

Verifying Probabilistic Programs with Separation Logic

Part 2

Joseph Tassarotti

New York University

Based on joint work with: Alejandro Aguirre, Philipp G. Haselwarter, Kwing Hei-Li, Markus de Medeiros, Puming Liu, Alex Bai, Simon Oddershede Gregersen, and Lars Birkedal

Course page: clutch-project.org/epit2026.html

Rocq code: github.com/logsem/clutch/tree/epit2026

Recap of Part 1

Goal: bridge probabilistic verification and modern separation logic.

Recap of Part 1

Goal: bridge probabilistic verification and modern separation logic.

Three approaches to probabilistic Hoare triples: distributional assertions, expectation transformers, and **liftings**.

Recap of Part 1

Goal: bridge probabilistic verification and modern separation logic.

Three approaches to probabilistic Hoare triples: distributional assertions, expectation transformers, and **liftings**.

Eris: error credits $\mathcal{L}(\epsilon)$ as a resource:

$\{\mathcal{L}(\epsilon) * P\} e \{Q\}$ means Q holds except with probability $\leq \epsilon$

- Splittable: $\mathcal{L}(\epsilon_1 + \epsilon_2) \dashv\vdash \mathcal{L}(\epsilon_1) * \mathcal{L}(\epsilon_2)$, bounded by $\mathcal{L}(1) \vdash \text{False}$
- Averaging rule supports redistributing credit across sample branches
- Resolves aHL limitations: error dependency, propagation, amortization

Example: Morris's Approximate Counter

Problem: count up to n events using fewer than $\log_2 n$ bits.

Example: Morris's Approximate Counter

Problem: count up to n events using fewer than $\log_2 n$ bits.

Idea (Morris, 1978): store an approximation of $\log_2 n$ rather than the count n itself.

That only needs $O(\log_2 \log_2 n)$ bits.

Example: Morris's Approximate Counter

Problem: count up to n events using fewer than $\log_2 n$ bits.

Idea (Morris, 1978): store an approximation of $\log_2 n$ rather than the count n itself.
That only needs $O(\log_2 \log_2 n)$ bits.

But how can you increment? Randomness!

Example: Morris's Approximate Counter

Problem: count up to n events using fewer than $\log_2 n$ bits.

Idea (Morris, 1978): store an approximation of $\log_2 n$ rather than the count n itself. That only needs $O(\log_2 \log_2 n)$ bits.

But how can you increment? Randomness!

$incr \triangleq \lambda c.$


 let $x = !c$ in

 if $\text{rand}(2^x - 1) = 0$ then $c \leftarrow x + 1$ else ()

$read \triangleq \lambda c. 2^{!c} - 1$

Key property: expected value of $read\ c$ is n .

Philippe Flajolet (1948–2011)



- A pioneer of analytic combinatorics for analysis of algorithms.
- 1985 paper analyzed many deep properties of Morris's counter.



- A pioneer of analytic combinatorics for analysis of algorithms.
- 1985 paper analyzed many deep properties of Morris's counter.

EPIT depuis 1973

1. Langages algébriques, Bonascre (J.-P. Crestin et M. Nivat), 1973
2. Complexité des algorithmes, Ile de Berder (Ph. Flajolet), 1974

Philippe Flajolet (1948-



EPIT depuis 197

1. Langages alg
2. Complexité

CIRM - BIBLIOTHÈQUE
 N° d'inventaire : L 37183
 Date : 14/03/2012

A FIRST COURSE IN STOCHASTIC PROCESSES

SECOND EDITION

SAMUEL KARLIN
 STANFORD UNIVERSITY
 AND
 THE WEIZMANN INSTITUTE OF SCIENCE

HOWARD M. TAYLOR
 CORNELL UNIVERSITY

Philippe Flajolet
 (1948 - 2011)



Mathématicien
 et Informaticien,
 Chercheur à l'INRIA



ACADEMIC PRESS New York San Francisco London

A Subsidiary of Harcourt Brace Jovanovich, Publishers



combinatorics for analysis

any deep properties of

t M. Nivat), 1973

Flajolet), 1974

Bounding the counter cell

In practice the cell holds a b -bit value, so we must rule out $x + 1 \geq 2^b$.

