

Weakest Precondition Reasoning for Expected Runtimes of Randomized Algorithms

BENJAMIN LUCIEN KAMINSKI, JOOST-PIETER KATOEN, and CHRISTOPH MATHEJA,
RWTH Aachen University
FEDERICO OLMEDO, Department of Computer Science, University of Chile

This article presents a wp-style calculus for obtaining bounds on the expected runtime of randomized algorithms. Its application includes determining the (possibly infinite) expected termination time of a randomized algorithm and proving positive almost-sure termination—does a program terminate with probability one in finite expected time? We provide several proof rules for bounding the runtime of loops, and prove the soundness of the approach with respect to a simple operational model. We show that our approach is a conservative extension of Nielson's approach for reasoning about the runtime of deterministic programs. We analyze the expected runtime of some example programs including the coupon collector's problem, a one-dimensional random walk and a randomized binary search.

CCS Concepts: • **Theory of computation** → **Probabilistic computation; Invariants; Pre- and post-conditions; Program verification; Logic and verification; Operational semantics; Hoare logic**; • **Software and its engineering** → **Formal software verification**;

Additional Key Words and Phrases: Probabilistic programs, expected runtime, positive almost-sure termination, weakest precondition, program verification.

ACM Reference format:

Benjamin Lucien Kaminski, Joost-Pieter Katoen, Christoph Matheja, and Federico Olmedo. 2018. Weakest Precondition Reasoning for Expected Runtimes of Randomized Algorithms. *J. ACM* 65, 5, Article 30 (August 2018), 68 pages.

<https://doi.org/10.1145/3208102>

1 INTRODUCTION

The Runtime of Randomized Algorithms. Since the early days of computing, randomization has been an important tool for algorithm design. It is used to obtain an efficient randomized algorithm, possibly at the cost of sacrificing correctness with low probability. The Rabin-Miller primality test and Freivalds' matrix multiplication are prime examples of this principle. Randomization is also used to accelerate existing deterministic algorithms. Randomly selecting the pivot in Hoare's quicksort algorithm lowers the quadratic worst-case runtime to $O(N \cdot \log N)$, where N is the

This work was supported by the Excellence Initiative of the German federal and state government as well as the Research Training Group 2236 UnRAVeL funded by the German Research Foundation.

Authors' addresses: B. L. Kaminski, J-P. Katoen, and C. Matheja, Lehrstuhl für Informatik 2, RWTH Aachen, Ahornstrasse 55, 52056 Aachen, Germany; emails: {benjamin.kaminski, katoen, matheja}@cs.rwth-aachen.de; F. Olmedo, Department of Computer Science, University of Chile, Av. Beauchef 851, Santiago, Chile; email: folmedo@dcc.uchile.cl.

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

© 2018 ACM 0004-5411/2018/08-ART30 \$15.00

<https://doi.org/10.1145/3208102>

size of the input. Furthermore, some problems inherently require randomized solutions, e.g., self-stabilization in anonymous distributed systems.

Randomized algorithms are conveniently described by probabilistic programming languages that (on top of the usual language constructs) offer the possibility of sampling values from a probability distribution. Sampling can be used in assignments as well as in Boolean guards. The interest in probabilistic computations has been rapidly growing recently. This is mainly due to their wide applicability [16]. Randomized algorithms are, e.g., used in security to describe cryptographic constructions and security notions. In machine learning, probabilistic programs are used to describe distribution functions that are analyzed using Bayesian inference.

The runtime of a randomized algorithm is affected by the outcomes of the samplings (aka coin tosses). Technically speaking, the runtime is a random variable, assuming value t_1 with probability p_1 , t_2 with probability p_2 , and so on. An important measure is the *average* or *expected* runtime $\sum_i p_i t_i$.¹ In that, one does not average with respect to a distribution of program inputs but rather with respect to the randomness that is inherent in the algorithm.

Expected Runtimes Are Intricate. Reasoning about the expected runtime of randomized algorithms is subtle and full of nuances. Let us illustrate this by discussing three phenomena of randomized algorithms:

- (1) They may have diverging runs but have a finite expected runtime.
- (2) They may terminate with probability one but have an infinite expected runtime.
- (3) Having a finite expected runtime is not preserved under sequencing.

(1) A single diverging run of an ordinary, i.e., nonprobabilistic, algorithm yields the program to have an infinite runtime. This is not true for randomized algorithms. They may admit arbitrarily long and even infinite runs while still having a finite expected runtime. The program

$$C_{geo} : \quad b := 1; \\ \quad \text{while}(b = 1)\{ \\ \quad \quad b := 1/2 \cdot \langle 0 \rangle + 1/2 \cdot \langle 1 \rangle \}$$

keeps flipping a fair coin until observing the first heads (represented by 0). The program C_{geo} admits arbitrarily long runs, since the probability of not seeing a heads in the first n trials is nonzero. This holds for every n . It even admits a nonterminating run (occurring with probability zero), namely, the one in which the outcome of all coin flips is tails. The runtime of the program C_{geo} , however, is geometrically distributed and therefore its expected runtime is finite: on average, it terminates after two loop iterations. It is worth mentioning that the decision problem “does a program terminate with probability one in finite expected time (on all inputs)?” is Π_3^0 -complete in the arithmetical hierarchy, and thus strictly harder than the universal halting problem for ordinary programs [27].

(2) The program C_{geo} terminates with probability one and its expected time until termination is finite. For ordinary sequential programs, termination always implies finite runtime. For probabilistic programs this is not always true. Consider the program

$$C_{rw} : \quad x := 10; \\ \quad \text{while}(x > 0)\{ \\ \quad \quad x := 1/2 \cdot \langle x-1 \rangle + 1/2 \cdot \langle x+1 \rangle \},$$

which models a one-dimensional random walk of a particle: starting from position 10, in each step the particle moves randomly one step to the left or one step to the right, until reaching position 0.

¹If the program does not terminate with some probability greater than zero, its expected runtime is infinity.

The particle reaches position 0 with probability one, but doing so requires infinitely many steps on average (cf. [26, Chapter 3.7.3]).

(3) In the classical setting, running two programs with finite runtime in a row yields a program with a finite runtime. For probabilistic programs this closure property breaks down. Consider the pair of programs

$$\begin{array}{ll}
 C_1 : & x := 1; b := 1; \\
 & \text{while}(b = 1)\{ \\
 & \quad b := 1/2 \cdot \langle 0 \rangle + 1/2 \cdot \langle 1 \rangle; \\
 & \quad x := 2x\}. \\
 C_2 : & \text{while}(x > 0)\{ \\
 & \quad x := x - 1\}
 \end{array}$$

As the loop in C_1 terminates on average in two iterations, it has a finite expected runtime. From any initial state in which x is nonnegative, C_2 makes x iterations, and thus its expected runtime is finite too. However, the program $C_1; C_2$ has an *infinite* expected runtime—even though it almost surely terminates; i.e. it terminates with probability one. This is intuitively due to the fact that the expected value of x after termination of C_1 is infinite and C_2 needs x steps to terminate.

Determining Expected Runtimes. Bounds on the expected runtime of randomized algorithms are typically obtained using classical probability theory, mostly with arguments relying on random variable expectations or martingales [15, 39]. These runtime bounds are nonetheless obtained partially following an ad hoc reasoning, which, moreover, usually takes for granted nontrivial relationships between the involved random variables. This constitutes a somewhat deficient proof methodology that often yields incomplete proof arguments. This article proposes an alternative using formal reasoning as typically used in deductive verification.

A naive approach toward a rigorous, formal reasoning about expected runtimes equips the program at hand with a runtime counter, say, rc . The variable rc is initially set to 0 and incremented for each basic operation that consumes time. The expected value of runtime counter rc is then obtained using existing verification techniques for randomized algorithms, such as Kozen’s probabilistic propositional dynamic logic (PPDL) [32] or the weakest pre-expectation calculus by McIver and Morgan [36]. This approach is, however, *unsound* as the expected value of counter rc may not coincide with the expected runtime.

Let us illustrate this by an example: applying the above principle to the certainly diverging program $\text{while}(\text{true})\{\text{skip}\}$ results (assuming skip consumes one time unit) in the program

$$\begin{array}{l}
 C_{div} : rc := 0 \\
 \quad \text{while}(\text{true})\{ \\
 \quad \quad \text{skip}; rc := rc + 1\}.
 \end{array}$$

PPDL and the weakest pre-expectation calculus would both yield zero as the expected value of rc : since no run of C_{div} terminates, the average value of rc upon termination is an empty sum, i.e. zero. However, the program’s expected runtime is infinite as every program run executes a statement that consumes time infinitely often. This is not just a corner case: in fact, the expected value of rc may assume an arbitrary positive integer value different from the actual expected runtime. Consider the program

$$\begin{array}{ll}
 C_{halfdiv} : rc := 0; \\
 \quad b := 1/2 \cdot \langle 0 \rangle + 1/2 \cdot \langle 1 \rangle; \\
 \quad \ell := 2k; & // \text{consume } 2k \text{ time units}
 \end{array}$$

```

while( $\ell > 0$ ){
     $\ell := \ell - 1$ ;  $rc := rc + 1$ ;
    while( $b = 1$ ){
        skip;  $rc := rc + 1$ 
    }
}

```

that first flips a fair coin to set the variable b to either 0 or 1. Afterward, the first while-loop performs $2k$ loop iterations before the second while-loop either terminates (if $b = 0$) or diverges (if $b = 1$). The variable rc counts the total number of loop iterations. Using the weakest pre-expectation calculus (or PPDL), one can show that for positive k , the expected value of rc for program $C_{halfdiv}$ is upper bounded by $\frac{1}{2} \cdot 2\lceil k \rceil + \frac{1}{2} \cdot 0 = \lceil k \rceil$, although the actual expected runtime of $C_{halfdiv}$ is infinite for any k .

The Approach of This Article. The inability of the weakest pre-expectation calculus and PPDL to soundly reason about expected runtimes, the manifold intricacies in the runtime analyses in the presence of randomization, and the deficiency of various ad hoc runtime analyses in the literature necessitate a more rigorous, systematic, and compelling approach for the runtime analysis of randomized algorithms. This article proposes a technique based on *formal program verification*. This technique derives runtime claims from first principles only. It consists of a wp-style calculus à la Dijkstra [12]. In a similar vein to Dijkstra’s predicate transformers, our wp-style calculus uses *runtime transformers*. The core of this calculus is the transformer ert :

$$ert[C](t)(\sigma)$$

gives the expected runtime of program C when started in initial state σ , under the assumption that t captures the expected runtime of the computation following C . In particular, the expected runtime of program C on initial state σ is given by $ert[C](\mathbf{0})(\sigma)$, where $\mathbf{0}$ is the constantly zero runtime.

Our calculus is defined over a simple probabilistic imperative language with recursive procedures. For most control structures, it is defined in a straightforward compositional manner. The action of the transformer on guarded loops and recursive procedures is given using fixed-point techniques. To avoid the tedious reasoning about such fixed points and enhance the calculus’s usability, we provide *invariant-based proof rules* that establish bounds on the expected runtime of loops and of recursive programs.

At the theoretical level, we validate our wp-style calculus in two ways: First, we show that the calculus corresponds to a simple, intuitive operational model of probabilistic programs based on Markov chains. Second, we show that the calculus is a conservative extension of Nielson’s approach for reasoning about the runtime of ordinary, nonprobabilistic programs [41]. On the more practical side, we use our calculus to perform a formal runtime analysis of a one-dimensional random walk, the coupon collector’s problem [37], and a Sherwood variant of the binary search algorithm [35].

We stress two important assets of our approach: First, our calculus for expected runtimes is amenable to a large degree of automation. For several cases, loop invariants can be synthesized from previously provided templates and our proof rules allow for mechanical verification of proposed invariants. An implementation of our calculus in the theorem prover Isabelle/HOL has recently been reported [23], certifying all the theoretical results in this article.

A second asset is that our calculus enables determining whether the expected runtime of a randomized algorithm (for all possible inputs) is finite or not. In combination with almost-sure termination, this is a very relevant property. To the best of our knowledge, this is the first formal verification framework that can handle both (universal) almost-sure termination (does a program

terminate with probability one on every input?) and (universal) positive almost-sure termination (does a program terminate with probability one in finite expected time on every inputs?).

Main Contributions of this Article. To summarize, this article presents a calculus for reasoning about the—finite or possibly infinite—expected runtime of randomized algorithms

- (1) with a set of invariant-based proof rules for obtaining bounds on the expected runtime of loops and recursive procedures,
- (2) which corresponds to a simple operational program model using Markov chains,
- (3) which conservatively extends a calculus [41] for runtime analysis of ordinary programs, and
- (4) which is applied to analyze the runtime of the coupon collector’s problem, a one-dimensional random walk, and a randomized binary search.

Organization of the Article. Section 2 defines our probabilistic programming language. Section 3 presents the transformer ert and studies its elementary properties. Section 4 presents proof rules for obtaining upper and lower bounds on the expected runtime of loops. Section 5 shows that the ert -transformer coincides with the expected runtime in a Markov chain that acts as an operational program model. Section 6 proves that the ert -transformer is a conservative extension of Nielson’s approach for obtaining upper bounds on deterministic programs. Section 7 extends the ert -calculus with recursion. Section 8 establishes a connection between transformer ert and the weakest precondition semantics of probabilistic programs. Section 9 discusses three case studies in detail. Section 10 summarizes related work, and Section 11 concludes.

The proofs of the results that do not proceed by a tedious induction on the program structure are provided in the article body. All inductive proofs as well as the detailed calculations for the case studies are provided in the appendix. The material presented in this article unifies and extends [28] and [44].

2 A PROBABILISTIC PROGRAMMING LANGUAGE FOR RANDOMIZED ALGORITHMS

In this section, we present the probabilistic programming language used throughout the article, together with its runtime model. For ease of presentation, we treat nonrecursive programs first and extend our calculus with recursion in Section 7. To model randomized algorithms, we employ a standard imperative language à la Dijkstra’s Guarded Command Language [12] with two distinguished features: We allow distribution expressions in assignments and guards to be probabilistic. For instance, we allow for probabilistic assignments like

$$y := \text{Unif}[1 \dots x],$$

which endows variable y with a uniform distribution in the interval $[1 \dots x]$. Notice that x is another program variable, so the resulting distribution of y depends on the current program state. We also allow for a program like

```
x := 0;
while (p · ⟨true⟩ + (1-p) · ⟨false⟩) {
  x := x + 1},
```

which uses a probabilistic loop guard to simulate a geometric distribution with success probability p ; i.e., the loop guard evaluates to true with probability p and to false with the remaining probability $1-p$.

Formally, the set of *probabilistic programs* pGCL is given by the grammar

$C ::=$	empty	empty program
	skip	effectless operation
	halt	immediate termination
	$x \approx \mu$	probabilistic assignment
	$C; C$	sequential composition
	if (ξ) { C } else { C }	probabilistic conditional choice
	while (ξ) { C }	probabilistic guarded while-loop

Here x represents a *program variable* in Var , μ a *distribution expression* in DExp , and ξ a distribution expression over the truth values, i.e., a *probabilistic guard*, in DExp . We assume distribution expressions in DExp to represent discrete probability distributions with a (possibly *infinite*) support of total probability mass 1. Let $p_1 \cdot \langle a_1 \rangle + \dots + p_n \cdot \langle a_n \rangle$ denote the distribution expression that assigns probability p_i to a_i . For instance, the distribution expression $1/2 \cdot \langle \text{true} \rangle + 1/2 \cdot \langle \text{false} \rangle$ represents the toss of a fair coin. Deterministic expressions over program variables such as $x - y$ or $x - y > 8$ are special instances of distribution expressions; they are understood as Dirac probability distributions that assign the total probability mass, i.e., one, to a single point.

To describe the effect of the different language constructs, we use some preliminaries. A *program state* σ is a mapping from program variables to values in Val . Let

$$\Sigma \triangleq \{ \sigma \mid \sigma : \text{Var} \rightarrow \text{Val} \}$$

be the set of all program states. We assume an interpretation function $\llbracket \cdot \rrbracket : \text{DExp} \rightarrow (\Sigma \rightarrow \mathcal{D}(\text{Val}))$ for distribution expressions, with $\mathcal{D}(\text{Val})$ being the set of discrete probability distributions over Val . For $\mu \in \text{DExp}$, $\llbracket \mu \rrbracket$ maps each program state to a probability distribution of values. Let $\llbracket \mu : v \rrbracket$ be a shorthand for the function mapping each program state σ to the probability that distribution $\llbracket \mu \rrbracket(\sigma)$ assigns to value v , i.e.,

$$\llbracket \mu : v \rrbracket(\sigma) \triangleq \text{Pr}_{\llbracket \mu \rrbracket(\sigma)}(v),$$

where Pr denotes the probability operator on distributions over values.

The effect of pGCL program constructs and their assumed timing is as follows:

- empty has no effect and its execution consumes no time.
- skip has also no effect but consumes, in contrast to empty, one unit of time.
- halt aborts any further program execution and consumes no time.
- $x \approx \mu$ is a probabilistic assignment that samples a value from $\llbracket \mu \rrbracket$ and assigns it to variable x . The sampling and assignment consume (altogether) one unit of time.
- $C_1; C_2$ is the sequential composition of programs C_1 and C_2 ; i.e. first C_1 is executed, then C_2 . The composition itself consumes no additional time.
- if (ξ) { C_1 } else { C_2 } is a probabilistic conditional branching: with probability $\llbracket \xi : \text{true} \rrbracket$ program C_1 is executed, whereas with probability $\llbracket \xi : \text{false} \rrbracket = 1 - \llbracket \xi : \text{true} \rrbracket$ program C_2 is executed. Evaluating (or more rigorously, sampling a value from) the probabilistic guard requires an additional unit of time.
- while (ξ) { C } is a probabilistic while loop: with probability $\llbracket \xi : \text{true} \rrbracket$ the loop body C is executed followed by a recursive execution of the loop, whereas with probability $\llbracket \xi : \text{false} \rrbracket$ the loop terminates immediately. As for conditionals, each evaluation of the guard consumes one unit of time.

Alternative Runtime Models. We stress that the above runtime model is a design decision for the sake of concreteness. All our development can be easily adapted to capture alternative models.

These include, for instance, the model where only the number of assignments in a program run or the model where only the number of loop iterations are of relevance. We can also capture more fine-grained models, where for instance the runtime of an assignment depends on some notion of *size* of the distribution expression being sampled.

Example 2.1 (Race between tortoise and hare). To illustrate the use of our programming language, consider the following program adopted from [9]:

```

 $h := 0; t := 30;$ 
while ( $h \leq t$ ) {
  if ( $\frac{1}{2} \cdot \langle \text{true} \rangle + \frac{1}{2} \cdot \langle \text{false} \rangle$ ) {
     $h := h + \text{Unif}[0 \dots 10]$ 
  }
  else {empty};
   $t := t + 1$ 
}

```

It models a race between a tortoise and a hare; the variables h and t represent their respective positions. The tortoise starts with a lead of 30 and in each round (i.e., in each loop iteration) advances one step forward. The hare with probability $\frac{1}{2}$ advances a random number of steps between 0 and 10 (governed by a uniform distribution) and with the remaining probability remains still. The race ends when the hare passes the tortoise.

Regarding the runtime, the program requires two units of time for the initial assignments. In every loop iteration, the program consumes either three or four units of time: it always takes one unit of time to evaluate the loop guard, evaluate the probabilistic conditional, and update variable t , respectively. If the probabilistic conditional is evaluated to true, an additional unit of time is consumed to update the value of variable h . \triangle

We conclude this section by fixing some notational conventions. To keep our program notation consistent with standard usage, we write $x := \mu$ instead of $x \approx \mu$ whenever μ represents a Dirac distribution given by a deterministic expression over program variables. For instance, in the program in Example 2.1 above, we shall write $t := t + 1$ instead of $t \approx t + 1$. Likewise, when ξ is a probabilistic guard given as a deterministic Boolean expression over program variables, we use $\llbracket \xi \rrbracket$ to denote $\llbracket \xi : \text{true} \rrbracket$ and $\llbracket \neg \xi \rrbracket$ to denote $\llbracket \xi : \text{false} \rrbracket$. For instance, we write $\llbracket b = 0 \rrbracket$ instead of $\llbracket b = 0 : \text{true} \rrbracket$.

3 A CALCULUS OF EXPECTED RUNTIMES

Our goal is to associate to any program C a function that maps each state σ to the average or expected runtime of executing C on initial state σ . To model these functions, we use the function space of *runtimes*:

$$\mathbb{T} \triangleq \{t \mid t : \Sigma \rightarrow \mathbb{R}_{\geq 0}^{\infty}\}.$$

Here, $\mathbb{R}_{\geq 0}^{\infty}$ represents the set of nonnegative real numbers extended with ∞ . Let \mathbf{k} denote the constant runtime $\lambda \sigma. k$ for $k \in \mathbb{R}_{\geq 0}^{\infty}$.

We express the expected runtime of programs in a continuation-passing style by means of the transformer

$$\text{ert}[\cdot] : \text{pGCL} \rightarrow (\mathbb{T} \rightarrow \mathbb{T}).$$

Concretely, $\text{ert}[C](t)(\sigma)$ gives the expected runtime of executing program C on initial state σ assuming that t captures the runtime of the computation that follows C . The function t is usually referred to as the *continuation* and we can think of it as being evaluated in the final states that are

Table 1. Rules for Defining the Expected Runtime Transformer ert

C	$\text{ert}[C](t)$
empty	t
skip	$\mathbf{1} + t$
halt	$\mathbf{0}$
$x := \mu$	$\mathbf{1} + \lambda\sigma. E_{\llbracket\mu\rrbracket(\sigma)}(\lambda v. t[x/v](\sigma))$
$C_1; C_2$	$\text{ert}[C_1](\text{ert}[C_2](t))$
if $(\xi) \{C_1\}$ else $\{C_2\}$	$\mathbf{1} + \llbracket\xi : \text{true}\rrbracket \cdot \text{ert}[C_1](t) + \llbracket\xi : \text{false}\rrbracket \cdot \text{ert}[C_2](t)$
while $(\xi) \{C'\}$	$\text{lfp } X \cdot \mathbf{1} + \llbracket\xi : \text{false}\rrbracket \cdot t + \llbracket\xi : \text{true}\rrbracket \cdot \text{ert}[C'](X)$

$\mathbf{1}$ is the constant runtime $\lambda\sigma. \mathbf{1}$; $E_\mu(h) \triangleq \sum_v \text{Pr}_\mu(v) \cdot h(v)$ represents the expected value of (random variable) h w.r.t. distribution μ ; $t[x/v] \triangleq \lambda\sigma. t(\sigma[x/v])$, where $\sigma[x/v]$ is the state obtained by updating in σ the value of x to v ; finally $\text{lfp } X. F(X)$ represents the least fixed point of transformer $F : \mathbb{T} \rightarrow \mathbb{T}$ with respect to the pointwise ordering on \mathbb{T} .

reached upon termination of C . Thus, $\text{ert}[C](\mathbf{0})(\sigma)$ gives the plain expected runtime of executing program C on initial state σ .

The transformer ert is defined by induction on the structure of C following the rules in Table 1. The rules correspond to the runtime model introduced in Section 2. That is, $\text{ert}[C](\mathbf{0})$ captures the expected number of assignments, guard evaluations, and skip statements. Most rules in Table 1 are self-explanatory. $\text{ert}[\text{empty}]$ behaves as the identity since empty does not modify the program state and its execution consumes no time. On the other hand, $\text{ert}[\text{skip}]$ adds one unit of time to any continuation since this is the time required by the execution of skip . $\text{ert}[\text{halt}]$ yields always the constant runtime $\mathbf{0}$ since halt aborts any subsequent program execution (making their runtime irrelevant) and consumes no time. The definition of ert on random assignments is more involved: $\text{ert}[x := \mu](t)(\sigma) = \mathbf{1} + \sum_v \text{Pr}_{\llbracket\mu\rrbracket(\sigma)}(v) \cdot t(\sigma[x/v])$ is obtained by adding one unit of time (due to the distribution sampling and assignment of its outcome) to the sum of the runtime of each possible subsequent execution, weighted according to their probabilities.

As for the composite statements, $\text{ert}[C_1; C_2]$ applies $\text{ert}[C_1]$ to the expected runtime obtained from the application of $\text{ert}[C_2]$ to the continuation t . $\text{ert}[\text{if } (\xi) \{C_1\} \text{ else } \{C_2\}]$ adds one unit of time (for the guard evaluation) to the weighted sum of the runtime of the two branches.

The ert of a loop is given as the least fixed point (with respect to the pointwise order in \mathbb{T}) of a runtime transformer $F : \mathbb{T} \rightarrow \mathbb{T}$, defined in terms of the runtime of the loop body. To guarantee the existence of such a fixed point, we use a standard argument (see, e.g., [50, Ch. 5]): we endow \mathbb{T} with the structure of an ω -complete partial order (ω -cpo for short) and we prove that F is continuous. Since the transformer F is used repeatedly in the rest of our development, we use the following notation:

Definition 3.1 (Characteristic functional of a loop). Given loop $\text{while } (\xi) \{C\}$ and runtime $t \in \mathbb{T}$, let

$$F_t^{(\xi, C)} : \mathbb{T} \rightarrow \mathbb{T}, \quad X \mapsto \mathbf{1} + \llbracket\xi : \text{false}\rrbracket \cdot t + \llbracket\xi : \text{true}\rrbracket \cdot \text{ert}[C](X)$$

be the *characteristic functional* of the loop with respect to (continuation) t . \triangle

When the loop is understood from the context, we simply write F_t for $F_t^{(\xi, C)}$. Using this definition, the ert of a loop can be recast as

$$\text{ert}[\text{while } (\xi) \{C\}](t) = \text{lfp } F_t^{(\xi, C)}.$$

The two steps that we take to guarantee the existence of $\text{lfp } F_t^{(\xi, C)}$ are as follows: First, we endow \mathbb{T} with the structure of an ω -cpo: runtimes are ordered pointwise, i.e., $t_1 \leq t_2$ if and only if $t_1(\sigma) \leq t_2(\sigma)$ for all $\sigma \in \Sigma$; the supremum of ω -chains is also defined pointwise, i.e., for $t_1 \leq t_2 \leq \dots$, $\sup_n t_n \triangleq \lambda\sigma. \sup_n t_n(\sigma)$; finally, the bottom element of the ω -cpo is the constant runtime 0. Second, we prove that the characteristic functional $F_t^{(\xi, C)}$ is continuous. The Kleene Fixed Point Theorem (Theorem A.3) ensures the existence of $\text{lfp } F_t^{(\xi, C)}$. It follows that the action of ert on loops is well defined. The continuity of $F_t^{(\xi, C)}$ follows immediately from the continuity of $\text{ert}[C]$, established below.

LEMMA 3.2 (CONTINUITY OF ERT). *For every program $C \in \text{pGCL}$, $\text{ert}[C] : \mathbb{T} \rightarrow \mathbb{T}$ is continuous; i.e., for every ω -chain of runtimes $t_0 \leq t_1 \leq \dots$ we have*

$$\text{ert}[C](\sup_n t_n) = \sup_n \text{ert}[C](t_n).$$

PROOF. By induction on the structure of C ; see Appendix B.1. □

We now illustrate the use of the ert transformer by analyzing the runtime of a program that simulates a truncated geometric distribution.

Example 3.3 (Truncated geometric distribution). Program C_{trunc} repeatedly flips a fair coin until observing the first heads (represented by `true` in the probabilistic guards) or completing the second unsuccessful trial:

$$\begin{aligned} C_{\text{trunc}} : & \text{ if } \left(\frac{1}{2} \cdot \langle \text{true} \rangle + \frac{1}{2} \cdot \langle \text{false} \rangle \right) \{ \\ & \quad \text{succ} := \text{true} \} \\ & \text{ else } \{ \\ & \quad \text{if } \left(\frac{1}{2} \cdot \langle \text{true} \rangle + \frac{1}{2} \cdot \langle \text{false} \rangle \right) \{ \\ & \quad \quad \text{succ} := \text{true} \} \\ & \quad \text{else } \{ \\ & \quad \quad \text{succ} := \text{false} \} \}. \end{aligned}$$

The calculation of the expected runtime $\text{ert}[C_{\text{trunc}}](\mathbf{0})$ of the program goes as follows:

$$\begin{aligned} \text{ert}[C_{\text{trunc}}](\mathbf{0}) &= \mathbf{1} + \frac{1}{2} \cdot \text{ert}[\text{succ} := \text{true}](\mathbf{0}) \\ &\quad + \frac{1}{2} \cdot \text{ert}[\text{if } (\dots) \{ \text{succ} := \text{true} \} \text{ else } \{ \text{succ} := \text{false} \}](\mathbf{0}) \\ &= \mathbf{1} + \frac{1}{2} \cdot \text{ert}[\text{succ} := \text{true}](\mathbf{0}) \\ &\quad + \frac{1}{2} \cdot \left(\mathbf{1} + \frac{1}{2} \cdot \text{ert}[\text{succ} := \text{true}](\mathbf{0}) + \frac{1}{2} \cdot \text{ert}[\text{succ} := \text{false}](\mathbf{0}) \right) \\ &= \mathbf{1} + \frac{1}{2} \cdot (\mathbf{1} + 0) + \frac{1}{2} \cdot \left(\mathbf{1} + \frac{1}{2} \cdot (\mathbf{1} + 0) + \frac{1}{2} \cdot (\mathbf{1} + 0) \right) = \frac{5}{2}. \end{aligned}$$

Therefore, the execution of C_{trunc} takes, on average, 2.5 units of time. △

Note that the calculation of the expected runtime in the example above is straightforward as the program is loop-free. Computing the runtime of loops requires the calculation of least fixed points, which is generally not feasible in practice. To circumvent this, in the next section we present proof rules based on loop invariants.

The ert transformer possesses several algebraic properties:

THEOREM 3.4 (BASIC PROPERTIES OF THE ert TRANSFORMER). *For any program $C \in \text{pGCL}$, any constant runtime k with $k \in \mathbb{R}_{\geq 0}$ and any runtimes $t, t' \in \mathbb{T}$, it holds:*

Monotonicity: $t \leq t'$ implies $\text{ert}[C](t) \leq \text{ert}[C](t')$;
Constant propagation: $\text{ert}[C](\mathbf{k} + t) = \mathbf{k} + \text{ert}[C](t)$, provided C is halt-free;
Preservation of ∞ : $\text{ert}[C](\infty) = \infty$, provided C is halt-free.

PROOF. Monotonicity follows from continuity (see Lemma 3.2). For the proof of constant propagation, see Appendix B.2. Finally, preservation of infinity follows from monotonicity and constant propagation since together they entail

$$\text{ert}[C](\infty) \geq \text{ert}[C](\mathbf{k}) = \mathbf{k} + \text{ert}[C](0) \geq \mathbf{k}$$

for every $k \in \mathbb{R}_{\geq 0}$. This immediately yields $\text{ert}[C](\infty) = \infty$. \square

The ert transformer is not linear in general, but it satisfies the weaker properties of subadditivity and subscalability, which are discussed in Section 8.

4 EXPECTED RUNTIME OF LOOPS

As mentioned earlier, reasoning about the runtime of loop-free programs involves mostly syntactic reasoning. The runtime of a loop, however, is given in terms of a least fixed point. It can thus be obtained by fixed-point iteration, but the fixed point need not be reached within a finite number of iterations. To overcome this problem, we study invariant-based proof rules for approximating the runtime of loops.

We present two families of proof rules that differ in the kind of invariants they build upon. In Section 4.1, we present a proof rule that rests on the presence of an invariant approximating the entire runtime of a loop in a global manner, while in Section 4.2 we present two proof rules that each rely on a parameterized invariant that approximates the runtime of a loop in an incremental fashion. In Section 4.3, we discuss how to tighten the runtime bounds yielded by any of these proof rules.

4.1 Proof Rule Based on Global Invariants

Our first proof rule allows for bounding the expected runtime of loops from above and rests on the notion of *upper invariants*:

Definition 4.1 (Upper invariants). Runtime $I \in \mathbb{T}$ is an *upper invariant of loop while* $(\xi) \{C\}$ with respect to t if and only if

$$F_t^{(\xi, C)}(I) \leq I. \quad \triangle$$

Such an upper invariant readily establishes an upper bound of the loop's runtime:

THEOREM 4.2 (UPPER BOUNDS FROM UPPER INVARIANTS). *If $I \in \mathbb{T}$ is an upper invariant of while $(\xi) \{C\}$ with respect to t , then*

$$\text{ert}[\text{while } (\xi) \{C\}](t) \leq I.$$

PROOF. The theorem follows by an application of Park's Theorem (Theorem A.4), which in our case, given that $F_t^{(\xi, C)}$ is continuous (see Lemma 3.2), states that

$$F_t^{(\xi, C)}(I) \leq I \text{ implies } \text{lfp } F_t^{(\xi, C)} \leq I.$$

The left-hand side of the implication is equivalent to I being an upper invariant, while the right-hand side is equivalent to $\text{ert}[\text{while } (\xi) \{C\}](f) \leq I$. \square

Example 4.3 (Geometric distribution). Consider the loop

$$C_{\text{geo}} : \text{while } (c = 1) \{ \\ c := 1/2 \cdot \langle 0 \rangle + 1/2 \cdot \langle 1 \rangle \}$$

whose runtime is geometrically distributed. Its characteristic functional with respect to continuation $t = 0$ is

$$F_0(X) = \mathbf{1} + \llbracket c \neq 1 \rrbracket \cdot \mathbf{0} + \llbracket c = 1 \rrbracket \cdot \text{ert}[c := 1/2 \cdot \langle 0 \rangle + 1/2 \cdot \langle 1 \rangle](X).$$

By the calculations below, we verify that $I = \mathbf{1} + \llbracket c = 1 \rrbracket \cdot \mathbf{4}$ is an upper invariant of the loop (with respect to $\mathbf{0}$):

$$\begin{aligned} F_0(I) &= \mathbf{1} + \llbracket c \neq 1 \rrbracket \cdot \mathbf{0} + \llbracket c = 1 \rrbracket \cdot \text{ert}[c := 1/2 \cdot \langle 0 \rangle + 1/2 \cdot \langle 1 \rangle](I) \\ &= \mathbf{1} + \llbracket c = 1 \rrbracket \cdot \left(\mathbf{1} + \frac{1}{2} \cdot I[c/0] + \frac{1}{2} \cdot I[c/1] \right) \\ &= \mathbf{1} + \llbracket c = 1 \rrbracket \cdot \left(\mathbf{1} + \frac{1}{2} \cdot \underbrace{(\mathbf{1} + \llbracket 0 = 1 \rrbracket \cdot \mathbf{4})}_{=1} + \frac{1}{2} \cdot \underbrace{(\mathbf{1} + \llbracket 1 = 1 \rrbracket \cdot \mathbf{4})}_{=5} \right) \\ &= \mathbf{1} + \llbracket c = 1 \rrbracket \cdot \mathbf{4} = I \leq I. \end{aligned}$$

By applying Theorem 4.2, we obtain

$$\text{ert}[C_{\text{geo}}](\mathbf{0}) \leq \mathbf{1} + \llbracket c = 1 \rrbracket \cdot \mathbf{4}.$$

In words, the expected runtime of C_{geo} is at most $1 + 4 = 5$ from any initial state where $c = 1$ and at most $1 + 0 = 1$ from any other state. \triangle

Notice that if the loop body is itself loop-free, as in the above example, verifying that some $I \in \mathbb{T}$ is an upper invariant is usually fairly easy. Inferring the invariant, in contrast, is one of the most involved parts of the verification effort.

