider the statement s in Fig. 10 which flips a fair coin N times and counts the number of heads. Using the logic, we prove that $c \sim B(N \cdot (N + 1)/2, 1/2)$ is a post-condition for s . We take the invariant:

$$c \sim B(j(j + 1)/2, 1/2)$$

The invariant holds initially, as $0 \sim B(0, 1/2)$. For the inductive case, we show:

$$\{c \sim B(0, 1/2)\} s_0 \{c \sim B((j + 1)(j + 2)/2, 1/2)\}$$

where s_0 represents the loop body, i.e. $x \stackrel{\$}{\leftarrow} B(j, 1/2); c \leftarrow c + x$. First, we apply the rule for sequence taking as intermediate assertion

$$c \sim B(j(j + 1)/2, 1/2) \wedge x \sim B(j, 1/2) \wedge \#\{x, c\}$$

¹² Consider the following program where $\text{Bern}(p)$ is the Bernoulli distribution with parameter p :

$$b \stackrel{\$}{\leftarrow} \text{Bern}(p); \text{ if } b \text{ then } x_1 \stackrel{\$}{\leftarrow} \text{Bern}(p_1); x_2 \stackrel{\$}{\leftarrow} \text{Bern}(p_2) \\ \text{ else } x_1 \stackrel{\$}{\leftarrow} \text{Bern}(p'_1); x_2 \stackrel{\$}{\leftarrow} \text{Bern}(p'_2)$$

Each branch establishes $\#\{x_1, x_2\}$, but this is not a valid post-condition for the conditional. There are similar examples using the binomial distribution.

$$\begin{array}{c}
\frac{}{\{\xi[x := e]\} \ x \leftarrow e \ \{\xi\}} \text{ [IL-ASSGN]} \\
\\
\frac{\{x\} \cap \text{FV}(E) \cap \text{FV}(e) = \emptyset}{\{\# E\} \ x \xleftarrow{\$} B(e, p) \ \{\#(E \cup \{x\}) \wedge x \sim B(e, p)\}} \text{ [IL-SAMPLE]} \\
\\
\frac{\{\xi\} \ s_1 \ \{\xi'\} \quad \{\xi'\} \ s_2 \ \{\xi''\}}{\{\xi\} \ s_1; s_2 \ \{\xi''\}} \text{ [IL-SEQ]} \\
\\
\frac{\{\xi\} \ s_1 \ \{\xi'\} \quad \{\xi\} \ s_2 \ \{\xi'\} \quad \xi \implies \text{det}(b)}{\{\xi\} \ \text{if } b \text{ then } s_1 \text{ else } s_2 \ \{\xi'\}} \text{ [IL-COND]} \\
\\
\frac{\{\xi\} \ s \ \{\xi\} \quad \xi \implies \text{det}(b) \quad \mathcal{C}_{\text{CTerm}}}{\{\xi\} \ \text{while } b \text{ do } s \ \{\xi\}} \text{ [IL-WHILE]}
\end{array}$$

Fig. 9. IL proof rules (selected)

```

proc sum () =
  var c:int, x:int;
  c  $\leftarrow$  0;
  for j  $\leftarrow$  1 to N do
    x  $\xleftarrow{\$}$  B(j, 1/2);
    c  $\leftarrow$  c + x;
  return c

```

Fig. 10. Sum of bin.

The first premise follows from the rule for random assignment and structural rules. The second premise follows from the rule for deterministic assignment and the rule of consequence, applying axioms about sums of binomial distributions.

We briefly comment on several limitations of IL. First, IL is restricted to programs with deterministic control flow, but this restriction could be partially relaxed by enriching IL with assertions for conditional independence. Such assertions are already expressible in the logic of ELLORA; adding conditional independence would significantly broaden the scope of the IL proof system and open the possibility to rely on axiomatizations of conditional independence (e.g., based on graphoids [36]). Second, the logic only supports sampling from binomial distributions. It is possible to enrich the language of assertions with clauses $c \sim g$ where g can model other distributions, like the uniform distribution or the Laplace distribution. The main design challenge is finding a core set of useful facts about these distributions. Enriching the logic and automating the analysis are interesting avenues for further work.

6.2 Embedding the Union Bound Logic

The program logic AHL [4] was recently introduced for estimating accuracy of randomized computations. One main application of AHL is proving accuracy of randomized algorithms, both in the offline and online settings—i.e. with adversary calls. AHL is based on the union bound, a basic tool from probability theory, and has judgments of the form $\models_{\beta} \{\Phi\} s \{\Psi\}$, where s is a statement, Φ and Ψ are first-order formulae over program variables, and β is a probability, i.e. $\beta \in [0, 1]$. A judgment $\models_{\beta} \{\Phi\} s \{\Psi\}$ is valid if for every memory m such that $\Phi(m)$, the probability of $\neg\Psi$ in $\llbracket s \rrbracket_m$ is upper bounded by β , i.e. $\Pr_{\llbracket s \rrbracket_m}[\neg\Psi] \leq \beta$.

Figure 11 presents some key rules of AHL, including a rule for sampling from the Laplace distribution \mathcal{L}_{ϵ} centered around e . The predicate $\mathcal{C}_{\text{CTerm}}(k)$ indicates that the loop terminates in at most k steps on any memory that satisfies the pre-condition. Moreover, β is a function of ϵ .

$$\begin{array}{c}
\frac{}{\models_{\beta} \{\top\} x \leftarrow \mathcal{L}_{\epsilon}(e) \{ |x - e| \leq \frac{1}{\epsilon} \log \frac{1}{\beta} \}} \text{ [AHL-SAMPLE]} \\
\\
\frac{\models_{\beta_1} \{\Phi\} s_1 \{\Theta\} \quad \models_{\beta_2} \{\Theta\} s_2 \{\Psi\}}{\models_{\beta_1 + \beta_2} \{\Phi\} s_1; s_2 \{\Psi\}} \text{ [AHL-SEQ]} \\
\\
\frac{\models_{\beta} \{\Phi\} c \{\Phi\} \quad \mathcal{C}_{\text{CTerm}}(k)}{\models_{k \cdot \beta} \{\Phi\} \text{ while } e \text{ do } c \{ \Phi \wedge \neg e \}} \text{ [AHL-WHILE]}
\end{array}$$

Fig. 11. AHL proof rules (selected)

AHL has a simple embedding into ELLORA.

Theorem 3 (Embedding of AHL). *If $\models_{\beta} \{\Phi\} s \{\Psi\}$ is derivable in AHL, then $\{\Box\Phi\} s \{\mathbb{E}[\mathbf{1}_{\neg\Psi}] \leq \beta\}$ is derivable in ELLORA.*

7 Case Studies: Verifying Randomized Algorithms

In this section, we will demonstrate ELLORA on a selection of examples; we present further examples in the supplemental material. Together, they exhibit a wide variety of different proof techniques and reasoning principles which are available in the ELLORA’s implementation.

Hypercube Routing. will begin with the *hypercube routing* algorithm [41,42]. Consider a network topology (the *hypercube*) where each node is labeled by a bitstring of length D and two nodes are connected by an edge if and only if the two corresponding labels differ in exactly one bit position.

In the network, there is initially one packet at each node, and each packet has a unique destination. The algorithm implements a routing strategy based on *bit fixing*: if the current position has bitstring i , and the target node has bitstring j , we compare the bits in i and j from left to right, moving along the edge that corrects the first differing bit. Valiant’s algorithm uses randomization to guarantee that the total number of steps grows *logarithmically* in the number of packets. In the first phase, each packet i select an intermediate destination $\rho(i)$ uniformly at random, and use bit fixing to reach $\rho(i)$. In the second phase, each packet use bit fixing to go from $\rho(i)$ to the destination j . We will focus on the first phase since the reasoning for the second phase is nearly identical. We can model the strategy with the code in Figure 18, using some syntactic sugar for the **for** loops.¹³ We assume that initially, the position of the packet i is at node

```

proc route (D T : int) :
  var  $\rho$ , pos, usedBy : node map;
  var nextE : edge;
  pos  $\leftarrow$  Map.init id  $2^D$ ;  $\rho \leftarrow$  Map.empty;
  for  $i \leftarrow 1$  to  $2^D$  do
     $\rho[i] \xleftarrow{\$}[1, 2^D]$ 
    for  $t \leftarrow 1$  to  $T$  do
      usedBy  $\leftarrow$  Map.empty;
      for  $i \leftarrow 1$  to  $2^D$  do
        if pos[i]  $\neq$   $\rho[i]$  then
          nextE  $\leftarrow$  getEdge pos[i]  $\rho[i]$ ;
          if usedBy[nextE] =  $\perp$  then
            // Mark edge used
            usedBy[nextE]  $\leftarrow$  i;
            // Move packet
            pos[i]  $\leftarrow$  dest nextE
      return (pos,  $\rho$ )

```

Fig. 12. Hypercube Routing

i (see Map.init). Then, we initialize the random intermediate destinations ρ . The remaining loop encodes the evaluation of the routing strategy iterated T time. The variable *usedBy* is a map that logs if an edge is already used by a packet, it is empty at the beginning of each iteration. For each packet, we try to move it across one edge along the path to its intermediate destination. The function

¹³ Recall that the number of node in a hypercube of dimension D is 2^D so each node can be identified by a number in $[1, 2^D]$.

getEdge returns the next edge to follow, following the bit-fixing scheme. If the packet can progress (its edge is not used), then its current position is updated and the edge is marked as used.