Bounding the counter cell

In practice the cell holds a b -bit value, so we must rule out $x + 1 \geq 2^b$.

Our language doesn't have bounded value, but we can simulate by adding an assertion that fails if the new exponent would overflow:

```
incr  $\triangleq$   $\lambda c.$  let  $x = !c$  in  
  if  $\text{rand}(2^x - 1) = 0$  then assert( $x + 1 < 2^b$ );  $c \leftarrow x + 1$  else ()
```

Bounding the counter cell

In practice the cell holds a b -bit value, so we must rule out $x + 1 \geq 2^b$.

Our language doesn't have bounded value, but we can simulate by adding an assertion that fails if the new exponent would overflow:

```
incr  $\triangleq$   $\lambda c.$  let  $x = !c$  in  
  if  $\text{rand}(2^x - 1) = 0$  then assert( $x + 1 < 2^b$ );  $c \leftarrow x + 1$  else ()
```

A verified spec for *incr* must **rule out** the assertion failing.

Spec 1: a state-tracking specification

$counter_val(c, k) \triangleq c \mapsto k.$

Spec 1: a state-tracking specification

$counter_val(c, k) \triangleq c \mapsto k.$

$$\{counter_val(c, k) * k + 1 < 2^b\}$$

$incr\ c$

$$\{counter_val(c, k) \vee counter_val(c, k + 1)\}$$

Spec 1: a state-tracking specification

$counter_val(c, k) \triangleq c \mapsto k.$

$$\{counter_val(c, k) * k + 1 < 2^b\}$$

incr c

$$\{counter_val(c, k) \vee counter_val(c, k + 1)\}$$

Weak, worst-case analysis: after 2^b increments cannot call again.

Lost the $O(\log \log n)$ benefit.

Spec 2: a coarse specification

Hide the value inside the predicate:

$$\mathit{counter}(c) \triangleq \exists k. \mathit{counter_val}(c, k)$$

Spec 2: a coarse specification

Hide the value inside the predicate:

$$\mathit{counter}(c) \triangleq \exists k. \mathit{counter_val}(c, k)$$

Charge a flat error budget per call:

$$\left\{ \mathit{counter}(c) * \epsilon \left(1/2^{2^b-1} \right) \right\}$$

incr c

$$\{ \mathit{counter}(c) \}$$

Spec 2: a coarse specification

Hide the value inside the predicate:

$$\mathit{counter}(c) \triangleq \exists k. \mathit{counter_val}(c, k)$$

Charge a flat error budget per call:

$$\left\{ \mathit{counter}(c) * \frac{1}{2^{2^b-1}} \right\}$$

incr c

$$\{\mathit{counter}(c)\}$$

The budget is the worst-case failure probability: at the boundary state $k = 2^b - 1$, any successful coin overflows.

Spec 2: a coarse specification

Hide the value inside the predicate:

$$\text{counter}(c) \triangleq \exists k. \text{counter_val}(c, k)$$

Charge a flat error budget per call:

$$\left\{ \text{counter}(c) * \frac{1}{2^{2^b-1}} \right\}$$

incr c

$$\{ \text{counter}(c) \}$$

The budget is the worst-case failure probability: at the boundary state $k = 2^b - 1$, any successful coin overflows.

Better, but still pessimistic: every call pays the boundary cost, even though early calls almost never overflow.

Spec 3: tight analysis with bounded-call budget

Compute the **exact** overflow probability under n calls from state k :

$$err_overflow(0, k) = 0$$

$$err_overflow(n + 1, k) = \frac{1}{2^k} \cdot e_{inc} + \left(1 - \frac{1}{2^k}\right) \cdot err_overflow(n, k)$$

where $e_{inc} = err_overflow(n, k + 1)$ if $k + 1 < 2^b$, else 1.