The invariant-based technique to reason about the runtime of loops presented in Theorem 4.2 is complete in the sense that there always exists an upper invariant that establishes the exact runtime of the loop at hand:

THEOREM 4.4. *There exists an upper invariant I of $\text{while } (\xi) \{C\}$ with respect to t such that $\text{ert}[\text{while } (\xi) \{C\}](t) = I$.*

PROOF. It suffices to show that $\text{ert}[\text{while } (\xi) \{C\}](t)$ itself is an upper invariant of the loop. Since $\text{ert}[\text{while } (\xi) \{C\}](t) = \text{lfp } F_t^{(\xi, C)}$, this amounts to showing that

$$F_t^{(\xi, C)} \left(\text{lfp } F_t^{(\xi, C)} \right) \leq \text{lfp } F_t^{(\xi, C)},$$

which holds by definition of lfp . \square

Intuitively, the proof of this theorem shows that $\text{ert}[\text{while } (\xi) \{C\}](t)$ itself is the tightest upper invariant that the loop admits.

4.2 Proof Rules Based on Parameterized Invariants

We now study a second family of proof rules that builds on the notion of ω -invariants for bounding the runtime of loops from both above and below.

Definition 4.5 (ω -Invariants). Let $n \in \mathbb{N}$. The parameterized runtime $I_n \in \mathbb{T}$ is a lower ω -invariant of loop $\text{while } (\xi) \{C\}$ with respect to t if and only if

$$I_0 \leq F_t^{(\xi, C)}(\mathbf{0}) \quad \text{and} \quad I_{n+1} \leq F_t^{(\xi, C)}(I_n), \quad \text{for all } n \geq 0.$$

Dually, I_n is an *upper ω -invariant* if and only if

$$F_t^{(\xi, C)}(\mathbf{0}) \leq I_0 \quad \text{and} \quad F_t^{(\xi, C)}(I_n) \leq I_{n+1}, \quad \text{for all } n \geq 0. \quad \triangle$$

Intuitively, a lower (upper) ω -invariant I_n represents a lower (upper) bound for the expected runtime of those program runs that finish within $n+1$ iterations, weighted according to their probabilities. Therefore, we can use the asymptotic behavior of I_n to approximate from below (above) the expected runtime of the entire loop.

THEOREM 4.6 (BOUNDS FROM ω -INVARIANTS).

- (1) *If I_n is a lower ω -invariant of loop while $(\xi) \{C\}$ with respect to t and the limit $\lim_{n \rightarrow \infty} I_n$ exists², then*

$$\text{ert}[\text{while } (\xi) \{C\}](t) \geq \lim_{n \rightarrow \infty} I_n.$$

- (2) *If I_n is an upper ω -invariant of loop while $(\xi) \{C\}$ with respect to t and the limit $\lim_{n \rightarrow \infty} I_n$ exists, then*

$$\text{ert}[\text{while } (\xi) \{C\}](t) \leq \lim_{n \rightarrow \infty} I_n.$$

PROOF. We prove only the case of lower ω -invariants since the other case follows by a dual argument. Let F_t be the characteristic functional of the loop with respect to t and let F_t^n denote the n -fold composition of F_t with itself, i.e., the function $F_t \circ \dots \circ F_t$ (n times). By the Kleene Fixed Point Theorem (Theorem A.3 in Appendix A), $\text{ert}[\text{while } (\xi) \{C\}](t) = \sup_n F_t^n(0)$, and since $F_t^0(0) \leq F_t^1(0) \leq \dots$ forms an ω -chain, by the Monotone Sequence Theorem (Theorem A.2 in Appendix A), $\sup_n F_t^n(0) = \lim_{n \rightarrow \infty} F_t^n(0)$. Then the result follows from showing that $F_t^{n+1}(0) \geq I_n$. We prove this by induction on n . The base case (i.e., $n = 1$) holds, since $F_t^1(0) \geq I_0$ holds by the assumption that I_n is a lower ω -invariant. For the inductive case we reason as follows:

$$F_t^{n+2}(0) = F_t(F_t^{n+1}(0)) \geq F_t(I_n) \geq I_{n+1}.$$

Here the first inequality follows by I.H. and monotonicity of F_t (recall that $\text{ert}[C]$ is monotonic by Theorem 3.4), while the second inequality holds by the assumption that I_n is a lower ω -invariant. \square

Example 4.7 (Lower bounds for C_{geo}). Reconsider loop C_{geo} from Example 4.3. We use Theorem 4.6 (1) to show that $\mathbf{1} + \llbracket c = 1 \rrbracket \cdot \mathbf{4}$ is also a lower bound of its runtime. Let us first show that $I_n = \mathbf{1} + \llbracket c = 1 \rrbracket \cdot (\mathbf{4} - 3/2^n)$ is a lower ω -invariant of the loop with respect to $\mathbf{0}$:

$$\begin{aligned} F_0(\mathbf{0}) &= \mathbf{1} + \llbracket c \neq 1 \rrbracket \cdot \mathbf{0} + \llbracket c = 1 \rrbracket \cdot \text{ert}[c \approx 1/2\langle 0 \rangle + 1/2\langle 1 \rangle](\mathbf{0}) \\ &= \mathbf{1} + \llbracket c = 1 \rrbracket \cdot \left(\mathbf{1} + \frac{1}{2} \cdot \mathbf{0}[c/0] + \frac{1}{2} \cdot \mathbf{0}[c/1] \right) \\ &= \mathbf{1} + \llbracket c = 1 \rrbracket \cdot \mathbf{1} = \mathbf{1} + \llbracket c = 1 \rrbracket \cdot (\mathbf{4} - 3/2^0) = I_0 \geq I_0 \end{aligned}$$

$$\begin{aligned} F_0(I_n) &= \mathbf{1} + \llbracket c \neq 1 \rrbracket \cdot \mathbf{0} + \llbracket c = 1 \rrbracket \cdot \text{ert}[c \approx 1/2\langle 0 \rangle + 1/2\langle 1 \rangle](I_n) \\ &= \mathbf{1} + \llbracket c = 1 \rrbracket \cdot \left(\mathbf{1} + \frac{1}{2} \cdot I_n[c/0] + \frac{1}{2} \cdot I_n[c/1] \right) \\ &= \mathbf{1} + \llbracket c = 1 \rrbracket \cdot \left(\mathbf{1} + \frac{1}{2} \cdot (\mathbf{1} + \mathbf{0}) + \frac{1}{2} \cdot \left(\mathbf{1} + \left(\mathbf{4} - \frac{3}{2^n} \right) \right) \right) \\ &= \mathbf{1} + \llbracket c = 1 \rrbracket \cdot \left(\mathbf{4} - \frac{3}{2^{n+1}} \right) = I_{n+1} \geq I_{n+1}. \end{aligned}$$

²The limit $\lim_{n \rightarrow \infty} I_n$ is to be understood pointwise on $\mathbb{R}_{\geq 0}^{\infty}$, i.e., $\lim_{n \rightarrow \infty} I_n = \lambda \sigma \cdot \lim_{n \rightarrow \infty} I_n(\sigma)$ and $\lim_{n \rightarrow \infty} I_n(\sigma) = \infty$ is considered a valid value.

Then from Theorem 4.6 (1) we obtain

$$\text{ert}[C_{\text{geo}}](\mathbf{0}) \geq \lim_{n \rightarrow \infty} \left(\mathbf{1} + \llbracket c = 1 \rrbracket \cdot \left(4 - \frac{3}{2^n} \right) \right) = \mathbf{1} + \llbracket c = 1 \rrbracket \cdot 4.$$

Combining this result with the upper bound $\text{ert}[C_{\text{geo}}](\mathbf{0}) \leq \mathbf{1} + \llbracket c = 1 \rrbracket \cdot 4$ we had established in Example 4.3, we conclude that $\mathbf{1} + \llbracket c = 1 \rrbracket \cdot 4$ is the *exact* runtime of C_{geo} . Observe, however, that the above calculations show that I_n is both a lower and an upper ω -invariant (exact equalities $F_0(\mathbf{0}) = I_0$ and $F_0(I_n) = I_{n+1}$ hold). Then we can apply Theorem 4.6 (1) and 4.6 (2) simultaneously to derive the exact runtime without having to resort to the result from Example 4.3.

Invariant Synthesis. In order to obtain invariant $I_n = \mathbf{1} + \llbracket c = 1 \rrbracket \cdot (4 - \frac{3}{2^n})$, we used template $I_n = \mathbf{1} + \llbracket c = 1 \rrbracket \cdot a_n$ and observed that under this template the definition of lower ω -invariant reduces to $a_0 \leq 1$, $a_{n+1} \leq 2 + \frac{1}{2}a_n$, with solution $a_n = 4 - \frac{3}{2^n}$. \triangle

Example 4.8 (Almost-sure termination at infinite expected runtime). Recall the program from the introduction that had an infinite expected runtime, despite being the concatenation of two programs with finite expected runtime:

```
C:  1: x := 1; b := 1;
    2: while (b = 1) { b := 1/2⟨0⟩ + 1/2⟨1⟩; x := 2x };
    3: while (x > 0) { x := x - 1 }
```

Now we apply Theorem 4.6 to formally analyze its runtime. We show that $\text{ert}[C](\mathbf{0}) \geq \infty$. In the course of doing so, we use C_i to denote the i th line of C . Since

$$\text{ert}[C](\mathbf{0}) = \text{ert}[C_1](\text{ert}[C_2](\text{ert}[C_3](\mathbf{0}))),$$

we start by analyzing the runtime of the loop C_3 . Using the lower ω -invariant

$$J_n = \mathbf{1} + \llbracket 0 < x < n \rrbracket \cdot 2x + \llbracket x \geq n \rrbracket \cdot (2n - 1),$$

we conclude that $\text{ert}[C_3](\mathbf{0}) \geq \mathbf{1} + \llbracket x > 0 \rrbracket \cdot 2x = \lim_{n \rightarrow \infty} J_n$. Next we show that

$$\begin{aligned} \text{ert}[C_2](\mathbf{1} + \llbracket x > 0 \rrbracket \cdot 2x) &\geq \mathbf{1} + \llbracket b \neq 1 \rrbracket \cdot (\mathbf{1} + \llbracket x > 0 \rrbracket \cdot 2x) \\ &\quad + \llbracket b = 1 \rrbracket \cdot (7 + \llbracket x > 0 \rrbracket \cdot \infty) \end{aligned}$$

by means of the lower ω -invariant

$$I_n = \mathbf{1} + \llbracket b \neq 1 \rrbracket \cdot (\mathbf{1} + \llbracket x > 0 \rrbracket \cdot 2x) + \llbracket b = 1 \rrbracket \cdot \left(7 - \frac{5}{2^n} + n \cdot \llbracket x > 0 \rrbracket \cdot 2x \right).$$

Let F be the characteristic functional of loop C_2 with respect to $\mathbf{1} + \llbracket x > 0 \rrbracket \cdot 2x$. The calculations to establish that I_n is a lower ω -invariant go as follows:

$$\begin{aligned} F(\mathbf{0}) &= \mathbf{1} + \llbracket b \neq 1 \rrbracket \cdot (\mathbf{1} + \llbracket x > 0 \rrbracket \cdot 2x) \\ &\quad + \llbracket b = 1 \rrbracket \cdot \left(\mathbf{1} + \frac{1}{2} \cdot (\mathbf{1} + \mathbf{0}[x, b/2x, 0]) + \frac{1}{2} \cdot (\mathbf{1} + \mathbf{0}[x, b/2x, 1]) \right) \\ &= \mathbf{1} + \llbracket b \neq 1 \rrbracket \cdot (\mathbf{1} + \llbracket x > 0 \rrbracket \cdot 2x) + \llbracket b = 1 \rrbracket \cdot \left(\mathbf{1} + \frac{1}{2} \cdot \mathbf{1} + \frac{1}{2} \cdot \mathbf{1} \right) \\ &= \mathbf{1} + \llbracket b \neq 1 \rrbracket \cdot (\mathbf{1} + \llbracket x > 0 \rrbracket \cdot 2x) + \llbracket b = 1 \rrbracket \cdot 2 = I_0 \geq I_0 \end{aligned}$$

$$\begin{aligned}
F(I_n) &= \mathbf{1} + \llbracket b \neq 1 \rrbracket \cdot (\mathbf{1} + \llbracket x > 0 \rrbracket \cdot 2x) \\
&\quad + \llbracket b = 1 \rrbracket \cdot \left(\mathbf{1} + \frac{1}{2} \cdot (\mathbf{1} + I_n[x, b/2x, 0]) + \frac{1}{2} \cdot (\mathbf{1} + I_n[x, b/2x, 1]) \right) \\
&= \mathbf{1} + \llbracket b \neq 1 \rrbracket \cdot (\mathbf{1} + \llbracket x > 0 \rrbracket \cdot 2x) \\
&\quad + \llbracket b = 1 \rrbracket \cdot \left(\mathbf{1} + \frac{1}{2} \cdot (\mathbf{3} + \llbracket 2x > 0 \rrbracket \cdot 4x) + \frac{1}{2} \cdot \left(\mathbf{9} - \frac{5}{2^n} + n \cdot \llbracket 2x > 0 \rrbracket \cdot 4x \right) \right) \\
&= \mathbf{1} + \llbracket b \neq 1 \rrbracket \cdot (\mathbf{1} + \llbracket x > 0 \rrbracket \cdot 2x) \\
&\quad + \llbracket b = 1 \rrbracket \cdot \left(\mathbf{7} - \frac{5}{2^{n+1}} + (n+1) \cdot \llbracket x > 0 \rrbracket \cdot 2x \right) = I_{n+1} \geq I_{n+1}.
\end{aligned}$$

Now we can complete the runtime analysis of program C :

$$\begin{aligned}
\text{ert}[C](\mathbf{0}) &= \text{ert}[C_1](\text{ert}[C_2](\text{ert}[C_3](\mathbf{0}))) \\
&\geq \text{ert}[C_1](\mathbf{1} + \llbracket b \neq 1 \rrbracket \cdot (\mathbf{1} + \llbracket x > 0 \rrbracket \cdot 2x) + \llbracket b = 1 \rrbracket \cdot (\mathbf{7} + \llbracket x > 0 \rrbracket \cdot \infty)) \\
&= \text{ert}[x := 1] \left(\text{ert}[b := 1] \left(\mathbf{1} + \llbracket b \neq 1 \rrbracket \cdot (\mathbf{1} + \llbracket x > 0 \rrbracket \cdot 2x) \right. \right. \\
&\quad \left. \left. + \llbracket b = 1 \rrbracket \cdot (\mathbf{7} + \llbracket x > 0 \rrbracket \cdot \infty) \right) \right) \\
&= \text{ert}[x := 1](\mathbf{8} + \llbracket x > 0 \rrbracket \cdot \infty) = \mathbf{8} + \infty = \infty.
\end{aligned}$$

Overall, we obtain that the expected runtime of program C is infinite even though it terminates with probability one. Notice furthermore that both subprograms C_2 and C_3 have finite expected runtimes since

$$\text{ert}[C_2](\mathbf{0}) = \mathbf{1} + \llbracket b = 1 \rrbracket \cdot \mathbf{4} \quad \text{and} \quad \text{ert}[C_3](\mathbf{0}) = \mathbf{1} + \llbracket x > 0 \rrbracket \cdot 2x.$$

Invariant Synthesis. In order to synthesize the ω -invariant I_n of loop C_2 , we propose the template $I_n = \mathbf{1} + \llbracket b \neq 1 \rrbracket \cdot (\mathbf{1} + \llbracket x > 0 \rrbracket \cdot 2x) + \llbracket b = 1 \rrbracket \cdot (a_n + b_n \cdot \llbracket x > 0 \rrbracket \cdot 2x)$ and from the definition of lower ω -invariants we obtain $a_0 \leq 2$, $a_{n+1} \leq 7/2 + 1/2 \cdot a_n$ and $b_0 \leq 0$, $b_{n+1} \leq 1 + b_n$. These recurrences admit solutions $a_n = 7^{-5/2^n}$ and $b_n = n$. \triangle

Just like the proof rule based on upper invariants, the proof rules based on ω -invariants are complete too: given loop $\text{while } (\xi) \{C\}$ and runtime t , it is enough to consider the ω -invariant $I_n = F_t^{n+1}(0)$, where F_t^n is defined as in the proof of Theorem 4.6 to yield the exact runtime $\text{ert}[\text{while } (\xi) \{C\}](t)$ from an application of Theorem 4.6. We formally capture this result by means of the following theorem:

THEOREM 4.9. *There exists a sequence I_n that is both a lower and an upper ω -invariant of $\text{while } (\xi) \{C\}$ with respect to t , such that*

$$\text{ert}[\text{while } (\xi) \{C\}](t) = \lim_{n \rightarrow \infty} I_n.$$

Theorem 4.9 together with Theorem 4.4 shows that the set of invariant-based proof rules presented in this section are complete. Next we study how to refine invariants to make the bounds that these proof rules yield more precise.

4.3 Refinement of Bounds

An important property of both upper and lower bounds of the runtime of loops is that they can be easily refined by repeated application of the characteristic functional. This works as follows. If u is an upper bound of $\text{ert}[\text{while } (\xi) \{C\}](t)$ and $F_t^{\xi, C}(u) \leq u$, then $F_t^{\xi, C}(u)$ is also an upper bound

at least as precise as u . Dually, if l is a lower bound and $F_t^{(\xi, C)}(l) \geq l$, then $F_t^{(\xi, C)}(l)$ is also a lower bound at least as precise as l . Formally, we have the following theorem:³

THEOREM 4.10 (REFINEMENT OF BOUNDS).

- (1) If $\text{ert}[\text{while } (\xi) \{C\}](t) \leq u$ and $F_t^{(\xi, C)}(u) \leq u$, then $\text{ert}[\text{while } (\xi) \{C\}](t) \leq F_t^{(\xi, C)}(u) \leq u$.
- (2) If $l \leq \text{ert}[\text{while } (\xi) \{C\}](t)$ and $l \leq F_t^{(\xi, C)}(l)$, then $l \leq F_t^{(\xi, C)}(l) \leq \text{ert}[\text{while } (\xi) \{C\}](t)$.

PROOF. We prove the first case only, as the proof for lower bounds is analogous. If u is an upper bound of $\text{ert}[\text{while } (\xi) \{C\}](t)$, then $\text{lfp } F_t^{(\xi, C)} \leq u$. By the monotonicity of $F_t^{(\xi, C)}$ (recall that ert is monotonic by Theorem 3.4) and from $F_t^{(\xi, C)}(u) \leq u$ we obtain

$$\text{ert}[\text{while } (\xi) \{C\}](t) = \text{lfp } F_t^{(\xi, C)} = F_t^{(\xi, C)}(\text{lfp } F_t^{(\xi, C)}) \leq F_t^{(\xi, C)}(u) \leq u,$$

which means that $F_t^{(\xi, C)}(u)$ is also an upper bound, possibly tighter than u . \square

Notice that if I is an upper invariant of $\text{while } (\xi) \{C\}$, then I fulfills all necessary conditions of Theorem 4.10. In practice, Theorem 4.10 provides a means of iteratively improving the precision of bounds yielded by Theorems 4.2 and 4.6. For instance, for upper invariant I we have

$$\text{ert}[\text{while } (\xi) \{C\}](t) \leq \dots \leq F_t^{(\xi, C)}(F_t^{(\xi, C)}(I)) \leq F_t^{(\xi, C)}(I) \leq I.$$

If I_n is an upper (lower, respectively) ω -invariant, applying Theorem 4.10 requires checking that $F_t^{(\xi, C)}(I) \leq L$ (and $F_t^{(\xi, C)}(L) \geq L$, respectively), where $L = \lim_{n \rightarrow \infty} I_n$. This proof obligation can be discharged by showing that I_n forms an ω -chain, i.e., that $I_n \leq I_{n+1}$ for all $n \in \mathbb{N}$.

This concludes our presentation of the ert -calculus for pGCL programs. In the next sections, we validate the calculus twofold. In Section 5, we show that the calculus corresponds to an intuitive operational model of pGCL programs based on Markov chains. Section 6 shows in detail that the ert -calculus is a conservative extension of Nielson's approach—basically an adaptation of Hoare triples—for obtaining upper bounds on the runtime of ordinary, i.e., nonprobabilistic, programs.

5 AN OPERATIONAL MODEL FOR EXPECTED RUNTIMES

We prove the soundness of the expected runtime transformer with respect to a simple operational model for our probabilistic programming language. This model is defined in terms of a reward Markov chain. We first briefly recall all essential notions. For a more comprehensive treatment, see [2, Ch. 10].

A *discrete-time Markov chain* (MC, for short) is a tuple $\mathcal{M} = (\mathcal{S}, \mathbf{P}, s_{\text{init}}, \text{rew})$, where \mathcal{S} is a countable, nonempty set of *control states*; $\mathbf{P} : \mathcal{S} \times \mathcal{S} \rightarrow [0, 1]$ is a *transition probability function* such that for all states $s \in \mathcal{S}$,

$$\sum_{s' \in \mathcal{S}} \mathbf{P}(s, s') = 1;$$

$s_{\text{init}} \in \mathcal{S}$ is the *initial state*; and $\text{rew} : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ is a *reward function*. Intuitively, for each pair of states s, s' , the transition probability function specifies the probability $\mathbf{P}(s, s')$ to take a transition from s to s' . We often write $s \xrightarrow{p} s'$ instead of $\mathbf{P}(s, s') = p$. Our operational model captures the runtime of probabilistic programs in terms of rewards. Thus, computing expected rewards of MCs along a path is essential. Formally, a *path* in an MC \mathcal{M} is a finite sequence $\pi = s_1 \dots s_n$ such

³A reader familiar with abstract interpretation will recognize this as applying a widening or narrowing step.

that $\mathbf{P}(s_i, s_{i+1}) > 0$ for each $1 \leq i < n$. The set of all paths in \mathcal{M} (starting in state s) is denoted by $\text{Paths}(\mathcal{M})$ ($\text{Paths}(\mathcal{M}, s)$). Moreover, given a set of target states \mathcal{T} , we define the set of all paths starting in state s that reach a state in \mathcal{T} as $\text{Paths}(\mathcal{M}, s, \mathcal{T}) = \text{Paths}(\mathcal{M}, s) \cap (\mathcal{S} \setminus \mathcal{T})^* \mathcal{T}$. For a path $\pi = s_1 \dots s_n$, the cumulative reward of π and the probability of π are given by

$$\text{rew}(\pi) \triangleq \sum_{k=1}^n \text{rew}(s_k) \quad \text{and} \quad \mathbf{P}(\pi) \triangleq \prod_{k=1}^{n-1} \mathbf{P}(s_k, s_{k+1}).$$

With these notions readily available, we can define the expected reward of an MC \mathcal{M} eventually reaching a set of target states from its initial state.

Definition 5.1 (Expected rewards over MCs). Let $\mathcal{M} = (\mathcal{S}, \mathbf{P}, s_{\text{init}}, \text{rew})$ be a Markov chain and $\mathcal{T} \subseteq \mathcal{S}$ a nonempty set of target states. The expected reward of \mathcal{M} eventually reaching \mathcal{T} from s_{init} is given by

$$\text{ExpRew}^{\mathcal{M}}(\mathcal{T}) \triangleq \sum_{\pi \in \text{Paths}(\mathcal{M}, s_{\text{init}}, \mathcal{T})} \mathbf{P}(\pi) \cdot \text{rew}(\pi)$$

if \mathcal{T} is reached almost surely from s_{init} , i.e., if

$$\sum_{\pi \in \text{Paths}(\mathcal{M}, s_{\text{init}}, \mathcal{T})} \mathbf{P}(\pi) = 1.$$

Otherwise, we set $\text{ExpRew}^{\mathcal{M}}(\mathcal{T}) \triangleq \infty$.

If $\mathcal{T} = \{s\}$ is a singleton, we often write $\text{ExpRew}^{\mathcal{M}}(s)$ instead of $\text{ExpRew}^{\mathcal{M}}(\{s\})$. We now turn to our operational model of probabilistic programs. For simplicity, we assume a canonical labeling for each program statement $C \in \text{pGCL}$. We collect the labels used in a given program C in Lab_* , including a special symbol \downarrow to denote successful program termination. These labels will—together with the set of program states Σ —form the state space of our operational model. Furthermore, we employ the following auxiliary functions between program statements and labels:

- $\text{init} : \text{pGCL} \rightarrow \text{Lab}_*$ provides the label corresponding to the beginning of a program.
- $\text{stmt} : \text{Lab}_* \rightarrow (\text{pGCL} \cup \{\downarrow\})$ yields the statement associated with a program label.
- $\text{first}, \text{second} : \text{Lab}_* \rightarrow \text{Lab}_*$ give the first and second successor of a program label. If no such successor exists, we define $\text{first}(\ell) = \downarrow$ and $\text{second}(\ell) = \downarrow$, respectively.

Example 5.2. Consider the following program C where each statement is annotated with its canonical program label:

$$\begin{aligned} C : & \text{ while } \left([1/2 \cdot \langle \text{true} \rangle + 1/2 \cdot \langle \text{false} \rangle]^1 \right) \{ \\ & \quad [succ := \text{true}]^2 \\ & \quad \}; \\ & \quad [succ := \text{false}]^3. \end{aligned}$$

Thus, $\text{Lab}_* = \{\downarrow, 1, 2, 3\}$. The definition of the auxiliary functions is straightforward:

$$\begin{array}{llll} \text{init}(C) = 1 & \text{stmt}(1) = \text{while } (\dots) \{succ := \text{true}\} & \text{first}(1) = 2 & \text{second}(1) = 3 \\ & \text{stmt}(2) = succ := \text{true} & \text{first}(2) = 1 & \text{second}(2) = \downarrow \\ & \text{stmt}(3) = succ := \text{false} & \text{first}(3) = \downarrow & \text{second}(3) = \downarrow \\ & \text{stmt}(\downarrow) = \downarrow & \text{first}(\downarrow) = \downarrow & \text{second}(\downarrow) = \downarrow. \end{array}$$

Definition 5.3 (MC of a pGCL program). For initial state $\sigma_0 \in \Sigma$ and $t \in \mathbb{T}$, the operational Markov chain of $C \in \text{pGCL}$ is given by $\mathcal{M}_{\sigma_0}^t[C] = (\mathcal{S}, \mathbf{P}, s_{\text{init}}, \text{rew})$, where:

$$\begin{array}{c}
 \frac{\text{stmt}(\ell) = \downarrow}{\langle \ell, \sigma \rangle \xrightarrow{1} \langle \text{sink} \rangle} \text{ [terminated]} \qquad \frac{}{\langle \text{sink} \rangle \xrightarrow{1} \langle \text{sink} \rangle} \text{ [sink]} \\
 \frac{\text{stmt}(\ell) = \text{empty} \quad \text{first}(\ell) = \ell'}{\langle \ell, \sigma \rangle \xrightarrow{1} \langle \ell', \sigma \rangle} \text{ [empty]} \qquad \frac{\text{stmt}(\ell) = \text{skip} \quad \text{first}(\ell) = \ell'}{\langle \ell, \sigma \rangle \xrightarrow{1} \langle \ell', \sigma \rangle} \text{ [skip]} \\
 \frac{\text{stmt}(\ell) = \text{halt}}{\langle \ell, \sigma \rangle \xrightarrow{1} \langle \text{sink} \rangle} \text{ [halt]} \qquad \frac{\text{stmt}(\ell) = x : \approx \mu \quad \llbracket \mu : v \rrbracket(\sigma) = p > 0 \quad \text{first}(\ell) = \ell'}{\langle \ell, \sigma \rangle \xrightarrow{p} \langle \ell', \sigma[x/v] \rangle} \text{ [pr-assgn]} \\
 \frac{\text{stmt}(\ell) = \text{if } (\xi) \{C_1\} \text{ else } \{C_2\} \quad \llbracket \xi : \text{true} \rrbracket(\sigma) = p > 0 \quad \text{first}(\ell) = \ell'}{\langle \ell, \sigma \rangle \xrightarrow{p} \langle \ell', \sigma \rangle} \text{ [if-true]} \\
 \frac{\text{stmt}(\ell) = \text{if } (\xi) \{C_1\} \text{ else } \{C_2\} \quad \llbracket \xi : \text{false} \rrbracket(\sigma) = p > 0 \quad \text{second}(\ell) = \ell'}{\langle \ell, \sigma \rangle \xrightarrow{p} \langle \ell', \sigma \rangle} \text{ [if-false]} \\
 \frac{\text{stmt}(\ell) = \text{while } (\xi) \{C\} \quad \llbracket \xi : \text{true} \rrbracket(\sigma) = p > 0 \quad \text{first}(\ell) = \ell'}{\langle \ell, \sigma \rangle \xrightarrow{p} \langle \ell', \sigma \rangle} \text{ [while-true]} \\
 \frac{\text{stmt}(\ell) = \text{while } (\xi) \{C\} \quad \llbracket \xi : \text{false} \rrbracket(\sigma) = p > 0 \quad \text{second}(\ell) = \ell'}{\ell, \sigma \xrightarrow{p} \ell', \sigma} \text{ [while-false]}
 \end{array}$$

Fig. 1. Rules for the transition probability function of operational MCs.

 Table 2. The Reward Function $\text{rew} : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ of Operational MCs

s	$\text{stmt}(\ell)$	$\text{rew}(s)$
$\langle \ell, \sigma \rangle$	\downarrow	$t(\sigma)$
$\langle \ell, \sigma \rangle$	skip, $x : \approx \mu$, if $(\xi) \{C_1\}$ else $\{C_2\}$, or while $(\xi) \{C\}$	1
$\langle \ell, \sigma \rangle$	empty, halt	0
$\langle \text{sink} \rangle$		0

- $\mathcal{S} = \{\langle \ell, \sigma \rangle \mid \ell \in \text{Lab}_*, \sigma \in \Sigma\} \cup \{\langle \text{sink} \rangle\}$,
- the transition probability function \mathbf{P} is given by the rules in Figure 1,
- $s_{\text{init}} = \langle \text{init}(C), \sigma_0 \rangle$, and
- the reward function $\text{rew} : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ is defined according to Table 2. \triangle

Most of the rules in Figure 1 defining the transition probability function of a program's MC are self-explanatory. As an example, consider the following rule for while loops:

$$\frac{\text{stmt}(\ell) = \text{while } (\xi) \{C\} \quad \llbracket \xi : \text{true} \rrbracket(\sigma) = p > 0 \quad \text{first}(\ell) = \ell'}{\langle \ell, \sigma \rangle \xrightarrow{p} \langle \ell', \sigma \rangle} \text{ [while-true]}.$$

To apply this rule, we first have to determine the probability p of the loop's guard being true. If $p > 0$, then the MC contains a transition to move with probability p from its current state to the first statement contained in the loop's body, i.e., $\text{first}(\ell)$. Analogously, by the rule [while-false], there is a transition to leave the loop, i.e., move to $\text{second}(\ell)$, with probability $1-p$, the likelihood of the guard being false.




Fig. 2. The operational MC $\text{ExpRew}_{\sigma}^t \llbracket C_{\text{trunc}} \rrbracket (\langle \text{sink} \rangle)$ corresponding to the truncated geometric distribution. For each state, the corresponding reward is provided in gray.

Let us explain the difference between state $\langle \text{sink} \rangle$ and states of the form $\langle \downarrow, \sigma \rangle$. The state $\langle \text{sink} \rangle$ is reached after either successful program termination or premature halting; states of the form $\langle \downarrow, \sigma \rangle$ are reached only upon successful program termination. After halting a program run, the MC immediately evolves to the sink state (see rule [halt] in Figure 1). Note that the continuation $t \in \mathbb{T}$ of a program contributes to the cumulative reward to reach the sink state only upon states of the form $\langle \downarrow, \sigma \rangle$ (see first row of Table 2).

Example 5.4 (MC for truncated geometric distribution). Recall the probabilistic program C_{trunc} from Example 3.3 together with its canonical labeling:

$$\begin{aligned}
 C_{\text{trunc}} : & \text{ if } \left([1/2 \cdot \langle \text{true} \rangle + 1/2 \cdot \langle \text{false} \rangle]^1 \right) \{ [succ := \text{true}]^2 \} \text{ else } \{ \\
 & \text{ if } \left([1/2 \cdot \langle \text{true} \rangle + 1/2 \cdot \langle \text{false} \rangle \right]^3 \{ [succ := \text{true}]^4 \} \\
 & \text{ else } \{ [succ := \text{false}]^5 \} \\
 & \}.
 \end{aligned}$$

Figure 2 depicts the MC $\mathcal{M}_{\sigma}^t \llbracket C_{\text{trunc}} \rrbracket$ for an initial program state $\sigma \in \Sigma$ and an arbitrary continuation $t \in \mathbb{T}$. Here labeled edges denote the value of the transition probability function for the respective states, while the reward of each state is provided in gray next to the state. To improve readability, edge labels are omitted if the probability of a transition equals one.

A brief inspection of Figure 2 reveals that $\mathcal{M}_{\sigma}^t \llbracket C_{\text{trunc}} \rrbracket$ contains three finite paths reaching $\langle \text{sink} \rangle$ from initial state $\langle 1, \sigma \rangle$:

$$\begin{aligned}
 \pi_1 &= \langle 1, \sigma \rangle \langle 2, \sigma \rangle \langle \downarrow, \sigma[succ/\text{true}] \rangle \langle \text{sink} \rangle, \\
 \pi_2 &= \langle 1, \sigma \rangle \langle 3, \sigma \rangle \langle 4, \sigma \rangle \langle \downarrow, \sigma[succ/\text{true}] \rangle \langle \text{sink} \rangle, \text{ and} \\
 \pi_3 &= \langle 1, \sigma \rangle \langle 3, \sigma \rangle \langle 5, \sigma \rangle \langle \downarrow, \sigma[succ/\text{false}] \rangle \langle \text{sink} \rangle.
 \end{aligned}$$

These paths correspond to the results of the two probabilistic guards in C . Hence, the expected reward of $\mathcal{M}_{\sigma}^t \llbracket C_{\text{trunc}} \rrbracket$ eventually reaching $\langle \text{sink} \rangle$ is given by

$$\begin{aligned}
\text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket C_{\text{trunc}} \rrbracket} (\langle \langle \text{sink} \rangle \rangle) &= \sum_{\pi \in \text{Paths}(\mathcal{M}_\sigma^t \llbracket C_{\text{trunc}} \rrbracket, \langle 1, \sigma \rangle, \langle \text{sink} \rangle)} \mathbf{P}(\pi) \cdot \text{rew}(\pi) \\
&= \mathbf{P}(\pi_1) \cdot \text{rew}(\pi_1) + \mathbf{P}(\pi_2) \cdot \text{rew}(\pi_2) + \mathbf{P}(\pi_3) \cdot \text{rew}(\pi_3) \\
&= \left(\frac{1}{2} \cdot 1 \cdot 1\right) \cdot (1 + 1 + t(\sigma[\text{succ}/\text{true}])) \\
&\quad + \left(\frac{1}{2} \cdot \frac{1}{2} \cdot 1 \cdot 1\right) \cdot (1 + 1 + 1 + t(\sigma[\text{succ}/\text{true}])) \\
&\quad + \left(\frac{1}{2} \cdot \frac{1}{2} \cdot 1 \cdot 1\right) \cdot (1 + 1 + 1 + t(\sigma[\text{succ}/\text{false}])) \\
&= \frac{5}{2} + \frac{3}{4} \cdot t(\sigma[\text{succ}/\text{true}]) + \frac{1}{4} \cdot t(\sigma[\text{succ}/\text{false}]),
\end{aligned}$$

as the probability of reaching $\langle \text{sink} \rangle$ (through these three paths) sums up to one.

Observe that for $t = \mathbf{0}$, the expected reward $\text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket C_{\text{trunc}} \rrbracket} (\langle \langle \text{sink} \rangle \rangle)$ and the expected runtime $\text{ert}[C](t)(\sigma)$ (cf. Example 3.3) coincide, both yielding $5/2$. \triangle

As the previous example suggests, the expected reward $\text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket C \rrbracket} (\langle \langle \text{sink} \rangle \rangle)$ to reach the sink state in the MC $\mathcal{M}_\sigma^t \llbracket C \rrbracket$ of program C coincides with the expected runtime $\text{ert}[C](t)(\sigma)$ as given by the ert-transformer. Formally:

THEOREM 5.5 (SOUNDNESS OF THE ERT TRANSFORMER). *Let $C \in \text{pGCL}$. Then for each initial program state $\sigma \in \Sigma$ and continuation $t \in \mathbb{T}$,*

$$\text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket C \rrbracket} (\langle \langle \text{sink} \rangle \rangle) = \text{ert}[C](t)(\sigma).$$

PROOF. By induction on the structure of pGCL programs. See Appendix B.3. \square

Hence, the expected runtime transformer ert is sound with respect to our operational program model based on Markov chains.