We show that if the number of timesteps T is $4D + 1$, then all packets reach their intermediate destination in at most T steps, except with a small probability 2^{-2D} of failure. That is, the number of timesteps grows linearly in D , logarithmic in the number of packets. This is formalized in our system as:

$$\{T = 4D + 1\} \text{ route } \{\Pr[\exists i. \text{pos}[i] \neq \rho[i]] \leq 2^{-2D}\}$$

```

proc coupon ( $N : \text{int}$ ) :
  var int  $\text{cp}[N]$ ,  $t[N]$ ;
  var int  $X \leftarrow 0$ ;
  for  $p \leftarrow 1$  to  $N$  do
     $\text{ct} \leftarrow 0$ ;
     $\text{cur} \xleftarrow{\$} [1, N]$ ;
    while  $\text{cp}[\text{cur}] = 1$  do
       $\text{ct} \leftarrow \text{ct} + 1$ ;
       $\text{cur} \xleftarrow{\$} [1, N]$ ;
     $t[p] \leftarrow \text{ct}$ ;
     $\text{cp}[\text{cur}] \leftarrow 1$ ;
     $X \leftarrow X + t[p]$ ;
  return  $X$ 

```

Fig. 13. Coupon collector

Modeling Infinite Processes. Our second example is the *coupon collector* process. The algorithm draws a uniformly random coupon (we have N coupon) on each day, terminating when it has drawn at least one of each kind of coupon. The code of the algorithm is displayed in Fig. 13; the array cp records of the coupons seen so far, t holds the number of steps taken before seeing a new coupon, and X tracks of the total number of steps. Our goal is to bound the average number of iterations. This is formalized in our logic as:

$$\{\mathcal{L}\} \text{ coupon } \left\{ \mathbb{E}[X] = \sum_{i \in [1, N]} \left(\frac{N}{N-i+1} \right) \right\}.$$

```

proc pwInd ( $N : \text{int}$ ) :
  var bool  $X[2^N]$ ,  $B[N]$ ;
  for  $i \leftarrow 1$  to  $N$  do
     $B[i] \xleftarrow{\$} \text{Ber}(1/2)$ ;
  for  $j \leftarrow 1$  to  $2^N$  do
     $X[j] \leftarrow 0$ ;
    for  $k \leftarrow 1$  to  $N$  do
      if  $k \in \text{bits}(j)$  then
         $X[j] \leftarrow X[j] \oplus B[k]$ 
  return  $X$ 

```

Fig. 14. Pairwise Independence

Limited Randomness. *Pairwise independence* says that if we see the result of X_i , we do not gain information about all other variables X_k . However, if we see the result of *two* variables X_i, X_j , we may gain information about X_k . There are many constructions in the algorithms literature that grow a small number

of independent bits into more pairwise independent bits. Figure 14 gives one procedure, where \oplus is exclusive-or, and $\text{bits}(j)$ is the set of positions set to 1 in the binary expansion of j . The proof uses the following fact, which we fully verify: for a uniformly distributed Boolean random variable Y , and a random variable Z of any type,

$$Y \# Z \Rightarrow Y \oplus f(Z) \# g(Z) \quad (1)$$

for any two Boolean functions f, g . Then, note that $x[i] = \bigoplus_{\{j \in \text{bits}(i)\}} \mathbb{B}[j]$ where the big XOR operator ranges over the indices j where the bit representation of i has bit j set. For any two $i, k \in [1, \dots, 2^N]$ distinct, there is a bit position in $[1, \dots, N]$ where i and k differ; call this position r and suppose it is set in i but not in k . By rewriting,

$$x[i] = \mathbb{B}[r] \oplus \bigoplus_{\{j \in \text{bits}(i) \setminus r\}} \mathbb{B}[j] \quad \text{and} \quad x[k] = \bigoplus_{\{j \in \text{bits}(k) \setminus r\}} \mathbb{B}[j].$$

Since $\mathbb{B}[j]$ are all independent, $x[i] \# x[k]$ follows from Eq. (1) taking Z to be the distribution on tuples $\langle \mathbb{B}[1], \dots, \mathbb{B}[N] \rangle$ excluding $\mathbb{B}[r]$. This verifies pairwise independence:

$$\{\mathcal{L}\} \text{ pwInd}(N) \{\mathcal{L} \wedge \forall i, k \in [2^N]. i \neq k \Rightarrow x[i] \# x[k]\}.$$

Adversarial Programs. Pseudorandom functions (PRF) and pseudorandom permutations (PRP) are two idealized primitives that play a central role in the design of symmetric-key systems. Although the most natural assumption to make about a blockcipher is that it behaves as a pseudorandom permutation, most commonly the security of such a system is analyzed by replacing the blockcipher with a perfectly random function. The PRP/PRF Switching Lemma [22,6] fills the gap: given a bound for the security of a blockcipher as a pseudorandom function, it gives a bound for its security as a pseudorandom permutation.

Lemma 4 (PRP/PRF switching lemma). *Let A be an adversary with blackbox access to an oracle O implementing either a random permutation on $\{0, 1\}^l$ or a random function from $\{0, 1\}^l$ to $\{0, 1\}^l$. Then the probability that the adversary A distinguishes between the two oracles in at most q calls is bounded by*

$$\left| \Pr_{PRP}[b \wedge |H| \leq q] - \Pr_{PRF}[b \wedge |H| \leq q] \right| \leq \frac{q(q-1)}{2^{l+1}},$$

where H is a map storing each adversary call and $|H|$ is its size.

Proving this lemma can be done using the Fundamental Lemma of Game-Playing, and bounding the probability of *bad* in the program from Fig. 15. We focus on the latter. Here we apply the [ADV] rule of ELLORA with the invariant $\forall k, \Pr[\text{bad} \wedge |H| \leq k] \leq \frac{k(k-1)}{2^{l+1}}$ where $|H|$ is the size of the map H , i.e. the number of adversary call. Intuitively, the invariant says that at each call to the oracle the probability that *bad* has been set before and that the number of adversary call is less than k is bounded by a polynomial in k .

The invariant is d -closed and true before the adversary call, since at that point $\Pr[\text{bad}] = 0$. Then we need to prove that the oracle preserves the invariant, which can be done easily using the precondition calculus ([PC] rule).

```

var H: ( $\{0, 1\}^l$ ,  $\{0, 1\}^l$ ) map;
proc orcl (q:  $\{0, 1\}^l$ ):
  var a:  $\{0, 1\}^l$ ;
  if q  $\notin$  H then
    a  $\xleftarrow{\$}$   $\{0, 1\}^l$ ;
    bad  $\leftarrow$  bad || a  $\in$  codom(H);
    H[q]  $\leftarrow$  a;
  return H[q];

proc main():
  var b: bool;
  bad  $\leftarrow$  false;
  H  $\leftarrow$  [];
  b  $\leftarrow$  A();
  return b;

```

Fig. 15. PRP/PRF game

8 Implementation and Mechanization

We have built a prototype implementation of ELLORA within EASYCRYPT [5,2], a theorem prover originally designed for verifying cryptographic protocols. EASYCRYPT provides a convenient environment for constructing proofs in various Hoare logics, supporting interactive, tactic-based proofs for manipulating assertions and allowing users to invoke external tools, like SMT-solvers, to discharge proof obligations. EASYCRYPT provides a mature set of libraries for both data structures (sets, maps, lists, arrays, etc.) and mathematical theorems (algebra, real analysis, etc.), which we extended with theorems from probability theory.

Example	LC FPLC	
hypercube	100	1140
coupon	27	184
vertex-cover	30	61
pairwise-indep	30	231
private-sums	22	80
poly-id-test	22	32
random-walk	16	42
dice-sampling	10	64
matrix-prod-test	20	75

Table 1. Benchmarks

branch of EASYCRYPT, including a general library on probabilistic independence.