Spec 3: tight analysis with bounded-call budget

Compute the **exact** overflow probability under n calls from state k :

$$err_overflow(0, k) = 0$$

$$err_overflow(n + 1, k) = \frac{1}{2^k} \cdot e_{inc} + \left(1 - \frac{1}{2^k}\right) \cdot err_overflow(n, k)$$

where $e_{inc} = err_overflow(n, k + 1)$ if $k + 1 < 2^b$, else 1.

Bake the remaining budget into the representation predicate:

$$counter_budget(c, n) \triangleq \exists k. counter_val(c, k) * \neg(err_overflow(n, k))$$

Spec 3: tight analysis with bounded-call budget

Compute the **exact** overflow probability under n calls from state k :

$$\begin{aligned}err_overflow(0, k) &= 0 \\err_overflow(n + 1, k) &= \frac{1}{2^k} \cdot e_{inc} + \left(1 - \frac{1}{2^k}\right) \cdot err_overflow(n, k)\end{aligned}$$

where $e_{inc} = err_overflow(n, k + 1)$ if $k + 1 < 2^b$, else 1.


Bake the remaining budget into the representation predicate:

$$counter_budget(c, n) \triangleq \exists k. counter_val(c, k) * \neg(err_overflow(n, k))$$

$$\{counter_budget(c, n + 1)\} incr\ c\ \{counter_budget(c, n)\}$$

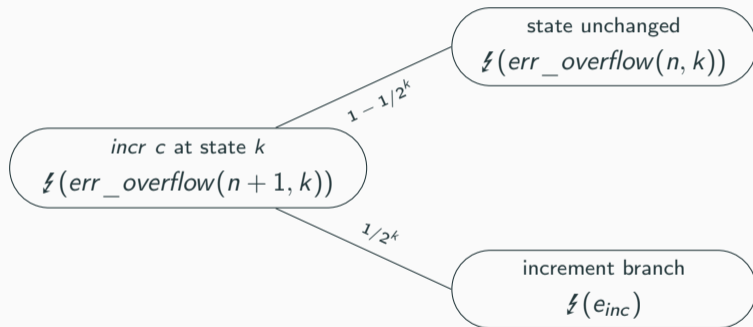
Proof sketch: average across the branches

Initial: state k , budget $\zeta(\text{err_overflow}(n+1, k))$. Split across the outcomes of $\text{rand}(2^k - 1)$ by the averaging rule:



Proof sketch: average across the branches

Initial: state k , budget $\zeta(\text{err_overflow}(n+1, k))$. Split across the outcomes of $\text{rand}(2^k - 1)$ by the averaging rule:



Averaging condition holds by definition of $\text{err_overflow}(n+1, k)$:

$$\frac{1}{2^k} \cdot e_{inc} + \left(1 - \frac{1}{2^k}\right) \cdot \text{err_overflow}(n, k) = \text{err_overflow}(n+1, k)$$

Proof sketch 2: Reason about the assert

If we increment, then before the assert we have e_{inc} where $e_{inc} = err_overflow(n, k + 1)$ if $k + 1 < 2^b$, else 1.

If $k + 1 < 2^b$ the assert passes.

Proof sketch 2: Reason about the assert

If we increment, then before the assert we have e_{inc} where $e_{inc} = err_overflow(n, k + 1)$ if $k + 1 < 2^b$, else 1.

If $k + 1 < 2^b$ the assert passes.

If $k + 1 = 2^b$, the assert would fail!

But then $e_{inc} = 1$ and $\neg(1) \vdash \text{False}$

Almost-Sure Termination with Eris_t

Background: Partial vs. Total Correctness

As described so far, Hoare triples have expressed **partial correctness**:

$\{P\} e \{Q\} \equiv$ if e starts in a state satisfying P and **terminates**, the final state satisfies Q .

Trivially satisfied by non-terminating programs.

Background: Partial vs. Total Correctness

As described so far, Hoare triples have expressed **partial correctness**:

$$\{P\} e \{Q\} \equiv \text{if } e \text{ starts in a state satisfying } P \text{ and } \mathbf{\textit{terminates}}, \text{ the final state satisfies } Q.$$

Trivially satisfied by non-terminating programs.