6 RUNTIME OF DETERMINISTIC PROGRAMS

Nielson [41, 42] has extended Hoare logic [21] so as to obtain a formal verification framework to reason about upper bounds on the runtime of deterministic programs, i.e., programs containing neither probabilistic guards nor probabilistic assignments. Whereas Hoare logic originally focuses on partial correctness, Nielson considered the variant for total correctness (see, e.g., [34]). We show in this section that applying ert to deterministic programs results in the tightest upper bound on the runtime obtained in Nielson's approach.

The language GCL of deterministic programs considered in [42] is given by the grammar

$$C ::= \text{skip} \mid x := E \mid C; C \mid \text{if } (B) \{C\} \text{ else } \{C\} \mid \text{while } (B) \{C\}.$$

Here E is a *deterministic* expression and B is a *deterministic* guard; i.e., $\llbracket E \rrbracket(\sigma)$ and $\llbracket B \rrbracket(\sigma)$ are Dirac distributions for each $\sigma \in \Sigma$. For simplicity, we slightly abuse notation and write $\llbracket E \rrbracket(\sigma)$ to denote the unique value $v \in \text{Val}$ such that $\llbracket E : v \rrbracket(\sigma) = 1$. Analogously, $\llbracket B \rrbracket(\sigma)$ denotes the unique value $b \in \{\text{true}, \text{false}\}$ such that $\llbracket B : b \rrbracket(\sigma) = 1$.

Nielson's calculus [41, 42] aims at verifying total program correctness while in addition establishing upper bounds on the program runtime. A correctness property in this extended calculus is of the form

$$\{P\} C \{E \Downarrow Q\},$$

where $C \in \text{GCL}$, E is a deterministic expression over the program variables, and P, Q are assertions expressed in first-order logic. The symbol \Downarrow is just a separator between the postcondition Q and the (bound on the) runtime E . Intuitively, the triple $\{P\} C \{E \Downarrow Q\}$ is valid, written $\models_E \{P\} C \{E \Downarrow Q\}$, if and only if there exists a natural number k such that from each initial state

$$\begin{array}{c}
\frac{}{\{P\} \text{ skip } \{1 \Downarrow P\}} \text{[skip]} \quad \frac{}{\{Q[x/\llbracket E \rrbracket]\} x := E \{1 \Downarrow Q\}} \text{[assgn]} \\
\\
\frac{\{P \wedge E'_2 = u\} C_1 \{E_1 \Downarrow Q \wedge E_2 \leq u\} \quad \{Q\} C_2 \{E_2 \Downarrow R\}}{\{P\} C_1; C_2 \{E_1 + E'_2 \Downarrow R\}} \text{[seq]} \\
\text{where } u \text{ is a fresh logical variable} \\
\\
\frac{\{P \wedge B\} C_1 \{E \Downarrow Q\} \quad \{P \wedge \neg B\} C_2 \{E \Downarrow Q\}}{\{P\} \text{ if } (B) \{C_1\} \text{ else } \{C_2\} \{E \Downarrow Q\}} \text{[if]} \\
\\
\frac{\{P(z+1) \wedge E' = u\} C \{E_1 \Downarrow P(z) \wedge E \leq u\}}{\{\exists z. P(z)\} \text{ while } (B) \{C\} \{E \Downarrow P(0)\}} \text{[while]} \\
\text{where } z \in \mathbb{N}, P(z+1) \Rightarrow B \wedge E \geq E_1 + E', P(0) \Rightarrow \neg B \wedge E \geq 1 \\
\text{and } u \text{ is a fresh logical variable} \\
\\
\frac{\{P'\} C \{E' \Downarrow Q'\}}{\{P\} C \{E \Downarrow Q\}} \text{[cons]} \\
\text{where } P \Rightarrow P' \wedge E' \leq k \cdot E \text{ for some } k \in \mathbb{N} \text{ and } Q' \Rightarrow Q
\end{array}$$

Fig. 3. Nielson's [41] inference system for order of magnitude of runtime of deterministic programs.

σ satisfying precondition P , program C terminates after at most $k \cdot \llbracket E \rrbracket(\sigma)$ steps in a state satisfying postcondition Q . Note that E is evaluated in the *initial* state σ .

The inference rules in Figure 3 are taken verbatim from [42] except for minor changes to match our notation. Let us briefly explain the inference rules. Most of the inference rules are self-explanatory extensions of the standard extension of Hoare calculus for total correctness of deterministic programs [34], which is obtained by omitting the gray parts. The runtime of `skip` and $x := E$ is one time unit. Since guard evaluations are assumed to consume no time in this calculus, any upper bound on the runtime of both branches of a conditional is also an upper bound on the runtime of the conditional itself; cf. rule [if]. The rule of consequence allows one to increase an already proven upper bound on the runtime by an arbitrary constant factor. The runtime of the sequential composition of C_1 and C_2 is, intuitively, the sum of their runtimes E_1 and E_2 . However, runtimes are expressions that are evaluated in the initial state. Thus, the runtime of C_2 has to be expressed in the initial state of $C_1; C_2$. Technically, this is achieved by adding a fresh (and hence universally quantified) variable u that is an upper bound on E_2 and at the same time equals a new expression E'_2 in the precondition of $C_1; C_2$. The runtime of $C_1; C_2$ is then given by $E_1 + E'_2$. The same principle is applied to each loop iteration. Here, the runtime of the loop body is given by E_1 and the runtime E' of the remaining z loop iterations is expressed in the initial state by using a fresh variable u . Any upper bound of $E \geq E_1 + E'$ bounds the runtime of z loop iterations from above.

For deterministic program $C \in \text{GCL}$, let $\vdash_E \{P\} C \{E \Downarrow Q\}$ denote that the correctness property $\{P\} C \{E \Downarrow Q\}$ is provable in Nielson's calculus. Analogously, provability of a total correctness property $\{P\} C \{\Downarrow Q\}$ in the standard Hoare calculus is denoted by $\vdash \{P\} C \{\Downarrow Q\}$. Our first result concerning Nielson's calculus asserts that a correctness proof of C in standard Hoare logic and the ert of C can be combined into a proof in Nielson's proof system.

THEOREM 6.1 (SOUNDNESS OF ert W.R.T. NIELSON'S CALCULUS). *For all deterministic programs $C \in \text{GCL}$ and assertions P, Q , we have*

$$\vdash \{P\} C \{\Downarrow Q\} \text{ implies } \vdash_E \{P\} C \{\text{ert}[C](\mathbf{0}) \Downarrow Q\}.$$

PROOF. By structural induction on C . See Appendix B.4 for a detailed proof. \square

Hence, our notion of ert is sound with respect to Nielson’s proof system. The next theorem asserts that no tighter bound can be derived in Nielson’s calculus. We cannot get a more precise relationship due to different runtime models: while guard evaluations are assumed to consume no time in Nielson’s logic, we assume each guard evaluation to consume one unit of time.

THEOREM 6.2 (COMPLETENESS OF ert W.R.T. NIELSON’S CALCULUS). *For all deterministic programs $C \in \text{GCL}$, assertions P, Q , and deterministic expressions E :*

$$\vdash_E \{P\} C \{E \Downarrow Q\} \text{ implies } \text{ert}[C](\mathbf{0})(\sigma) \leq k \cdot \llbracket E \rrbracket(\sigma),$$

for some $k \in \mathbb{N}$ and all program states $\sigma \in \Sigma$ satisfying P .

PROOF. By induction on C ’s structure; see Appendix B.5 for a detailed proof. \square

Theorem 6.1 together with Theorem 6.2 shows that ert is a conservative extension of Nielson’s approach for reasoning about the runtime of deterministic programs. In particular, given a correctness proof of a deterministic program C in Hoare logic, it suffices to compute $\text{ert}[C](\mathbf{0})$ in order to obtain a corresponding proof in Nielson’s proof system.

7 RECURSION

We now extend the results of the previous sections to be able to reason about the expected runtime of *recursive* randomized algorithms. For achieving that, we extend the transformer ert to allow for recursive procedures. We also prove that this extension preserves all properties studied earlier (such as continuity, etc.) and present a set of proof rules to effectively reason about the runtime of recursive procedure calls. As for an operational semantics, we show that probabilistic programs with recursive procedures can be interpreted as *push-down* Markov chains and we show that the result from Theorem 5.5 relating the wp-based and the operational approach remains valid for recursive programs.

7.1 A Recursive Probabilistic Language

As a first step, we extend pGCL with recursive programs by incorporating procedure calls. For simplicity, we assume the presence of only a single procedure, say, P . We defer the treatment of multiple (possibly mutually recursive) procedures to Section 7.5. The syntax of our *probabilistic Recursive Guarded Command Language* (pRGCL) thus extends the syntax of pGCL (see Section 2) by an atomic statement for procedure calls:

$$\text{call } P.$$

We assume that the procedure P manipulates the global program state and we thus dispense with local variables, parameters, and return statements for passing information across procedure calls. The *declaration* of P consists of a *procedure body* (much like a loop body), which is written in pRGCL syntax. In particular, the body of a procedure can itself contain procedure calls and thus procedures can recursively invoke themselves. We denote by

$$P \triangleright C$$

that procedure P has body $C \in \text{pRGCL}$. A pRGCL *program* is a pair $\langle C, \mathcal{D} \rangle$, where $C \in \text{pRGCL}$ is the “main” command and $\mathcal{D} : \{P\} \rightarrow \text{pRGCL}$ is the declaration of P .⁴

⁴The declaration of P is a mapping from a singleton and not the mere body of P because this minimizes the changes to accommodate the subsequent treatment of multiple procedures.

Table 3. Inductive Definition for the Runtime Transformer $\text{ert}[\cdot, \mathcal{D}]$ for Recursive Programs

C	$\text{ert}[C, \mathcal{D}](t)$
empty	t
skip	$\mathbf{1} + t$
halt	$\mathbf{0}$
$x := \mu$	$\mathbf{1} + \lambda\sigma. \mathbf{E}_{\llbracket \mu \rrbracket(\sigma)}(\lambda v. t[x/v](\sigma))$
$C_1; C_2$	$\text{ert}[C_1, \mathcal{D}](\text{ert}[C_2, \mathcal{D}](t))$
if $(\xi) \{C_1\}$ else $\{C_2\}$	$\mathbf{1} + \llbracket \xi : \text{true} \rrbracket \cdot \text{ert}[C_1, \mathcal{D}](t) + \llbracket \xi : \text{false} \rrbracket \cdot \text{ert}[C_2, \mathcal{D}](t)$
while $(\xi) \{C'\}$	$\text{lfp } X. \mathbf{1} + \llbracket \xi : \text{false} \rrbracket \cdot t + \llbracket \xi : \text{true} \rrbracket \cdot \text{ert}[C', \mathcal{D}](X)$
call P	$(\text{lfp } \eta. \mathbf{1} \oplus \text{ert}[\mathcal{D}(P)]_{\eta}^{\#})(t)$

Example 7.1 (Faulty recursive factorial). Consider the following recursive procedure for (erroneously) computing the factorial of a natural number stored in x :

$$\begin{aligned}
P_{\text{fact}} \triangleright & \text{if } (x \leq 0) \{y := 1\} \text{ else } \{ \\
& \quad \text{if } (5/6 \cdot \langle \text{true} \rangle + 1/6 \cdot \langle \text{false} \rangle) \{ \\
& \quad \quad x := x - 1; \text{ call } P_{\text{fact}}; x := x + 1 \\
& \quad \} \text{ else } \{ \\
& \quad \quad x := x - 2; \text{ call } P_{\text{fact}}; x := x + 2 \\
& \quad \}; \\
& \quad y := y \cdot x \\
& \}.
\end{aligned}$$

In each recursive call, variable x is decreased by either one or two, with probability $5/6$ and $1/6$, respectively. Therefore, some factors might be missing in the computation of the factorial of x . \triangle

Concerning pRGCL's runtime model, we assume that invoking a procedure (i.e., the procedure call itself) consumes one unit of time. The overall runtime of a procedure call is then one plus the runtime of the procedure's body.

7.2 The ert Transformer for Recursive Randomized Algorithms

Since declarations are part of recursive programs, the ert transformer on pRGCL now has shape

$$\text{ert}[C, \mathcal{D}] : \mathbb{T} \rightarrow \mathbb{T}.$$

As a consequence, the rules in Table 1 giving the inductive definition of the transformer must be slightly adapted so that it also propagates declarations. The adaptation for all constructs (except procedure calls) is quite straightforward; see Table 3.

It remains to define the runtime of procedure calls. In the same way as for loops, the action of the transformer on procedure calls is defined using fixed-point techniques. Procedure calls require, however, higher-order fixed points and the use of a subsidiary transformer:

$$\text{ert}[C]_{\eta}^{\#} : \mathbb{T} \rightarrow \mathbb{T}.$$

This transformer behaves exactly as the ert transformer (see Table 1) for all programs except procedure calls. For those, it reverts to a provided *runtime environment* η in

$$\text{RtEnv} \triangleq \{ \eta \mid \eta : \mathbb{T} \rightarrow \mathbb{T} \text{ is continuous} \},$$

instead of the procedure declaration. We can think of such an environment η as a lookup table, in which the transformer $\text{ert}[C]_{\eta}^{\#}$ can look up how to handle the effects of procedure calls. With this in mind, we define

$$\text{ert}[\text{call } P]_{\eta}^{\#}(t) \triangleq \eta(t).$$

Using the subsidiary transformer $\text{ert}[\cdot]_{\eta}^{\#}$, we can (implicitly) define the action of $\text{ert}[\cdot, \mathcal{D}]$ on procedure calls using the equation

$$\text{ert}[\text{call } P, \mathcal{D}] = \underline{\mathbf{1}} \oplus \text{ert}[\mathcal{D}(P)]_{\text{ert}[\text{call } P, \mathcal{D}]}^{\#}.$$

Here, $\underline{\mathbf{1}} \triangleq \lambda t. \mathbf{1}$ represents the constantly $\mathbf{1}$ runtime transformer and “ \oplus ” the pointwise sum between runtime transformers; i.e., for $\gamma_1, \gamma_2 : \mathbb{T} \rightarrow \mathbb{T}$, we let $(\gamma_1 \oplus \gamma_2)(t) \triangleq \gamma_1(t) + \gamma_2(t)$. This equation immediately leads to the (fixed point) characterization

$$\text{ert}[\text{call } P, \mathcal{D}] \triangleq \text{lfp } \eta. \underline{\mathbf{1}} \oplus \text{ert}[\mathcal{D}(P)]_{\eta}^{\#} \quad (1)$$

as in Table 3. To guarantee the existence of the above fixed point, we follow the same strategy as for loops: we endow the set of runtime environments with the structure of an ω -cpo ($\text{RtEnv}, \sqsubseteq$) and we prove that the environment transformer $\lambda \eta. \underline{\mathbf{1}} \oplus \text{ert}[\mathcal{D}(P)]_{\eta}^{\#}$ is continuous. Kleene’s Fixed Point Theorem ensures that the transformer’s fixed point is well defined; for details, see Appendix B.6.

As a next step, we show that all the transformer properties for programs with loops in Section 3 remain valid for recursive programs.

THEOREM 7.2 (BASIC PROPERTIES OF ERT FOR RECURSIVE PROGRAMS). *For any recursive program $\langle C, \mathcal{D} \rangle \in \text{pRGCL}$, any constant runtime k with $k \in \mathbb{R}_{\geq 0}$, any runtimes $t, t' \in \mathbb{T}$, and any ω -chain $t_0 \leq t_1 \leq \dots$ of runtimes in \mathbb{T} :*

Continuity:	$\sup_n \text{ert}[C, \mathcal{D}](t_n) = \text{ert}[C, \mathcal{D}](\sup_n t_n)$;
Monotonicity:	$t \leq t' \Rightarrow \text{ert}[C, \mathcal{D}](t) \leq \text{ert}[C, \mathcal{D}](t')$;
Constant propagation:	$\text{ert}[C, \mathcal{D}](\mathbf{k} + t) = \mathbf{k} + \text{ert}[C, \mathcal{D}](t)$, <i>provided C is halt-free</i> ;
Preservation of ∞ :	$\text{ert}[C, \mathcal{D}](\infty) = \infty$, <i>provided C is halt-free</i> .

PROOF. For the proof of continuity see Appendix B.7. Monotonicity follows from continuity. Preservation of constants is proved in Appendix B.8. Preservation of infinity follows by the same argument as for nonrecursive programs (see the proof of Theorem 3.4). \square

7.3 Proof Rules for Recursion

As for while loops, the runtime of recursive procedures is defined using fixed points. However, reasoning about recursive procedures is typically more involved than reasoning about loops since recursive procedures involve *higher-order* fixed points [19]—the runtime of a loop requires the fixed point of a runtime transformer, while the runtime of a recursive procedure requires the fixed point of an environment transformer. To alleviate this, we propose a set of proof rules for approximating the runtime of (recursive) procedures avoiding the use of any fixed points.

Loosely speaking, the proof rules say that in order to prove that a procedure call satisfies a given runtime specification, it suffices to show that the procedure’s body satisfies the specification, assuming that the recursive calls in the body do so too. To formally state the rules, we require the

notion of *constructive derivability*. Given logical formulae A and B , we write $A \Vdash B$ to denote that B can be derived assuming A . In particular, we will consider claims of the form

$$\text{ert}[\text{call } P](t_1) \bowtie g_1 \Vdash \text{ert}[C](t_2) \bowtie g_2,$$

where $\bowtie \in \{\leq, \geq\}$; t_1, g_1 give the runtime specification of $\text{call } P$; and t_2, g_2 give the runtime specification of C . Notice that in such a claim, we omit any procedure declaration \mathcal{D} as the derivation is independent of P 's body. That is, whenever we encounter a procedure call $\text{call } P$ in C , we may replace it by an upper or lower bound g_1 (depending on \bowtie) without taking procedure declarations into account. Formally, the statement $\text{ert}[\text{call } P](t_1) \bowtie g_1 \Vdash \text{ert}[C](t_2) \bowtie g_2$ can thus also be understood as

$$\forall \eta \text{ with } \eta(t_1) \bowtie g_1 : \text{ert}[C]_{\eta}^{\sharp}(t_2) \bowtie g_2,$$

where $\eta(t_1)$ plays the role of $\text{ert}[\text{call } P](t_1)$.

THEOREM 7.3 (RUNTIME BOUNDS FOR PROCEDURE CALLS). *Let \mathcal{D} be the declaration of procedure P , $u \in \mathbb{T}$, and let $(l_n)_{n \in \mathbb{N}}$, $(u_n)_{n \in \mathbb{N}}$ be two sequences of runtimes in \mathbb{T} .*

(1) *If*

$$\text{ert}[\text{call } P](t) \leq \mathbf{1} + u \Vdash \text{ert}[\mathcal{D}(P)](t) \leq u,$$

then

$$\text{ert}[\text{call } P, \mathcal{D}](t) \leq \mathbf{1} + u.$$

(2) *If $\lim_{n \rightarrow \infty} u_n$ exists, $u_0 = \mathbf{0}$, and for all $n \geq 0$,*

$$\text{ert}[\text{call } P](t) \leq \mathbf{1} + u_n \Vdash \text{ert}[\mathcal{D}(P)](t) \leq u_{n+1},$$

then

$$\text{ert}[\text{call } P, \mathcal{D}](t) \leq \mathbf{1} + \lim_{n \rightarrow \infty} u_n.$$

(3) *If $\lim_{n \rightarrow \infty} l_n$ exists, $l_0 = \mathbf{0}$, and for all $n \geq 0$,*

$$\mathbf{1} + l_n \leq \text{ert}[\text{call } P](t) \Vdash l_{n+1} \leq \text{ert}[\mathcal{D}(P)](t),$$

then

$$\mathbf{1} + \lim_{n \rightarrow \infty} l_n \leq \text{ert}[\text{call } P, \mathcal{D}](t).$$

These proof rules can be seen as a direct counterpart of the proof rules for reasoning about the runtime of loops studied in Section 4: rule (1) is the counterpart of the rule based on loop upper invariants (Theorem 4.2), while rules (2) and (3) are the counterpart of the rules based on ω -invariants (Theorem 4.6). The above proof rules are inspired by the traditional proof rule used for reasoning about functional correctness of ordinary recursive programs; for the traditional weakest precondition transformer wp , this rule says that if

$$\text{wp}[\text{call } P](Q) \Rightarrow R \Vdash \text{wp}[\mathcal{D}(P)](Q) \Rightarrow R,$$

then

$$\text{wp}[\text{call } P, \mathcal{D}](Q) \Rightarrow R,$$

R and Q being the pre- and postcondition of $\text{call } P$, respectively [19]. Compared to our proof rule (1) in Theorem 7.3, the partial order “ \leq ” over runtimes is replaced by the partial order “ \Rightarrow ” over predicates, and the shifting of one unit of time occurs because the runtime of a procedure call is one plus the runtime of its body, whereas the semantics of a procedure call fully agrees with the semantics of its body. A detailed soundness proof for our rules is provided in Appendix B.9.

Example 7.4 (Proving exact expected runtimes). Consider the pRGCL procedure P_{geo} that is given by declaration \mathcal{D} :

$$P_{\text{geo}} \triangleright \text{if } (1/2 \cdot \langle \text{true} \rangle + 1/2 \cdot \langle \text{false} \rangle) \{ \text{call } P_{\text{geo}} \} \text{ else } \{ \text{skip} \}.$$

We prove that the exact expected runtime of calling procedure P_{geo} is five units of time, i.e., $\text{ert}[\text{call } P_{\text{geo}}, \mathcal{D}](\mathbf{0}) = 5$, by using simultaneously rules (2) and (3) from Theorem 7.3. To this end, we propose a sequence of runtimes t_n as follows:

$$t_n = \begin{cases} \mathbf{0}, & \text{if } n = 0 \\ \mathbf{4} - \frac{1}{2^{n-2}}, & \text{if } n > 0. \end{cases}$$

Clearly, $\lim_{n \rightarrow \infty} t_n = 4$. To apply Theorem 7.3, it thus suffices to discharge that we have

$$\text{ert}[\text{call } P_{\text{geo}}](\mathbf{0}) = \mathbf{1} + t_n \Vdash \text{ert}[\mathcal{D}(P_{\text{geo}})](\mathbf{0}) = t_{n+1}$$

for all natural numbers n . This amounts to the following calculations:

$$\begin{aligned} & \text{ert}[\mathcal{D}(P_{\text{geo}})](\mathbf{0}) \\ &= \mathbf{1} + 1/2 \cdot \text{ert}[\text{call } P_{\text{geo}}](\mathbf{0}) + 1/2 \cdot \text{ert}[\text{skip}](\mathbf{0}) \\ &= \mathbf{1} + 1/2 \cdot (\mathbf{1} + t_n) + 1/2 \cdot \text{ert}[\text{skip}](\mathbf{0}) && \text{(claim :ert}[\text{call } P_{\text{geo}}](\mathbf{0}) = \mathbf{1} + t_n) \\ &= \mathbf{1} + 1/2 \cdot (\mathbf{1} + t_n) + 1/2 \cdot (\mathbf{1} + \mathbf{0}) && \text{(Table 3)} \\ &= \mathbf{2} + 1/2 \cdot t_n. \end{aligned}$$

For $n = 0$, this means that

$$\text{ert}[\mathcal{D}(P_{\text{geo}})](\mathbf{0}) = \mathbf{2} + 1/2 \cdot t_0 = \mathbf{2} + \mathbf{0} = \mathbf{4} - \frac{1}{2^{1-2}} = t_1.$$

Further, for $n > 0$, we obtain

$$\begin{aligned} & \text{ert}[\mathcal{D}(P_{\text{geo}})](\mathbf{0}) = \mathbf{2} + 1/2 \cdot t_n \\ &= \mathbf{2} + 1/2 \cdot \left(\mathbf{4} - \frac{1}{2^{n-2}} \right) \\ &= \mathbf{4} - \frac{1}{2^{(n+1)-2}} \\ &= t_{n+1}. \end{aligned}$$

Hence, Theorem 7.3 yields $\text{ert}[\text{call } P_{\text{geo}}, \mathcal{D}](\mathbf{0}) = \mathbf{1} + \lim_{n \rightarrow \infty} t_n = 5$. \triangle

7.4 Finite Approximations of Recursive Programs

Equation (1) characterizes the runtime of recursive procedures in terms of fixed points. We next present an alternative characterization, where the runtime of procedures is given as the limit or asymptotic runtime of their finite approximations. The result will be essential in a subsequent section for extending Theorem 5.5 to recursive programs.

The notion of “finite approximation” to a procedure is materialized by its finite inlinings. Formally, the n -th inlining $\text{call}_n^{\mathcal{D}} P$ of a procedure P with respect to declaration \mathcal{D} is defined inductively, by clauses

$$\begin{aligned} \text{call}_0^{\mathcal{D}} P &\triangleq \text{halt} \\ \text{call}_{n+1}^{\mathcal{D}} P &\triangleq \text{skip}; \mathcal{D}(P)[\text{call } P / \text{call}_n^{\mathcal{D}} P]. \end{aligned}$$

Here, $\mathcal{D}(P)[\text{call } P/\text{call}_n^{\mathcal{D}} P]$ denotes the syntactic replacement of every occurrence of $\text{call } P$ in $\mathcal{D}(P)$ with $\text{call}_n^{\mathcal{D}} P$ and the initial skip statement is used to simulate the runtime of the procedure call itself (as both consume one unit of time).⁵ As so defined, the family of (call-free) programs $\text{call}_n^{\mathcal{D}} P$ define a sequence of approximations to $\text{call } P$, where $\text{call}_0^{\mathcal{D}} P$ is the “poorest” approximation, while the larger the n , the more precise the approximation becomes. Observe that, in general, $\text{call}_{n+1}^{\mathcal{D}} P$ mimics the exact behavior—and runtime—of $\text{call } P$ for all executions that finish after at most n recursive calls. The runtime of $\text{call } P$ can then be obtained as the limit of the runtimes of $\text{call}_n^{\mathcal{D}} P$.

THEOREM 7.5 (RECURSIVE PROCEDURES AS LIMIT OF FINITE APPROXIMATIONS). *For any runtime $t \in \mathbb{T}$,*

$$\text{ert}[\text{call } P, \mathcal{D}](t) = \sup_n \text{ert}[\text{call}_n^{\mathcal{D}} P](t).^6$$

PROOF. Recall that $\text{ert}[\text{call } P, \mathcal{D}] = \text{lfp } \eta. F(\eta)$, where $F(\eta) = \underline{1} \oplus \text{ert}[\mathcal{D}(P)]_{\eta}^{\sharp}$. Since F is continuous (see Appendix B.6), we can apply Kleene’s Fixed Point Theorem to express $\text{lfp } \eta. F(\eta)$ as $\sup_n F^n(\perp_{\text{RtEnv}})$, where F^n denotes the composition of F with itself n times and $\perp_{\text{RtEnv}} = \lambda t. \mathbf{0}$. The theorem then follows from

$$\forall n. F^n(\perp_{\text{RtEnv}}) = \text{ert}[\text{call}_n^{\mathcal{D}} P],$$

which is proven by induction on n . The base case is immediate. For the inductive case we rely on a result that follows from the definition of transformer $\text{ert}[C]_{\eta}^{\sharp}$:

$$\text{ert}[C]_{\text{ert}[C']}^{\sharp} = \text{ert}[C[\text{call } P/C']] \quad \forall C \in \text{pRGCL}, C' \in \text{pGCL}.$$

Using this result, we reason as follows:

$$\begin{aligned} F^{n+1}(\perp_{\text{RtEnv}}) &= \underline{1} \oplus \text{ert}[\mathcal{D}(P)]_{F^n(\perp_{\text{RtEnv}})}^{\sharp} && \text{(definition } F^{n+1}) \\ &= \underline{1} \oplus \text{ert}[\mathcal{D}(P)]_{\text{ert}[\text{call}_n^{\mathcal{D}} P]}^{\sharp} && \text{(I.H.)} \\ &= \underline{1} \oplus \text{ert}[\mathcal{D}(P)[\text{call } P/\text{call}_n^{\mathcal{D}} P]] && \text{(auxiliary result)} \\ &= \text{ert}[\text{skip}; \mathcal{D}(P)[\text{call } P/\text{call}_n^{\mathcal{D}} P]] && \text{(Table 1)} \\ &= \text{ert}[\text{call}_{n+1}^{\mathcal{D}} P] && \text{(definition } \text{call}_{n+1}^{\mathcal{D}} P). \quad \square \end{aligned}$$

7.5 Mutual Recursion

For the sake of simplicity, we have so far assumed the presence of a single procedure. Notwithstanding, our ert-calculus can be readily extended to handle multiple procedures. Say we want to handle m (possibly mutually recursive) procedures P_1, \dots, P_m with declaration $\mathcal{D} : \{P_1, \dots, P_m\} \rightarrow \text{pRGCL}$. A runtime environment is now a tuple $\eta = (\eta_1, \dots, \eta_m)$, where η_i is meant to provide the behavior of procedure P_i in $\text{ert}[\cdot]_{\eta}^{\sharp}$, i.e., $\text{ert}[\text{call } P_i]_{\eta}^{\sharp} = \eta_i$. The action of ert on procedure calls is defined simultaneously as

$$\left(\text{ert}[\text{call } P_1, \mathcal{D}], \dots, \text{ert}[\text{call } P_m, \mathcal{D}] \right) \triangleq \text{lfp } \eta. \left(\underline{1} \oplus \text{ert}[\mathcal{D}(P_1)]_{\eta}^{\sharp}, \dots, \underline{1} \oplus \text{ert}[\mathcal{D}(P_m)]_{\eta}^{\sharp} \right).$$

For determining the least fixed point above, environments are compared componentwise, i.e., $(\eta_1, \dots, \eta_m) \sqsubseteq (\nu_1, \dots, \nu_m)$ if and only if $\eta_i \sqsubseteq \nu_i$ for all $i = 1 \dots m$, where “ \sqsubseteq ” is the partial order over RtEnv defined in Appendix B.6. Technically, this corresponds to taking as underlying

⁵The formal definition of the syntactic replacement proceeds by a routine induction on the structure of $\mathcal{D}(P)$.

⁶Strictly speaking, $\text{call}_n^{\mathcal{D}} P$ is a pGCL-program. We therefore prefer to write $\text{ert}[\text{call}_n^{\mathcal{D}} P](t)$ rather than $\text{ert}[\text{call}_n^{\mathcal{D}} P, \mathcal{D}](t)$.

ω -cpo in the fixed point above the product ω -cpo

$$(\text{RtEnv}, \sqsubseteq) \times \overset{m \text{ times}}{\dots} \times (\text{RtEnv}, \sqsubseteq).$$

The proof rules from Theorem 7.3 are also easily adapted to reason about multiple procedures. For instance, rule (1) now says that if for all $i = 1 \dots m$:

$$\text{ert}[\text{call } P_1](t_1) \leq \mathbf{1} + u_1, \dots, \text{ert}[\text{call } P_m](t_m) \leq \mathbf{1} + u_m \Vdash \text{ert}[\mathcal{D}(P_i)](t_i) \leq u_i,$$

then also for all $i = 1 \dots m$,

$$\text{ert}[\text{call } P_i, \mathcal{D}](t_i) \leq \mathbf{1} + u_i.$$

The rule reasons about all the procedures simultaneously. Roughly speaking, the rule premise requires deriving the runtime specification for the body of each procedure P_i , assuming the corresponding specification for all procedure calls in it. The rule conclusion establishes the runtime specification of the set of procedures altogether. The other two rules from Theorem 7.3 admit a similar adaptation.

Theorem 7.5 characterizing (the expected runtime of) recursive procedures as the limit of their n th inlinings naturally extends to multiple procedures. We only need to adapt the definition of the n -inlining $\text{call}_n^{\mathcal{D}} P_i$ of procedure P_i so as to inline the calls of all procedures:

$$\text{call}_{n+1}^{\mathcal{D}} P_i \triangleq \text{skip}; \mathcal{D}(P_i)[\text{call } P_1 / \text{call}_n^{\mathcal{D}} P_1, \dots, \text{call } P_m / \text{call}_n^{\mathcal{D}} P_m].$$

7.6 Operational Model

To incorporate recursion into our operational (Markov chain) model of programs, we need to keep track of the procedures that have been invoked during the current program execution but have not yet terminated. This knowledge is necessary upon the termination of a procedure to determine whether this termination represents the “entire” program termination or the execution is to be continued with the remainder statements of the caller.

This kind of book-keeping of procedure calls is done by extending MCs to *pushdown Markov chains* (PMC) [24] whose verification has been studied by Brázdil, Esparza, and Kucera [6, 7, 33]. Formally, a PMC is a tuple $\mathcal{P} = (\mathcal{S}, \Gamma, \gamma_0, \mathbf{P}, s_{\text{init}}, \text{rew})$, where \mathcal{S} is a countable, nonempty set of *control states*, Γ is a finite *stack alphabet*, $\gamma_0 \in \Gamma$ is a special *bottom-of-stack* symbol,

$$\mathbf{P} : \mathcal{S} \times \Gamma \times \mathcal{S} \rightarrow [0, 1] \times (\Gamma \setminus \{\gamma_0\})^*$$

is a *pushdown transition probability function*, $s_{\text{init}} \in \mathcal{S}$ is the *initial control state*, and $\text{rew} : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ is a *reward function*. We assume that the topmost symbol of a stack $\alpha = \gamma \cdot \alpha' \in \Gamma \cdot \Gamma^*$ corresponds to the leftmost symbol γ in α .

In contrast to standard Markov chains, each transition of a PMC additionally depends on the topmost stack symbol, which is popped from the stack upon each transition. Moreover, a PMC may push zero or more stack symbols—with the exception of bottom-of-stack symbol γ_0 —onto the stack. For convenience, we restrict ourselves to PMCs pushing at most one new symbol onto the stack at once. Thus, we write

- $s \xrightarrow{p, \text{push}(\gamma)} s'$ instead of $\mathbf{P}(s, \alpha, s') = (p, \gamma \cdot \alpha)$,
- $s \xrightarrow{p, \text{pop}(\gamma)} s'$ instead of $\mathbf{P}(s, \gamma, s') = (p, \varepsilon)$,
- $s \xrightarrow{p} s'$ instead of $\mathbf{P}(s, \gamma, s') = (p, \gamma)$, and
- $s \xrightarrow{p, \text{empty}} s'$ instead of $\mathbf{P}(s, \gamma_0, s') = (p, \gamma_0)$




Fig. 4. Illustration of PMC from Example 7.6.

to denote that PMC \mathcal{P} moves with probability p from state s to s' and pushes $\gamma \in \Gamma$ on the stack, pops γ from the stack, ignores the stack, and has an empty stack, respectively. Again, transition probabilities $p = 1$ are usually omitted from figures.

A state-stack pair $(s, \alpha) \in \mathcal{S} \times \Gamma^*$ is called a *configuration*. Then, all notions defined over states of MCs introduced in Section 5, such as paths, cumulative reward, and expected reward, can be lifted to corresponding notions defined over configurations of PMCs. In particular, a *path* in a PMC is a finite sequence $\pi = (s_1, \alpha_1) \dots (s_n, \alpha_n)$ such that for each $1 \leq i < n$ with $\alpha_i = \gamma \cdot \alpha'$, we have $\mathbf{P}(s_i, \gamma, s_{i+1}) = (p, \beta)$ with $p > 0$ and $\alpha_{i+1} = \beta \cdot \alpha'$. The set of all paths in a PMC \mathcal{P} starting in configuration (s, α) (reaching a goal state in $\mathcal{T} \subseteq \mathcal{S}$) is given by $\text{Paths}(\mathcal{P}, s, \alpha)$ ($\text{Paths}(\mathcal{P}, s, \alpha, \mathcal{T})$). Then the expected reward of a PMC \mathcal{P} eventually reaching a nonempty set \mathcal{T} of target configurations from its initial configuration is defined as

$$\text{ExpRew}^{\mathcal{P}}(\mathcal{T}) \triangleq \sum_{\pi \in \text{Paths}(\mathcal{M}, s_{\text{init}}, \gamma_0, \mathcal{T})} \mathbf{P}(\pi) \cdot \text{rew}(\pi)$$

if \mathcal{T} is reached almost surely from initial configuration $(s_{\text{init}}, \gamma_0)$. Otherwise, let $\text{ExpRew}^{\mathcal{P}}(\mathcal{T}) \triangleq \infty$.