We used the implementation for verifying many examples from the literature, including all the programs presented in § 7 as well as some additional examples (such as polynomial identity test, private running sums, properties about random walks, etc.). The verified proofs bear a strong resemblance to the existing, paper proofs. Independently of this work, ELLORA has been used to formalize the main theorem about a randomized gossip-based protocol for distributed systems [26, Theorem 2.1]. Some libraries developed in the scope of ELLORA have been incorporated into the main

A New Library for Probabilistic Independence. In order to support assertions of the concrete program logic, we enhanced the standard libraries of EASYCRYPT, notably the ones dealing with big operators and sub-distributions. Like all EASYCRYPT libraries, they are written in a foundational style, i.e. they are defined instead of axiomatized. A large part of our libraries are proved formally from first principles. However, some results, such as concentration bounds, are currently declared as axioms.

Our formalization of probabilistic independence deserves special mention. We formalized two different (but logically equivalent) notions of independence. The first is in terms of products of probabilities, and is based on heterogeneous lists.

Since ELLORA (like EASYCRYPT) has no support for heterogeneous lists, we use a smart encoding based on second-order predicates. The second definition is more abstract, in terms of product and marginal distributions. While the first definition is easier to use when reasoning about randomized algorithms, the second definition is more suited for proving mathematical facts. We prove the two definitions equivalent, and formalize a collection of related theorems.

Mechanized Meta-Theory. The proofs of soundness and relative completeness of the abstract logic, without adversary calls, and the syntactical termination arguments have been mechanized in the Coq proof assistant. The development is available in supplemental material.

9 Related Work

More on Assertion-Based Techniques. The earliest assertion-based system is due to Ramshaw [37], who proposes a program logic where assertions can be formulas involving *frequencies*, essentially probabilities on sub-distributions. Ramshaw’s logic allows assertions to be combined with operators like \oplus , similar to our approach. [18] presents a Hoare-style logic with general assertions on the distribution, allowing expected values and probabilities. However, his **while** rule is based on a semantic condition on the guarded loop body, which is less desirable for verification because it requires reasoning about the semantics of programs. [8] give decidability results for a probabilistic Hoare logic without **while** loops. We are not aware of any existing system that supports assertions about general expected values; existing works also restrict to Boolean distributions. [38] formalize a Hoare logic for probabilistic programs but unlike our work, their assertions are interpreted on *distributions* rather than sub-distributions. For conditionals, their semantics rescales the distribution of states that enter each branch. However, their assertion language is limited and they impose strong restrictions on loops.

Other Approaches. Researchers have proposed many other approaches to verify probabilistic program. For instance, verification of Markov transition systems goes back to at least [17,40]; our condition for ensuring almost-sure termination in loops is directly inspired by their work. Automated methods include model checking (see e.g., [1,25,29]) and abstract interpretation (see e.g., [32,12]). Techniques for reasoning about higher-order (functional) probabilistic languages are an active subject of research (see e.g., [7,13,14]). For analyzing probabilistic loops, in particular, there are tools for reasoning about running time. There are also automated systems for synthesizing invariants [11,3]. [9,10] use a martingale method to compute the expected time of the coupon collector process for $N = 5$ —fixing N lets them focus on a program where the outer **while** loop is fully unrolled. Martingales are also used by [15] for analyzing probabilistic termination. Finally, there are approaches involving symbolic execution; [39] use a mix of static and dynamic analysis to check probabilistic programs from the approximate computing literature.

10 Conclusion and Perspectives

We introduced an expressive program logic for probabilistic programs, and showed that assertion-based systems are suited for practical verification of probabilistic programs. Owing to their richer assertions, program logics are a more suitable foundation for specialized reasoning principles than expectation-based systems. As evidence, our program logic can be smoothly extended with custom reasoning for probabilistic independence and union bounds. Future work includes proving better accuracy bounds for differentially private algorithms, and exploring further integration of ELLORA into EASYCRYPT.

Acknowledgments. We thank the reviewers for their helpful comments. This work benefited from discussions with Dexter Kozen, Annabelle McIver, and Carroll Morgan. This work was partially supported by ERC Grant #679127, and NSF grants 1513694 and 1718220.

References

1. Baier, C.: Probabilistic model checking. In: Dependable Software Systems Engineering, NATO Science for Peace and Security Series - D: Information and Communication Security, vol. 45, pp. 1–23. IOS Press (2016), <http://dx.doi.org/10.3233/978-1-61499-627-9-1>
2. Barthe, G., Dupressoir, F., Grégoire, B., Kunz, C., Schmidt, B., Strub, P.: EasyCrypt: A tutorial. In: Foundations of Security Analysis and Design (FOSAD). Lecture Notes in Computer Science, vol. 8604, pp. 146–166. Springer-Verlag (2013)
3. Barthe, G., Espitau, T., Ferrer Fioriti, L.M., Hsu, J.: Synthesizing probabilistic invariants via Doob’s decomposition. In: International Conference on Computer Aided Verification (CAV), Toronto, Ontario (2016), <https://arxiv.org/abs/1605.02765>
4. Barthe, G., Gaboardi, M., Grégoire, B., Hsu, J., Strub, P.Y.: A program logic for union bounds. In: International Colloquium on Automata, Languages and Programming (ICALP), Rome, Italy (2016), <http://arxiv.org/abs/1602.05681>
5. Barthe, G., Grégoire, B., Heraud, S., Béguelin, S.Z.: Computer-aided security proofs for the working cryptographer. In: IACR International Cryptology Conference (CRYPTO), Santa Barbara, California. Lecture Notes in Computer Science, vol. 6841, pp. 71–90. Springer-Verlag (2011)
6. Bellare, M., Rogaway, P.: The security of triple encryption and a framework for code-based game-playing proofs. In: IACR International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT), Saint Petersburg, Russia. pp. 409–426 (2006), http://dx.doi.org/10.1007/11761679_25
7. Bizjak, A., Birkedal, L.: Step-indexed logical relations for probability. In: International Conference on Foundations of Software Science and Computation Structures (FoSSaCS), London, England. pp. 279–294. Springer-Verlag (2015), <https://arxiv.org/abs/1501.02623>
8. Chadha, R., Cruz-Filipe, L., Mateus, P., Sernadas, A.: Reasoning about probabilistic sequential programs. Theoretical Computer Science 379(1–2), 142–165 (2007)

9. Chakarov, A., Sankaranarayanan, S.: Probabilistic program analysis with martingales. In: International Conference on Computer Aided Verification (CAV), Saint Petersburg, Russia. Lecture Notes in Computer Science, vol. 8044, pp. 511–526. Springer-Verlag (2013)
10. Chakarov, A., Sankaranarayanan, S.: Expectation invariants as fixed points of probabilistic programs. In: International Symposium on Static Analysis (SAS), Perpignan, France. Lecture Notes in Computer Science, vol. 8723, pp. 85–100. Springer-Verlag (2014)
11. Chatterjee, K., Fu, H., Novotný, P., Hasheminezhad, R.: Algorithmic analysis of qualitative and quantitative termination problems for affine probabilistic programs. In: ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages (POPL), Saint Petersburg, Florida. pp. 327–342 (2016), <http://doi.acm.org/10.1145/2837614.2837639>
12. Cousot, P., Monerau, M.: Probabilistic abstract interpretation. In: European Symposium on Programming (ESOP), Tallinn, Estonia. Lecture Notes in Computer Science, vol. 7211, pp. 169–193. Springer-Verlag (2012)
13. Crubillé, R., Dal Lago, U.: On probabilistic applicative bisimulation and call-by-value λ -calculi. In: European Symposium on Programming (ESOP), Grenoble, France. Lecture Notes in Computer Science, vol. 8410, pp. 209–228. Springer-Verlag (2014), <https://arxiv.org/abs/1501.02623>
14. Dal Lago, U., Sangiorgi, D., Alberti, M.: On coinductive equivalences for higher-order probabilistic functional programs. In: ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages (POPL), San Diego, California. pp. 297–308 (2014), <https://arxiv.org/abs/1311.1722>
15. Fioriti, L.M.F., Hermanns, H.: Probabilistic termination: Soundness, completeness, and compositionality. In: ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages (POPL), Mumbai, India. pp. 489–501 (2015)
16. Gretz, F., Katoen, J.P., McIver, A.: Operational versus weakest pre-expectation semantics for the probabilistic guarded command language. *Performance Evaluation* 73, 110–132 (2014)
17. Hart, S., Sharir, M., Pnueli, A.: Termination of probabilistic concurrent programs. *ACM Transactions on Programming Languages and Systems* 5(3), 356–380 (1983)
18. den Hartog, J.: Probabilistic extensions of semantical models. Ph.D. thesis, Vrije Universiteit Amsterdam (2002)
19. Hurd, J.: Formal verification of probabilistic algorithms. Tech. Rep. UCAM-CL-TR-566, University of Cambridge, Computer Laboratory (2003)
20. Hurd, J.: Verification of the Miller-Rabin probabilistic primality test. *Journal of Logic and Algebraic Programming* 56(1–2), 3–21 (2003), [http://dx.doi.org/10.1016/S1567-8326\(02\)00065-6](http://dx.doi.org/10.1016/S1567-8326(02)00065-6)
21. Hurd, J., McIver, A., Morgan, C.: Probabilistic guarded commands mechanized in HOL. *Theoretical Computer Science* 346(1), 96–112 (2005)
22. Impagliazzo, R., Rudich, S.: Limits on the provable consequences of one-way permutations. In: ACM SIGACT Symposium on Theory of Computing (STOC), Seattle, Washington. pp. 44–61 (1989), <http://doi.acm.org/10.1145/73007.73012>
23. Kaminski, B.L., Katoen, J.P., Matheja, C.: Inferring covariances for probabilistic programs. In: International Conference on Quantitative Evaluation of Systems (QEST). pp. 191–206. Springer-Verlag (2016)
24. Kaminski, B., Katoen, J.P., Matheja, C., Olmedo, F.: Weakest Precondition Reasoning for Expected Run-Times of Probabilistic Programs. In: European Symposium on Programming (ESOP), Eindhoven, The Netherlands (Jan 2016)