Example 7.6. Consider the PMC $(\mathcal{S}, \Gamma, \gamma_0, \mathbf{P}, s_{\text{init}}, \text{rew})$ with states $\mathcal{S} = \{1, 2, 3, 4, 5, \downarrow, \text{sink}\}$, stack alphabet $\Gamma = \{\gamma_0, 4, 5, \downarrow\}$, and initial state $s_{\text{init}} = 1$. The pushdown transition probability function \mathbf{P} is depicted in Figure 4. Moreover, the reward rew is 1 for states 1,2,3,4,5 and 0 for states \downarrow and sink , respectively.

Starting in configuration $(1, \gamma_0)$, this PMC first randomly chooses whether to move to state 2 (with probability $3/4$) or state 3 (with $1/4$). Say the left transition in Figure 4 is chosen; we move to 3. After that 4 is pushed on the stack and we move back to state 1. Thus, the current configuration is $(1, 4 \cdot \gamma_0)$. Now, assume the right transition is chosen; we move to 2 and subsequently to \downarrow . Since the topmost stack symbol is 4, we move to state 4; the current configuration is $(4, \gamma_0)$. Then 5 is pushed onto the stack and we are in state 1 again. If state \downarrow is eventually reached with an empty stack, we end up in the sink state sink . \triangle

Coming back to probabilistic programs, we now show how the expected runtime of pRGCL programs is related to the expected reward of PMCs. As in Section 5, we assume a canonical labeling of programs and employ auxiliary functions `stmt`, `first`, and `second` to denote the program statement, the first successor, and the second successor of a label, respectively. If no such successor exists, these functions yield \downarrow . Furthermore, in addition to the initial label of a pRGCL program, we need the initial label of each procedure. To that end, let

$$\text{init} : \text{pRGCL} \cup \{P_1, \dots, P_m\} \rightarrow \text{Lab}_*,$$

where P_1, \dots, P_m are the available procedures and Lab_* denotes the set of labels used in a program.

Example 7.7. Consider the pRGCL program $\langle C, \mathcal{D} \rangle$, where $\mathcal{D}(P) = C$ and C is the labeled recursive probabilistic program

$$C : \text{if } \left(\left[\frac{3}{4} \cdot \langle \text{true} \rangle + \frac{1}{4} \cdot \langle \text{false} \rangle \right]^1 \right) \{ [\text{skip}]^2 \} \text{ else } \{ [\text{call } P]^3; [\text{call } P]^4; [\text{call } P]^5 \}.$$

Moreover, the auxiliary functions from above are given by

$$\text{init}(C) = \text{init}(P) = 1, \quad \text{first}(1) = 2, \quad \text{second}(1) = 3, \quad \text{first}(3) = 4, \quad \text{first}(4) = 5.$$

In every other case, first and second are mapped to \downarrow . △

Using these notions, the Markov chain interpretation of pGCL programs in Section 5 naturally extends to a *pushdown* Markov chain interpretation of pRGCL programs. Intuitively, each program label may be pushed onto the stack. Whenever a procedure terminates, such a label is popped from the stack and the execution is continued at the popped label. Thus, apart from a new rule for procedure calls, the rules determining the transition probability function are exactly the same as in Figure 1, where the stack is ignored entirely. The only exception is the termination rule [terminated]: Our PMC first tries to pop a label—the return address—from the stack and continues execution at the popped label. If no such label is on the stack, the program has successfully terminated. Consequently, we obtain the following two new rules:

$$\frac{\text{stmt}(\ell) = \downarrow}{\langle \ell, \sigma \rangle \xrightarrow{\text{pop}(\ell')} \langle \ell', \sigma \rangle} [\text{return}] \quad \frac{\text{stmt}(\ell) = \downarrow}{\langle \ell, \sigma \rangle \xrightarrow{\text{empty}} \langle \text{sink} \rangle} [\text{terminated}].$$

In addition to that, a new rule for procedure calls is needed. This rule first pushes the return address of the current procedure, i.e., a procedure's direct successor, onto the stack and then executes the first statement of the called procedure:

$$\frac{\text{stmt}(\ell) = \text{call } P_i \quad \text{first}(\ell) = \ell'}{\langle \ell, \sigma \rangle \xrightarrow{\text{push}(\ell')} \langle \text{init}(P_i), \sigma \rangle} [\text{call}].$$

Since we assume procedure calls to consume one unit of time, the reward of states containing a program label corresponding to a procedure call is set to 1. The rewards of other states are the same as for nonrecursive programs. Formally:

Definition 7.8 (PMC of recursive programs). Given a pRGCL program $\langle C, \mathcal{D} \rangle$, an initial state $\sigma_0 \in \Sigma$, and a continuation $t \in \mathbb{T}$, the PMC of C is given by $\mathcal{P}_\sigma^t \llbracket C, \mathcal{D} \rrbracket = (\mathcal{S}, \Gamma, \gamma_0, \mathbf{P}, s_{\text{init}}, \text{rew})$, where

- $\mathcal{S} = \{ \langle \ell, \sigma \rangle \mid \ell \in \text{Lab}_*, \sigma \in \Sigma \} \cup \{ \langle \text{sink} \rangle \}$;
- the transition probability function \mathbf{P} is given by the rules in Figure 1 (without [terminated]) and the three aforementioned rules [return], [terminated], and [call];
- $s_{\text{init}} = \langle \text{init}(C), \sigma_0 \rangle$; and
- the reward function $\text{rew} : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ is defined according to Table 4. △

Example 7.9. The PMC \mathcal{P} considered in Example 7.6 and depicted in Figure 4 is the operational PMC of the pRGCL program $\langle C, \mathcal{D} \rangle$, where \mathcal{D} is given by $P \triangleright C$ and

$$C : \text{if } \left(\left[\frac{3}{4} \cdot \langle \text{true} \rangle + \frac{1}{4} \cdot \langle \text{false} \rangle \right]^1 \right) \{ [\text{skip}]^2 \} \text{ else } \{ [\text{call } P]^3; [\text{call } P]^4; [\text{call } P]^5 \}. \quad \Delta$$

As for pGCL programs and operational Markov chains (cf. Theorem 5.5), we obtain a correspondence between pRGCL programs and operational pushdown Markov chains.

THEOREM 7.10 (CORRESPONDENCE THEOREM). *Let $\langle C, \mathcal{D} \rangle$ be a pRGCL program and $t \in \mathbb{T}$. Then for each $\sigma \in \Sigma$, we have*

$$\text{ExpRew}^{\mathcal{P}_\sigma^t \llbracket C, \mathcal{D} \rrbracket} (\langle \text{sink} \rangle) = \text{ert}[C, \mathcal{D}](t)(\sigma).$$

Table 4. Definition of the Reward Function $rew : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ of Operational PMCs for Recursive pRGCL Programs

s	$\text{stmt}(\ell)$	$rew(s)$
$\langle \ell, \sigma \rangle$	\downarrow	$t(\sigma)$
$\langle \ell, \sigma \rangle$	skip, $x : \approx \mu$, if $(\xi) \{C_1\}$ else $\{C_2\}$, while $(\xi) \{C\}$, or call P	1
$\langle \ell, \sigma \rangle$	empty, halt	0
$\langle \text{sink} \rangle$		0

PROOF. By induction on the structure of pRGCL programs. See Appendix B.10. \square

Example 7.11. We exploit the correspondence between pRGCL programs and operational PMCs to analyze the expected runtime of the PMC from Example 7.9. More precisely, we show that C terminates, on average, after at most 10 units of time. This follows, in turn, from showing that

$$\text{ert}[\text{call } P, \mathcal{D}](\mathbf{0}) \leq \mathbf{11}$$

as our runtime model assumes that the runtime of a procedure call (P in this case) is one unit of time more than the runtime of its body (C in this case). To prove the above runtime specification, we apply the first proof rule from Theorem 7.3, taking $u = \mathbf{10}$. We have to derive $\text{ert}[C](0) \leq \mathbf{10}$ assuming for the recursive calls that $\text{ert}[\text{call } P](0) \leq \mathbf{11}$.

$$\begin{aligned}
 \text{ert}[C](\mathbf{0}) &= \mathbf{1} + \frac{3}{4} \cdot \mathbf{1} + \frac{1}{4} \cdot \text{ert}[\text{call } P; \text{call } P; \text{call } P](\mathbf{0}) \\
 &= \mathbf{1} + \frac{3}{4} \cdot \mathbf{1} + \frac{1}{4} \cdot \text{ert}[\text{call } P](\text{ert}[\text{call } P](\text{ert}[\text{call } P](\mathbf{0}))) \\
 &\leq \mathbf{1} + \frac{3}{4} \cdot \mathbf{1} + \frac{1}{4} \cdot \text{ert}[\text{call } P](\text{ert}[\text{call } P](\mathbf{11})) && \text{(assumption)} \\
 &= \mathbf{1} + \frac{3}{4} \cdot \mathbf{1} + \frac{1}{4} \cdot (\mathbf{11} + \text{ert}[\text{call } P; \text{call } P](\mathbf{0})) && \text{(Thm. 7.2, const. prop.)} \\
 &\leq \mathbf{1} + \frac{3}{4} \cdot \mathbf{1} + \frac{1}{4} \cdot (\mathbf{11} + \mathbf{11} + \mathbf{11}) && \text{(repeat previous two steps twice)} \\
 &= \mathbf{10}. && \text{(algebra)}
 \end{aligned}$$

By Theorem 7.10, we can then conclude that the expected runtime of C , or equivalently, the expected reward of the PMC in Figure 4 of reaching $\langle \text{sink} \rangle$, is bounded by 10. \triangle

In the previous example, we exploited the correspondence between ert and operational PMCs to manually compute an upper bound on the expected reward of a PMC. For PMCs with finitely many control states, however, the converse direction is usually more desirable. Expected rewards of finite-state PMCs can be computed by solving linear recurrences [6]. The following result is a direct consequence of their approach.

COROLLARY 7.12. *Let $\langle C, \mathcal{D} \rangle$ be a pRGCL program whose PMC has finitely many control states. Then the expected runtime $\text{ert}[C, \mathcal{D}](\mathbf{0})$ is computable in polynomial space.*

8 RELATION TO EXPECTATION TRANSFORMERS

Expectation transformers are the probabilistic counterpart of Dijkstra's predicate transformers approach to program semantics. We now establish a connection between expectation transformers and our runtime transformers and exploit this connection to derive further algebraic properties of our runtime transformer.

8.1 Expectation Transformers Semantics of Programs

In his seminal work, Dijkstra [12] introduced the predicate transformer wp (standing for **w**eakest **p**recondition) to reason about the semantics of a simple imperative language. Intuitively, for a program C and predicate—or postcondition— Q ,

$$\text{wp}[C](Q)$$

gives the weakest predicate—or precondition—that must hold in the initial state of C so that the execution of C terminates in a final state satisfying Q . Kozen [32] extended the transformer to probabilistic computations in terms of a PDDL, and later on McIver and Morgan [36] showed how to handle programs that combine both probabilistic and nondeterministic behavior.

The wp -semantics over probabilistic programs generalizes Dijkstra's original wp -semantics over ordinary programs twofold: First, instead of being predicates over program states, pre- and postconditions are now (nonnegative) real-valued functions over program states. Second, instead of merely evaluating a (Boolean-valued) postcondition in the final state(s) of a program, we now *measure* the expected value of a (real-valued) postcondition with respect to the distribution of final states.⁷ Formally, for a probabilistic program C and postcondition $f : \Sigma \rightarrow \mathbb{R}_{\geq 0}$, we let

$$\text{wp}[C](f) \triangleq \lambda \sigma. \mathbb{E}_{\llbracket C \rrbracket(\sigma)}(f),$$

where $\llbracket C \rrbracket(\sigma)$ denotes the distribution of final states from executing C in initial state σ and $\mathbb{E}_{\llbracket C \rrbracket(\sigma)}(f)$ denotes the expected value of f with respect to the distribution of final states $\llbracket C \rrbracket(\sigma)$. Consider, for instance, the program from Example 3.3 that simulates a truncated geometric distribution:

$$\begin{aligned} C_{\text{trunc}} : & \text{ if } \left(\frac{1}{2} \cdot \langle \text{true} \rangle + \frac{1}{2} \cdot \langle \text{false} \rangle \right) \{ \text{succ} := \text{true} \} \text{ else } \{ \\ & \quad \text{ if } \left(\frac{1}{2} \cdot \langle \text{true} \rangle + \frac{1}{2} \cdot \langle \text{false} \rangle \right) \{ \text{succ} := \text{true} \} \\ & \quad \text{ else } \{ \text{succ} := \text{false} \} \\ & \}. \end{aligned}$$

For this program we have

$$\text{wp}[C_{\text{trunc}}](f) = \lambda \sigma. \frac{3}{4} \cdot f(\sigma[\text{succ}/\text{true}]) + \frac{1}{4} \cdot f(\sigma[\text{succ}/\text{false}]).$$

Observe that, in particular, if $[A]$ denotes the indicator function of a predicate A over program states, $\text{wp}[C]([A])(\sigma)$ gives the probability of (terminating and) establishing A after executing C from state σ . For instance, we can determine the probability that, from an initial state σ , C_{trunc} terminates in a final state where $\text{succ}=\text{true}$ by

$$\text{wp}[C_{\text{trunc}}]([\text{succ}=\text{true}])(\sigma) = \frac{3}{4} \cdot 1 + \frac{1}{4} \cdot 0 = \frac{3}{4}.$$

Formally, the transformer wp operates on unbounded, so-called *expectations* [36] in

$$\mathbb{E} \triangleq \left\{ f \mid f : \Sigma \rightarrow \mathbb{R}_{\geq 0}^{\infty} \right\},$$

and thus has type

$$\text{wp}[\cdot] : \text{pGCL} \rightarrow (\mathbb{E} \rightarrow \mathbb{E}).$$

We include infinity in the range of expectations because even if $f(\sigma) < \infty$ for all $\sigma \in \Sigma$, the expected value $\text{wp}[C](f)(\sigma)$ can be infinite for some initial state σ , and we need a homogeneous treatment of the domain (i.e., postexpectations) and range (i.e., preexpectations) of wp . Observe,

⁷Strictly speaking, we consider *subdistributions* of final states, where the missing mass captures the probability of nontermination.

Table 5. Rules for Defining the Weakest Pre-Expectation Transformer wp

C	$\text{wp}[C](f)$
empty	f
skip	f
halt	$\mathbf{0}$
$x := \mu$	$\lambda\sigma. \mathbb{E}_{\ \mu\ (\sigma)}(\lambda v. f[x/v](\sigma))$
$C_1; C_2$	$\text{wp}[C_1](\text{wp}[C_2](f))$
if $(\xi) \{C_1\}$ else $\{C_2\}$	$\llbracket \xi : \text{true} \rrbracket \cdot \text{wp}[C_1](f) + \llbracket \xi : \text{false} \rrbracket \cdot \text{wp}[C_2](f)$
while $(\xi) \{C'\}$	$\text{lfp } X. \llbracket \xi : \text{false} \rrbracket \cdot f + \llbracket \xi : \text{true} \rrbracket \cdot \text{wp}[C'](X)$

$\mathbb{E}_\eta(h) \triangleq \sum_v \text{Pr}_\eta(v) \cdot h(v)$ represents the expected value of (random variable) h with respect to distribution η ; $f[x/v] \triangleq \lambda\sigma. f(\sigma[x/v])$, where $\sigma[x/v]$ is the state obtained by updating in σ the value of x to v ; finally, $\text{lfp } X. F(X)$ represents the least fixed point of transformer $F : \mathbb{E} \rightarrow \mathbb{E}$ with respect to the pointwise ordering on \mathbb{E} .

moreover, that the set of expectations \mathbb{E} coincides with the set of runtimes \mathbb{T} , but we prefer to distinguish the two sets because they are to represent different objects.

In a similar vein to the (runtime) transformer $\text{ert}[C]$, (expectation) transformer $\text{wp}[C]$ can be defined by induction on the structure of C , following the rules in Table 5. Most of the rules are self-explanatory and akin to those in Table 1 for the $\text{ert}[C]$ transformer.

8.2 Relation between Expectation and Runtime Transformers

We will now establish the relation between expectation and runtime transformers.

To gain some intuition about this relationship, recall the inductive definitions of wp (cf. Table 5) and ert (cf. Table 1). For every atomic program statement C , we observe that $\text{ert}[C](t)$, where $t \in \mathbb{T}$ is some postruntime, can be decomposed into the time consumed by executing C , i.e., $\text{ert}[C](0)$, and the expected value of the postruntime t , i.e., $\text{wp}[C](t)$. Thus, the expected runtime of program C itself is independent of postruntime t , but t might depend on the expected values of program variables manipulated by C . For example, for a probabilistic assignment, we have

$$\text{ert}[x := \mu](t) = \underbrace{\mathbf{1}}_{=\text{ert}[x := \mu](0)} + \underbrace{\lambda\sigma. \mathbb{E}_{\|\mu\|(\sigma)}(\lambda v. t[x/v](\sigma))}_{=\text{wp}[x := \mu](t)}.$$

This decomposition property for expectation and runtime transformers also holds for composed programs including loops. The precise relationship between expectation and runtime transformers is given by the following theorem.

THEOREM 8.1 (CONNECTION BETWEEN ert AND wp). *For every pGCL program C and runtime $t \in \mathbb{T}$,*

$$\text{ert}[C](t) = \text{ert}[C](\mathbf{0}) + \text{wp}[C](t).$$

More generally, for every two runtimes $t, t' \in \mathbb{T}$,

$$\text{ert}[C](t + t') = \text{ert}[C](t) + \text{wp}[C](t').$$

PROOF. By induction on the program structure. See Appendix B.11 for details. \square

Combining Theorem 8.1 with the linearity of transformer wp (see, e.g., [36]), we can easily establish the subadditivity and subscaling of ert briefly mentioned in Section 3.

COROLLARY 8.2. For all pGCL program C , runtimes $t, t' \in \mathbb{T}$, and constant $r \in \mathbb{R}_{\geq 0}$,

$$\begin{aligned} \text{Subadditivity:} \quad & \text{ert}[C](t + t') \leq \text{ert}[C](t) + \text{ert}[C](t'); \\ \text{Subscaling:} \quad & \text{ert}[C](r \cdot t) \geq \min\{1, r\} \cdot \text{ert}[C](t); \\ & \text{ert}[C](r \cdot t) \leq \max\{1, r\} \cdot \text{ert}[C](t). \end{aligned}$$

PROOF. As for the subadditivity, we have

$$\begin{aligned} \text{ert}[C](t + t') &= \text{ert}[C](t) + \text{wp}[C](t') && \text{(Theorem 8.1)} \\ &= \text{ert}[C](t) + \text{ert}[C](t') - \text{ert}[C](\mathbf{0}) && \text{(Theorem 8.1)} \\ &\leq \text{ert}[C](t) + \text{ert}[C](t') && (\text{ert}[C](\mathbf{0}) \geq \mathbf{0}). \end{aligned}$$

The subscaling follows from the reasoning below (we establish only one inequality; the proof argument for the other one is analogous):

$$\begin{aligned} \text{ert}[C](r \cdot t) &= 1 \cdot \text{ert}[C](\mathbf{0}) + \text{wp}[C](r \cdot t) && \text{(Theorem 8.1)} \\ &= \underbrace{1}_{\geq \min\{1, r\}} \cdot \text{ert}[C](\mathbf{0}) + \underbrace{r}_{\geq \min\{1, r\}} \cdot \text{wp}[C](t) && (\text{wp}[C] \text{ linear}) \\ &\geq \min\{1, r\} \cdot (\text{ert}[C](\mathbf{0}) + \text{wp}[C](t)) && \text{(algebra)} \\ &= \min\{1, r\} \cdot \text{ert}[C](t) && \text{(Theorem 8.1)} \quad \square \end{aligned}$$

For the sake of clarity, we have presented the results considering only pGCL programs and left out recursion. All the results are extensible to recursive programs in pRGCL as shown in [44] and its full version [43].

9 CASE STUDIES

In this section, we use the ert-calculus to analyze the runtime of three well-known randomized algorithms: we use global and incremental invariants for nonrecursive probabilistic programs to analyze the *coupon collector's problem* and a fair *one-dimensional (symmetric) random walk*, respectively. Furthermore, our proof rules for reasoning about recursive probabilistic programs are employed to reason about a *randomized binary search* algorithm.

9.1 The Coupon Collector's Problem

To illustrate the use of our proof rule for loops based on global invariants, we apply the ert-calculus to the coupon collector's problem. This problem arises from the following scenario:⁸ suppose each box of cereal contains one of N different types of coupons, and once a consumer has collected a coupon of each type, he or she can trade them for a prize. The problem is to determine the average number of cereal boxes a coupon collector has to buy in order to collect at least one coupon of each type. It is assumed that each coupon type occurs with the same probability in the cereal boxes.

We can model the coupon collector's problem by a probabilistic program C_{cp} as follows:

$$\begin{aligned} C_{cp} &: \quad cp := [0, \dots, 0]; i := 1; x := N; \\ C_{out} &: \quad \text{while } (x > 0) \{ \end{aligned}$$

⁸The problem formulation presented here is taken from [37].

$$\begin{aligned}
C_{in} : \quad & \text{while } (cp[i] \neq 0) \{ \\
& \quad i \approx \text{Unif}[1 \dots N] \\
& \quad ; \\
& \quad cp[i] := 1; x := x - 1 \\
& \quad \}.
\end{aligned}$$

Here, we use an array cp as a shortcut for N distinct program variables. The array cp (of size N) is initialized with 0s, and whenever we obtain the first coupon of type i , we set $cp[i]$ to 1. The outer loop C_{out} is iterated N times, and in each iteration of the outer loop we collect a new—unseen—coupon type. The collection of the new coupon type is performed by the inner loop C_{in} .

In order to analyze the runtime of program C_{cp} , we need to find a suitable invariant for the outer loop C_{out} . To this end, we propose the following global upper invariant I :

$$\begin{aligned}
I \triangleq & 1 + \sum_{\ell=0}^{\infty} [x > \ell] \cdot \left(4 + 2 \cdot \sum_{k=0}^{\infty} \left(\frac{\#col + \ell}{N} \right)^k \right) \\
& - 2 \cdot [cp[i] = 0] \cdot [x > 0] \cdot \sum_{k=0}^{\infty} \left(\frac{\#col}{N} \right)^k,
\end{aligned}$$

where $\#col \triangleq \sum_{i=1}^N [cp[i] \neq 0]$ denotes the number of coupons that have already been collected. This invariant was essentially obtained by performing a few fixed-point iterations to approximate the least fixed point that determines $\text{ert}[C_{out}](\mathbf{0})$ and then generalizing the result.

What Is the Intuition Underlying This Invariant? At least one time unit is consumed by the outer loop, because the loop guard is evaluated at least once. Variable ℓ represents the number of iterations of the outer loop. Since x is decremented in every iteration, C_{out} performs x iterations; all summands in I are thus zero if $x \leq \ell$. How long, then, is the expected runtime of every loop iteration? At least four units of time are required to evaluate the guard of the outer loop and the guard of the inner loop for the first time, and to perform the two assignments at the end of the loop body of C_{out} , respectively. Every iteration of the inner loop requires two additional units of time to evaluate the loop guard and perform an assignment. The expected number of iterations performed by the inner loop is expressed by the second sum, where k intuitively represents the number of times we enter the body of the inner loop. However, in the very first iteration of the outer loop, we have $[cp[i] = 0]$. In other words, no iteration of the inner loop is performed; i.e., the inner sum represents runtime that is never actually consumed by the program. We thus subtract this illegally added runtime. Furthermore, notice that the probability of performing the k -t iteration of the inner loop depends on the number of already collected coupons, i.e., the number $\#col$ of initially collected coupons plus ℓ , due to the fact that one coupon is collected in every iteration of the outer loop.

A detailed verification that I is indeed an upper global invariant is found in Appendix C.2 (this invariant verification further requires exhibiting a suitable invariant for the inner loop C_{in}).

The Expected Runtime of C_{cp} . We are now in a position to compute an upper bound for $\text{ert}[C_{cp}](\mathbf{0})$ using the previously proposed invariant I :

$$\begin{aligned}
\text{ert}[C_{cp}](\mathbf{0}) &= \text{ert}[cp := [0, \dots, 0]; i := 1; x := N; C_{out}](0) \\
&= 3 + (\text{ert}[C_{out}](\mathbf{0})) [x/N, i/1, cp[1]/0, \dots, cp[N]/0] \\
&\leq 3 + I[x/N, i/1, cp[1]/0, \dots, cp[N]/0]. \qquad (\text{ert}[C_{out}](0) \leq I)
\end{aligned}$$

Inserting our invariant I in the computation of $\text{ert}[C_{cp}](\mathbf{0})$ then yields

$$\begin{aligned} \text{ert}[C_{cp}](0) &\leq 4 + [N > 0] \cdot \left(4N + 2 \sum_{\ell=1}^{N-1} \sum_{k=0}^{\infty} \left(\frac{\ell}{N} \right)^k \right) && \left(\left(\sum_{\ell=0}^{\infty} [x > \ell] \right) [x/N] = \sum_{\ell=0}^{N-1} 1 \right) \\ &= 4 + [N > 0] \cdot \left(4N + 2 \sum_{\ell=1}^{N-1} \frac{N}{\ell} \right) && \left(\text{geom. series and} \right. \\ &= 4 + [N > 0] \cdot 2N \cdot (2 + \mathcal{H}_{N-1}), && \left. \text{sum reordering} \right) \end{aligned}$$

where $\mathcal{H}_{N-1} \triangleq 0 + 1/1 + 1/2 + 1/3 + \dots + 1/N-1$ denotes the $(N-1)$ -th harmonic number. Since the harmonic numbers approach asymptotically to the natural logarithm, we conclude that the coupon collector algorithm C_{cp} runs in expected time $\mathcal{O}(N \cdot \log(N))$.

9.2 One-Dimensional Random Walk

The probabilistic program

```

Crw : x := 10;
      while (x > 0) {
C :   x := 1/2 · ⟨x-1⟩ + 1/2 · ⟨x+1⟩
      }

```

models a one-dimensional walk of a particle that starts at position $x = 10$ and moves with equal probability to the left or to the right in each turn. The random walk terminates when the particle reaches position $x = 0$. It can be shown that the program terminates with probability one [25] but requires, on average, an infinite amount of time to do so. We apply the ert-calculus to formally derive this.

It is easy to see that $\text{ert}[C_{rw}](0) \leq \infty$, so we concentrate on proving that ∞ is a lower bound on the expected runtime of C_{rw} . To that end, we derive a lower ω -invariant I_n of loop while $(x > 0)$ {C} with respect to continuation $\mathbf{0}$.

Invariant Synthesis. To obtain a suitable invariant I_n , we first propose the following template:

$$I_n = 1 + \sum_{k=0}^n [x > k] \cdot a_{n,k}.$$

We always need one unit of time, because the loop guard is evaluated at least once. Furthermore, parameter n corresponds to the number of loop iterations of loop while $(x > 0)$ {C}. Intuitively, we then consider (a lower bound of) the expected time $a_{n,k}$ such that, after n loop iterations entering the loop body, the position of the particle is k . Since the random walk does not terminate within n loop iterations, we know that it suffices to consider these runtimes for $0 \leq k \leq n$.

As a next step, we show that for all nonnegative constants $a_{n,k}$, I_n is a lower ω -invariant of the loop with respect to continuation $\mathbf{0}$ whenever these $a_{n,k}$ satisfy the recurrence relation

$$\begin{aligned} a_{0,0} &= 1, \\ a_{n+1,0} &= 2 + \frac{1}{2} \cdot (a_{n,0} + a_{n,1}), \\ a_{n+1,k} &= \frac{1}{2} \cdot (a_{n,k-1} + a_{n,k+1}) \text{ for all } 1 \leq k \leq n+1, \\ a_{n,k} &= 0 \text{ for all } k > n, \end{aligned}$$

for $n \geq 0$. Intuitively, if the target position k after 0 further loop iterations is zero, then the expected runtime after one iteration is one due to the guard evaluation that terminates the loop. Similarly, if $k = 0$ but we have to perform $n + 1$ loop iterations, we require at least two units of time to evaluate

the guard and execute the probabilistic assignment in the loop body. The remaining expected runtime is then the weighted sum of $a_{n,k-1}$ (i.e., the position is increased in the last step and it suffices to reach position $k-1$) and $a_{n,k+1}$ (i.e., the position is decreased in the last iteration and it suffices to reach position $k+1$). For the case $1 \leq k < n$, we analogously express the expected time to reach position k in terms of the expected time to reach distance $k-1$ and $k+1$, respectively. Since we are only interested in a lower bound, it suffices to omit the two additional units of time. Furthermore, if $k > n$, then we abort, because the random walk cannot terminate within n loop iterations.

Now, one suitable solution for $a_{n,k}$ is

$$a_{n,k} = \frac{1}{2^n} \left[-\binom{n}{\lfloor \frac{n-k}{2} \rfloor} + 2 \sum_{i=0}^{n-k} 2^i \binom{n-i}{\lfloor \frac{n-i-k}{2} \rfloor} \right].$$

This solution was obtained with the help of Mathematica. A formal proof showing that I_n (with the above coefficients $a_{n,k}$) is indeed a lower ω -invariant with respect to continuation 0 is found in Appendix C.1.

The Expected Runtime of C_{rw} . Now Theorem 4.6 and the fact that $\lim_{n \rightarrow \infty} a_{n,0} = \infty$ yield

$$\text{ert}[\text{while } (x > 0) \{C\}](0) \geq \lim_{n \rightarrow \infty} I_n \geq \lim_{n \rightarrow \infty} 1 + [x > 0] \cdot a_{n,0} = 1 + [x > 0] \cdot \infty.$$

The calculations for the aforementioned steps can be found in Appendix C.1. Altogether we have

$$\begin{aligned} \text{ert}[C_{rw}](\mathbf{0}) &= \text{ert}[x := 10](\text{ert}[\text{while } (x > 0) \{C\}](\mathbf{0})) \\ &\geq \text{ert}[x := 10](1 + [x > 0] \cdot \infty) \\ &= 1 + (1 + [x > 0] \cdot \infty)[x/10] \\ &= 1 + (1 + 1 \cdot \infty) = \infty. \end{aligned}$$

Thus, $\text{ert}[C_{rw}](\mathbf{0}) \geq \infty$. As the reverse inequality trivially holds, the expected runtime of the one-dimensional random walk is infinite, i.e., $\text{ert}[C_{rw}](\mathbf{0}) = \infty$.

9.3 Randomized Binary Search

As our third case study, we show the applicability of our approach to randomized algorithms by analyzing a probabilistic, so-called Sherwood [35], variant of the classical recursive binary search algorithm. The main difference with the classical version is that in each recursive call the pivot element is picked uniformly at random from the remaining array, aligning, this way, the worst, best, and average case of the algorithm's runtime.

The Sherwood binary search algorithm searches for the value val in array $a[\text{left}.. \text{right}]$. It is encoded by procedure B with declaration \mathcal{D} presented in Figure 5. We use the random assignment $mid \approx \text{Unif}(\text{left} \dots \text{right})$ to model the random selection of the pivot element. For simplicity, we assume that the random assignment is performed in constant time 1 if $\text{left} \leq \text{right}$ and that it diverges (thus has runtime ∞) if $\text{left} > \text{right}$.

As the runtime heavily depends on the input data, we restrict our input and perform a runtime analysis for those inputs where val does *not* occur in the array, which constitutes the worst case for the classical binary search algorithm. Under this assumption, we can distinguish two cases: val is either smaller than every element in the array or larger than all of them.

For the first case, the expected runtime is bounded from above by $1+u$, where

$$u = [\text{left} > \text{right}] \cdot \infty + 3 + [\text{left} < \text{right}] \cdot (5 \cdot H_{\text{right}-\text{left}+1} - 1/5),$$

and H_k is the k th harmonic number. Ignoring the actual values of the constants, the intuition for the above expression is this: If $\text{left} > \text{right}$, the uniform sampling is assumed to diverge and thus we

$$u = [left > right] \cdot \infty + 3 + [left < right] \cdot \left(5 + \sum_{i=left}^{right} \left(\frac{[\min(i+1, right) < right]}{right - left + 1} \cdot \left(5 \cdot H_{right - \min(i+1, right) < right + 1} - 5/2 \right) \right) \right)$$

```

B ▷ 1: mid := Unif(left, right);
      2 + [left < right] · (2 + [a[mid] < val] · (3
      + [min(mid + 1, right) < right] · (5 · Hright - min(mid+1, right)+1 - 5/2)
      + [a[mid] > val] · (...))
2: if (left < right){
      3 + [a[mid] < val] · u[left/min(mid + 1, right)] + [a[mid] > val] · (...
3:   if (a[mid] < val){
      2 + u[left/min(mid + 1, right)]
4:     left := min(mid + 1, right);
      1 + u
5:     call B
      0
6:   } else {
      2 + [a[mid] > val] · (...
7:     if (a[mid] > val){
      2 + u[right/max(mid - 1, left)]
8:       right := max(mid - 1, left);
      1 + u
9:       call B
      0
10:    } else { 1 skip 0 } 0
11:  } 0
12: } else { 1 skip 0 } 0

```

Fig. 5. Declaration (boldface black) of the randomized binary search procedure B together with the runtime analysis (lightface gray) for the case that every value occurring in $a[left..right]$ is smaller than val . We write $j \mathbf{C} h$ for $\text{ert}[C, \mathcal{D}](h) \leq j \cdot H_k$ stands for the k th harmonic number.

get infinite runtime. If $left = right$, $O(1)$ steps are performed, but the procedure is not recursively called. Finally, if $left < right$, the test for $a[mid] < val$ will by assumption (val is smaller than every value in the array a) fail and the test $a[mid] > val$ will succeed. The binary search procedure is then called recursively on the randomly selected new interval and this takes on average $O(H_{right-left+1})$ units of time.

For formally showing

$$\text{ert}[\text{call } B, \mathcal{D}](0) \leq 1 + u$$

by application of Theorem 7.3 (1), we have to establish

$$\text{ert}[\text{call } B, \mathcal{D}](0) \leq 1 + u \quad \Vdash \quad \text{ert}[\mathcal{D}(B), \mathcal{D}](0) \leq u.$$

The details of this derivation are provided in Figure 5. In the annotations of Line 9, we use the assumption $\text{ert}[\text{call } B, \mathcal{D}](0) \leq 1 + u$. All other annotations are straightforward applications of the rules in Table 1 while keeping in mind that val is smaller than every element in the array. The top-most annotation containing the sum is equal to u by algebraic reasoning. All in all, we have proven $\text{ert}[\mathcal{D}(B), \mathcal{D}](0) \leq u$ assuming $\text{ert}[\text{call } B, \mathcal{D}](0) \leq 1 + u$. By Theorem 7.3 (1), we now know that

$\text{ert}[\text{call } B, \mathcal{D}](0) \leq 1 + u$ for all initial states in which val is smaller than every element in the array.

Similarly, we can show that when val is greater than every element in the array, the expected runtime is upper bounded by $1 + u$, with

$$u = [\text{left} > \text{right}] \cdot \infty + 3 + [\text{left} < \text{right}] \cdot (6 \cdot H_{\text{right}-\text{left}+1} - 3).$$

The verification for this case is analogous and therefore omitted here.

Combining the two cases, we conclude that when the sought-after value does not occur in the array, the algorithm terminates in expected time in $O(\log n)$, where $n = \text{right} - \text{left} + 1$ is the size of the array, since $H_k \in \Theta(\log k)$.

10 RELATED WORK

Resource Analysis of Deterministic Programs. Several works apply wp-style or Floyd-Hoare-style reasoning to study quantitative aspects of classical algorithms. Nielson [41, 42] provides a Hoare logic for determining upper bounds on the runtime of deterministic programs. Our approach applied to such programs yields the tightest upper bound on the runtime that can be derived using Nielson's approach. Arthan et al. [1] provide a general framework for sound and complete Hoare-style logics and show that an instance of their theory can be used to obtain upper bounds on the runtime of while-programs. Hickey and Cohen [20] automate the average-case analysis of deterministic programs by generating a system of recurrence equations derived from a program whose efficiency is to be analyzed. They build on Kozen's seminal work [31] on the semantics of probabilistic programs. Berghammer and Müller-Olm [5] show how Hoare-style reasoning can be extended to obtain bounds on the closeness of results obtained using approximate algorithms to the optimal solution. Deriving space and time consumption of deterministic programs has also been considered by Hehner [17]. Alternative approaches for analyzing resource consumption in deterministic programs include, among others, type-checking [22], abstract interpretation [47], and worst-case execution time analysis [49].

Runtime Analysis of Probabilistic Programs. Classical techniques to analyze the runtime of randomized algorithms include probabilistic recurrence relations [29] and martingale theory. Formal reasoning about probabilistic programs goes back to Kozen [31] and has been developed further by Hehner [18] and McIver and Morgan [36]. A general abstract interpretation framework for the analysis of probabilistic programs has been given by Cousot and Monerau [11]. They capture our approach as an abstraction of a low-level trace semantics, but their work does not provide any proof rules. The work by Celiku and McIver [8] is perhaps the closest to our article. They provide a wp-calculus for obtaining performance properties of probabilistic programs, including upper bounds on expected runtimes. Their focus is on refinement. They neither provide a soundness result of their approach nor consider lower bounds. We believe that our transformer is simpler to work with in practice too. Monniaux [38] exploits abstract interpretation to automatically prove the probabilistic termination of programs using exponential bounds on the tail of the distribution. His analysis can be used to prove the soundness of experimental statistical methods to determine the average runtime of probabilistic programs. Brázdil et al. [7] study the runtime of probabilistic programs with unbounded recursion by considering probabilistic pushdown automata (pPDAs). They show (using Martingale theory) that for every pPDA, the probability of performing a long run decreases exponentially (polynomially) in the length of the run, if and only if the pPDA has a finite (infinite) expected runtime. As opposed to our program verification technique, [7] considers reasoning at the operational level. Fioriti and Hermanns [14] proposed a typing scheme for deciding almost-sure termination. They showed, among others, that if a program is well typed,

then it almost surely terminates. This result does not cover positive almost-sure termination; that is, their approach cannot be used to obtain that the runtime is infinite.

Automated Expected Runtime Analysis. Chatterjee et al. [10] recently presented a linear-time algorithm to derive (logarithmic, linear, or almost-linear) bounds on expected runtimes of randomized algorithms whose runtime is described by a set of recurrence relations. The key idea is to overapproximate terms in a recurrence relation through integral and Taylor expansion enabling one to obtain bounds by comparing the leading terms of pseudo-polynomials. This technique derives bounds for randomized univariate and separable bivariate recurrence relations and is applicable to classical algorithms such as quick-select, Sherwood randomized search, and the coupon collector.

The ert-calculus developed in this article provides a solid basis for the automated analysis of expected runtimes. This is witnessed by some recent follow-up works. Ngo et al. [40] combined the principles of the ert-calculus with existing automated amortized resource analysis techniques. This results in an automated approach to derive upper bounds (as symbolic polynomials) on the expected runtime of probabilistic programs. The correctness is shown by proving that the technique provides upper bounds on the expected runtime of a program as defined by our ert-calculus. In essence, their technique assumes potential functions to be a linear combination of base functions and derives using inference rules akin to those in this article of a system of in-equations that is solved by an LP solver. Experimental results show that this works for an interesting class of sample programs.

Batz et al. [4] showed how the ert-calculus can be used to obtain exact expected sampling times of Bayesian networks in a fully automated fashion. The key idea here is that loops in the probabilistic programs describing such networks are statistically independent, enabling obtaining closed-form symbolic expressions for their expected runtime. An experimental evaluation on Bayesian network benchmarks demonstrates that ill-conditioned networks—resulting in very large simulation times—can be automatically inferred within less than a second.

11 CONCLUSION

We have presented a wp-style calculus for reasoning about the expected runtime and positive almost-sure termination of randomized algorithms. Our main contribution consists of several sound and complete proof rules for obtaining upper as well as lower bounds on the expected runtime of loops. We applied these rules to analyze the expected runtime of a variety of example algorithms including the well-known coupon collector problem. While finding invariants is, in general, a challenging task, we were able to find correct invariants by considering a few loop unrollings most of the time. Hence, we believe that our proof rules are natural and widely applicable, and provide a viable alternative to existing techniques to determine the expected runtime. The approach is a conservative extension of Nielson's approach for reasoning about the runtime of deterministic programs and our calculus is sound with respect to a simple operational model defined in terms of (push-down) Markov chains. Toward automation of our approach, an important step is to develop techniques for automated loop-invariant synthesis; initial approaches for probabilistic programs can be found in, e.g., [3, 13, 30].

APPENDIXES

A RECAP OF AUXILIARY RESULTS

For the sake of self-containment, we recall here some well-known theorems that we use to establish our main results.

THEOREM A.1 (LEBESGUE'S MONOTONE CONVERGENCE THEOREM). *Let $(f_n)_{n \in \mathbb{N}}$ be a sequence of functions of type $A \rightarrow \mathbb{R}_{\geq 0}^{\infty}$ such that $f_n(a) \leq f_{n+1}(a)$ for all $a \in A$. Then, for every probability distribution ν over A , we have*

$$E_{\nu}(\sup_n f_n) = \sup_n E_{\nu}(f_n),$$

where $E_{\nu}(f)$ denotes the expected value of f with respect to ν and $\sup_n f_n$ is taken pointwise.

PROOF. See, e.g., [46, Ch. 21]. □

THEOREM A.2 (MONOTONE SEQUENCE THEOREM). *If $(a_n)_{n \in \mathbb{N}}$ is a monotonically increasing sequence in $\mathbb{R}_{\geq 0}^{\infty}$, then*

$$\sup_n a_n = \lim_{n \rightarrow \infty} a_n.$$

PROOF. See, e.g., [45, Ch. 2]. □

THEOREM A.3 (KLEENE'S FIXED-POINT THEOREM). *Let (A, \leq) be an ω -cpo with bottom element \perp and let $F : A \rightarrow A$ be continuous.⁹ Then F admits a least fixed-point $\text{lfp}(F)$, which can moreover be obtained as*

$$\text{lfp}(F) = \sup_n F^n(\perp).$$

Here, F^n denotes the composition of F with itself n times, i.e., $F^0 = \text{id}$ and $F^{n+1} = F \circ F^n$.

PROOF. See, e.g., [48, Ch. 1]. □

THEOREM A.4 (PARK'S THEOREM). *Let (A, \leq) be an ω -cpo with bottom element and let $F : A \rightarrow A$ be continuous. Then, for every $a \in A$,*

$$F(a) \leq a \Rightarrow \text{lfp}(F) \leq a.$$

PROOF. See [48, Ch. 1]. □

LEMMA A.5 (DIAGONALIZATION OF DOUBLY INDEXED CHAINS). *Let $a_{n,m}$ be elements of an ω -cpo (A, \leq) such that $a_{n,m} \leq a_{n',m'}$ whenever $n \leq n'$ and $m \leq m'$. Then,*

$$\sup_n (\sup_m a_{n,m}) = \sup_m (\sup_n a_{n,m}) = \sup_i a_{i,i}.$$

PROOF. See [50, Ch. 8]. □

LEMMA A.6 (RELATIONAL FIXED-POINT FUSION). *Let (\mathcal{D}_1, \leq_1) , (\mathcal{D}_2, \leq_2) and (\mathcal{D}, \leq) be ω -cpo's with bottom elements \perp_1 , \perp_2 , and \perp , respectively. Moreover, let*

$$F_1 : \mathcal{D}_1 \rightarrow \mathcal{D}_1 \quad F_2 : \mathcal{D}_2 \rightarrow \mathcal{D}_2 \quad f_1 : \mathcal{D}_1 \rightarrow \mathcal{D} \quad f_2 : \mathcal{D}_2 \rightarrow \mathcal{D}$$

be continuous and $h_1, h_2 : \mathcal{D} \rightarrow \mathcal{D}$. If

- (1) $\forall d_1. f_1(F_1(d_1)) \leq h_1(f_1(d_1))$ and $\forall d_2. f_2(F_2(d_2)) \leq h_2(f_2(d_2))$,
- (2) $f_1(\perp_1) \leq f_2(\text{lfp } F_2)$ and $f_2(\perp_2) \leq f_1(\text{lfp } F_1)$, and
- (3) $h_1(f_2(\text{lfp } F_2)) \leq f_2(\text{lfp } F_2)$ and $h_2(f_1(\text{lfp } F_1)) \leq f_1(\text{lfp } F_1)$,

then

$$f_1(\text{lfp } F_1) = f_2(\text{lfp } F_2).$$

PROOF. See [43, App. A.6]. □

⁹A function $F : A \rightarrow A$ is said to be *continuous* if and only if it preserves suprema of ω -chains; i.e., for every ω -chain $a_0 \leq a_1 \leq \dots$ in A we have $\sup_n F(a_n) = F(\sup_n a_n)$.

B OMITTED PROOFS

B.1 Proof of Lemma 3.2 (Continuity of ert)

We prove that for every program C , the transformer $\text{ert}[C]$ is continuous; i.e., for every ω -chain of runtimes $t_0 \leq t_1 \leq \dots$,

$$\text{ert}[C](\sup_n t_n) = \sup_n \text{ert}[C](t_n),$$

by induction on the structure of C . We consider only the cases of assignments, conditional choices, and while-loops as the proof argument for the remaining language constructs is straightforward.

Assignment. The proof relies on Lebesgue's Monotone Convergence Theorem (LMCT) recalled in Theorem A.1. We have:

$$\begin{aligned} \text{ert}[x \approx \mu](\sup_n t_n) &= \mathbf{1} + \lambda\sigma. \mathbb{E}_{\|\mu\|(\sigma)}(\lambda v. (\sup_n t_n)[x/v](\sigma)) && \text{(Table 1)} \\ &= \mathbf{1} + \lambda\sigma. \mathbb{E}_{\|\mu\|(\sigma)}(\sup_n \lambda v. t_n[x/v](\sigma)) \\ &= \mathbf{1} + \lambda\sigma. \sup_n \mathbb{E}_{\|\mu\|(\sigma)}(\lambda v. t_n[x/v](\sigma)) && \text{(LMCT)} \\ &= \sup_n \mathbf{1} + \lambda\sigma. \mathbb{E}_{\|\mu\|(\sigma)}(\lambda v. t_n[x/v](\sigma)) && \text{(1 is constant)} \\ &= \sup_n \text{ert}[x \approx \mu](t_n). && \text{(Table 1)} \end{aligned}$$

Conditional Choice. The proof relies on a Monotone Sequence Theorem (MST) recalled in Theorem A.2. We have:

$$\begin{aligned} \text{ert}[\text{if } (\xi) \{C_1\} \text{ else } \{C_2\}](\sup_n t_n) &= \mathbf{1} + \llbracket \xi : \text{true} \rrbracket \cdot \text{ert}[C_1](\sup_n t_n) + \llbracket \xi : \text{false} \rrbracket \cdot \text{ert}[C_2](\sup_n t_n) && \text{(Table 1)} \\ &= \mathbf{1} + \llbracket \xi : \text{true} \rrbracket \cdot \sup_n \text{ert}[C_1](t_n) + \llbracket \xi : \text{false} \rrbracket \cdot \sup_n \text{ert}[C_2](t_n) && \text{(I.H. on } C_1, C_2) \\ &= \mathbf{1} + \llbracket \xi : \text{true} \rrbracket \cdot \lim_{n \rightarrow \infty} \text{ert}[C_1](t_n) + \llbracket \xi : \text{false} \rrbracket \cdot \lim_{n \rightarrow \infty} \text{ert}[C_2](t_n) && \text{(MST)} \\ &= \lim_{n \rightarrow \infty} \mathbf{1} + \llbracket \xi : \text{true} \rrbracket \cdot \text{ert}[C_1](t_n) + \llbracket \xi : \text{false} \rrbracket \cdot \text{ert}[C_2](t_n) \\ &= \sup_n \mathbf{1} + \llbracket \xi : \text{true} \rrbracket \cdot \text{ert}[C_1](t_n) + \llbracket \xi : \text{false} \rrbracket \cdot \text{ert}[C_2](t_n) && \text{(MST)} \\ &= \sup_n \text{ert}[\text{if } (\xi) \{C_1\} \text{ else } \{C_2\}](t_n). && \text{(Table 1)} \end{aligned}$$

While-Loop. Let $F_t(X) = \mathbf{1} + \llbracket \xi : \text{false} \rrbracket \cdot t + \llbracket \xi : \text{true} \rrbracket \cdot \text{ert}[C'](X)$ be the characteristic functional of loop while $(\xi) \{C'\}$. The proof relies on three facts about F_t and the lfp operator:

- (1) $F_{\sup_n t_n} = \sup_n F_{t_n}$, which follows from a straightforward reasoning.
- (2) $\sup_n F_{t_n}$ is continuous (in $\mathbb{T} \rightarrow \mathbb{T}$), which follows from the fact that $\langle F_{t_n} \rangle$ forms an ω -chain of continuous transformers (since by I.H. $\text{ert}[C']$ is continuous) and continuous functions are closed under supremums.
- (3) $\text{lfp} : [\mathbb{T} \rightarrow \mathbb{T}] \rightarrow \mathbb{T}$ is itself continuous when restricted to the set of continuous transformers in $\mathbb{T} \rightarrow \mathbb{T}$, denoted $[\mathbb{T} \rightarrow \mathbb{T}]$ [48, Proposition 12].

We then have

$$\begin{aligned} \text{ert}[\text{while } (\xi) \{C'\}](\sup_n t_n) &= \text{lfp}(F_{\sup_n t_n}) && \text{(Table 1)} \\ &= \text{lfp}(\sup_n F_{t_n}) && \text{(Fact(1))} \\ &= \sup_n \text{lfp}(F_{t_n}) && \text{(Facts(2)and(3))} \\ &= \sup_n \text{ert}[\text{while } (\xi) \{C'\}](t_n). && \text{(Table 1)} \quad \square \end{aligned}$$

B.2 Proof of Theorem 3.4 (Constant Propagation)

For any halting-free program $C \in \text{pGCL}$, any constant $k \in \mathbb{R}_{\geq 0}$, and any runtime $t \in \mathbb{T}$, we prove

$$\text{ert}[C](\mathbf{k} + t) = \mathbf{k} + \text{ert}[C](t)$$

by induction on the structure of C . We consider only the cases of assignments and while-loops as the proof argument for the remaining language constructs is straightforward.

Assignment. For $x \approx \mu$, the proof relies on the linearity of the expected value operator; i.e., $E_\nu(\mathbf{k} + t) = k + E_\nu(t)$ provided distribution ν has total mass 1. We have

$$\begin{aligned} \text{ert}[x \approx \mu](\mathbf{k} + t) &= \mathbf{1} + \lambda\sigma. E_{\llbracket \mu \rrbracket(\sigma)}(\lambda\nu. (\mathbf{k} + t)[x/\nu](\sigma)) && \text{(Table 1)} \\ &= \mathbf{1} + \lambda\sigma. E_{\llbracket \mu \rrbracket(\sigma)}(\lambda\nu. k + t[x/\nu](\sigma)) && (\mathbf{k}[x/\nu](\sigma) = k) \\ &= \mathbf{1} + \mathbf{k} + \lambda\sigma. E_{\llbracket \mu \rrbracket(\sigma)}(\lambda\nu. t[x/\nu](\sigma)) && \text{(linearity of E)} \\ &= \mathbf{k} + \text{ert}[x \approx \mu](t). && \text{(Table 1)} \end{aligned}$$

While-Loop. Let $F_t(X) = \mathbf{1} + \llbracket \xi : \text{false} \rrbracket \cdot t + \llbracket \xi : \text{true} \rrbracket \cdot \text{ert}[C'](X)$ be the characteristic functional of loop while $(\xi) \{C'\}$. We have to show

$$\text{lfp } F_{\mathbf{k}+t} = \mathbf{k} + \text{lfp } F_t,$$

which is equivalent to the pair of inequalities

$$\text{lfp } F_{\mathbf{k}+t} \leq \mathbf{k} + \text{lfp } F_t \quad \text{and} \quad \text{lfp } F_t \leq \text{lfp } F_{\mathbf{k}+t} - \mathbf{k}.$$

These inequalities follow, in turn, from equalities

$$F_{\mathbf{k}+t}(\mathbf{k} + \text{lfp } F_t) = \mathbf{k} + \text{lfp } F_t \quad \text{and} \quad F_t(\text{lfp } F_{\mathbf{k}+t} - \mathbf{k}) = \text{lfp } F_{\mathbf{k}+t} - \mathbf{k}.$$

(This is because lfp gives the *least* fixed point of a transformer and then $F(x) = x$ implies $\text{lfp } F \leq x$.) Let us now discharge each of the above equalities:

$$\begin{aligned} F_{\mathbf{k}+t}(\mathbf{k} + \text{lfp } F_t) &= \mathbf{1} + \llbracket \xi : \text{false} \rrbracket \cdot (\mathbf{k} + t) + \llbracket \xi : \text{true} \rrbracket \cdot \text{ert}[C'](\mathbf{k} + \text{lfp } F_t) && \text{(def. } F_{\mathbf{k}+t}) \\ &= \mathbf{1} + \llbracket \xi : \text{false} \rrbracket \cdot (\mathbf{k} + t) + \llbracket \xi : \text{true} \rrbracket \cdot (\mathbf{k} + \text{ert}[C'](\text{lfp } F_t)) && \text{(I.H. on } C') \\ &= \mathbf{1} + \mathbf{k} + \llbracket \xi : \text{false} \rrbracket \cdot t + \llbracket \xi : \text{true} \rrbracket \cdot \text{ert}[C'](\text{lfp } F_t) \\ &= \mathbf{k} + F_t(\text{lfp } F_t) && \text{(def. } F_t) \\ &= \mathbf{k} + \text{lfp } F_t. && \text{(def. lfp)} \end{aligned}$$

$$\begin{aligned} F_t(\text{lfp } F_{\mathbf{k}+t} - \mathbf{k}) &= \mathbf{1} + \llbracket \xi : \text{false} \rrbracket \cdot t + \llbracket \xi : \text{true} \rrbracket \cdot \text{ert}[C'](\text{lfp } F_{\mathbf{k}+t} - \mathbf{k}) && \text{(def. } F_t) \\ &= \mathbf{1} + \llbracket \xi : \text{false} \rrbracket \cdot t + \llbracket \xi : \text{true} \rrbracket \cdot (\text{ert}[C'](\text{lfp } F_{\mathbf{k}+t} - \mathbf{k}) + 2\mathbf{k} - 2\mathbf{k}) \\ &= \mathbf{1} + \llbracket \xi : \text{false} \rrbracket \cdot t + \llbracket \xi : \text{true} \rrbracket \cdot (\text{ert}[C'](\text{lfp } F_{\mathbf{k}+t} - \mathbf{k} + 2\mathbf{k}) - 2\mathbf{k}) && \text{(I.H. on } C') \\ &= \mathbf{1} + \llbracket \xi : \text{false} \rrbracket \cdot t + \llbracket \xi : \text{true} \rrbracket \cdot (\text{ert}[C'](\text{lfp } F_{\mathbf{k}+t} + \mathbf{k}) - 2\mathbf{k}) \\ &= \mathbf{1} + \llbracket \xi : \text{false} \rrbracket \cdot t + \llbracket \xi : \text{true} \rrbracket \cdot (\text{ert}[C'](\text{lfp } F_{\mathbf{k}+t}) + \mathbf{k} - 2\mathbf{k}) && \text{(I.H. on } C') \\ &= \mathbf{1} + \llbracket \xi : \text{false} \rrbracket \cdot t + \llbracket \xi : \text{true} \rrbracket \cdot (\text{ert}[C'](\text{lfp } F_{\mathbf{k}+t}) - \mathbf{k}) \\ &= \mathbf{1} + \llbracket \xi : \text{false} \rrbracket \cdot (\mathbf{k} + t) + \llbracket \xi : \text{true} \rrbracket \cdot \text{ert}[C'](\text{lfp } F_{\mathbf{k}+t}) - \mathbf{k} \\ &= F_{\mathbf{k}+t}(\text{lfp } F_{\mathbf{k}+t}) - \mathbf{k} && \text{(def. } F_{\mathbf{k}+t}) \\ &= \text{lfp } F_{\mathbf{k}+t} - \mathbf{k}. && \text{(def. lfp)} \quad \square \end{aligned}$$

B.3 Proof of Theorem 5.5 (Soundness of the ert-Transformer)


Before we prove the soundness of ert with respect to the simple operational model of our probabilistic programming language introduced in Section 5, some preparation is needed to deal with sequential composition and while-loops. In particular, we use the following decomposition lemma.

LEMMA B.1. *Let $C_1, C_2 \in \text{pGCL}$, $t \in \mathbb{T}$, and $\sigma \in \Sigma$. Then*

$$\text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket C_1; C_2 \rrbracket} (\langle \text{sink} \rangle) = \text{ExpRew}^{\mathcal{M}_\sigma^u \llbracket C_1 \rrbracket} (\langle \text{sink} \rangle),$$

where $u \triangleq \text{ExpRew}^{\lambda\rho \cdot \mathcal{M}_\rho^t \llbracket C_2 \rrbracket} (\langle \text{sink} \rangle)$.

PROOF. The MC $\mathcal{M}_\sigma^t \llbracket C_1; C_2 \rrbracket$ is of the following form



Here, $\text{stmt}(\ell_1) = C_1$ and $\text{stmt}(\ell_2) = C_2$. Hence, every path starting in $\langle \ell_1, \sigma \rangle$ either eventually reaches $\langle \ell_2, \sigma' \rangle$, for some $\sigma' \in \Sigma$, or diverges, i.e., never reaches $\langle \text{sink} \rangle$. Since $\langle \ell_2, \sigma' \rangle$ is the initial state of the MC $\mathcal{M}_\sigma^t \llbracket C_2 \rrbracket$, we can transform $\mathcal{M}_\sigma^t \llbracket C_1; C_2 \rrbracket$ into an MC $\mathcal{M}_\sigma^u \llbracket C_1 \rrbracket$ with the same expected reward by setting

$$u \triangleq \text{ExpRew}^{\lambda\rho \cdot \mathcal{M}_\rho^t \llbracket C_2 \rrbracket} (\langle \text{sink} \rangle). \quad \square$$

Furthermore, we have to consider bounded while-loops that are obtained by successively unrolling a while-loop up to a finite number of executions of the loop body.

Definition B.2 (Bounded while-loops). Let $\xi \in \text{DExp}$ and $C \in \text{pGCL}$. Then the *bounded while loops* of while are given by

$$\begin{aligned}
 \text{while}^{<0}(\xi) \{C\} &\triangleq \text{halt}, \text{ and} \\
 \text{while}^{<k+1}(\xi) \{C\} &\triangleq \text{if } (\xi) \{C; \text{while}^{<k}(\xi) \{C\}\} \text{ else } \{\text{empty}\}.
 \end{aligned}$$

As for ordinary programs, the runtime of a while loop can be expressed in terms of the runtime of bounded while loops.

LEMMA B.3. *Let $\xi \in \text{DExp}$, $C \in \text{pGCL}$, and $t \in \mathbb{T}$. Then,*

$$\sup_{k \in \mathbb{N}} \text{ert}[\text{while}^{<k}(\xi) \{C\}](t) = \text{ert}[\text{while}(\xi) \{C\}](t).$$

PROOF. Let $F_t(X)$ be the characteristic functional corresponding to $\text{while}(\xi) \{C\}$ as defined in Definition 3.1. Assume, for the moment, that for each $k \in \mathbb{N}$, we have $\text{ert}[\text{while}^{<k}(\xi) \{C\}](t) = F_t^k(\mathbf{0})$. Then, using Kleene's Fixed Point Theorem, we can establish that

$$\sup_{k \in \mathbb{N}} \text{ert}[\text{while}^{<k}(\xi) \{C\}](t) = \sup_{k \in \mathbb{N}} F_t^k(\mathbf{0}) = \text{lfp } X.F_t(X) = \text{ert}[\text{while}(\xi) \{C\}](t).$$

Hence, it suffices to show that $\text{ert}[\text{while}^{<k}(\xi) \{C\}](t) = F_t^k(\mathbf{0})$ for each $k \in \mathbb{N}$. This can be established by a straightforward induction on k . \square

Intuitively, this means that the expected runtime of a loop and the runtime of its finite approximations by bounded loops coincide in the limit. This also holds for the expected reward of the MC of a loop and the expected reward of its finite approximations.

LEMMA B.4. *Let $\xi \in \text{DExp}$, $C \in \text{pGCL}$, $t \in \mathbb{T}$, and $\sigma \in \Sigma$. Then*

$$\sup_{k \in \mathbb{N}} \text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket \text{while}^{<k}(\xi) \{C\} \rrbracket} (\langle \text{sin}\kappa \rangle) = \text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket \text{while}(\xi) \{C\} \rrbracket} (\langle \text{sin}\kappa \rangle).$$

PROOF. First, observe that every path in the operational MC $\mathcal{M}_\sigma^t \llbracket \text{while}^{<k}(\xi) \{C\} \rrbracket$ either terminates or halts after k loop iterations. Since the operational MC $\mathcal{M}_\sigma^t \llbracket \text{while}(\xi) \{C\} \rrbracket$ does not prematurely stop after k loop iterations, it includes the above paths. Thus, we obtain the inequality

$$\sup_{k \in \mathbb{N}} \text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket \text{while}^{<k}(\xi) \{C\} \rrbracket} (\langle \text{sin}\kappa \rangle) \leq \text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket \text{while}(\xi) \{C\} \rrbracket} (\langle \text{sin}\kappa \rangle).$$

To prove the converse inequality, observe that every infinite path in the MC $\mathcal{M}_\sigma^t \llbracket \text{while}(\xi) \{C\} \rrbracket$ yields either reward 0 or reward ∞ . We distinguish two cases:

- (1) All paths in the operational MC $\mathcal{M}_\sigma^t \llbracket \text{while}(\xi) \{C\} \rrbracket$ that collect positive reward are finite. Then, for each of these paths, there exists some natural number $k \geq 0$ such that the operational MC $\mathcal{M}_\sigma^t \llbracket \text{while}^{<k}(\xi) \{C\} \rrbracket$ includes this path. By taking the supremum of these numbers k , we include every path of $\mathcal{M}_\sigma^t \llbracket \text{while}(\xi) \{C\} \rrbracket$ with positive reward. Hence,

$$\sup_{k \in \mathbb{N}} \text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket \text{while}^{<k}(\xi) \{C\} \rrbracket} (\langle \text{sin}\kappa \rangle) \geq \text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket \text{while}(\xi) \{C\} \rrbracket} (\langle \text{sin}\kappa \rangle).$$

- (2) There exists an infinite path in the operational MC $\mathcal{M}_\sigma^t \llbracket \text{while}(\xi) \{C\} \rrbracket$ that yields positive reward. By the above observation, this path yields reward ∞ . Then, for all natural numbers k , the operational MC $\mathcal{M}_\sigma^t \llbracket \text{while}^{<k}(\xi) \{C\} \rrbracket$ contains a prefix of this path that yields positive reward proportional to k . By taking the supremum of these numbers k , we end up with an infinite reward. Hence,

$$\sup_{k \in \mathbb{N}} \text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket \text{while}^{<k}(\xi) \{C\} \rrbracket} (\langle \text{sin}\kappa \rangle) = \infty = \text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket \text{while}(\xi) \{C\} \rrbracket} (\langle \text{sin}\kappa \rangle). \quad \square$$


We are now in a position to show the soundness of ert with respect to the operational semantics.

THEOREM 5.5 (SOUNDNESS OF THE ERT TRANSFORMER). *Let $C \in \text{pGCL}$, and $t \in \mathbb{T}$. Then, for each $\sigma \in \Sigma$, we have*

$$\text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket C \rrbracket} (\langle \text{sin}\kappa \rangle) = \text{ert}[C](t)(\sigma).$$

PROOF. We prove the claim by induction on the structure of pGCL programs $C \in \text{pGCL}$. The base cases $C = \text{empty}$, $C = \text{skip}$, and $C = \text{halt}$ are straightforward.

The Probabilistic Assignment $C = x := \mu$. For some $n \in \mathbb{N}$, the MC $\mathcal{M}_\sigma^t \llbracket x := \mu \rrbracket$ is of the following form—where $\text{stmt}(\ell) = x := \mu$:



Hence, all of its prefix-free paths reaching $\langle \text{sink} \rangle$ are of the form

$$\pi_i = \langle \ell, \sigma \rangle \langle \downarrow, \sigma[x/v_i] \rangle \langle \text{sink} \rangle \dots,$$

with $\Pr\{\pi_i\} = p_i$ for each $v_i \in \text{Val}$ with $\llbracket \mu \rrbracket(\sigma)(v_i) = p_i > 0$ and

$$\sum_{k=1}^n p_k = 1.$$

Moreover, $\text{rew}(\pi_i) = 1 + t(\sigma[x/v_i])$. Thus, we have

$$\begin{aligned} & \text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket x := \mu \rrbracket} (\langle \text{sink} \rangle) \\ &= \sum_{1 \leq i \leq n} \Pr\{\pi_i\} \cdot \text{rew}(\pi_i) \\ &= \sum_{1 \leq i \leq n} \Pr\{\pi_i\} \cdot (1 + t(\sigma[x/v_i])) \\ &= 1 + \sum_{1 \leq i \leq n} \Pr\{\pi_i\} \cdot t(\sigma[x/v_i]) \\ &= 1 + \sum_{1 \leq i \leq n} \llbracket \mu \rrbracket(\sigma)(v_i) \cdot t(\sigma[x/v_i]) \\ &= 1 + \mathbb{E}_{\llbracket \mu \rrbracket(\sigma)} (\lambda v_i \bullet t(\sigma[x/v_i])) \\ &= \text{ert}[x := \mu](t)(\sigma). \end{aligned}$$

Induction Hypothesis. For all (substatements) $C' \in \text{pGCL}$ of C , $t \in \mathbb{T}$ and $\sigma \in \Sigma$,

$$\text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket C' \rrbracket} (\langle \text{sink} \rangle) = \text{ert}[C'](t)(\sigma).$$

For the induction step, we consider sequential composition and while-loops. The case of conditional statements is straightforward.

Sequential Composition. $C = C_1; C_2$.

$$\begin{aligned} & \text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket C_1; C_2 \rrbracket} (\langle \text{sink} \rangle) \\ &= \text{ExpRew}^{\mathcal{M}_\sigma^{\text{ExpRew}^{\lambda \rho \bullet \mathcal{M}_\rho^t \llbracket C_2 \rrbracket} (\langle \text{sink} \rangle)} \llbracket C_1 \rrbracket} (\langle \text{sink} \rangle) \quad (\text{Lemma B.1}) \\ &= \text{ExpRew}^{\mathcal{M}_\sigma^{\lambda \rho \bullet \text{ert}[C_2](t)(\rho)} \llbracket C_1 \rrbracket \langle \text{sink} \rangle} (\langle \text{sink} \rangle) \quad (\text{I.H. on } C_2) \\ &= \text{ert}[C_1](\text{ert}[C_2](t)(\sigma)) \quad (\text{I.H. on } C_1) \\ &= \text{ert}[C_1; C_2](t)(\sigma). \end{aligned}$$

While-Loop. Let $C = \text{while}$. For any natural number $k \geq 1$ and $\sigma \in \Sigma$, we have

$$\begin{aligned} & \text{ert}[\text{while}^{<k}(\xi)\{C'\}](t)(\sigma) \\ &= \text{ert}[\text{if}(\xi)\{C'; \text{while}^{<k-1}(\xi)\{C'\}\} \text{else}\{\text{empty}\}](t)(\sigma) \\ &= 1 + \llbracket \xi : \text{true} \rrbracket(\sigma) \cdot \text{ert}[C'; \text{while}^{<k-1}(\xi)\{C'\}](t)(\sigma) \\ &\quad + \llbracket \xi : \text{false} \rrbracket(\sigma) \cdot \text{ert}[\text{empty}](t)(\sigma) \\ &= 1 + \llbracket \xi : \text{true} \rrbracket(\sigma) \cdot \text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket C'; \text{while}^{<k-1}(\xi)\{C'\} \rrbracket} (\langle \text{sink} \rangle) \quad (\text{I.H.}) \\ &\quad + \llbracket \xi : \text{false} \rrbracket(\sigma) \cdot \text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket \text{empty} \rrbracket} (\langle \text{sink} \rangle) \\ &= \text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket \text{while}^{<k}(\xi)\{C'\} \rrbracket} (\langle \text{sink} \rangle). \end{aligned}$$

Moreover, for $k = 0$, we have

$$\begin{aligned}
& \text{ert}[\text{while}^{<0}(\xi) \{C'\}](t)(\sigma) \\
&= \text{ert}[\text{halt}](t)(\sigma) && \text{Definition B.2} \\
&= \text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket \text{halt} \rrbracket} (\langle \text{sin}\mathcal{K} \rangle) && \text{(already shown in base case)} \\
&= \text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket \text{while}^{<0}(\xi) \{C'\} \rrbracket} (\langle \text{sin}\mathcal{K} \rangle). && \text{(Definition B.2)}
\end{aligned}$$

Putting both cases together, we can establish that

$$\begin{aligned}
& \text{ert}[\text{while}(\xi) \{C'\}](t)(\sigma) \\
&= \sup_{k \in \mathbb{N}} \text{ert}[\text{while}^{<k}(\xi) \{C'\}](t)(\sigma) && \text{(Lemma B.3)} \\
&= \sup_{k \in \mathbb{N}} \text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket \text{while}^{<k}(\xi) \{C'\} \rrbracket} (\langle \text{sin}\mathcal{K} \rangle) \\
&= \text{ExpRew}^{\mathcal{M}_\sigma^t \llbracket \text{while}(\xi) \{C'\} \rrbracket} (\langle \text{sin}\mathcal{K} \rangle). && \text{(Lemma B.4)} \quad \square
\end{aligned}$$

B.4 Proof of Theorem 6.1 (Soundness of ert w.r.t. Nielson's Proof System)

The proof relies on three auxiliary results that are presented first. The first lemma shows a standard relationship between while-loops and their unrollings.