25. Katoen, J.P.: The probabilistic model-checking landscape. In: IEEE Symposium on Logic in Computer Science (LICS), New York, New York (2016)
26. Kempe, D., Dobra, A., Gehrke, J.: Gossip-based computation of aggregate information. pp. 482–491 (2003), <http://dx.doi.org/10.1109/SFCS.2003.1238221>
27. Kozen, D.: Semantics of probabilistic programs. vol. 22, pp. 328–350. Elsevier (1981), <https://www.sciencedirect.com/science/article/pii/0022000081900362>
28. Kozen, D.: A probabilistic PDL. *Journal of Computer and System Sciences* 30(2), 162–178 (1985)
29. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: Verification of probabilistic real-time systems. In: International Conference on Computer Aided Verification (CAV), Snowbird, Utah. *Lecture Notes in Computer Science*, vol. 6806, pp. 585–591. Springer-Verlag (2011)
30. McIver, A., Morgan, C.: Abstraction, Refinement, and Proof for Probabilistic Systems. *Monographs in Computer Science*, Springer-Verlag (2005)
31. McIver, A., Morgan, C., Kaminski, B.L., Katoen, J.P.: A new rule for almost-certain termination. *Proceedings of the ACM on Programming Languages* 1(POPL) (2018), <https://arxiv.org/abs/1612.01091>, appeared at ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages (POPL), Los Angeles, California
32. Monniaux, D.: Abstract interpretation of probabilistic semantics. In: International Symposium on Static Analysis (SAS), Santa Barbara, California. *Lecture Notes in Computer Science*, vol. 1824, pp. 322–339. Springer-Verlag (2000)
33. Morgan, C.: Proof rules for probabilistic loops. In: BCS–FACS Conference on Refinement, Bath, England (1996)
34. Morgan, C., McIver, A., Seidel, K.: Probabilistic predicate transformers. *ACM Transactions on Programming Languages and Systems* 18(3), 325–353 (1996)
35. Olmedo, F., Kaminski, B.L., Katoen, J.P., Matheja, C.: Reasoning about recursive probabilistic programs. In: IEEE Symposium on Logic in Computer Science (LICS), New York, New York. pp. 672–681 (2016)
36. Pearl, J., Paz, A.: Graphoids: Graph-based logic for reasoning about relevance relations. In: ECAI. pp. 357–363 (1986)
37. Ramshaw, L.H.: Formalizing the Analysis of Algorithms. Ph.D. thesis, Computer Science (1979)
38. Rand, R., Zdancewic, S.: VPHL: a verified partial-correctness logic for probabilistic programs. In: Conference on the Mathematical Foundations of Programming Semantics (MFPS), Nijmegen, The Netherlands (2015)
39. Sampson, A., Panchekha, P., Mytkowicz, T., McKinley, K.S., Grossman, D., Ceze, L.: Expressing and verifying probabilistic assertions. In: ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Edinburgh, Scotland. p. 14 (2014)
40. Sharir, M., Pnueli, A., Hart, S.: Verification of probabilistic programs. *SIAM Journal on Computing* 13(2), 292–314 (1984)
41. Valiant, L.G.: A scheme for fast parallel communication. *SIAM Journal on Computing* 11(2), 350–361 (1982)
42. Valiant, L.G., Brebner, G.J.: Universal schemes for parallel communication. In: ACM SIGACT Symposium on Theory of Computing (STOC), Milwaukee, Wisconsin. pp. 263–277 (1981), <http://doi.acm.org/10.1145/800076.802479>

A Soundness of ELLORA

Before presenting the proof of soundness, we will introduce two technical lemmas needed for the loop rules. Intuitively, d - and t -closed assertions are preserved in the limit of general and lossless loops, respectively.

Proposition 5.

1. If η is d -closed and is s.t.

$$\forall \mu. \mu \models \eta \implies \llbracket \mathbf{if } b \mathbf{ then } s \rrbracket_\mu \models \eta,$$

then $\forall \mu. \mu \models \eta \implies \llbracket \mathbf{while } b \mathbf{ do } s \rrbracket_\mu \models \eta$.

2. If η is t -closed and is s.t.

$$\forall \mu. \mu \models \eta \implies \llbracket \mathbf{if } b \mathbf{ then } s \rrbracket_\mu \models \eta,$$

then $\forall \mu. \mu \models \eta \implies \llbracket \mathbf{while } b \mathbf{ do } s \rrbracket_\mu \models \eta$, provided that $\mathbf{while } b \mathbf{ do } s$ is lossless.

Proof. We only treat the first case; the second is similar. Let η be a d -closed assertion s.t. for any sub-distribution μ , if $\eta(\mu)$, then $\llbracket \mathbf{if } b \mathbf{ then } s \rrbracket_\mu \models \eta$. We prove by induction on n that for any sub-distribution μ such that $\mu \models \eta$, we have $\llbracket (\mathbf{if } b \mathbf{ then } s)^n \rrbracket_\mu \models \eta$. By downward closedness of η , we have $\llbracket (\mathbf{if } b \mathbf{ then } s)^n \rrbracket_{\neg b} \models \eta$. We conclude by t -closedness of η . \square

Lemma 6. Let η be an assertion and s be a command. Then $\{\eta[\llbracket s \rrbracket]\} s \{\eta\}$.

Proof. Let $\mu \models \eta[\llbracket s \rrbracket]$. By definition of $\eta[\llbracket s \rrbracket]$, this amounts to have $\llbracket s \rrbracket_\mu \models \eta$, which exactly gives the expected result. \square

Lemma 7. Let μ be a sub-distribution and s be a command. Then, $|\llbracket s \rrbracket_\mu| \leq |\mu|$.

Proof. We have

$$\begin{aligned} |\llbracket s \rrbracket_\mu| &= |\mathbb{E}_{m \sim \mu} [\llbracket s \rrbracket_m]| = \sum_m \mu(m) \cdot \overbrace{|\llbracket s \rrbracket_m|}^{\leq 1} \\ &\leq \sum_m \mu(m) = |\mu| \square \end{aligned}$$

Lemma 8. Let μ be a sub-distribution s.t. $\mu \triangleq \sum_{i \in I} \lambda_i \mu_i$ where I is a finite set, all the λ_i 's are in $[0, 1]$ and all the μ_i 's are sub-distributions. Then, for any command s ,

$$\llbracket s \rrbracket_\mu = \sum_{i \in I} \lambda_i \cdot \llbracket s \rrbracket_{\mu_i}$$

Proof. Immediate consequence of the linearity of \mathbb{E} . \square

Lemma 9. *Let s be a lossless command and μ be a sub-distribution. Then $\mu_{|\overline{\text{mod}(s)}} = \llbracket s \rrbracket_{\mu_{|\overline{\text{mod}(s)}}}$.*

Proof. The proof is done by a direct induction of s . \square

Lemma 10. *The rules of ELLORA are sound.*

Proof. We use the notation of the rules.