LEMMA B.5. *For all pGCL programs C , $t \in \mathbb{T}$ and deterministic guards B , we have*

$$\text{ert}[\text{while}(B) \{C\}](t) = \text{ert}[\text{if}(B) \{C; \text{while}(B) \{C\}\} \text{else} \{\text{empty}\}](t).$$

PROOF. Let $F_t(X)$ be the characteristic functional corresponding to $\text{while}(B) \{C\}$ as introduced in Definition 3.1. Then,

$$\begin{aligned}
& \text{ert}[\text{while}(B) \{C\}](t) \\
&= \text{lfp } F_t \\
&= F_t(\text{lfp } F_t) && \text{(Def. lfp)} \\
&= 1 + \llbracket B : \text{true} \rrbracket \cdot \text{ert}[C](\text{lfp } F_t) + \llbracket B : \text{false} \rrbracket \cdot t && \text{(Definition 3.1)} \\
&= 1 + \llbracket B : \text{true} \rrbracket \cdot \text{ert}[C](\text{ert}[\text{while}(B) \{C\}](t)) + \llbracket B : \text{false} \rrbracket \cdot t \\
&= 1 + \llbracket B : \text{true} \rrbracket \cdot \text{ert}[C](\text{ert}[\text{while}(B) \{C\}](t)) + \llbracket B : \text{false} \rrbracket \cdot \text{ert}[\text{empty}](t) \\
&= 1 + \llbracket B : \text{true} \rrbracket \cdot \text{ert}[C; \text{while}(B) \{C\}](t) + \llbracket B : \text{false} \rrbracket \cdot \text{ert}[\text{empty}](t) \\
&= \text{ert}[\text{if}(B) \{C; \text{while}(B) \{C\}\} \text{else} \{\text{empty}\}](t). \quad \square
\end{aligned}$$

Moreover, we observe that the runtime of two sequentially composed *terminating* deterministic programs C_1, C_2 can be decomposed into the sum of the individual runtimes of C_1 and C_2 . To that end, we need the following. The MC $\mathcal{M}_\sigma^0 \llbracket C \rrbracket$ of a deterministic program C and a program state $\sigma \in \Sigma$ (cf. Definition 5.3) reduces to a labeled transition system. In particular, if a state $\langle \downarrow, \sigma \rangle$ indicating successful termination is reachable from the initial state of $\mathcal{M}_\sigma^0 \llbracket C \rrbracket$, it is unique. Hence, we capture the effect of a deterministic program by a partial function $\mathbb{C} \llbracket \cdot \rrbracket (\cdot) : \text{GCL} \times \Sigma \rightarrow \Sigma$ mapping each deterministic program $C \in \text{GCL}$ and each program state $\sigma \in \Sigma$ to a program state $\sigma' \in \Sigma$ if and only if there exists a state $\langle \downarrow, \sigma' \rangle$ that is reachable in the MC $\mathcal{M}_\sigma^0 \llbracket C \rrbracket$ from its initial state $\langle \text{init}(C), \sigma \rangle$. Otherwise, $\mathbb{C} \llbracket C \rrbracket (\sigma)$ is undefined.

LEMMA B.6. *Let $C_1 \in \text{GCL}$ terminate on program state $\sigma \in \Sigma$ and $C_2 \in \text{GCL}$ terminate on $\mathbb{C} \llbracket C_1 \rrbracket (\sigma)$. Then,*

$$\text{ert}[C_1; C_2](\mathbf{0})(\sigma) = \text{ert}[C_1](\mathbf{0})(\sigma) + \text{ert}[C_2](\mathbf{0})(\mathbb{C} \llbracket C_1 \rrbracket (\sigma)).$$

PROOF. Immediate by inspection of the MC $\mathcal{M}_\sigma^0 \llbracket C_1; C_2 \rrbracket$ and Theorem 5.5. \square

Our third lemma extends the previous decomposition result to while-loops.

LEMMA B.7. *Let B be a deterministic guard and $C \in \text{GCL}$. For all program states $\sigma \in \Sigma$ with $\llbracket B \rrbracket(\sigma) = \text{true}$ such that $\text{while}(B) \{C\}$ terminates on σ , we have*

$$\text{ert}[\text{while}(B) \{C\}](\mathbf{0})(\sigma) = 1 + \text{ert}[C](\mathbf{0})(\sigma) + \text{ert}[\text{while}(B) \{C\}](\mathbf{0})(\mathbb{C}\llbracket C \rrbracket(\sigma)).$$

PROOF.

$$\begin{aligned} & \text{ert}[\text{while}(B) \{C\}](\mathbf{0})(\sigma) \\ &= \text{ert}[\text{if}(B) \{C; \text{while}(B) \{C\}\} \text{ else } \{\text{empty}\}](\mathbf{0})(\sigma) && \text{(Lemma B.5)} \\ &= 1 + \llbracket B \rrbracket(\sigma) \cdot \text{ert}[C; \text{while}(B) \{C\}](\mathbf{0})(\sigma) + \llbracket \neg B \rrbracket(\sigma) \cdot 0 \\ &= 1 + \text{ert}[C; \text{while}(B) \{C\}](\mathbf{0})(\sigma) && (\llbracket B \rrbracket(\sigma) = \text{true}) \\ &= 1 + \text{ert}[C](\mathbf{0})(\sigma) + \text{ert}[\text{while}(B) \{C\}](\mathbf{0})(\mathbb{C}\llbracket C \rrbracket(\sigma)). && \text{(Lemma B.6)} \quad \square \end{aligned}$$

We are now in a position to prove the soundness of ert with respect to Nielson's proof system.

THEOREM 6.1 (SOUNDNESS OF ert FOR DETERMINISTIC PROGRAMS). *For all deterministic programs $C \in \text{GCL}$ and assertions P, Q , we have*

$$\vdash \{P\} C \{ \Downarrow Q \} \quad \text{implies} \quad \vdash_E \{P\} C \{ \text{ert}[C](\mathbf{0}) \Downarrow Q \}.$$

PROOF. By induction on the structure of GCL-program C . The base cases $C = \text{skip}$ and $C = x := E$ are immediate, because

$$\text{ert}[\text{skip}](\mathbf{0}) = \text{ert}[x := E](\mathbf{0}) = 1$$

and $\{P\} \text{skip} \{1 \Downarrow P\}$ as well as $\{P\} x := E \{1 \Downarrow P\}$ are axioms.

Induction Hypothesis. Assume that for each substatement C' of C and each pair of assertions P, Q , we have

$$\vdash \{P\} C' \{ \Downarrow Q \} \text{ implies } \vdash_E \{P\} C' \{ \text{ert}[C'](\mathbf{0}) \Downarrow Q \}.$$

For the induction step, we have to consider sequential composition, conditionals, and loops.

Sequential Composition $C' = C_1; C_2$. Assume that

$$\vdash \{P\} C_1; C_2 \{ \Downarrow Q \}.$$

Then, there exists an assertion R such that

$$\vdash \{P\} C_1 \{ \Downarrow R \} \text{ and } \vdash \{R\} C_2 \{ \Downarrow Q \}.$$

By induction hypothesis, we know that

$$\vdash_E \{P\} C_1 \{ \text{ert}[C_1](\mathbf{0}) \Downarrow R \} \text{ and } \vdash_E \{R\} C_2 \{ \text{ert}[C_2](\mathbf{0}) \Downarrow Q \}.$$

Now, let

$$E'_2 \triangleq \text{ert}[C_1; C_2](\mathbf{0}) - \text{ert}[C_1](\mathbf{0})$$

and consider the triple

$$\{P \wedge E'_2 = u\} C_1 \{ \text{ert}[C_1](\mathbf{0}) \Downarrow R \wedge \text{ert}[C_2](\mathbf{0}) \leq u \},$$

where u is a fresh logical variable. Since u does not occur in P , each state $\sigma \in \Sigma$ satisfying P also satisfies $P \wedge E'_2 = u$; the latter is denoted by $\sigma \models P \wedge E'_2 = u$. Then,

$$\begin{aligned}
& \sigma \models P \wedge E'_2 = u \text{ and } \mathbb{C}[\![C_1]\!](\sigma) \models R \text{ and } \text{ert}[C_2](\mathbf{0})(\mathbb{C}[\![C_1]\!](\sigma)) \leq u \\
& \Leftrightarrow \sigma \models P \wedge E'_2 = u \text{ and } \mathbb{C}[\![C_1]\!](\sigma) \models R \text{ and} \\
& \quad \text{ert}[C_2](\mathbf{0})(\mathbb{C}[\![C_1]\!](\sigma)) \leq \text{ert}[C_1; C_2](\mathbf{0})(\sigma) - \text{ert}[C_1](\mathbf{0})(\sigma) \quad (\text{Definition of } E'_2) \\
& \Leftrightarrow \sigma \models P \wedge E'_2 = u \text{ and } \mathbb{C}[\![C_1]\!](\sigma) \models R \text{ and} \\
& \quad \text{ert}[C_1](\mathbf{0})(\sigma) + \text{ert}[C_2](\mathbf{0})(\mathbb{C}[\![C_1]\!](\sigma)) \leq \text{ert}[C_1; C_2](\mathbf{0})(\sigma) \quad (\text{Lemma B.6}) \\
& \Rightarrow \models_E \{ P \wedge E'_2 = u \} C_1 \{ \text{ert}[C_1](\mathbf{0}) \Downarrow R \wedge \text{ert}[C_2](\mathbf{0}) \leq u \}.
\end{aligned}$$

Since both

$$\{ P \wedge E'_2 = u \} C_1 \{ \text{ert}[C_1](\mathbf{0}) \Downarrow R \wedge \text{ert}[C_2](\mathbf{0}) \leq u \}$$

and

$$\{ R \} C_2 \{ \text{ert}[C_2](\mathbf{0}) \Downarrow Q \}$$

are valid triples, we may apply the rule of sequential composition to conclude

$$\begin{aligned}
& \vdash_E \{ P \} C_1; C_2 \{ \text{ert}[C_1](\mathbf{0}) + E'_2 \Downarrow Q \} \\
& \Leftrightarrow \vdash_E \{ P \} C_1; C_2 \{ \text{ert}[C_1; C_2](\mathbf{0}) \Downarrow Q \}. \quad (\text{Definition of } E'_2)
\end{aligned}$$

Conditionals. $C' = \text{if } (B) \{C_1\} \text{ else } \{C_2\}$. Assume that

$$\vdash \{ P \} \text{if } (B) \{C_1\} \text{ else } \{C_2\} \{ \Downarrow Q \}$$

is a provable triple in Hoare logic. Then, also,

$$\vdash \{ P \wedge B \} C_1 \{ \Downarrow Q \} \text{ and } \vdash \{ P \wedge \neg B \} C_2 \{ \Downarrow Q \}.$$

By induction hypothesis, it follows that

$$\vdash_E \{ P \wedge B \} C_1 \{ \text{ert}[C_1](\mathbf{0}) \Downarrow Q \} \text{ and } \vdash_E \{ P \wedge \neg B \} C_2 \{ \text{ert}[C_2](\mathbf{0}) \Downarrow Q \}$$

are provable triples in Nielson's logic. Now, let

$$\begin{aligned}
E & \triangleq \text{ert}[\text{if } (B) \{C_1\} \text{ else } \{C_2\}](\mathbf{0}) \\
& = 1 + \llbracket B \rrbracket \cdot \text{ert}[C_1](\mathbf{0}) + \llbracket \neg B \rrbracket \cdot \text{ert}[C_2](\mathbf{0}).
\end{aligned}$$

Then $E \geq \text{ert}[C_1](\mathbf{0})$ and $E \geq \text{ert}[C_2](\mathbf{0})$ hold for all states satisfying the precondition P . Hence, we can apply the rule of consequence to conclude

$$\vdash_E \{ P \wedge B \} C_1 \{ E \Downarrow Q \} \text{ and } \vdash_E \{ P \wedge \neg B \} C_2 \{ E \Downarrow Q \}.$$

Now, applying the rule for conditionals yields

$$\vdash_E \{ P \} \text{if } (B) \{C_1\} \text{ else } \{C_2\} \{ E \Downarrow Q \}.$$

Loops. $C' = \text{while } (B) \{C_1\}$. Assume that

$$\vdash \{ P \} \text{while } (B) \{C_1\} \{ \Downarrow Q \}$$

is a provable triple. Then, there exists an assertion $R(z)$ such that $P \Rightarrow \exists z. R(z)$, $R(0) \Rightarrow Q$ and

$$\vdash \{ \exists z. R(z) \} \text{while } (B) \{C_1\} \{ \Downarrow R(0) \}.$$

By the while-rule of Hoare logic for total correctness, we have

$$\vdash \{ R(z+1) \} C_1 \{ \Downarrow R(z) \}. \quad (*)$$

The induction hypothesis yields

$$\vdash_E \{ R(z+1) \} C_1 \{ E_1 \Downarrow R(z) \},$$

where $E_1 = \text{ert}[C_1](0)$. Now, let

$$E' \triangleq \text{ert}[\text{while } (B) \{C_1\}](0) - \text{ert}[C_1](0)$$

and

$$E \triangleq \text{ert}[\text{while } (B) \{C_1\}](0).$$

Our goal is to apply the while-rule in order to show

$$\vdash_E \{ \exists z. R(z) \} \text{ while } (B) \{C_1\} \{ E \Downarrow R(0) \}.$$

We first check the side conditions of this rule for our choice of E' ; i.e., we show $R(0) \Rightarrow \neg B \wedge E \geq 1$ as well as $R(z+1) \Rightarrow B \wedge E \geq E_1 + E'$.

If $R(0)$ is valid, then $\neg B$ is valid due to $(*)$ and

$$\llbracket E \rrbracket(\sigma) = \text{ert}[\text{while } (B) \{C_1\}](0)(\sigma) = 1 \geq 1.$$

Furthermore, if $R(z+1)$ is valid for some $z \in \mathbb{N}$, then B is valid by $(*)$ and for each program state $\sigma \in \Sigma$, we have

$$\begin{aligned} \llbracket E \rrbracket(\sigma) &= \text{ert}[\text{while } (B) \{C_1\}](0)(\sigma) \\ &= \text{ert}[\text{while } (B) \{C_1\}](0)(\sigma) - \text{ert}[C_1](0)(\sigma) + \text{ert}[C_1](0)(\sigma) \\ &= \llbracket E' \rrbracket(\sigma) + \llbracket E_1 \rrbracket(\sigma). \end{aligned}$$

Hence, the side conditions of the while-rule hold. In order to apply the rule, we also have to show validity of the triple

$$\{ R(z+1) + E' = u \} C_1 \{ \text{ert}[C_1](0) \Downarrow R(z) \wedge E \leq u \},$$

where u is a fresh logical variable. Since u does not occur in P , we know that for each state $\sigma \in \Sigma$ with $\sigma \models P$, we also have $\sigma \models R(z+1) \wedge E' = u$. Then

$$\begin{aligned} &\sigma \models R(z+1) \wedge E' = u \text{ and } \mathbb{C}\llbracket C_1 \rrbracket(\sigma) \models R(z) \text{ and } E(\mathbb{C}\llbracket C_1 \rrbracket(\sigma)) \leq u \\ \Leftrightarrow &\sigma \models R(z+1) \wedge E' = u \text{ and } \mathbb{C}\llbracket C_1 \rrbracket(\sigma) \models R(z) \text{ and } E(\mathbb{C}\llbracket C_1 \rrbracket(\sigma)) \leq E'(\sigma) \quad (\text{Def. of } u) \\ \Leftrightarrow &\sigma \models R(z+1) \wedge E' = u \text{ and } \mathbb{C}\llbracket C_1 \rrbracket(\sigma) \models R(z) \text{ and} \\ &E(\mathbb{C}\llbracket C_1 \rrbracket(\sigma)) \leq \text{ert}[\text{while } (B) \{C_1\}](0)(\sigma) - \text{ert}[C_1](0)(\sigma) \quad (\text{Definition of } E') \\ \Leftrightarrow &\sigma \models R(z+1) \wedge E' = u \text{ and } \mathbb{C}\llbracket C_1 \rrbracket(\sigma) \models R(z) \text{ and} \\ &E(\mathbb{C}\llbracket C_1 \rrbracket(\sigma)) + \text{ert}[C_1](0)(\sigma) \leq \text{ert}[\text{while } (B) \{C_1\}](0)(\sigma) \\ \Rightarrow &\vdash_E \{ R(z+1) \wedge E' = u \} C_1 \{ \text{ert}[C_1](0) \Downarrow R(z) \wedge E \leq u \}. \quad (\text{Lemma B.7}) \end{aligned}$$

Thus, we may apply the while-rule to conclude

$$\vdash_E \{ \exists z. R(z) \} \text{ while } (B) \{C_1\} \{ \text{ert}[\text{while } (B) \{C_1\}](0) \Downarrow R(0) \}.$$

By assumption, the implications $P \Rightarrow \exists z. R(z)$ as well as $R(0) \Rightarrow Q$ are valid, i.e.,

$$\vdash_E \{ P \} \text{ while } (B) \{C_1\} \{ \text{ert}[\text{while } (B) \{C_1\}](0) \Downarrow Q \}. \quad \square$$

B.5 Proof of Theorem 6.2 (Completeness of ert w.r.t. Nielson's Proof System)

We show the following claim by induction on the structure of program statements: For all assertions P, Q , deterministic program $C \in \text{GCL}$, and arithmetic expressions E ,

$$\vdash_E \{ P \} C \{ E \Downarrow Q \}$$

implies that there exists a natural number $k \in \mathbb{N}$ such that for all program states $\sigma \in \Sigma$, we have

$$\text{ert}[C](0)(\sigma) \leq k \cdot \llbracket E \rrbracket(\sigma).$$

The Effectless Program $C = \text{skip}$. Assume

$$\vdash_E \{P\} \text{skip} \{E \Downarrow Q\}.$$

Then there exists an assertion R such that

$$\vdash_E \{R\} \text{skip} \{1 \Downarrow R\}$$

and $P \Rightarrow R \wedge 1 \leq k \cdot E$ and $R \Rightarrow Q$ are valid for some $k \in \mathbb{N}$. Hence, there exists a $k \in \mathbb{N}$ such that $\text{ert}[\text{skip}](\mathbf{0})(\sigma) = 1 \leq k \cdot \llbracket E \rrbracket(\sigma)$ for each $\sigma \in \Sigma$.

The Assignment $C = x := E$. Analogous to skip.

The Sequential Composition $C = C_1; C_2$. Assume

$$\vdash_E \{P\} C_1; C_2 \{E \Downarrow Q\}.$$

Then there exists an assertion R and arithmetic expressions E_1, E_2, E'_2 such that

$$\vdash_E \{P \wedge E'_2 = u\} C_1 \{E_1 \Downarrow R \wedge E_2 \leq u\} \quad \text{and} \quad \vdash_E \{R\} C_2 \{E_2 \Downarrow Q\}$$

for some fresh logical variable u . By the induction hypothesis, there exist natural numbers $k_1, k_2 \in \mathbb{N}$ such that for all program states $\sigma \in \Sigma$, we have

$$\text{ert}[C_1](\mathbf{0})(\sigma) \leq k_1 \cdot \llbracket E_1 \rrbracket(\sigma) \quad \text{and} \quad \text{ert}[C_2](\mathbf{0})(\sigma) \leq k_2 \cdot \llbracket E_2 \rrbracket(\sigma).$$

By setting $k = \max\{k_1, k_2\}$, we obtain

$$\text{ert}[C_1](\mathbf{0})(\sigma) \leq k \cdot \llbracket E_1 \rrbracket(\sigma) \quad \text{and} \quad \text{ert}[C_2](\mathbf{0})(\sigma) \leq k \cdot \llbracket E_2 \rrbracket(\sigma). \quad (*)$$

In particular, this means that

$$\text{ert}[C_2](\mathbf{0})(\mathbb{C}\llbracket C_1 \rrbracket(\sigma)) \leq k \cdot \llbracket E_2 \rrbracket(\mathbb{C}\llbracket C_1 \rrbracket(\sigma)) \leq k \cdot \llbracket E'_2 \rrbracket(\sigma). \quad (\dagger)$$

Hence,

$$\begin{aligned} & \text{ert}[C_1; C_2](\mathbf{0})(\sigma) \\ &= \text{ert}[C_1](\mathbf{0})(\sigma) + \text{ert}[C_2](\mathbf{0})(\mathbb{C}\llbracket C_1 \rrbracket(\sigma)) && \text{(Lemma B.6)} \\ &\leq k \cdot \llbracket E_1 \rrbracket(\sigma) + \text{ert}[C_2](\mathbf{0})(\mathbb{C}\llbracket C_1 \rrbracket(\sigma)) && \text{(by *)} \\ &\leq k \cdot \llbracket E_1 \rrbracket(\sigma) + k \cdot \llbracket E'_2 \rrbracket(\sigma) && \text{(by } \dagger) \\ &= k \cdot (\llbracket E_1 \rrbracket(\sigma) + \llbracket E'_2 \rrbracket(\sigma)). \end{aligned}$$

Conditionals $C' = \text{if } (B) \{C_1\} \text{ else } \{C_2\}$. Assume

$$\vdash_E \{P\} \text{if } (B) \{C_1\} \text{ else } \{C_2\} \{E \Downarrow Q\}.$$

Then

$$\vdash_E \{P \wedge B\} C_1 \{E \Downarrow Q\} \quad \text{and} \quad \vdash_E \{P \wedge \neg B\} C_2 \{E \Downarrow Q\}.$$

By induction hypothesis, there exist natural numbers $k_1, k_2 \in \mathbb{N}$ such that for each $\sigma \in \Sigma$ and $i \in \{1, 2\}$,

$$\text{ert}[C_i](\mathbf{0})(\sigma) \leq k_i \cdot \llbracket E \rrbracket(\sigma).$$

By setting $k = \max\{k_1, k_2\}$, we obtain

$$\text{ert}[C_i](\mathbf{0})(\sigma) \leq k \cdot \llbracket E \rrbracket(\sigma). \quad (**)$$

Thus,

$$\begin{aligned}
& \text{ert}[\text{if } (B) \{C_1\} \text{ else } \{C_2\}](0)(\sigma) \\
&= 1 + \llbracket B \rrbracket(\sigma) \cdot \text{ert}[C_1](0)(\sigma) + \llbracket \neg B \rrbracket(\sigma) \cdot \text{ert}[C_2](0)(\sigma) && \text{(Table 1)} \\
&\leq k + \llbracket B \rrbracket(\sigma) \cdot k \cdot \llbracket E \rrbracket(\sigma) + \llbracket \neg B \rrbracket(\sigma) \cdot k \cdot \llbracket E \rrbracket(\sigma) && \text{(by **)} \\
&\leq (3 \cdot k) \cdot \llbracket E \rrbracket(\sigma).
\end{aligned}$$

Loops. $C' = \text{while } (B) \{C_1\}$. Assume

$$\vdash_E \{P\} \text{ while } (B) \{C_1\} \{E \Downarrow Q\}.$$

Then there exists an assertion $R(z)$ such that $P \Rightarrow \exists z. R(z)$ and $R(0) \Rightarrow Q$ are valid. Furthermore, there exists $z \in \mathbb{N}$ such that

$$\vdash_E \{R(z+1) \wedge E' = u\} C_1 \{E_1 \Downarrow R(z) \wedge E \leq u\} \quad (\ddagger)$$

for some fresh logical variable u . Additionally, for each $z \in \mathbb{N}$, the side conditions

$$R(z+1) \Rightarrow B \wedge E \geq E_1 + E' \text{ as well as } R(0) \Rightarrow \neg B \wedge E \geq 1 \quad (\dagger\dagger)$$

are valid. Our proof obligation is to show for some $k \in \mathbb{N}$ and all program states $\sigma \in \Sigma$ satisfying P that

$$\text{ert}[\text{while } (B) \{C_1\}](0)(\sigma) \leq k \cdot \llbracket E \rrbracket(\sigma). \quad (\spadesuit)$$

By induction hypothesis, there exists a $k' \in \mathbb{N}$ such that for each $\sigma \in \Sigma$, we have

$$1 \leq \text{ert}[C_1](0)(\sigma) \leq k' \cdot \llbracket E_1 \rrbracket(\sigma). \quad (\clubsuit)$$

We show by complete induction over $z \in \mathbb{N}$ that for all $\sigma \in \Sigma$ with $\sigma \models R(z)$ and

$$\vdash_E \{R(z)\} \text{ while } (B) \{C_1\} \{E \Downarrow R(0)\},$$

we have $\text{ert}[\text{while } (B) \{C_1\}](0)(\sigma) \leq (k'+1) \cdot \llbracket E \rrbracket(\sigma)$.

For the base case $z=0$, the side condition $R(0) \Rightarrow \neg B \wedge E \geq 1$ yields

$$\begin{aligned}
& (k'+1) \cdot \llbracket E \rrbracket(\sigma) \\
&\geq 1 && \text{(by } \dagger\dagger) \\
&= 1 + \llbracket B \rrbracket(\sigma) \cdot \text{ert}[C_1](0)(\sigma) + \llbracket \neg B \rrbracket(\sigma) \cdot 0 \\
&= \text{ert}[\text{while } (B) \{C_1\}](0)(\sigma). && \text{(Table 1)}
\end{aligned}$$

Now, for the induction step, assume $\sigma \models R(z+1)$. Then the side condition $R(z+1) \Rightarrow B \wedge E \geq E_1 + E'$ yields

$$\begin{aligned}
& (k'+1) \cdot \llbracket E \rrbracket(\sigma) \\
&\geq (k'+1) \cdot (\llbracket E_1 \rrbracket(\sigma) + \llbracket E' \rrbracket(\sigma)) && \text{(by } \dagger\dagger) \\
&\geq (k'+1) \cdot \llbracket E_1 \rrbracket(\sigma) + (k'+1) \cdot \llbracket E \rrbracket(\mathbb{C}\llbracket C_1 \rrbracket(\sigma)) && \text{(Postcondition of } \ddagger) \\
&\geq (k'+1) \cdot \llbracket E_1 \rrbracket(\sigma) + \text{ert}[\text{while } (B) \{C_1\}](0)(\mathbb{C}\llbracket C_1 \rrbracket(\sigma)) && \text{(I.H., } \mathbb{C}\llbracket C_1 \rrbracket(\sigma) \models R(z)) \\
&\geq \llbracket E_1 \rrbracket(\sigma) + \text{ert}[C_1](0)(\sigma) + \text{ert}[\text{while } (B) \{C_1\}](0)(\mathbb{C}\llbracket C_1 \rrbracket(\sigma)) && \text{(by } \clubsuit) \\
&\geq 1 + \text{ert}[C_1](0)(\sigma) + \text{ert}[\text{while } (B) \{C_1\}](0)(\mathbb{C}\llbracket C_1 \rrbracket(\sigma)) \\
&= \text{ert}[\text{while } (B) \{C_1\}](0)(\sigma), && \text{(Lemma B.7)}
\end{aligned}$$

which completes the inner induction. Moreover, (\spadesuit) follows immediately by setting $k = k'+1$. \square

B.6 Well-Definedness of ert on Procedure Calls

We prove that the least fixed point

$$\text{lfp } \eta. \underline{\mathbf{1}} \oplus \text{ert}[\mathcal{D}(P)]_{\eta}^{\#} \quad (2)$$

defining $\text{ert}[\text{call } P, \mathcal{D}]$ is well defined. To this end we take the following four steps:

- (1) We endow the set RtEnv of runtime environments with the structure of an ω -cpo (Proposition B.8).
- (2) We show that transformer $\text{ert}[C]_{\eta}^{\#}(t)$ is continuous with respect to t for every $C \in \text{pRGCL}$ (Lemma B.9).
- (3) We show that the transformer $\text{ert}[C]_{\eta}^{\#}(t)$ is also continuous with respect to η (Lemma B.10).
- (4) We conclude that $\lambda \eta. \underline{\mathbf{1}} \oplus \text{ert}[C]_{\eta}^{\#}$ is continuous and by the Kleene Fixed-Point Theorem (Theorem A.3) the least fixed point in Equation (2) is well defined.

Let “ \sqsubseteq ” denote the pointwise order between runtime environments; i.e., for $\eta_1, \eta_2 \in \text{RtEnv}$, $\eta_1 \sqsubseteq \eta_2$ if and only if $\eta_1(t) \leq \eta_2(t)$ for every $t \in \mathbb{T}$.

PROPOSITION B.8. *The pair $(\text{RtEnv}, \sqsubseteq)$ is an ω -cpo with bottom element $\perp_{\text{RtEnv}} \triangleq \lambda t : \mathbb{T}. \mathbf{0}$, where the supremum of an ω -chain $\eta_0 \sqsubseteq \eta_1 \sqsubseteq \dots$ is given pointwise, i.e., $(\sup_n \eta_n)(t) \triangleq \sup_n \eta_n(t)$.*

LEMMA B.9 (CONTINUITY OF $\text{ert}[C]_{\eta}^{\#}(t)$ W.R.T. t). *Let $t_0 \leq t_1 \leq \dots$ be an ω -chain in \mathbb{T} . Then for every $C \in \text{pRGCL}$ and $\eta \in \text{RtEnv}$,*

$$\text{ert}[C]_{\eta}^{\#}(\sup_n t_n) = \sup_n \text{ert}[C]_{\eta}^{\#}(t_n).$$

PROOF. By induction on the structure of C . Except for procedure calls, all program constructs use the same proof argument as for the continuity of plain transformer $\text{ert}[\cdot]$ (see the proof of Lemma 3.2). For procedure calls, the statement follows immediately from the continuity of η since

$$\text{ert}[\text{call } P]_{\eta}^{\#}(\sup_n t_n) = \eta(\sup_n t_n) = \sup_n \eta(t_n) = \sup_n \text{ert}[\text{call } P]_{\eta}^{\#}(t_n). \quad \square$$

LEMMA B.10 (CONTINUITY OF $\text{ert}[C]_{\eta}^{\#}(t)$ W.R.T. η). *Let $\eta_0 \sqsubseteq \eta_1 \sqsubseteq \dots$ be an ω -chain in RtEnv . Then, for every $C \in \text{pRGCL}$,*

$$\text{ert}[C]_{\sup_n \eta_n}^{\#} = \sup_n \text{ert}[C]_{\eta_n}^{\#}.$$

PROOF. By induction on the structure of C . For the four basic instructions `empty`, `skip`, `halt`, and `$x := \mu$` the proof is straightforward since the action of the transformer is independent of the runtime environment (i.e., constant functions are always continuous). We omit the case of `while`-loops since they can be readily simulated by recursive procedures. For the remaining program constructs we reason as follows:

Sequential Composition. We use the fact that $\text{ert}[C_1]_{\sup_m \eta_m}^{\#}(f)$ is continuous in f (by Lemma B.9) and Lemma A.5 to “merge” a pair of nested supremums into a single supremum:

$$\begin{aligned} \text{ert}[C_1; C_2]_{\sup_n \eta_n}^{\#}(t) &= \text{ert}[C_1]_{\sup_m \eta_m}^{\#}(\text{ert}[C_2]_{\sup_n \eta_n}^{\#}(t)) && \text{(Table 3)} \\ &= \text{ert}[C_1]_{\sup_m \eta_m}^{\#}(\sup_n \text{ert}[C_2]_{\eta_n}^{\#}(t)) && \text{(I.H. on } C_2) \\ &= \sup_n \text{ert}[C_1]_{\sup_m \eta_m}^{\#}(\text{ert}[C_2]_{\eta_n}^{\#}(t)) && \text{(Lemma B.9)} \\ &= \sup_n \sup_m \text{ert}[C_1]_{\eta_m}^{\#}(\text{ert}[C_2]_{\eta_n}^{\#}(t)) && \text{(I.H. on } C_1) \\ &= \sup_i \text{ert}[C_1]_{\eta_i}^{\#}(\text{ert}[C_2]_{\eta_i}^{\#}(t)) && \text{(Lemma A.5)} \\ &= \sup_i \text{ert}[C_1; C_2]_{\eta_i}^{\#}(t). && \text{(Table 3)} \end{aligned}$$

Conditional Choice. The idea here is to substitute $\sup_n \text{ert}[C_1]_{\eta_n}^\#$ with $\lim_{n \rightarrow \infty} \text{ert}[C_1]_{\eta_n}^\#$ using the Monotone Sequence Theorem. This is possible because by I.H., $\text{ert}[C_1]_{\eta_n}^\#$ is (continuous and thus) monotonic in η_n and $\eta_0 \sqsubseteq \eta_1 \sqsubseteq \dots$; therefore, $(\text{ert}[C_1]_{\eta_n}^\#(t))_{n \in \mathbb{N}}$ defines a monotonic sequence:

$$\begin{aligned}
& \text{ert}[\text{if } (\xi) \{C_1\} \text{ else } \{C_2\}]_{\sup_n \eta_n}^\#(t) \\
&= \llbracket \xi : \text{true} \rrbracket \cdot \text{ert}[C_1]_{\sup_n \eta_n}^\#(t) + \llbracket \xi : \text{false} \rrbracket \cdot \text{ert}[C_2]_{\sup_n \eta_n}^\#(t) && \text{(Table 3)} \\
&= \llbracket \xi : \text{true} \rrbracket \cdot \sup_n \text{ert}[C_1]_{\eta_n}^\#(t) + \llbracket \xi : \text{false} \rrbracket \cdot \sup_n \text{ert}[C_2]_{\eta_n}^\#(t) && \text{(I.H. on } C_1, C_2) \\
&= \llbracket \xi : \text{true} \rrbracket \cdot \lim_{n \rightarrow \infty} \text{ert}[C_1]_{\eta_n}^\#(t) + \llbracket \xi : \text{false} \rrbracket \cdot \lim_{n \rightarrow \infty} \text{ert}[C_2]_{\eta_n}^\#(t) && \text{(Theorem A.2)} \\
&= \lim_{n \rightarrow \infty} (\llbracket \xi : \text{true} \rrbracket \cdot \text{ert}[C_1]_{\eta_n}^\#(t) + \llbracket \xi : \text{false} \rrbracket \cdot \text{ert}[C_2]_{\eta_n}^\#(t)) && \text{(algebra of limits)} \\
&= \sup_n (\llbracket \xi : \text{true} \rrbracket \cdot \text{ert}[C_1]_{\eta_n}^\#(t) + \llbracket \xi : \text{false} \rrbracket \cdot \text{ert}[C_2]_{\eta_n}^\#(t)) && \text{(Theorem A.2)} \\
&= \sup_n \text{ert}[\text{if } (\xi) \{C_1\} \text{ else } \{C_2\}]_{\eta_n}^\#. && \text{(Table 3)}
\end{aligned}$$

Procedure Call. The reasoning here is straightforward:

$$\begin{aligned}
\text{ert}[\text{call } P]_{\sup_n \eta_n}^\#(t) &= (\sup_n \eta_n)(t) && \text{(Table 3)} \\
&= \sup_n \eta_n(t) && \text{(def. } \sup_n \eta_n) \\
&= \sup_n \text{ert}[\text{call } P]_{\eta_n}^\#(t). && \text{(Table 3)}
\end{aligned}$$

□

B.7 Proof of Theorem 7.2 (Continuity of ert for Recursive Programs)

We prove that for pRGCL program $\langle C, \mathcal{D} \rangle$ and ω -chain of runtimes $t_0 \leq t_1 \leq \dots$,

$$\sup_n \text{ert}[C, \mathcal{D}](f_n) = \text{ert}[C, \mathcal{D}](\sup_n f_n).$$

The proof proceeds by induction on the structure of C . We consider only the case of procedure calls since all other language constructs follow the same reasoning as in the proof of Lemma 3.2 (see Appendix B.1), the only difference being that the transformer now propagates declarations. For $\text{ert}[\text{call } P, \mathcal{D}]$, let $F(\eta) = \mathbf{1} \oplus \text{ert}[\mathcal{D}(P)]_{\eta}^\#$. Since F is continuous (see Appendix B.6), we can apply the Kleene Fixed-Point Theorem to characterize $\text{lfp } \eta \cdot F(\eta)$ as $\sup_m F^m(\perp_{\text{RtEnv}})$, and we have

$$\begin{aligned}
\text{ert}[\text{call } P, \mathcal{D}](\sup_n t_n) &= (\text{lfp } \eta \cdot F(\eta))(\sup_n t_n) && \text{(Equation (1))} \\
&= \sup_m F^m(\perp_{\text{RtEnv}})(\sup_n t_n) && \text{(Theorem A.3)} \\
&= \sup_m \sup_n F^m(\perp_{\text{RtEnv}})(t_n) && (F^m(\perp_{\text{RtEnv}}) \text{ continuous)} \\
&= \sup_n \sup_m F^m(\perp_{\text{RtEnv}})(t_n) && \text{(Theorem A.5)} \\
&= \sup_n (\text{lfp } \eta \cdot F(\eta))(t_n) && \text{(Theorem A.3)} \\
&= \sup_n \text{ert}[\text{call } P, \mathcal{D}](t_n). && \text{(Equation (1))}
\end{aligned}$$

We are only left to prove that $F^m(\perp_{\text{RtEnv}})$ is continuous for all $m \in \mathbb{N}$. We prove this by induction on m . The base case is immediate since $F^0(\perp_{\text{RtEnv}}) = \perp_{\text{RtEnv}}$ and \perp_{RtEnv} is continuous. For the inductive case we have $F^{m+1}(\perp_{\text{RtEnv}}) = F(F^m(\perp_{\text{RtEnv}}))$. The continuity of $F^{m+1}(\perp_{\text{RtEnv}})$ follows from the I.H. and the fact that F preserves continuity; i.e., η continuous implies $F(\eta)$ continuous. □

B.8 Proof of Theorem 7.2 (Constant Propagation of ert for Recursive Programs)

We prove that for halt-free pRGCL program $\langle C, \mathcal{D} \rangle$ and constant runtime \mathbf{k} with $k \in \mathbb{R}_{\geq 0}$,

$$\text{ert}[C, \mathcal{D}](\mathbf{k} + t) = \mathbf{k} + \text{ert}[C, \mathcal{D}](t). \quad (3)$$

The proof relies on two subsidiary results that we present below.

LEMMA B.11. For halt-free pRGCL program $\langle C, \mathcal{D} \rangle$ and constant runtime \mathbf{k} with $k \in \mathbb{R}_{\geq 0}$,

$$\text{ert}[C, \mathcal{D}](\mathbf{k}) \geq \mathbf{k}.$$

PROOF. By induction on the structure of C . Except for the case of procedure calls, all other program constructs pose no difficulty. For a procedure call, we must do a case distinction on whether the procedure terminates almost surely or not. This requires extending the weakest precondition transformer wp to probabilistic recursive programs. A detailed proof containing this case analysis is provided in [43, App. 6]. \square

For stating the second auxiliary result, we require the notion of “constant separable” runtime environment. We say that $\eta \in \text{RtEnv}$ is *constant separable into* $v \in \text{RtEnv}$ if and only if for all $k \in \mathbb{R}_{\geq 0}$ and $t \in \mathbb{T}$, $\eta(\mathbf{k} + t) = \mathbf{k} + v(t)$.

LEMMA B.12. Let η be a runtime environment constant separable¹⁰ into v . Then, for all halt-free $C \in \text{pRGCL}$ and constant runtime \mathbf{k} with $k \in \mathbb{R}_{\geq 0}$,

$$\text{ert}[C]_{\eta}^{\#}(\mathbf{k} + t) = \mathbf{k} + \text{ert}[C]_{v}^{\#}(t).$$

PROOF. By a routine induction on the structure of C . \square

Now we have all prerequisites to establish Equation (3). By letting $F(\eta) = \underline{1} \oplus \text{ert}[\mathcal{D}(P)]_{\eta}^{\#}$, we can recast the equation as $(\text{lfp } F)(\mathbf{k} + t) = \mathbf{k} + (\text{lfp } F)(t)$, or equivalently,

$$(\lambda \eta^{\star}. \lambda t^{\star}. \eta^{\star}(\mathbf{k} + t^{\star}))(\text{lfp } F) = (\lambda \eta^{\star}. \lambda t^{\star}. \mathbf{k} + \eta^{\star}(t^{\star}))(\text{lfp } F).$$

To prove this equation, we apply the Lemma A.6 with instantiations

$$\begin{aligned} F_1 &= F_2 = F \\ f_1 &= \lambda \eta^{\star}. \lambda t^{\star}. \eta^{\star}(\mathbf{k} + t^{\star}) \\ f_2 &= \lambda \eta^{\star}. \lambda t^{\star}. \mathbf{k} + \eta^{\star}(t^{\star}) \\ h_1 &= \lambda \eta^{\star}. \lambda t^{\star}. \mathbf{1} + \text{ert}[\mathcal{D}(P)]_{\lambda t^{\star}. \eta^{\star}(t^{\star} - \mathbf{k})}^{\#}(\mathbf{k} + t^{\star}) \\ h_2 &= \lambda \eta^{\star}. \lambda t^{\star}. \mathbf{k} + \mathbf{1} + \text{ert}[\mathcal{D}(P)]_{\lambda t^{\star}. \eta^{\star}(t^{\star}) - \mathbf{k}}^{\#}(t^{\star}) \end{aligned}$$

and underlying ω -cpos $(\mathcal{D}_1, \leq_1) = (\mathcal{D}_2, \leq_2) = (\mathcal{D}, \leq) = (\text{RtEnv}, \sqsubseteq)$ and bottom elements $\perp_1 = \perp_2 = \perp = \perp_{\text{RtEnv}}$. The application of Lemma A.6 requires the continuity of F , which follows from Lemma B.10; the continuity of f_1 and f_2 , which holds because runtime environments are continuous by definition; and finally the monotonicity of h_1 and h_2 . This latter fact, together with the fact that h_1 and h_2 are effectively well defined (i.e., have type $\text{RtEnv} \rightarrow \text{RtEnv}$), can be proved with an inductive argument (on the structure of $\mathcal{D}(P)$).

We are left to discharge hypotheses 1 through 3 of Lemma A.6. A simple unfolding of the involved functions yields $f_1(F(\eta)) \sqsubseteq h_1(f_1(\eta))$ and $f_2(F(\eta)) \sqsubseteq h_2(f_2(\eta))$ for all $\eta \in \text{RtEnv}$; this establishes hypothesis 1. As for hypothesis 2, $f_1(\perp_{\text{RtEnv}}) \sqsubseteq f_2(\text{lfp } F)$ holds because $f_1(\perp_{\text{RtEnv}}) = \perp_{\text{RtEnv}}$

¹⁰For the definition of *constant separable* runtime environment, see the paragraph above Lemma B.12.

and $f_2(\perp_{\text{RtEnv}}) \sqsubseteq f_1(\text{lfp } F)$ reduces to $\mathbf{k} \leq \text{ert}[\text{call } P, \mathcal{D}](\mathbf{k} + t)$, which holds in view of the monotonicity of ert and the auxiliary Lemma B.11. Finally, to discharge hypothesis 3, we let $\eta(t') = \mathbf{k} + \text{ert}[\text{call } P, \mathcal{D}](t' - \mathbf{k})$ and reason as follows:

$$\begin{aligned}
& h_1(f_2(\text{lfp } F))(t) \leq f_2(\text{lfp } F)(t) \\
& \iff \mathbf{1} + \text{ert}[\mathcal{D}(P)]_{\eta}^{\#}(\mathbf{k} + t) \leq \mathbf{k} + \text{ert}[\text{call } P, \mathcal{D}](t) && \text{(definition } h_1, f_2, F) \\
& \iff \mathbf{1} + \mathbf{k} + \text{ert}[\mathcal{D}(P)]_{\text{ert}[\text{call } P, \mathcal{D}]}^{\#}(t) \leq \mathbf{k} + \text{ert}[\text{call } P, \mathcal{D}](t) && \begin{array}{l} (\eta \text{ constant separable into}) \\ \text{(ert}[\text{call } P, \mathcal{D}]; \text{Lemma B.12)} \end{array} \\
& \iff \mathbf{k} + F(\text{ert}[\text{call } P, \mathcal{D}](t)) \leq \mathbf{k} + \text{ert}[\text{call } P, \mathcal{D}](t) && \text{(definition } F) \\
& \iff \mathbf{k} + F(\text{lfp } F)(t) \leq \mathbf{k} + (\text{lfp } F)(t) && \text{(Equation (1))} \\
& \iff \mathbf{k} + (\text{lfp } F)(t) \leq \mathbf{k} + (\text{lfp } F)(t) && \text{(definition } \text{lfp}) \\
& \iff \text{true} && \text{("} \leq \text{" is a partial order)}
\end{aligned}$$

To prove the other part of hypothesis 3, i.e., $h_2(f_1(\text{lfp } F))(t) \leq f_1(\text{lfp } F)(t)$, we follow a similar reasoning. \square

B.9 Proof of Theorem 7.3 (Proof Rules for Expected Runtimes of Recursive Programs)

We show that proof rules (1) and (2) from Theorem 7.3 reproduced below are sound with respect to the ert -calculus in pRGCL . Proof rule (3) follows the same argument as (2):

$$\begin{aligned}
& \frac{\text{ert}[\text{call } P](t) \leq \mathbf{1} + u \Vdash \text{ert}[\mathcal{D}(P)](t) \leq u}{\text{ert}[\text{call } P, \mathcal{D}](t) \leq \mathbf{1} + u} \quad (1) \\
& \quad \quad \quad u_0 = \mathbf{0} \\
& \frac{\text{ert}[\text{call } P](t) \leq \mathbf{1} + \text{ert}[\mathcal{D}(P)](t) \leq u_{n+1}}{\text{ert}[\text{call } P, \mathcal{D}](t) \leq \mathbf{1} + \lim_{n \rightarrow \infty} u_n}. \quad (2)
\end{aligned}$$

To prove the rules sound, we make use of the following result:

PROPOSITION B.13. *The derivability assertion*

$$\text{ert}[\text{call } P](t_1) \leq u_1 \Vdash \text{ert}[C](t_2) \leq u_2$$

implies that for every runtime environment η ,

$$\eta(t_1) \leq u_1 \Rightarrow \text{ert}[C]_{\eta}^{\#}(t_2) \leq u_2.$$

Soundness of Rule (1). Let runtime environment η^* map t to u and all other runtimes to (the constant runtime) ∞ . Then,

$$\begin{aligned}
& \text{ert}[\text{call } P, \mathcal{D}](t) \leq \mathbf{1} + u \\
& \iff (\text{lfp } \eta \cdot \mathbf{1} \oplus \text{ert}[\mathcal{D}(P)]_{\eta}^{\#})(t) \leq \mathbf{1} + u && \text{(Equation (1))} \\
& \iff \text{lfp } \eta \cdot \mathbf{1} \oplus \text{ert}[\mathcal{D}(P)]_{\eta}^{\#} \sqsubseteq \mathbf{1} \oplus \eta^* && \text{(definition } \eta^*, \sqsubseteq) \\
& \iff \mathbf{1} \oplus \text{ert}[\mathcal{D}(P)]_{\mathbf{1} \oplus \eta^*}^{\#} \sqsubseteq \mathbf{1} \oplus \eta^* && \text{(Theorem A.4, B.10)} \\
& \iff \mathbf{1} + \text{ert}[\mathcal{D}(P)]_{\mathbf{1} \oplus \eta^*}^{\#}(t) \leq \mathbf{1} + u && \text{(definition } \eta^*, \sqsubseteq) \\
& \iff \text{ert}[\mathcal{D}(P)]_{\mathbf{1} \oplus \eta^*}^{\#}(t) \leq u && \text{(algebra)} \\
& \iff (\mathbf{1} \oplus \eta^*)(t) \leq \mathbf{1} + u && \text{(Proposition B.13, rule premise)} \\
& \iff \text{true}. && \text{(definition } \eta^*)
\end{aligned}$$

Soundness of Rule (2). Let $F(\eta) = \mathbf{1} \oplus \text{ert}[\mathcal{D}(P)]_{\eta}^{\#}$. Since F is continuous (see Appendix B.6), we can apply the Kleene Fixed-Point Theorem (Theorem A.3) to characterize $\text{lfp } \eta \cdot F(\eta)$ as $\sup_n F^n(\perp_{\text{RtEnv}})$. This yields

$$\text{ert}[\text{call } P, \mathcal{D}](t) = \text{lfp } \eta \cdot F(\eta) = \sup_n F^n(\perp_{\text{RtEnv}})(t).$$

Moreover, the sequence $(F^n(\perp_{\text{RtEnv}})(t))_{n \in \mathbb{N}}$ is monotonic. Then by the Monotone Sequence Theorem (Theorem A.2),

$$\sup_n F^n(\perp_{\text{RtEnv}})(t) = \lim_{n \rightarrow \infty} F^n(\perp_{\text{RtEnv}})(t).$$

Combining the above two equations, we obtain

$$\begin{aligned} \text{ert}[\text{call } P, \mathcal{D}](t) &\leq \mathbf{1} + \lim_{n \rightarrow \infty} u_n \\ &\iff \lim_{n \rightarrow \infty} F^n(\perp_{\text{RtEnv}})(t) \leq \mathbf{1} + \lim_{n \rightarrow \infty} u_n \\ &\iff \forall n \cdot F^n(\perp_{\text{RtEnv}})(t) \leq \mathbf{1} + u_n. \end{aligned}$$

We prove the last statement by induction on n . The base case $F^0(\perp_{\text{RtEnv}})(t) \leq \mathbf{1} + u_0$ is immediate since $F^0(\perp_{\text{RtEnv}})(t) = \perp_{\text{RtEnv}}(t) = 0$. For the inductive case we have

$$\begin{aligned} F^{n+1}(\perp_{\text{RtEnv}})(t) &\leq \mathbf{1} + u_{n+1} \\ &\iff \mathbf{1} + \text{ert}[\mathcal{D}(P)]_{F^n(\perp_{\text{RtEnv}})}^{\#}(t) \leq \mathbf{1} + u_{n+1} && \text{(definition } F^{n+1}(\perp_{\text{RtEnv}})) \\ &\iff \text{ert}[\mathcal{D}(P)]_{F^n(\perp_{\text{RtEnv}})}^{\#}(t) \leq u_{n+1} && \text{(algebra)} \\ &\iff F^n(\perp_{\text{RtEnv}})(t) \leq \mathbf{1} + u_n && \text{(Proposition B.13, rule premise)} \\ &\iff \text{true}. && \text{(I.H.) } \square \end{aligned}$$

B.10 Proof of Theorem 7.10 (Soundness of ert w.r.t. Pushdown Markov Chains)

Let us fix procedures P_1, \dots, P_k and a declaration \mathcal{D} . As in the proof of Theorem 5.5, we prove the soundness of ert with respect to operational PMCs by induction on the structure of pRGCL-programs. Similar to bounded while-loops (cf. Definition B.2), this proof relies on expressing the runtime of procedure calls in terms of inlined call-free programs. Thus, recall from Section 7.4 the definition of bounded procedure calls $\text{call}_n^{\mathcal{D}} P$ (for multiple procedures):

Definition B.14 (Bounded procedure calls). The bounded procedure calls of $\text{call } P_i, 1 \leq i \leq k$, are given by

$$\begin{aligned} \text{call}_0^{\mathcal{D}} P_i &\triangleq \text{halt} \\ \text{call}_{n+1}^{\mathcal{D}} P_i &\triangleq \text{skip}; \mathcal{D}(P_i) \left[\text{call } P_1 / \text{call}_n^{\mathcal{D}} P_1, \dots, \text{call } P_k / \text{call}_n^{\mathcal{D}} P_k \right]. \end{aligned}$$

The main motivation to consider bounded procedure calls is to use them as runtime environments. This usage is formally expressed by the following auxiliary result.

LEMMA B.15. *For each $C \in \text{pRGCL}$, $1 \leq i \leq k$, and $t \in \mathbb{T}$, we have*

$$\text{ert}[C, \mathcal{D}](t) = \sup_{n \in \mathbb{N}} \text{ert}[C, \mathcal{D}]_{(\text{ert}[\text{call}_n^{\mathcal{D}} P_1], \dots, \text{ert}[\text{call}_n^{\mathcal{D}} P_k])}^{\#}(t).$$

PROOF. By induction on the structure of pRGCL programs. In all cases except for procedure calls, the claims follows immediately from the definition of $\text{ert}\cdot$ and $\text{ert}[\cdot]_{\eta}^{\#}(\cdot)$ and the induction

hypothesis for compound statements. For a procedure call $\text{call } P_i$, we have

$$\begin{aligned}
 & \text{ert}[\text{call } P_i, \mathcal{D}](t) \\
 &= \sup_{n \in \mathbb{N}} \text{ert} \left[\text{call}_n^{\mathcal{D}} P_i \right] (t) \quad (\text{Theorem 7.5}) \\
 &= \sup_{n \in \mathbb{N}} \text{ert} [\text{call } P_i, \mathcal{D}]_{(\text{ert}[\text{call}_n^{\mathcal{D}} P_1], \dots, \text{ert}[\text{call}_n^{\mathcal{D}} P_k])}^{\#} (t) \quad (\text{ert}[C]_{\text{ert}[C']}^{\#} = \text{ert}[C[\text{call } P/C']]).
 \end{aligned}$$

□

Now, consider a PMC ${}^n\langle \mathcal{P}_\sigma^f \llbracket C, \mathcal{D} \rrbracket \rangle$ that behaves exactly the same as the operational PMC $\mathcal{P}_\sigma^f \llbracket C, \mathcal{D} \rrbracket$ but counts the number of symbols currently on the stack. Moreover, if this number is exactly n and $\mathcal{P}_\sigma^f \llbracket C, \mathcal{D} \rrbracket$ would perform another push onto the stack, ${}^n\langle \mathcal{P}_\sigma^f \llbracket C, \mathcal{D} \rrbracket \rangle$ immediately moves to $\langle \text{sink} \rangle$ —without entering a state of the form $\langle \downarrow, \sigma' \rangle$ indicating successful termination first and without collecting reward for the procedure call. Intuitively, this means that the modified PMC *halts* its execution after encountering $n + 1$ procedure calls without returns. It is evident that:

$$\text{LEMMA B.16. } \text{ExpRew}^{\mathcal{P}_\sigma^f \llbracket C, \mathcal{D} \rrbracket} (\langle \text{sink} \rangle) = \sup_{n \in \mathbb{N}} \text{ExpRew}^{{}^n\langle \mathcal{P}_\sigma^f \llbracket C, \mathcal{D} \rrbracket \rangle} (\langle \text{sink} \rangle).$$

PROOF SKETCH. ${}^n\langle \mathcal{P}_\sigma^f \llbracket C, \mathcal{D} \rrbracket \rangle$ exhibits a partial behavior of $\mathcal{P}_\sigma^f \llbracket C, \mathcal{D} \rrbracket$ in the sense that every path of ${}^n\langle \mathcal{P}_\sigma^f \llbracket C, \mathcal{D} \rrbracket \rangle$ eventually reaching $\langle \text{sink} \rangle$ is—up to renaming—also a path of $\mathcal{P}_\sigma^f \llbracket C, \mathcal{D} \rrbracket$. Hence,

$$\text{ExpRew}^{\mathcal{P}_\sigma^f \llbracket C, \mathcal{D} \rrbracket} (\langle \text{sink} \rangle) \leq \sup_{n \in \mathbb{N}} \text{ExpRew}^{{}^n\langle \mathcal{P}_\sigma^f \llbracket C, \mathcal{D} \rrbracket \rangle} (\langle \text{sink} \rangle).$$

Conversely, every finite path π of $\mathcal{P}_\sigma^f \llbracket C, \mathcal{D} \rrbracket$ eventually reaching $\langle \text{sink} \rangle$ can be implemented with finite stack size. Therefore, there exists an $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$ the path π of $\mathcal{P}_\sigma^f \llbracket C, \mathcal{D} \rrbracket$ is also a path of ${}^n\langle \mathcal{P}_\sigma^f \llbracket C, \mathcal{D} \rrbracket \rangle$. Thus,

$$\text{ExpRew}^{\mathcal{P}_\sigma^f \llbracket C, \mathcal{D} \rrbracket} (\langle \text{sink} \rangle) \geq \sup_{n \in \mathbb{N}} \text{ExpRew}^{{}^n\langle \mathcal{P}_\sigma^f \llbracket C, \mathcal{D} \rrbracket \rangle} (\langle \text{sink} \rangle). \quad \square$$

Assume for the moment the following result:

LEMMA B.17. *For all $n \in \mathbb{N}$ it holds that*

$$\lambda \sigma. \text{ExpRew}^{{}^n\langle \mathcal{P}_\sigma^t \llbracket C, \mathcal{D} \rrbracket \rangle} (\langle \text{sink} \rangle) = \text{ert}[C, \mathcal{D}]_{(\text{ert}[\text{call}_n^{\mathcal{D}} P_1], \dots, \text{ert}[\text{call}_n^{\mathcal{D}} P_k])}^{\#} (t).$$

With these auxiliary results readily available, we are in a position to show the transformer ert for pRGCL-programs to be sound with respect to operational PMCs.

THEOREM 7.10 (CORRESPONDENCE THEOREM). *Let $\langle C, \mathcal{D} \rangle$ be a pRGCL program, $t \in \mathbb{T}$. Then for each $\sigma \in \Sigma$, we have*

$$\text{ExpRew}^{\mathcal{P}_\sigma^t \llbracket C, \mathcal{D} \rrbracket} (\langle \text{sink} \rangle) = \text{ert}[C, \mathcal{D}](t)(\sigma).$$

PROOF.

$$\begin{aligned}
 & \text{ExpRew}^{\mathcal{P}_\sigma^t \llbracket C, \mathcal{D} \rrbracket} (\langle \text{sink} \rangle) \\
 &= \sup_{n \in \mathbb{N}} \text{ExpRew}^{{}^n\langle \mathcal{P}_\sigma^t \llbracket C, \mathcal{D} \rrbracket \rangle} (\langle \text{sink} \rangle) \quad (\text{Lemma B.16}) \\
 &= \sup_{n \in \mathbb{N}} \text{ert}[C, \mathcal{D}]_{(\text{ert}[\text{call}_n^{\mathcal{D}} P_1], \dots, \text{ert}[\text{call}_n^{\mathcal{D}} P_k])}^{\#} (t) \quad (\text{Lemma B.17}) \\
 &= \text{ert}[C, \mathcal{D}](t). \quad (\text{Lemma B.15.}) \quad \square
 \end{aligned}$$

It remains to prove Lemma B.17.

PROOF OF LEMMA B.17. By induction on n .

The Base Case $n = 0$. By Definition B.14, we have $\text{call}_0^{\mathcal{D}} P_i = \text{halt}$ for each procedure P_i . We proceed by induction on the structure of pRGCL programs to show that

$$\lambda\sigma. \text{ExpRew}^{0\langle \mathcal{P}_\sigma^t \llbracket C, \mathcal{D} \rrbracket \rangle} (\langle \text{sink} \rangle) = \text{ert} [C, \mathcal{D}]_{(\text{halt}, \dots, \text{halt})}^\# (t).$$

We first consider the base case of procedure calls $C = \text{call } P_i$. Since $n = 0$, no single push onto the stack may be performed without the PMC ${}^n\langle \mathcal{P}_\sigma^t \llbracket C, \mathcal{D} \rrbracket \rangle$ immediately moving to $\langle \text{sink} \rangle$. Thus, every path of the operational PMC ${}^n\langle \mathcal{P}_\sigma^t \llbracket \text{call } P_i, \mathcal{D} \rrbracket \rangle$ is of the form

$$(\langle \text{init}(\text{call } P_i), \sigma \rangle, \gamma_0) \xrightarrow{1} (\langle \text{sink} \rangle, \gamma_0) \xrightarrow{1} \dots \xrightarrow{1} (\langle \text{sink} \rangle, \gamma_0),$$

collecting zero reward. Then

$$\begin{aligned} & \text{ExpRew}^{0\langle \mathcal{P}_\sigma^t \llbracket \text{call } P_i, \mathcal{D} \rrbracket \rangle} (\langle \text{sink} \rangle) \\ &= 0 \\ &= \text{ert} [\text{halt}, \mathcal{D}]_{(\text{halt}, \dots, \text{halt})}^\# (t)(\sigma) \\ &= \text{ert} [\text{call}_0^{\mathcal{D}} P_i, \mathcal{D}]_{(\text{halt}, \dots, \text{halt})}^\# (t)(\sigma). \end{aligned}$$

All other base cases are analogous to the soundness proof for pGCL-programs presented in the proof of Theorem 5.5, because $\text{ert} [C, \mathcal{D}]_{(\text{halt}, \dots, \text{halt})}^\# (t) = \text{ert}[C](t)$ holds for each call-free program C . The same holds for the compound statements—sequential composition, conditionals, and loops—by using the inductive hypothesis on C .

Inductive Hypothesis on n . For an arbitrary but fixed $n \in \mathbb{N}$ we assume that for all pRGCL-programs C , we have

$$\lambda\sigma. \text{ExpRew}^{n\langle \mathcal{P}_\sigma^t \llbracket C, \mathcal{D} \rrbracket \rangle} (\langle \text{sink} \rangle) = \text{ert} [C, \mathcal{D}]_{(\text{call}_n^{\mathcal{D}} P_1, \dots, \text{call}_n^{\mathcal{D}} P_k)}^\# (t).$$

Inductive Step $n \mapsto n+1$. As in the base case $n = 0$, the proof proceeds by structural induction on C , where each case except for procedure calls is analogous to the soundness proof for pGCL-programs (cf. Theorem 5.5). We thus concentrate on the treatment of procedure calls, a base case of our structural induction. By definition of the transition relation Δ of ${}^n\langle \mathcal{P}_\sigma^t \llbracket C, \mathcal{D} \rrbracket \rangle$, we observe that

$$\text{ExpRew}^{n+1\langle \mathcal{P}_\sigma^t \llbracket \text{call } P_i, \mathcal{D} \rrbracket \rangle} (\langle \text{sink} \rangle) = 1 + \text{ExpRew}^{n\langle \mathcal{P}_\sigma^t \llbracket \mathcal{D}(P_i), \mathcal{D} \rrbracket \rangle} (\langle \text{sink} \rangle).$$

In other words, the expected reward of a procedure call equals the reward of the call itself plus the reward of executing the procedure's body. We then obtain the desired result as follows:

$$\begin{aligned} & \lambda\sigma. \text{ExpRew}^{n+1\langle \mathcal{P}_\sigma^t \llbracket \text{call } P_i, \mathcal{D} \rrbracket \rangle} (\langle \text{sink} \rangle) \\ &= \lambda\sigma. 1 + \text{ExpRew}^{n\langle \mathcal{P}_\sigma^t \llbracket \mathcal{D}(P_i), \mathcal{D} \rrbracket \rangle} (\langle \text{sink} \rangle) && \text{(previous observation)} \\ &= \mathbf{1} + \lambda\sigma. \text{ExpRew}^{n\langle \mathcal{P}_\sigma^t \llbracket \mathcal{D}(P_i), \mathcal{D} \rrbracket \rangle} (\langle \text{sink} \rangle) \\ &= \mathbf{1} + \text{ert} [\mathcal{D}(P_i), \mathcal{D}]_{(\text{call}_n^{\mathcal{D}} P_1, \dots, \text{call}_n^{\mathcal{D}} P_k)}^\# (t) && \text{(I.H. on } n) \\ &= \text{ert} [\text{skip}; \mathcal{D}(P_i), \mathcal{D}]_{(\text{call}_n^{\mathcal{D}} P_1, \dots, \text{call}_n^{\mathcal{D}} P_k)}^\# (t) \\ &= \text{ert} [\text{call } P_i, \mathcal{D}]_{(\text{call}_{n+1}^{\mathcal{D}} P_1, \dots, \text{call}_{n+1}^{\mathcal{D}} P_k)}^\# (t). \end{aligned}$$

This completes the proof of Lemma B.17 as well as Theorem 7.10. \square

B.11 Proof of Theorem 8.1 (Connection between ert and wp)

We prove that for every program $C \in \text{pGCL}$,

$$\text{ert}[C](t + t') = \text{ert}[C](t) + \text{wp}[C](t')$$

by induction on the structure of C . We consider only the cases of compound statements and assignments as the proof argument for the remaining basic instructions is straightforward.

Assignment.

$$\begin{aligned} \text{ert}[x := \mu](t + t') &= \mathbf{1} + \lambda\sigma. \mathbb{E}_{\llbracket \mu \rrbracket(\sigma)}(\lambda v. (t + t')[x/v](\sigma)) && \text{(Table 1)} \\ &= \mathbf{1} + \lambda\sigma. \mathbb{E}_{\llbracket \mu \rrbracket(\sigma)}(\lambda v. t[x/v](\sigma)) + \lambda\sigma. \mathbb{E}_{\llbracket \mu \rrbracket(\sigma)}(\lambda v. t'[x/v](\sigma)) && \text{(E}_\eta(\cdot)\text{linear)} \\ &= \text{ert}[x := \mu](t) + \text{wp}[x := \mu](t') && \text{(Tables 1, 5)} \end{aligned}$$

Conditionals.

$$\begin{aligned} \text{ert}[\text{if } (\xi) \{C_1\} \text{ else } \{C_2\}](t + t') &= \mathbf{1} + \llbracket \xi : \text{true} \rrbracket \cdot \text{ert}[C_1](t + t') + \llbracket \xi : \text{false} \rrbracket \cdot \text{ert}[C_2](t + t') && \text{(Table 1)} \\ &= \mathbf{1} + \llbracket \xi : \text{true} \rrbracket \cdot (\text{ert}[C_1](t) + \text{wp}[C_1](t')) \\ &\quad + \llbracket \xi : \text{false} \rrbracket \cdot (\text{ert}[C_2](t) + \text{wp}[C_2](t')) && \text{(I.H. on } C_1, C_2\text{)} \\ &= \mathbf{1} + \llbracket \xi : \text{true} \rrbracket \cdot \text{ert}[C_1](t) + \llbracket \xi : \text{false} \rrbracket \cdot \text{ert}[C_2](t) \\ &\quad + \llbracket \xi : \text{true} \rrbracket \cdot \text{wp}[C_1](t') + \llbracket \xi : \text{false} \rrbracket \cdot \text{wp}[C_2](t') \\ &= \text{ert}[\text{if } (\xi) \{C_1\} \text{ else } \{C_2\}](t) + \text{wp}[\text{if } (\xi) \{C_1\} \text{ else } \{C_2\}](t') && \text{(Tables 1, 5)} \end{aligned}$$

Sequential Composition.

$$\begin{aligned} \text{ert}[C_1; C_2](t + t') &= \text{ert}[C_1](\text{ert}[C_2](t + t')) && \text{(Table 1)} \\ &= \text{ert}[C_1](\text{ert}[C_2](t) + \text{wp}[C_2](t')) && \text{(I.H. on } C_2\text{)} \\ &= \text{ert}[C_1](\text{ert}[C_2](t)) + \text{wp}[C_1](\text{wp}[C_2](t')) && \text{(I.H. on } C_1\text{)} \\ &= \text{ert}[C_1; C_2](t) + \text{wp}[C_1; C_2](t') && \text{(Tables 1, 5)} \end{aligned}$$

Loops. Let $F_t(X) \triangleq \mathbf{1} + \llbracket \xi : \text{false} \rrbracket \cdot t + \llbracket \xi : \text{true} \rrbracket \cdot \text{ert}[C'](X)$ and $G_t(X) \triangleq \llbracket \xi : \text{false} \rrbracket \cdot t + \llbracket \xi : \text{true} \rrbracket \cdot \text{wp}[C'](X)$. We have to show that

$$\text{lfp } X. F_{t+t'}(X) = \text{lfp } X. F_t(X) + \text{lfp } X. G_{t'}(X),$$

which, following the same argument for the soundness of the second rule in Theorem 7.3 (see Appendix B.9), becomes

$$\lim_{n \rightarrow \infty} F_{t+t'}^n(\mathbf{0}) = \lim_{n \rightarrow \infty} F_t^n(\mathbf{0}) + G_{t'}^n(\mathbf{0}),$$

where $F_{t+t'}^n$ denotes the composition of $F_{t+t'}$ with itself n times (and likewise for $G_{t'}^n$). To conclude, we prove by induction on n that

$$\forall n. F_{t+t'}^n(\mathbf{0}) = F_t^n(\mathbf{0}) + G_{t'}^n(\mathbf{0}).$$

The base case is immediate since it reduces to $\mathbf{0} = \mathbf{0} + \mathbf{0}$. For the inductive case, we derive

$$\begin{aligned}
F_{t+t'}^{n+1}(\mathbf{0}) &= \mathbf{1} + \llbracket \xi : \text{false} \rrbracket \cdot (t + t') + \llbracket \xi : \text{true} \rrbracket \cdot \text{ert}[C'](F_{t+t'}^n(\mathbf{0})) && (\text{def. } F_{t+t'}^{n+1}) \\
&= \mathbf{1} + \llbracket \xi : \text{false} \rrbracket \cdot (t + t') + \llbracket \xi : \text{true} \rrbracket \cdot \text{ert}[C'](F_t^n(\mathbf{0}) + G_{t'}^n(\mathbf{0})) && (\text{I.H. on } n) \\
&= \mathbf{1} + \llbracket \xi : \text{false} \rrbracket \cdot (t + t') \\
&\quad + \llbracket \xi : \text{true} \rrbracket \cdot (\text{ert}[C'](F_t^n(\mathbf{0})) + \text{wp}[C'](G_{t'}^n(\mathbf{0}))) && (\text{I.H. on } C') \\
&= \mathbf{1} + \llbracket \xi : \text{false} \rrbracket \cdot t + \llbracket \xi : \text{true} \rrbracket \cdot \text{ert}[C'](F_t^n(\mathbf{0})) \\
&\quad + \llbracket \xi : \text{false} \rrbracket \cdot t' + \llbracket \xi : \text{true} \rrbracket \cdot \text{wp}[C'](G_{t'}^n(\mathbf{0})) \\
&= F_t^{n+1}(\mathbf{0}) + G_{t'}^{n+1}(\mathbf{0}). && (\text{def. } F_t^{n+1}, G_{t'}^{n+1}) \quad \square
\end{aligned}$$

C OMITTED CALCULATIONS

C.1 Invariant Verification for the Random Walk

First, we verify that

$$I_n = 1 + \sum_{k=0}^n \llbracket x > k \rrbracket \cdot a_{n,k}$$

is a lower ω -invariant of the loop with respect to $\mathbf{0}$ if for all $n \geq 0$,

$$a_{0,0} = 1 \tag{4}$$

$$a_{n+1,0} = 2 + \frac{1}{2} \cdot (a_{n,0} + a_{n,1}) \tag{5}$$

$$a_{n+1,k} = \frac{1}{2} \cdot (a_{n,k-1} + a_{n,k+1}), \text{ for all } 1 \leq k \leq n+1 \tag{6}$$

$$a_{n,k} = 0, \text{ for all } k > n. \tag{7}$$

Let F be the characteristic functional of the loop with respect to $\mathbf{0}$. Then, for the first condition $F(\mathbf{0}) \geq I_0$, we have

$$\begin{aligned}
F(\mathbf{0}) &= 1 + \llbracket x \leq 0 \rrbracket \cdot 0 + \llbracket x > 0 \rrbracket \cdot \text{ert}[C](0) \\
&= 1 + \llbracket x > 0 \rrbracket \cdot \left(1 + \frac{1}{2} \cdot 0[x/x - 1] + \frac{1}{2} \cdot 0[x/x + 1] \right) \\
&= 1 + \llbracket x > 0 \rrbracket \cdot 1 \\
&= 1 + \llbracket x > 0 \rrbracket \cdot a_{0,0} = I_0. && (\text{Equation (4)})
\end{aligned}$$

For the second condition $F(I_n) \geq I_{n+1}$, consider

$$\begin{aligned}
F(I_n) &= 1 + \llbracket x \leq 0 \rrbracket \cdot 0 + \llbracket x > 0 \rrbracket \cdot \text{ert}[C](I_n) \\
&= 1 + \llbracket x > 0 \rrbracket \cdot \left(1 + \frac{1}{2} \cdot I_n[x/x - 1] + \frac{1}{2} \cdot I_n[x/x + 1] \right) \\
&= 1 + \llbracket x > 0 \rrbracket \cdot \left(2 + \frac{1}{2} \cdot \sum_{k=0}^n \llbracket x-1 > k \rrbracket \cdot a_{n,k} + \frac{1}{2} \cdot \sum_{k=0}^n \llbracket x+1 > k \rrbracket \cdot a_{n,k} \right) \\
&= 1 + \llbracket x > 0 \rrbracket \cdot \left(2 + \frac{1}{2} \cdot \sum_{k=1}^{n+1} \llbracket x > k \rrbracket \cdot a_{n,k-1} + \frac{1}{2} \cdot \sum_{k=-1}^{n-1} \llbracket x > k \rrbracket \cdot a_{n,k+1} \right)
\end{aligned}$$

$$\begin{aligned}
 &= 1 + \llbracket x > 0 \rrbracket \cdot \left(2 + \frac{1}{2} \cdot \llbracket x > -1 \rrbracket \cdot a_{n,0} + \frac{1}{2} \cdot \llbracket x > 0 \rrbracket \cdot a_{n,1} \right. \\
 &\quad \left. + \frac{1}{2} \cdot \sum_{k=1}^{n-1} \llbracket x > k \rrbracket \cdot (a_{n,k-1} + a_{n,k+1}) \right. \\
 &\quad \left. + \frac{1}{2} \cdot \llbracket x > n \rrbracket \cdot a_{n,n-1} + \frac{1}{2} \cdot \llbracket x > n+1 \rrbracket \cdot a_{n,n} \right)
 \end{aligned}$$