- [SKIP] — immediate since $\llbracket \mathbf{skip} \rrbracket_{\mu} = \mu$.
- [ABORT] — immediate by definition of $\Box \eta$ and since

$$\text{supp}(\llbracket \mathbf{abort} \rrbracket_{\mu}) = \text{supp}(\mathcal{K}) = \emptyset.$$

- [ASSGN] & [SAMPLE] — immediate consequence of Lemma 6.
- [SEQ] — let $\mu \models \eta_1$. Then, by the first premise, $\llbracket s_1 \rrbracket_{\mu} \models \eta_2$. Hence, by the second premise, $\llbracket s_2 \rrbracket_{\llbracket s_1 \rrbracket_{\mu}} \models \eta_3$.
- [CONSEQ] — if $\mu \models \eta_0$, then $\mu \models \eta_1$ by the first premise. Hence, $\llbracket s \rrbracket_{\mu} \models \eta_2$ by the second premise and $\llbracket s \rrbracket_{\mu} \models \eta_3$ by the third premise.
- [SPLIT] — let $\mu \models \eta_1 \oplus \eta_2$. Then, there exists μ_1, μ_2 s.t. $\mu = \mu_1 + \mu_2$ and $\mu_i \models \eta_i$ for $i \in \{1, 2\}$. By the two premises of the rule, we have $\llbracket s \rrbracket_{\mu_i} \models \eta'_i$ for $i \in \{1, 2\}$. Now, by Lemma 8, we have $\llbracket s \rrbracket_{\mu} = \llbracket s \rrbracket_{\mu_1} + \llbracket s \rrbracket_{\mu_2}$. By taking resp. $\llbracket s \rrbracket_{\mu_1}$ and $\llbracket s \rrbracket_{\mu_2}$ for the witnesses of $\eta'_1 \oplus \eta'_2$, we obtain that $\llbracket s \rrbracket_{\mu} \models \eta'_1 \oplus \eta'_2$.
- [COND] — we first prove that

$$\{\eta_1 \wedge \Box e\} \text{ if } e \text{ then } s_1 \text{ else } s_2 \{\eta'_1\}. \quad (2)$$

Let $\mu \models \eta_1 \wedge \Box e$. Then, for $m \in \text{supp}(\mu)$, we know that $\llbracket e \rrbracket_m = \top$. It follows that

$$\begin{aligned} \llbracket \text{if } e \text{ then } s_1 \text{ else } s_2 \rrbracket_{\mu} &= \mathbb{E}_{m \sim \mu} \underbrace{\llbracket \text{if } e \text{ then } s_1 \text{ else } s_2 \rrbracket_m}_{= \llbracket s_1 \rrbracket_m} \\ &= \llbracket s_1 \rrbracket_{\mu}. \end{aligned}$$

However, by the first premise, $\llbracket s_1 \rrbracket_{\mu} \models \eta'_1$. Hence,

$$\llbracket \text{if } e \text{ then } s_1 \text{ else } s_2 \rrbracket_{\mu} \models \eta'_1,$$

concluding the proof of Eq. (2). By a similar reasoning, we have:

$$\{\eta_2 \wedge \Box \neg e\} \text{ if } e \text{ then } s_1 \text{ else } s_2 \{\eta'_2\}. \quad (3)$$

We conclude the proof of soundness of [COND] from the one of [SPLIT] applied to Eq. (2) and Eq. (3).

- [CALL] — immediate consequence of [SEQ] and [ASSGN].
- [FRAME] — let $\mu \models \eta$. Let $\mu \models \eta$. To show that $\llbracket s \rrbracket_{\mu} \models \eta$, from the definition of $\text{separated}(\eta, \text{mod}(s))$, it is sufficient to show that $|\mu| = |\llbracket s \rrbracket_{\mu}|$ and $\mu_{|\overline{\text{mod}(s)}} = \llbracket s \rrbracket_{\mu_{|\overline{\text{mod}(s)}}}$. The first one is a consequence of the losslessness of s , the second one is a direct application of Lemma 9.

- [WHILE] — from the first premise, by induction on n and using [SEQ], we know that for any n , the following holds:

$$\{\eta\} \text{ (if } b \text{ then } s)^n \{\eta\}. \quad (4)$$

Now, let $\mu \models \eta$. First, from the definition of $\llbracket \text{while } b \text{ do } s \rrbracket_\mu$, we know that

$$\llbracket \text{while } b \text{ do } s \rrbracket_\mu \models \Box \neg b.$$

It remains to prove $\llbracket \text{while } b \text{ do } s \rrbracket_\mu \models \eta$. In the case of certain termination, we know the existence of a k s.t.

$$\llbracket \text{while } b \text{ do } s \rrbracket = \llbracket (\text{if } b \text{ then } s \text{ else })^k \rrbracket$$

and we conclude by Eq. (4). For the cases of almost surely termination and almost termination, we conclude for Eq. (4) and Proposition 5 of the paper.

- [ADV] — the proof is done by induction on the body of the external procedure which is of the form $s_1; c_1^?; \dots; s_n; c_n^?$ where the s_i 's do not contain calls and the $c_i^?$'s are potential calls to the oracles. From the [FRAME] rule, we know that the s_i 's preserve the invariant η — noting that the losslessness of the adversary implies the losslessness of the s_i 's. Likewise, from the last premise of the [ADV] rule, we know that calls to the oracles also preserve the invariant. Hence, by multiple application of the [SEQ] rule, we obtain that the adversary body maintains η . \square

B Semantics of Assertions

The semantics of assertions is given Fig. 16.

C Precondition Calculus

Figure 17 contains the full definition of the precondition calculus.

D Soundness of the Syntactic Rules

D.1 Certain Termination

Proposition 11 (Soundness of rule [WHILE- $\mathcal{C}_{\text{TERM}}$]). *Let μ be a sub-distributions such that $\mathcal{C}_{\text{Term}}$ is valid. Then,*

$$\llbracket \text{while } b \text{ do } s \rrbracket_\mu \models \eta \wedge \Box \neg b.$$

Proof. Given μ satisfying $\mathcal{C}_{\text{Term}}$, we first claim that there exists a decreasing function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that $\Box(\tilde{e} \leq f(n) \vee \neg b)$ holds at each iteration of the

$$\begin{array}{ll}
\llbracket v \rrbracket_m^\rho & \triangleq \rho(v) \\
\llbracket \mathbf{1}_\phi \rrbracket_m^\rho & \triangleq \mathbf{1}_{\llbracket \phi \rrbracket_m^\rho} \\
\llbracket \mathbb{E}_{v \sim g}[\tilde{e}] \rrbracket_m^\rho & \triangleq \mathbb{E}_{w \sim \llbracket g \rrbracket_m^\rho} [\llbracket \tilde{e} \rrbracket_m^{\rho[v:=w]}] \\
\llbracket o(\tilde{e}) \rrbracket_m^\rho & \triangleq o(\llbracket \tilde{e} \rrbracket_m^\rho) \\
\hline
\llbracket \tilde{e}_1 \bowtie \tilde{e}_2 \rrbracket_m^\rho & \triangleq \llbracket \tilde{e}_1 \rrbracket_m^\rho \bowtie \llbracket \tilde{e}_2 \rrbracket_m^\rho \\
\llbracket FO(\phi) \rrbracket_m^\rho & \triangleq FO(\llbracket \phi \rrbracket_m^\rho) \\
\hline
\llbracket \mathbb{E}[\tilde{e}] \rrbracket_\mu^\rho & \triangleq \mathbb{E}_{m \sim \mu} [\llbracket \tilde{e} \rrbracket_m^\rho] \\
\llbracket o(p) \rrbracket_\mu^\rho & \triangleq o(\llbracket p \rrbracket_\mu^\rho) \\
\hline
\llbracket p_1 \bowtie p_2 \rrbracket_\mu^\rho & \triangleq \llbracket p_1 \rrbracket_\mu^\rho \bowtie \llbracket p_2 \rrbracket_\mu^\rho \\
\llbracket \eta_1 \oplus \eta_2 \rrbracket_\mu^\rho & \triangleq \exists \mu_1, \mu_2. \mu = \mu_1 + \mu_2 \wedge \llbracket \eta_1 \rrbracket_{\mu_1}^\rho \wedge \llbracket \eta_2 \rrbracket_{\mu_2}^\rho \\
\llbracket FO(\eta) \rrbracket_\mu^\rho & \triangleq FO(\llbracket \eta \rrbracket_\mu^\rho)
\end{array}$$

Fig. 16. Semantics of assertions

$$\begin{array}{ll}
\text{pe}(s, o(p)) & \triangleq o(\text{pe}(s, p)) \quad \text{where } o \in \mathbf{Ops} \\
\text{pe}(\mathbf{skip}, \mathbb{E}[\tilde{e}]) & \triangleq \mathbb{E}[\tilde{e}] \\
\text{pe}(s_1; s_2, \mathbb{E}[\tilde{e}]) & \triangleq \text{pe}(s_1, \text{pe}(s_2, \mathbb{E}[\tilde{e}])) \\
\text{pe}(x \leftarrow e, \mathbb{E}[\tilde{e}]) & \triangleq \mathbb{E}[\tilde{e}][x := e] \\
\text{pe}(x \stackrel{s}{\leftarrow} g, \mathbb{E}[\tilde{e}]) & \triangleq \mathcal{P}_x^g(\mathbb{E}[\tilde{e}]) \\
\text{pe}(\mathbf{if } e \mathbf{ then } s_1 \mathbf{ else } s_2, \mathbb{E}[\tilde{e}]) & \triangleq \text{pe}(s_1, \mathbb{E}[\tilde{e}])|_e + \text{pe}(s_2, \mathbb{E}[\tilde{e}])|_{\neg e} \\
\text{pe}(\mathbf{abort}, \mathbb{E}[\tilde{e}]) & \triangleq 0 \\
\hline
\text{pc}(s, p_1 \bowtie p_2) & \triangleq \text{pc}(s, p_1) \bowtie \text{pc}(s, p_2) \\
\text{pc}(s, FO(\eta)) & \triangleq FO(\text{pc}(s, \eta))
\end{array}$$

Fig. 17. Precondition calculus

loop. Indeed, we remark that the statement $\exists k \Box(\tilde{e} \leq k)$ holds at first iteration by the precondition hypothesis. Let:

$$f(n) = \begin{cases} k - n & \text{if } n \leq k \\ 0 & \text{otherwise.} \end{cases}$$

Then unrolling the loop and using the semantical [SEQ] rule ensure by induction the claimed domination. The termination of the loop arises then naturally from the exit condition $\Box(\tilde{e} = 0 \Rightarrow \neg b)$. \square

D.2 Almost-Sure Termination

The main challenge for proving the soundness of the is proving termination; from there, we can conclude by rule [WHILE- $\mathcal{C}_{\text{ASTERM}}$]. Our arguments use basic notations and theorems from the theory of Markov processes.

Proposition 12 (Soundness of rule [WHILE- $\mathcal{C}_{\text{ASTERM}}$]). *Let η be a t -closed assertion. For any sub-distributions μ , such that $\mathcal{C}_{\text{ASTerm}}$ is valid. Then*

$$\llbracket \text{while } b \text{ do } s \rrbracket \mu \models \eta.$$

Proof. As indicated, by soundness of the semantic while rule for almost-sure termination, and the premises, it suffices to prove almost-sure termination. The sketch of the proof is the following:

1. We follow the behavior of the variant by seeing it as a random variable on the space of state.
2. We first introduce a Markov chain that reaches a particular state (the state zero) with probability 1.
3. We then stochastically dominates the variant by the latter chain.
4. In particular, this shows that the probability of the variant to eventually reach 0 is 1, ensuring almost-sure termination.

Step one: Variant as a Random variable.


We consider the integer-valued variant \tilde{e} as a random variable over the space of states. Let the family $(\tilde{e}_i)_i$ represents the value taken by \tilde{e} after the i -th iteration of the loop. Then we have, using the semantical [SEQ] rule and the

- The sequence is uniformly bounded by K .
- The probability of decreasing is bounded below by ϵ :

$$\forall i \in \mathbb{N}. \Pr[\tilde{e}_i > \tilde{e}_{i+1} \mid \tilde{e}_i, \dots, \tilde{e}_0] = \epsilon_i \geq \epsilon > 0.$$

Step two: Modelization with a Markov chain.

First, we can assume that $K > 0$ since if $K = 0$, the loop terminates immediately. Consider the following finite Markov chain $(X_i)_i$, over the state space $S = \{0, \dots, K\}$, following the transitions:



This chain models the following behavior: with probability ϵ the value decrease by one, while with probability $1 - \epsilon$ it jumps to K . Since zero is the only connected component of the underlying graph, the probability of terminating in the state zero is 1 by a standard result on Markov chains.

Step three: Stochastic domination.

By construction, there exists a natural coupling between (\tilde{X}_i) and $(\tilde{e}_i)_i$, since the probability of the event \tilde{e}_i decrease is greater than the probability of the event \tilde{X}_i decrease. Since we impose that $X_0 = \tilde{e}_0$ (initial position), a simple induction over i ensure that we have, for this coupling: $\Pr[\tilde{e}_i > 0] \leq \Pr[X_i > 0]$, so $\lim_{i \rightarrow \infty} \Pr[\tilde{e}_i = 0] \geq 1 - \lim_{i \rightarrow \infty} \Pr[X_i > 0] = 1$ as desired. \square

E Further Examples

E.1 Randomization for Approximation: Vertex Cover

We begin with a classical application of randomization: *approximation algorithms* for computationally hard problems. For problems that take long time to solve in the worst case, we can sometimes devise efficient algorithms that find a solution that is “nearly” as good as the true solution.

Our first example illustrates a famous approximation algorithm for the *vertex cover* problem. The input is a graph described by vertices V and edges E . The goal is to output a vertex cover: a subset $C \subseteq V$ such that each edge has at least one endpoint in C , and such that C is as small as possible.

It is known that this problem is NP-complete, but there is simple randomized algorithm that returns a vertex cover that is at on average at most twice the size of the optimal vertex cover. The algorithm proceeds by maintaining a current cover (initially empty) and considering each edge in order. If neither endpoint is in current cover, the algorithm adds one of the two endpoints uniformly at random. The ELLORA program is:

```

proc VC (E : set<edge>) :
var  set<node> C ← ∅;
for (e1, e2) in E do
  if (e1 ∉ C) ∧ (e2 ∉ C) then
    b  $\stackrel{\$}{\leftarrow}$  {0,1};
    C ← (b ? e1 : e2) ∪ C;
  fi
end

```

Here, we represent edges as a finite set of pairs of nodes. We loop through the edges, adding one point of each uncovered edge to the cover C uniformly at random. The operator $b ? e_1 : e_2$ returns e_1 if b is true, and e_2 if not.

To prove the approximation guarantee, we first assume that we have a set of nodes C^* . We only assume that C^* is a valid vertex cover; i.e., each edge has at least one endpoint in C^* . Then, we use the following loop invariant:

$$\mathbb{E}[\text{size}(C \setminus C^*)] \leq \mathbb{E}[\text{size}(C \cap C^*)]. \quad (5)$$

Given the loop invariant, we can prove the conclusion by letting C^* be the cover OPT of minimal size, and reasoning about intersections and differences of sets.

Clearly the invariant is initially true. To see why the invariant is preserved, let e be the current edge, with both endpoints out of C . Since C^* is a vertex cover, it has at least one endpoint of e . Since our algorithm includes an endpoint of e uniformly at random, the probability we choose a vertex not in C^* is at most $1/2$, so the expectation on the left in Eq. (5) increases by at most $1/2$. If e is not covered in C but is covered by C^* , there is at least a $1/2$ probability that we increase the intersection $C \cap C^*$, so the right side in Eq. (5) increases by at least $1/2$. Thus, the invariant is preserved, and we can prove

$$\{\text{isVC}(C^*, E)\} \text{ VC}(E) \{ \mathbb{E}[\text{size}(C \setminus C^*)] \leq \mathbb{E}[\text{size}(C \cap C^*)] \}$$

and by reasoning on intersection and difference of sets,

$$\{\text{isVC}(C^*, E)\} \text{ VC}(E) \{ \mathbb{E}[\text{size}(C)] \leq 2 \cdot \mathbb{E}[\text{size}(C^*)] \}.$$

E.2 Random Walks: Termination and Reachability

A canonical example of a random process is a *random walk*. There are many variations, but the basic scenario describes a numeric position changing over time, where the position depends on the position at the previous timestep, influenced by random noise drawn from some distribution.

To demonstrate how to verify interesting facts about random processes, we will model a one-dimensional random walk on the natural numbers. We start at position 0, and repeatedly update our position according to the following rules. From 0, we always make a step to 1. From non-zero positions, we flip a fair coin that is biased to come up heads with probability $p > 0$. If heads, we increase our position by 1; if tails, we decrease by 1.

We will prove two facts about this random walk. First, for any natural number T , the probability of eventually reaching T is 1. Second, when we reach T , we must first pass through all intermediate points $1, \dots, T-1$. In ELLORA, we can express the random walk with the following code.

```
var bool visited[T];
var int pos = 0;
visited[0] ← true;
while pos ≠ T do:
  c  $\stackrel{\$}{\leftarrow}$  Ber(p);
```

```

    pos ← pos + ((pos = 0) ∨ c) ? 1 : -1;
    visited[pos] ← true;
end

```

In order to verify this example, we will use the probabilistic while rule [WHILE-ASTERM]. First we establish almost-sure termination by finding an appropriate termination measure: the distance between our current position, and τ . Indeed, this measure is bounded by τ , and has a non-negative (at least $1/2$) probability of decreasing each iteration. Therefore, the loop terminates almost-surely, and thus our random walk eventually reaches any point τ with probability 1.