We distribute $\llbracket x > 0 \rrbracket$ and use the fact that $\llbracket x > 0 \rrbracket \cdot \llbracket x > a \rrbracket = \llbracket x > \max\{0, a\} \rrbracket$ to obtain

$$\begin{aligned}
 &= 1 + \llbracket x > 0 \rrbracket \cdot \left(2 + \frac{1}{2} \cdot (a_{n,0} + a_{n,1}) \right) + \frac{1}{2} \cdot \sum_{k=1}^{n-1} \llbracket x > k \rrbracket \cdot (a_{n,k-1} + a_{n,k+1}) \\
 &\quad + \frac{1}{2} \cdot \llbracket x > n \rrbracket \cdot a_{n,n-1} + \frac{1}{2} \cdot \llbracket x > n+1 \rrbracket \cdot a_{n,n} \\
 &= 1 + \llbracket x > 0 \rrbracket \cdot a_{n+1,0} + \frac{1}{2} \cdot \sum_{k=1}^{n-1} \llbracket x > k \rrbracket \cdot (a_{n,k-1} + a_{n,k+1}) \tag{Equation (5)} \\
 &\quad + \frac{1}{2} \cdot \llbracket x > n \rrbracket \cdot a_{n,n-1} + \frac{1}{2} \cdot \llbracket x > n+1 \rrbracket \cdot a_{n,n}
 \end{aligned}$$

$$\begin{aligned}
 &= 1 + \llbracket x > 0 \rrbracket \cdot a_{n+1,0} + \frac{1}{2} \cdot \sum_{k=1}^{n-1} \llbracket x > k \rrbracket \cdot (a_{n,k-1} + a_{n,k+1}) \tag{Equation (7)} \\
 &\quad + \frac{1}{2} \cdot \llbracket x > n \rrbracket \cdot (a_{n,n-1} + a_{n,n+1}) + \frac{1}{2} \cdot \llbracket x > n+1 \rrbracket \cdot (a_{n,n} + a_{n,n+2})
 \end{aligned}$$

$$\begin{aligned}
 &= 1 + \llbracket x > 0 \rrbracket \cdot a_{n+1,0} + \frac{1}{2} \cdot \sum_{k=1}^{n+1} \llbracket x > k \rrbracket \cdot (a_{n,k-1} + a_{n,k+1}) \\
 &= 1 + \llbracket x > 0 \rrbracket \cdot a_{n+1,0} + \sum_{k=1}^{n+1} \llbracket x > k \rrbracket \cdot a_{n+1,k} \tag{Equation (6)}
 \end{aligned}$$

$$= 1 + \sum_{k=0}^{n+1} \llbracket x > k \rrbracket \cdot a_{n+1,k} = I_{n+1}.$$

Now we show that

$$a_{n,k} = \frac{1}{2^n} \left[-\binom{n}{\lfloor \frac{n-k}{2} \rfloor} + 2 \sum_{i=0}^{n-k} 2^i \binom{n-i}{\lfloor \frac{n-i-k}{2} \rfloor} \right]$$

satisfies the recursion in Equations (4) to (7). Here, we assume that $\binom{n}{m}$ is 0 whenever $m < 0$. Equations (4) and (7) are immediate. For Equation (5), i.e., $a_{n+1,0} = 2 + 1/2 \cdot (a_{n,0} + a_{n,1})$, we make use of the identity

$$\binom{k}{\lfloor \frac{k+1}{2} \rfloor} = \binom{k}{\lfloor \frac{k}{2} \rfloor}, \tag{★}$$

which is shown by a simple case analysis on k being even or odd. Then

$$\begin{aligned}
 a_{n+1,0} &= \frac{1}{2^{n+1}} \left[-\binom{n+1}{\lfloor \frac{n+1}{2} \rfloor} + 2 \sum_{i=0}^{n+1} 2^i \binom{n+1-i}{\lfloor \frac{n+1-i}{2} \rfloor} \right] \tag{Def. $a_{n+1,0}$ } \\
 &= \frac{1}{2^{n+1}} \left[-\binom{n}{\lfloor \frac{n+1}{2} \rfloor} - \binom{n}{\lfloor \frac{n+1}{2} \rfloor - 1} \right] \quad \left(\binom{k+1}{\ell+1} = \binom{k}{\ell} + \binom{k}{\ell+1} \right)
 \end{aligned}$$

$$\begin{aligned}
& + 2 \sum_{i=0}^{n+1} 2^i \left(\binom{n-i}{\lfloor \frac{n+1-i}{2} \rfloor} + \binom{n-i}{\lfloor \frac{n+1-i}{2} \rfloor - 1} \right) \Big] \\
& = \frac{1}{2^{n+1}} \left[- \binom{n}{\lfloor \frac{n+1}{2} \rfloor} - \binom{n}{\lfloor \frac{n+1}{2} \rfloor - 1} \right. \\
& \quad \left. + 2^{n+2} + 2 \sum_{i=0}^n 2^i \left(\binom{n-i}{\lfloor \frac{n+1-i}{2} \rfloor} + \binom{n-i}{\lfloor \frac{n+1-i}{2} \rfloor - 1} \right) \right] \\
& = 2 + \frac{1}{2} \cdot \left(\frac{1}{2^n} \left[- \binom{n}{\lfloor \frac{n+1}{2} \rfloor} + 2 \sum_{i=0}^{n-1} 2^i \binom{n-i}{\lfloor \frac{n-i-1}{2} \rfloor} \right] \right. \\
& \quad \left. + \frac{1}{2^n} \left[- \binom{n}{\lfloor \frac{n+1}{2} \rfloor} + 2 \sum_{i=0}^n 2^i \binom{n-i}{\lfloor \frac{n+1-i}{2} \rfloor} \right] \right) \quad (\text{by } \binom{0}{-1} = 0 \text{ and } \lfloor \frac{n+1}{2} \rfloor - 1 = \lfloor \frac{n-1}{2} \rfloor) \\
& = 2 + \frac{1}{2} \left(a_{n,1} + \frac{1}{2^n} \left[- \binom{n}{\lfloor \frac{n+1}{2} \rfloor} + 2 \sum_{i=0}^n 2^i \binom{n-i}{\lfloor \frac{n+1-i}{2} \rfloor} \right] \right) \quad (\text{Def. } a_{n,1}) \\
& = 2 + \frac{1}{2} \left(a_{n,1} + \frac{1}{2^n} \left[- \binom{n}{\lfloor \frac{n}{2} \rfloor} + 2 \sum_{i=0}^n 2^i \binom{n-i}{\lfloor \frac{n-i}{2} \rfloor} \right] \right) \quad (\text{using } \star) \\
& = 2 + \frac{1}{2} (a_{n,1} + a_{n,0}). \quad (\text{Def. } a_{n,0})
\end{aligned}$$

It remains to show Equation (6), i.e., $a_{n+1,k} = \frac{1}{2} \cdot (a_{n,k-1} + a_{n,k+1})$, for all $1 \leq k \leq n+1$:

$$\begin{aligned}
a_{n+1,k} & = \frac{1}{2^{n+1}} \left[- \binom{n+1}{\lfloor \frac{n+1-k}{2} \rfloor} + 2 \sum_{i=0}^{n+1-k} 2^i \binom{n+1-i}{\lfloor \frac{n+1-i-k}{2} \rfloor} \right] \quad (\text{Def. } a_{n+1,k}) \\
& = \frac{1}{2^{n+1}} \left[- \binom{n+1}{\lfloor \frac{n+1-k}{2} \rfloor} + 2^{n+2-k} + 2 \sum_{i=0}^{n-k} 2^i \binom{n+1-i}{\lfloor \frac{n+1-i-k}{2} \rfloor} \right] \quad \left(\binom{m}{0} = 1 \right) \\
& = \frac{1}{2^{n+1}} \left[- \binom{n}{\lfloor \frac{n-(k-1)}{2} \rfloor} - \binom{n}{\lfloor \frac{n-(k+1)}{2} \rfloor} + 2^{n+2-k} \right. \\
& \quad \left. + 2 \sum_{i=0}^{n-k} 2^i \binom{n-i}{\lfloor \frac{n-i-(k-1)}{2} \rfloor} + 2 \sum_{i=0}^{n-k} 2^i \binom{n-i}{\lfloor \frac{n-i-(k+1)}{2} \rfloor} \right] \quad \left(\binom{k+1}{\ell+1} = \binom{k}{\ell} + \binom{k}{\ell+1} \right) \\
& = \frac{1}{2^{n+1}} \left[- \binom{n}{\lfloor \frac{n-(k-1)}{2} \rfloor} - \binom{n}{\lfloor \frac{n-(k+1)}{2} \rfloor} + 2^{n+2-k} \right. \\
& \quad \left. + 2 \sum_{i=0}^{n-k} 2^i \binom{n-i}{\lfloor \frac{n-i-(k-1)}{2} \rfloor} + 2 \sum_{i=0}^{n-(k+1)} 2^i \binom{n-i}{\lfloor \frac{n-i-(k+1)}{2} \rfloor} \right] \quad \left(\binom{m}{-1} = 0 \right) \\
& = \frac{1}{2} \left(a_{n,k+1} + \frac{1}{2^n} \left[- \binom{n}{\lfloor \frac{n-(k-1)}{2} \rfloor} + 2 \sum_{i=0}^{n-(k-1)} 2^i \binom{n-i}{\lfloor \frac{n-i-(k-1)}{2} \rfloor} \right] \right) \quad (\text{Def. } a_{n,k+1}) \\
& = \frac{1}{2} (a_{n,k+1} + a_{n,k-1}). \quad (\text{Def. } a_{n,k-1})
\end{aligned}$$

Finally, we prove that $\lim_{n \rightarrow \infty} a_{n,0} = \infty$. The crux of the proof is showing that for all $n \geq 2$, $a_{n,0} \geq 1 + \mathcal{H}_{\lfloor n/2 \rfloor}$, where \mathcal{H}_m denotes the m th Harmonic number, i.e.,

$$\mathcal{H}_m = \sum_{k=1}^m \frac{1}{k}.$$

The result then follows since $\lim_{m \rightarrow \infty} \mathcal{H}_m = \infty$. Calculations go as follows:

$$\begin{aligned} a_{n,0} &= \frac{1}{2^n} \left[-\binom{n}{\lfloor \frac{n}{2} \rfloor} + 2 \sum_{i=0}^n 2^i \binom{n-i}{\lfloor \frac{n-i}{2} \rfloor} \right] \\ &= \frac{1}{2^n} \left[-\binom{n}{\lfloor \frac{n}{2} \rfloor} + 2 \sum_{k=0}^n 2^{n-k} \binom{k}{\lfloor \frac{k}{2} \rfloor} \right] && \text{(Take } k = n - i \text{)} \\ &\geq \frac{1}{2^n} \left[-2^n + 2 \sum_{k=0}^n 2^{n-k} \binom{k}{\lfloor \frac{k}{2} \rfloor} \right] && \left(\binom{n}{\lfloor n/2 \rfloor} \leq 2^n \right) \\ &= -1 + 2 \sum_{k=0}^n 2^{-k} \binom{k}{\lfloor \frac{k}{2} \rfloor} \\ &\geq -1 + 2 \sum_{j=0}^{\lfloor n/2 \rfloor} 2^{-2j} \binom{2j}{j} && \text{(Keep only even } k \text{'s)} \\ &\geq 1 + 2 \sum_{j=1}^{\lfloor n/2 \rfloor} 2^{-2j} \binom{2j}{j} && \text{(Extract } j = 0 \text{ out of the sum)} \\ &\geq 1 + 2 \sum_{j=1}^{\lfloor n/2 \rfloor} 2^{-2j} \frac{2^{2j-1}}{\sqrt{j}} && \left(\text{Stirling approximation : } \binom{2j}{j} \geq \frac{2^{2j-1}}{\sqrt{j}} \right) \\ &= 1 + \sum_{j=1}^{\lfloor n/2 \rfloor} \frac{1}{\sqrt{j}} \\ &\geq 1 + \sum_{j=1}^{\lfloor n/2 \rfloor} \frac{1}{j} = 1 + \mathcal{H}_{\lfloor n/2 \rfloor}. && (\sqrt{j} \leq j) \end{aligned}$$

C.2 Invariant Verification for the Coupon Collector Algorithm

Recall our proposed invariant

$$\begin{aligned} I \triangleq & \mathbf{1} + \sum_{\ell=0}^{\infty} [x > \ell] \cdot \left(4 + 2 \cdot \sum_{k=0}^{\infty} \left(\frac{\#col + \ell}{N} \right)^k \right) \\ & - 2 \cdot [cp[i] = 0] \cdot [x > 0] \cdot \sum_{k=0}^{\infty} \left(\frac{\#col}{N} \right)^k. \end{aligned}$$

To prove this invariant correct, we have to show that $F(I) \leq I$, where $F(X)$ denotes the characteristic functional of the coupon collector's outer loop with respect to runtime 0. As such, $F(X)$ is given by

$$\begin{aligned} F(X) &= \mathbf{1} + [x \leq 0] \cdot 0 + [x > 0] \cdot \text{ert}[C_{in}; cp[i] := 1; x := x - 1](X) \\ &= \mathbf{1} + [x > 0] \cdot (2 + \text{ert}[C_{in}](X[x/x - 1, cp[i]/1])), \end{aligned}$$

where C_{in} corresponds to the inner loop of the coupon collector algorithm. Thus, in order to verify that $F(I) \leq I$, we need an upper invariant for the inner loop first. Note that this invariant cannot

be chosen with respect to continuation 0, but with respect to the invariant I . We will, however, derive an invariant of the inner loop with respect to an arbitrary continuation.

Invariant for the Inner Loop. Given an arbitrary runtime $f \in \mathbb{T}$, we propose that

$$\begin{aligned} J_f &\triangleq \mathbf{1} + [cp[i] = 0] \cdot f + [cp[i] \neq 0] \cdot \sum_{\ell=0}^{\infty} \left(\frac{\#col}{N}\right)^{\ell} \cdot \underbrace{\left(2 + \frac{1}{N} \cdot \sum_{j=1}^N [cp[j] = 0] \cdot f[i/j]\right)}_{= G_f} \\ &= \mathbf{1} + [cp[i] = 0] \cdot f + [cp[i] \neq 0] \cdot \sum_{\ell=0}^{\infty} \left(\frac{\#col}{N}\right)^{\ell} \cdot G_f \end{aligned}$$

is an invariant of the inner loop C_{in} with respect to f . Toward a correctness proof of our proposed invariant, we have to verify that $H_f(J_f) \leq J_f$, where $H_f(Y)$ is the characteristic functional of the inner loop with respect to runtime f . Thus,

$$\begin{aligned} H_f(J_f) &= \mathbf{1} + [cp[i] = 0] \cdot f + [cp[i] \neq 0] \cdot \text{ert}[i \approx \text{Unif}[1 \dots N]](J_f) \\ &= \mathbf{1} + [cp[i] = 0] \cdot f + [cp[i] \neq 0] \cdot \left(1 + \frac{1}{N} \cdot \sum_{k=1}^N J_f[i/k]\right) \\ &= \mathbf{1} + [cp[i] = 0] \cdot f + [cp[i] \neq 0] \quad (\text{Definition } J_f) \\ &\quad \cdot \left(1 + \frac{1}{N} \cdot \sum_{k=1}^N \left(1 + [cp[k] = 0] \cdot f[i/k] + [cp[k] \neq 0] \cdot \sum_{\ell=0}^{\infty} \left(\frac{\#col}{N}\right)^{\ell} \cdot G_f\right)\right) \\ &= \mathbf{1} + [cp[i] = 0] \cdot f + 2 \cdot [cp[i] \neq 0] \\ &\quad + \frac{[cp[i] \neq 0]}{N} \cdot \sum_{k=1}^N [cp[k] = 0] \cdot f[i/k] + \frac{[cp[i] \neq 0]}{N} \cdot \sum_{k=1}^N [cp[k] \neq 0] \cdot \sum_{\ell=0}^{\infty} \left(\frac{\#col}{N}\right)^{\ell} \cdot G_f \\ &= \mathbf{1} + [cp[i] = 0] \cdot f + 2 \cdot [cp[i] \neq 0] \quad (\text{Definition of } \#col) \\ &\quad + \frac{[cp[i] \neq 0]}{N} \cdot \sum_{k=1}^N [cp[k] = 0] \cdot f[i/k] + \frac{[cp[i] \neq 0]}{N} \cdot \#col \cdot \sum_{\ell=0}^{\infty} \left(\frac{\#col}{N}\right)^{\ell} \cdot G_f \\ &= \mathbf{1} + [cp[i] = 0] \cdot f + [cp[i] \neq 0] \cdot G_f + \frac{[cp[i] \neq 0]}{N} \cdot \#col \cdot \sum_{\ell=0}^{\infty} \left(\frac{\#col}{N}\right)^{\ell} \cdot G_f \\ &= \mathbf{1} + [cp[i] = 0] \cdot f + [cp[i] \neq 0] \cdot G_f + [cp[i] \neq 0] \cdot \sum_{\ell=1}^{\infty} \left(\frac{\#col}{N}\right)^{\ell} \cdot G_f \\ &= \mathbf{1} + [cp[i] = 0] \cdot f + [cp[i] \neq 0] \cdot \sum_{\ell=0}^{\infty} \left(\frac{\#col}{N}\right)^{\ell} \cdot G_f = J_f. \end{aligned}$$

Hence, J_f is an upper global invariant of the inner loop C_{in} . We are now in a position to verify that $F(I) \leq I$; i.e., I is an upper global invariant of the outer loop C_{out} .

Invariant Verification for the Outer Loop. In order to keep calculations readable, let

$$K_i^j \triangleq \sum_{\ell=i}^{\infty} [x > \ell + j] \cdot \left(4 + 2 \cdot \sum_{k=0}^{\infty} \left(\frac{\#col + j + \ell}{N}\right)^k\right).$$

Then our proposed invariant I can be rewritten as

$$I = 1 + K_0^0 - 2 \cdot [cp[i] = 0] \cdot [x > 0] \cdot \sum_{k=0}^{\infty} \left(\frac{\#col}{N} \right)^k.$$

The proof $F(I) \leq I$ relies on two properties:

LEMMA C.1.

$$N - \#col = \sum_{i=1}^N (1 - [cp[i] \neq 0]) = \sum_{i=1}^N [cp[i] = 0].$$

PROOF. Immediate by definition of $\#col$. □

LEMMA C.2.

$$[cp[i] = 0] \cdot I[x/x - 1, cp[i]/1] = [cp[i] = 0] \cdot (1 + K_0^1).$$

PROOF. Immediate by observing that

$$I_n = 1 + K_0^0 - 2 \cdot [cp[i] = 0] \cdot [x > 0] \cdot \sum_{k=0}^{\infty} \left(\frac{\#col}{N} \right)^k$$

and applying the respective substitutions. □

We are now in a position to verify that $F(I) \leq I$; i.e., our proposed invariant I is indeed an upper global invariant of the outer loop.

$$\begin{aligned} & F(I) \\ &= 1 + [x > 0] \cdot (2 + \text{ert}[C_{in}](I[x/x - 1, cp[i]/1])) && \text{(Definition of } F) \\ &\leq 1 + [x > 0] \cdot (2 + J_{I[x/x-1, cp[i]/1]}) && (\text{ert}[C_{in}](f) \leq J_f) \\ &= 1 + [x > 0] \cdot (3 + [cp[i] = 0] \cdot I[x/x - 1, cp[i]/1] && \text{(Definition of } J_f) \\ &\quad + [cp[i] \neq 0] \cdot \sum_{\ell=0}^{\infty} \left(\frac{\#col}{N} \right)^{\ell} \cdot G_{I[x/x-1, cp[i]/1]}) \\ &= 1 + [x > 0] \cdot (3 + [cp[i] = 0] \cdot I[x/x - 1, cp[i]/1] && \text{(Definition of } G_f) \\ &\quad + [cp[i] \neq 0] \cdot \sum_{\ell=0}^{\infty} \left(\frac{\#col}{N} \right)^{\ell} \cdot \left(2 + \sum_{j=1}^N \frac{[cp[j] = 0]}{N} \cdot I[x/x - 1, cp[i]/1, i/j] \right)) \\ &= 1 + [x > 0] \cdot (3 + [cp[i] = 0] \cdot (1 + K_0^1) && \text{(Lemma C.2)} \\ &\quad + [cp[i] \neq 0] \cdot \sum_{\ell=0}^{\infty} \left(\frac{\#col}{N} \right)^{\ell} \cdot \left(2 + \sum_{j=1}^N \frac{[cp[j] = 0]}{N} \cdot (1 + K_0^1) \right)) \\ &= 1 + [x > 0] \cdot (3 + [cp[i] = 0] \cdot (1 + K_1^0) && (K_0^1 = K_1^0) \\ &\quad + [cp[i] \neq 0] \cdot \sum_{\ell=0}^{\infty} \left(\frac{\#col}{N} \right)^{\ell} \cdot \left(2 + \sum_{j=1}^N \frac{[cp[j] = 0]}{N} \cdot (1 + K_1^0) \right)) \\ &= 1 + [x > 0] \cdot (3 + [cp[i] = 0] \cdot (1 + K_1^0) && \text{(Lemma C.1)} \\ &\quad + [cp[i] \neq 0] \cdot \sum_{\ell=0}^{\infty} \left(\frac{\#col}{N} \right)^{\ell} \cdot \left(2 + \left(1 - \frac{\#col}{N} \right) \cdot (1 + K_1^0) \right)) \end{aligned}$$

$$\begin{aligned}
&= 1 + [x > 0] \cdot (3 + [cp[i] = 0] \cdot (1 + K_1^0)) \\
&\quad + [cp[i] \neq 0] \cdot \sum_{\ell=0}^{\infty} \left(\frac{\#col}{N}\right)^{\ell} \cdot (3 + K_1^0 - \frac{\#col}{N} \cdot (1 + K_1^0)) \\
&= 1 + [x > 0] \cdot (4 + K_1^0) \\
&\quad + [cp[i] \neq 0] \cdot \sum_{\ell=1}^{\infty} \left(\frac{\#col}{N}\right)^{\ell} \cdot \left(3 + K_1^0 - \frac{\#col}{N} \cdot (1 + K_1^0)\right) \\
&\quad + [cp[i] \neq 0] \cdot \left(2 - \frac{\#col}{N} (1 + K_1^0)\right) \\
&= 1 + [x > 0] \cdot (4 + K_1^0) \\
&\quad + [cp[i] \neq 0] \cdot \sum_{\ell=1}^{\infty} \left(\frac{\#col}{N}\right)^{\ell} \cdot \left(\left(1 - \frac{\#col}{N}\right) \cdot (1 + K_1^0)\right) \\
&\quad + [cp[i] \neq 0] \cdot \left(2 - \frac{\#col}{N} (1 + K_1^0)\right) + 2 \cdot [cp[i] \neq 0] \cdot \sum_{\ell=1}^{\infty} \left(\frac{\#col}{N}\right)^{\ell} \\
&= 1 + [x > 0] \cdot (4 + K_1^0 + [cp[i] \neq 0] \cdot (1 + K_1^0) \cdot \frac{\#col}{N}) \\
&\quad + [cp[i] \neq 0] \cdot \left(2 - \frac{\#col}{N} (1 + K_1^0)\right) + 2 \cdot [cp[i] \neq 0] \cdot \sum_{\ell=1}^{\infty} \left(\frac{\#col}{N}\right)^{\ell} \\
&= 1 + [x > 0] \cdot \left(4 + K_1^0 + 2 \cdot [cp[i] \neq 0] + 2 \cdot [cp[i] \neq 0] \cdot \sum_{\ell=1}^{\infty} \left(\frac{\#col}{N}\right)^{\ell}\right) \\
&= 1 + [x > 0] \cdot \left(4 + K_1^0 + 2 \cdot [cp[i] \neq 0] \cdot \sum_{\ell=0}^{\infty} \left(\frac{\#col}{N}\right)^{\ell}\right) \\
&= 1 + K_1^0 + [x > 0] \cdot \left(4 + 2 \cdot [cp[i] \neq 0] \cdot \sum_{\ell=0}^{\infty} \left(\frac{\#col}{N}\right)^{\ell}\right) \\
&= 1 + K_1^0 + [x > 0] \cdot \left(4 + 2 \cdot \sum_{\ell=0}^{\infty} \left(\frac{\#col}{N}\right)^{\ell}\right) \\
&\quad - 2 \cdot [x > 0] \cdot [cp[i] = 0] \cdot \sum_{\ell=0}^{\infty} \left(\frac{\#col}{N}\right)^{\ell} \\
&= 1 + K_0^0 - 2 \cdot [x > 0] \cdot [cp[i] = 0] \cdot \sum_{\ell=0}^{\infty} \left(\frac{\#col}{N}\right)^{\ell} \quad (\text{Definition of } K_0^0) \\
&= I. \quad (\text{Definition of } I)
\end{aligned}$$

Hence, by Theorem 4.2, we know that $\text{ert}[C_{out}](0) \leq I$.

ACKNOWLEDGMENTS

We thank Gilles Barthe for bringing the coupon collector problem as a particularly intricate case study for formal verification of expected runtimes to our attention, Thomas Noll for bringing Nielson's Hoare logic to our attention, Johannes Hölzl for pointing out a flaw in the analysis of

the random walk as in [28], and several contributors from the MathOverflow community for their valuable insights in finding a closed form of the invariant for the random walk problem.

REFERENCES

- [1] Rob Arthan, Ursula Martin, Erik A. Mathiesen, and Paulo Oliva. 2009. A general framework for sound and complete Floyd-Hoare logics. *ACM Trans. Comput. Log.* 11, 1 (2009), 7:1–7:31.
- [2] Christel Baier and Joost-Pieter Katoen. 2008. *Principles of Model Checking*. MIT Press.
- [3] Gilles Barthe, Thomas Espitau, Luis María Ferrer Fioriti, and Justin Hsu. 2016. Synthesizing probabilistic invariants via doob’s decomposition. In *Computer-Aided Verification (CAV’16) (LNCS)*, Vol. 9779. Springer, 43–61.
- [4] Kevin Batz, Benjamin Lucien Kaminski, Joost-Pieter Katoen, and Christoph Matheja. 2018. How long, O Bayesian network, will I sample thee? In *European Symposium on Programming (ESOP’18) (LNCS)*, Vol. 10801. Springer.
- [5] Rudolf Berghammer and Markus Müller-Olm. 2004. Formal development and verification of approximation algorithms using auxiliary variables. In *Logic Based Program Synthesis and Transformation (LOPSTR’04) (LNCS)*, Vol. 3018. Springer, 59–74.
- [6] Tomáš Brázdil, Javier Esparza, Stefan Kiefer, and Antonín Kucera. 2013. Analyzing probabilistic pushdown automata. *Formal Methods Syst. Design* 43, 2 (2013), 124–163.
- [7] Tomáš Brázdil, Stefan Kiefer, Antonín Kucera, and Ivana Hutarová Vareková. 2015. Runtime analysis of probabilistic programs with unbounded recursion. *J. Comput. Syst. Sci.* 81, 1 (2015), 288–310.
- [8] Orieta Celiku and Annabelle McIver. 2005. Compositional specification and analysis of cost-based properties in probabilistic programs. In *Formal Methods (FM’05) (LNCS)*, Vol. 3582. Springer, 107–122.
- [9] Aleksandar Chakarov and Sriram Sankaranarayanan. 2013. Probabilistic program analysis with martingales. In *Computer Aided Verification (CAV’13) (LNCS)*, Vol. 8044. Springer, 511–526.
- [10] Krishnendu Chatterjee, Hongfei Fu, and Aniket Murhekar. 2017. Automated recurrence analysis for almost-linear expected-runtime bounds. In *Computer-Aided Verification (CAV’17) (LNCS)*, Vol. 10426. Springer, 118–139.
- [11] Patrick Cousot and Michael Monerau. 2012. Probabilistic abstract interpretation. In *European Symposium on Programming (ESOP’12) (LNCS)*, Vol. 7211. Springer, 169–193.
- [12] Edsger W. Dijkstra. 1976. *A Discipline of Programming*. Prentice Hall.
- [13] Yijun Feng, Lijun Zhang, David N. Jansen, Naijun Zhan, and Bican Xia. 2017. Finding polynomial loop invariants for probabilistic programs. In *Automated Technology for Verification and Analysis (ATVA’17) (LNCS)*, Vol. 10482. Springer, 400–416.
- [14] Luis María Ferrer Fioriti and Holger Hermanns. 2015. Probabilistic termination: Soundness, completeness, and compositionality. In *Principles of Programming Languages (POPL’15)*. ACM, 489–501.
- [15] Gudmund S. Frandsen. 1998. *Randomised Algorithms*. Lecture Notes, University of Aarhus, Denmark.
- [16] Andrew D. Gordon, Thomas A. Henzinger, Aditya V. Nori, and Sriram K. Rajamani. 2014. Probabilistic programming. In *Future of Software Engineering (FOSE’14)*. ACM, 167–181.
- [17] Eric C. R. Hehner. 1998. Formalization of time and space. *Formal Aspects Comput.* 10, 3 (1998), 290–306.
- [18] Eric C. R. Hehner. 2011. A probability perspective. *Formal Aspects Comput.* 23, 4 (2011), 391–419.
- [19] Wim H. Hesselink. 1993. Proof rules for recursive procedures. *Formal Aspects Comput.* 5, 6 (1993), 554–570.
- [20] Timothy Hickey and Jacques Cohen. 1988. Automating program analysis. *J. ACM* 35, 1 (1988), 185–220.
- [21] Charles A. R. Hoare. 1969. An axiomatic basis for computer programming. *Commun. ACM* 12, 10 (1969), 576–580.
- [22] Jan Hoffmann, Klaus Aehlig, and Martin Hofmann. 2012. Multivariate amortized resource analysis. *ACM Trans. Program. Lang. Syst.* 34, 3 (2012), 14:1–14:62.
- [23] Johannes Hölzl. 2016. Formalising semantics for expected running time of probabilistic programs. In *Interactive Theorem Proving (ITP’16) (LNCS)*, Vol. 9807. Springer, 475–482.
- [24] Juraj Hromkovic and Georg Schnitger. 2010. On probabilistic pushdown automata. *Inf. Comput.* 208, 8 (2010), 982–995.
- [25] Joe Hurd. 2002. A formal approach to probabilistic termination. In *Theorem Proving in Higher Order Logics (TPHOL’02)*. LNCS, Vol. 2410. Springer, 230–245.
- [26] Oliver C. Ibe. 2013. *Elements of Random Walk and Diffusion Processes*. John Wiley & Sons.
- [27] Benjamin Lucien Kaminski and Joost-Pieter Katoen. 2015. On the hardness of almost-sure termination. In *Mathematical Foundations of Computer Science (MFCS’15), Part I (LNCS)*, Vol. 9234. Springer, 307–318.
- [28] Benjamin Lucien Kaminski, Joost-Pieter Katoen, Christoph Matheja, and Federico Olmedo. 2016. Weakest precondition reasoning for expected run-times of probabilistic programs. In *European Symposium on Programming (ESOP’16) (LNCS)*, Vol. 9632. Springer, 364–389.
- [29] Richard M. Karp. 1994. Probabilistic recurrence relations. *J. ACM* 41, 6 (1994), 1136–1150.
- [30] Joost-Pieter Katoen, Annabelle McIver, Larissa Meinicke, and Carroll C. Morgan. 2010. Linear-invariant generation for probabilistic programs: - Automated support for proof-based methods. In *Static Analysis Symposium (SAS’10) (LNCS)*, Vol. 6337. Springer, 390–406.

- [31] Dexter Kozen. 1981. Semantics of probabilistic programs. *J. Comput. Syst. Sci.* 22, 3 (1981), 328–350.
- [32] Dexter Kozen. 1985. A probabilistic PDL. *J. Comput. Syst. Sci.* 30, 2 (1985), 162–178.
- [33] Antonin Kucera, Javier Esparza, and Richard Mayr. 2006. Model checking probabilistic pushdown automata. *Logical Methods Comput. Sci.* 2, 1 (2006), 12–21.
- [34] Zohar Manna and Amir Pnueli. 1974. Axiomatic approach to total correctness of programs. *Acta Inf.* 3 (1974), 243–263.
- [35] Jeffrey J. McConnell. 2008. *Analysis of Algorithms – An Active Learning Approach*. Jones and Bartlett Publishers.
- [36] Annabelle McIver and Carroll Morgan. 2004. *Abstraction, Refinement and Proof for Probabilistic Systems*. Springer.
- [37] Michael Mitzenmacher and Eli Upfal. 2005. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press.
- [38] David Monniaux. 2001. An abstract analysis of the probabilistic termination of programs. In *Symposium on Static Analysis (SAS'01) (LNCS)*, Vol. 2126. Springer, 111–126.
- [39] Rajeev Motwani and Prabhakar Raghavan. 1995. *Randomized Algorithms*. Cambridge University Press.
- [40] Van Chan Ngo, Quentin Carbonneaux, and Jan Hoffmann. 2018. Bounded expectations: Resource analysis for probabilistic programs. In *Programming Language Design and Implementation (PLDI'18)*. ACM.
- [41] Hanne Riis Nielson. 1987. A Hoare-like proof system for analysing the computation time of programs. *Sci. Comput. Program.* 9, 2 (1987), 107–136.
- [42] Hanne Riis Nielson and Flemming Nielson. 2007. *Semantics with Applications: An Appetizer*. Springer.
- [43] Federico Olmedo, Benjamin Lucien Kaminski, Joost-Pieter Katoen, and Christoph Matheja. 2016. Reasoning about recursive probabilistic programs. *CoRR* abs/1603.02922 (2016). Retrieved from <http://arxiv.org/abs/1603.02922>.
- [44] Federico Olmedo, Benjamin Lucien Kaminski, Joost-Pieter Katoen, and Christoph Matheja. 2016. Reasoning about recursive probabilistic programs. In *31st Annual ACM/IEEE Symposium on Logic in Computer Science (LICS'16)*. IEEE Computer Society, 672–681.
- [45] Saminathan Ponnusamy. 2011. *Foundations of Mathematical Analysis*. Springer Science & Business Media.
- [46] Eric Schechter. 1996. *Handbook of Analysis and Its Foundations*. Elsevier Science.
- [47] Alejandro Serrano, Pedro López-García, and Manuel V. Hermenegildo. 2014. Resource usage analysis of logic programs via abstract interpretation using sized types. *TPLP* 14, 4–5 (2014), 739–754.
- [48] Wolfgang Wechler. 1992. *Universal Algebra for Computer Scientists*. EATCS Monographs on Theoretical Computer Science, Vol. 25. Springer.
- [49] Reinhard Wilhelm, Jakob Engblom, Andreas Ermedahl, Niklas Holsti, Stephan Thesing, David B. Whalley, Guillem Bernat, Christian Ferdinand, Reinhold Heckmann, Tulika Mitra, Frank Mueller, Isabelle Puaut, Peter P.uschner, Jan Staschulat, and Per Stenström. 2008. The worst-case execution-time problem - Overview of methods and survey of tools. *ACM Trans. Embedded Comput. Syst.* 7, 3 (2008), 36:1–36:53.
- [50] Glynn Winskel. 1993. *The Formal Semantics of Programming Languages: An Introduction*. MIT Press.

Received May 2017; revised February 2018; accepted April 2018