To prove our second assertion—that we visit every point from 0 to τ —we use the following loop invariant for the while command:

$$\Box(\forall i. 0 \leq i \leq \text{pos} \Rightarrow \text{visited}[i] = \text{true}).$$

In other words, if we have reached position pos , then we must have already reached every position in $[0, \text{pos}]$. Since this invariant is t -closed, we may apply rule [WHILE-ASTERM] and the invariant holds at the end of the loop. With the assertion $\text{pos} = \tau$ at termination, we have enough to prove that each position is visited:

$$\{\mathcal{L}\} s \{\mathcal{L} \wedge \Box(\forall i \in [\tau]. \text{visited}[i] = \text{true})\}.$$

The losslessness post-condition indicates that the walk terminates almost-surely.

E.3 Amplification: Polynomial Identity Testing

A second use of randomness is in running independent trials of the same algorithm. This technique, known as *probability amplification*, runs a randomized algorithm several times in order to reduce the error probability. Roughly speaking, a single trial may have high error with some probability, but by repeating the trial it is unlikely that all of the trials yield poor results. By combining the results appropriately—e.g., with a majority vote for algorithms with binary outputs, or by selecting the best answer when we can check the quality of the solution—we can produce an output that is more accurate than a single run of the original algorithm.

An example in this vein is probabilistic *polynomial identity testing*. Given two multivariate polynomials $P(x_1, \dots, x_n)$ and $Q(x_1, \dots, x_n)$ over the finite field \mathbb{F}_q of q elements, we want to check whether $P = Q$, or equivalently, whether the polynomial $P - Q$ is zero or not. We will take n uniformly random samples $(v_i)_i$ from \mathbb{F}_q and check whether $(P - Q)(v_1, \dots, v_n) = 0$. We repeat the trial q times, rejecting if we see a sample where the difference is non-zero. In our system, this corresponds to the following program:

```

var bool res = true;
for i = 1 to q do:
  v ←  $\mathbb{S}$  Unif( $\mathbb{F}_q^n$ );
  res ← res ∧ ((P - Q)(v) = 0);
end

```

The proof uses an instance of the Schwartz-Zippel lemma due to Øystein Ore for finite fields, which upper bounds the probability of randomly picking a root of a polynomial over a finite field.

Lemma 13. *Let P a non-zero polynomial function over \mathbb{F}_q . If we sample the variables v_1, \dots, v_n uniformly at random from \mathbb{F}_q , then*

$$\Pr[P(v_1, \dots, v_n) = 0] \leq 1 - 1/q.$$

We encode this lemma as an axiom in our system:

$$P \neq 0 \wedge v \sim \mathbf{Unif}(\mathbb{F}_q^n) \Rightarrow \Pr[P(v) = 0] \leq 1 - 1/q.$$

With this fact, we can prove the following loop invariant:

$$P \neq 0 \Rightarrow \Pr[\text{res} = \text{true}] \leq (1 - 1/q)^i,$$

finally proving that

$$\{P \neq 0\} \text{ s } \{\Pr[\text{res} = \text{true}] \leq (1 - 1/q)^q \leq 1/e\}.$$

We have also verified Freivald’s algorithm, an amplification-based algorithm for checking matrix multiplication.

E.4 Concentration Bounds: Private Running Sums

Now, we turn to examples involving independence of random variables. Our first example is drawn from the *differential privacy* literature. Given a list of N integers, we add noise from a two-sided geometric distribution to each entry in order to protect privacy, and we calculate the *partial sums* of the noisy values: $x_1, x_1 + x_2, x_1 + x_2 + x_3$, and so on. We wish to measure how far the noisy sums deviate from the true sums.

In ELLORA, we can express this algorithm as follows:

```

var int X[N], noise[N], out[N];
var int acc = 0;
for i = 1 to N do
  noise[i]  $\stackrel{\$}{\leftarrow}$  twogeom( $\epsilon$ );
  out[i]  $\leftarrow$  acc + X[i] + noise[i];
  acc  $\leftarrow$  out[i];
end

```

The parameter ϵ to the noise distribution `twogeom` is a numeric parameter controlling the strength of the privacy guarantee, by changing the magnitude of the noise.

Our loop invariant tracks three pieces of information: (i) `noise[i]` is distributed according to `twogeom(ϵ)`; (ii) the array `noise` remains independent; and (iii) `out[i]` stores the noisy running sum:

$$\forall q \in [i]. \text{out}[q] - \sum_{p \in [q]} x[p] = \sum_{p \in [q]} \text{noise}[p].$$

To bound the error introduced by the noise, we need to bound $\left| \sum_{p \in [q]} \text{noise}[p] \right|$. Since we know that the elements in noise are all independent, we can apply a concentration bound to bound the probability of a large error in the p -th running sum, concluding:

$$\{\mathcal{L}\} \text{ s } \left\{ \forall p \in [\mathbb{N}]. \Pr \left[\left| \sum_{i \in [p]} x_i \right| \geq T \right] \leq Q(\epsilon, T) / \sqrt{\mathbb{N}} \right\}$$

for a particular function Q derived from the Berry-Esseen theorem.

E.5 Hypercube Routing in More Details

We will begin with the *hypercube routing* algorithm [41,42]. To set the stage, consider a network where each node is labeled by a bitstring of length D , and two nodes are connected by an edge if and only if the two corresponding labels differ in exactly one bit position. This network topology is known as a *hypercube*, a D -dimensional version of the standard cube; a simple example with $D = 3$ is in Fig. 18. In the network, there is initially one packet at each node, and each




Fig. 18. Hypercube path from 111 to 100 ($D = 3$)

packet has a unique destination. Our goal is to design a routing strategy that will move the packets from node to node, following the edges, until all packets reach their destination. Furthermore, the routing should be *oblivious*: to avoid coordination overhead, each packet must select a path without considering the behavior of the other packets.

To model the flow of packets, each packet's current position is a node in the graph. Time proceeds in a series of steps, and at each step at most one packet can traverse any single edge. If several packets try to use the same edge

simultaneously, one packet will be selected to move (for any selection strategy that selects *some* packet). The other packets wait at their current position; these packets are *delayed* and make no progress this step.

The routing strategy is based on *bit fixing*: if the current position has bitstring i , and the target node has bitstring j , we compare the bits in i and j from left to right, moving along the edge that corrects the first differing bit. For instance, if we are at 111 and we wish to reach 100, we will move along the edge corresponding to the second position: from 111 to 101, and then along the edge corresponding to the third position: from 101 to 100. See Fig. 18 for a picture.

While this strategy is simple and oblivious, there are permutations π that require a total number of steps growing linearly in the number of packets to route all packets. Valiant proposes a simple modification, so that the total number of steps grows *logarithmically* in the number of packets. In the first phase, each packet i will first select an intermediate destination $\rho(i)$ uniformly at random from all nodes, and use bit fixing to reach $\rho(i)$. In the second phase, each packet will use bit fixing to go from $\rho(i)$ to the destination $\pi(i)$. We will focus on the first phase, since the reasoning for the second phase is nearly identical. We can model the strategy with the following code, using some syntactic sugar for the **for** loops; recall that the number of node in a hypercube of dimension D is 2^D so each node can be identified by a number in $[1, 2^D]$.

```

proc route ( $D$   $T$  : int) :
  var  $\rho$ , pos, usedBy : node map;
  var nextE : edge;
  pos  $\leftarrow$  Map.init id  $2^D$ ;  $\rho \leftarrow$  Map.empty;
  for  $i \leftarrow 1$  to  $2^D$  do  $\rho[i] \leftarrow \mathbb{S}[1, 2^D]$ 
  for  $t \leftarrow 1$  to  $T$  do
    usedBy  $\leftarrow$  Map.empty;
    for  $i \leftarrow 1$  to  $2^D$  do
      if pos[ $i$ ]  $\neq$   $\rho[i]$  then
        nextE  $\leftarrow$  getEdge pos[ $i$ ]  $\rho[i]$ ;
        if usedBy[nextE] =  $\perp$  then
          usedBy[nextE]  $\leftarrow i$ ; // Mark edge used
          pos[ $i$ ]  $\leftarrow$  dest nextE // Move packet
  return (pos,  $\rho$ )

```

We assume that initially the position of the packet i is at node i (see Map.init). Then, we initialize the random intermediate destinations ρ . The remaining loop encodes the evaluation of the routing strategy iterated T time. The variable `usedBy` is a map that logs if an edge is already used by a packet, it is empty at the beginning of each iteration. For each packet, we try to move it across one edge along the path to its intermediate destination. The function `getEdge` returns the next edge to follow, following the bit-fixing scheme. If the packet can progress (its edge is not used), then its current position is updated and the edge is marked as used.

Our goal is to show that if the number of timesteps T is $4D + 1$, then all packets reach their intermediate destination in at most T steps, except with a small probability 2^{-2^D} of failure. That is, the number of timesteps grows linearly in D , logarithmic in the number of packets. At a high level, the analysis involves three steps.

The first step is to consider a single packet traveling from i to $\rho(i)$, and to calculate the average number of other packets that share at least one edge with i 's path P . Let $H^\rho(i, j)$ be 1 if the paths of packets i and j share at least one edge and 0 otherwise. Then, the load $R^\rho(i)$ on i 's path is the sum of $H^\rho(i, j)$ over all the other packets j . By using the fact that each intermediate destination in ρ is uniformly distributed, and by analyzing the number of packets beside i 's that use each edge on i 's path, we can upper bound the expected value of $R^\rho(i)$ by $D/2$. In the second step, we move from a bound on the expectation of $R^\rho(i)$ to a *high-probability bound*: we want to show that $R^\rho(i) < 3D$ holds except with a small probability of failure. The key tool is the *Chernoff bound*, which gives high probability bounds of sums of independent samples. The libraries in ELLORA include a mechanized proof of a generalized Chernoff bound, with the following statement (leaving off some of the parameters):

```
lemma Chernoff  $n(d:\text{mem } \mathbf{distr})(X:\text{int} \rightarrow \text{mem} \rightarrow \text{bool})$ :
  let  $X\ m = \Sigma_{i \in [1, n]} X\ i\ m$  in
   $\mathbb{E}[X\ d] \leq \mu \wedge \mathbf{indep}\ [0, n]\ (X\ i)\ d$ 
   $\Rightarrow \Pr[X > (1 + \delta)\mu\ d] \leq (e^\delta / (1 + \delta)^{1+\delta})^\mu$ 
```

Returning to the proof, we bounded the expectation of R^ρ in the previous step. To apply the Chernoff bound, we need to show that for any packet i , the expressions $H^\rho(i, j)$ are independent for all j . This is not exactly true, since $H^\rho(i, j_1)$ and $H^\rho(i, j_2)$ both depend on the (random) destination $\rho(i)$ of packet i . However, it suffices to show that these variables are independent if we fix the value of $\rho(i)$; then we can apply the Chernoff bound to upper bound R^ρ with high probability.

Finally, we can bound the total delay of any packet. This portion of the proof rests on an intricate loop invariant assigning an imaginary coin for each delay step to some packet that crosses P . By showing that each packet holds at most one coin, we can conclude that the i 's delay is at most the number $R^\rho(i)$ of crossing packets. With our high probability bound from the previous step, we can show that $T = 4D + 1$ timesteps is sufficient to route all packets i to $\rho(i)$, except with some small probability:

$$\{T = 4D + 1\} \text{ route } \{\Pr[\exists i. \text{pos}[i] \neq \rho[i]] \leq 2^{-2D}\}$$

E.6 Coupon Collector in More Details

Our second example is the *coupon collector* process. The algorithm draws a uniformly random coupon (we have N coupon) on each day, terminating when it has drawn at least one of each kind of coupon. The code of the algorithm is as follows.

```
proc coupon ( $N : \text{int}$ ) :
  var  $\text{int } \text{cp}[N],\ t[N]$ ;
  var  $\text{int } X \leftarrow 0$ ;
  for  $p \leftarrow 1$  to  $N$  do
     $\text{ct} \leftarrow 0$ ;
     $\text{cur} \xleftarrow{\$} [1, N]$ ;
    while  $\text{cp}[\text{cur}] = 1$  do
       $\text{ct} \leftarrow \text{ct} + 1$ ;
```

```

    cur  $\stackrel{\$}{\leftarrow}$  [1, N];
    t[p]  $\leftarrow$  ct;
    cp[cur]  $\leftarrow$  1;
    X  $\leftarrow$  X + t[p];
return X

```

The array `cp` keeps track of the coupons seen so far; initially $\text{cp}[i] = 0$. We divide the loop into a sequence of phases (the outer loop) where in each phase we repeatedly sample coupons and wait until we see a new coupon (the inner loop). We keep track of the number of steps we spend in each phase p in $t[p]$, and the total number of steps in X .

Our goal is to bound the average number of iterations. The code involves two nested loops, so we have two loop invariants. The inner loop has a probabilistic guard: at every iteration, there is a finite probability that the loop terminates (i.e., if we draw a new coupon), but there is no finite bound on the number of iterations we need to run. Since our desired loop invariant is not downwards closed, we must apply the rule for almost sure termination. We use a variant that is 1 if we have not seen a new coupon, and 0 if we have seen a new coupon. Note that each iteration, we have a strictly positive probability $\rho(p)$ of seeing a new coupon and decreasing the variant. Furthermore, the variant is bounded by 1, and the loop exits when the variant reaches 0. So, the side-condition $\mathcal{C}_{\text{ASTerm}}$ holds. For the inner loop, we prove that forall c the invariant η_i is preserved:

$$\bigvee \left\{ \begin{array}{l} (\Box(\text{cp}[\text{cur}] = 1 \Rightarrow c \leq \text{ct}) \wedge \Pr[\text{cp}[\text{cur}] = 0 \wedge c = \text{ct}] = (1 - \rho(p))^c \rho(p)) \\ \exists k \in [1, c). \left\{ \begin{array}{l} \Box(\text{cp}[\text{cur}] = 1 \Rightarrow \text{ct} = k) \wedge (\text{cp}[\text{cur}] = 0 \Rightarrow \text{ct} \leq k) \\ \Pr[\text{cp}[\text{cur}] = 1 \wedge \text{ct} = k] = (1 - \rho(p))^k. \end{array} \right. \end{array} \right.$$

Note that this is a t -closed formula; there is an existential in the second disjunction, but it has finite domain (for fixed c).

For intuition, for every natural number c there are two cases: Either we have already unrolled more than c iterations, or not. The first disjunction corresponds to the first case, since loops where the guard is true all have $\text{ct} \geq c$, and the probability of stopping at c iterations is $(1 - \rho(p))^c \rho(p)$ —we see c old coupons, and then a new one. Otherwise, we have the second disjunction. The integer k represents the current number of unfoldings of the loop. If the loop is continuing then $k = \text{ct}$. If the loop is terminated, it terminated before the current iteration: $\text{ct} < k$. Furthermore, the probability of continuing at iteration k is $(1 - \rho(p))^k$. At the end of the loop we have $\Box \text{cp}[\text{cur}] = 0$. So, by the first conjunct and some manipulations,

$$\forall c \in \mathbb{N}. \Pr[c = \text{ct}] = (1 - \rho(p))^c \rho(p)$$

holds when the inner loop exits, precisely describing the distribution of iterations ct as $\mathbf{Geom}(\rho(p))$ by definition.

The outer loop is easier to handle, since the loop has a fixed bound N on the number of iterations so we can use the rule for certain termination. For the loop

invariant, we take:

$$\eta_{out} \triangleq \begin{cases} \forall i \in [p-1]. t[i] \sim \text{Geom}(\rho(i)) \wedge \Box \left(X = \sum_{i \in [p-1]} t[i] \right) \\ \wedge \Box \left(\sum_{i \in [1, N]} \text{cp}[i] = p-1 \right) \\ \wedge \forall i \in [1, N]. \Box(\text{cp}[i] \in \{0, 1\}). \end{cases}$$

The first conjunct states that the previous waiting times follow a geometric distribution with parameter $\rho(i)$; this assertion follows from the previous reasoning on the inner loop. The second conjunct asserts that X holds the total waiting time so far. The final two conjuncts state that there are at most $p-1$ flags set in `cp`. Thus,

$$\{\mathcal{L}\} \text{ coupon } \left\{ \begin{array}{l} \forall i \in [1, N]. t[i] \sim \text{Geom}(\rho(i)) \\ \wedge \Box X = \sum_{i \in [1, N]} t[i] \end{array} \right\}$$

at the end of the outer loop. By applying linearity of expectations and a fact about the expectation of the geometric distribution, we can bound the expected running time:

$$\{\mathcal{L}\} \text{ coupon } \left\{ \mathbb{E}[X] = \sum_{i \in [1, N]} \left(\frac{N}{N-i+1} \right) \right\}.$$