Can also consider **total correctness** that requires termination:

$$[P] e [Q] \equiv \text{if } e \text{ starts in a state satisfying } P, \text{ then } e \mathbf{\textit{terminates}} \text{ and the final state satisfies } Q.$$

Total Correctness for Randomized Programs

Eris is similarly a **partial** logic.

We have developed a total version called Eris_t .

In a probabilistic setting, we ask for **almost-sure termination**: the program terminates with probability 1.

Total Soundness.

If $[\not\perp(\epsilon)] e [Q]$, then for any starting state, with probability $\geq 1 - \epsilon$, running e will terminate and end in a state satisfying Q .

Total Soundness.

If $[\not\downarrow(\epsilon)] e [Q]$, then for any starting state, with probability $\geq 1 - \epsilon$, running e will terminate and end in a state satisfying Q .

Prove almost-sure termination by **continuity**:

If for all $\epsilon > 0$, $[\not\downarrow(\epsilon)] e [Q]$, then e terminates almost-surely from any starting state.

Total Soundness.

If $[\not\downarrow(\epsilon)] e [Q]$, then for any starting state, with probability $\geq 1 - \epsilon$, running e will terminate and end in a state satisfying Q .

Prove almost-sure termination by **continuity**:

If for all $\epsilon > 0$, $[\not\downarrow(\epsilon)] e [Q]$, then e terminates almost-surely from any starting state.

Internalize this continuity as a rule:

$$\frac{\forall \epsilon > 0. [P * \not\downarrow(\epsilon)] e [Q]}{[P] e [Q]}$$

How do the rules change?

The partial logic supports the Löb induction rule we saw before.

Total version removes Löb induction. Must reason about recursion by **induction** on a well-founded relation.

Otherwise all the same rules.

The Trouble with Induction for Almost-Sure Termination

```
let rec coinToss () =  
  if rand(1) = 0 then ()  
  else coinToss ()
```

This terminates with probability 1.

The Trouble with Induction for Almost-Sure Termination

```
let rec coinToss () =  
  if rand(1) = 0 then ()  
  else coinToss ()
```

This terminates with probability 1.

But what is “getting smaller” on each recursive call?

Key Idea: Error Credit Induction

```
let rec coinToss () =  
  if rand(1) = 0 then ()  
  else coinToss ()
```

Let's suppose we start with $\zeta(\epsilon)$ for some $\epsilon > 0$.

Recall: $\zeta(1) \vdash \text{False}$

Key Idea: Error Credit Induction

```
let rec coinToss () =  
  if rand(1) = 0 then ()  
  else coinToss ()
```

Let's suppose we start with $\zeta(\epsilon)$ for some $\epsilon > 0$.

Recall: $\zeta(1) \vdash \text{False}$

Each iteration we can **double** the credits we have for recursive call.


Induct on number of rounds until getting $\zeta(1)$.

Error Credit Induction for coinToss

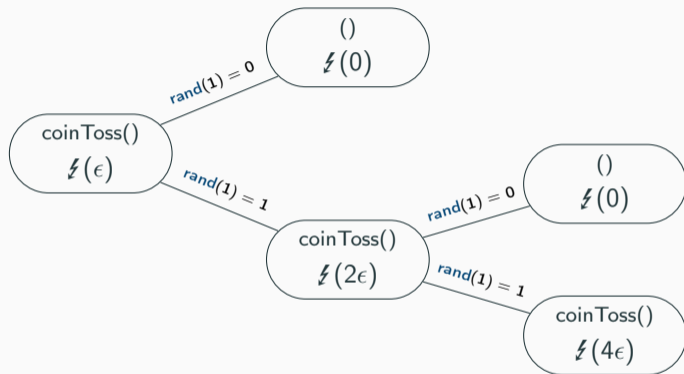
coinToss()

ϵ

Error Credit Induction for coinToss



Error Credit Induction for coinToss



No matter how small ϵ , eventually get $\epsilon(1)$.

Error Credit Induction Rule

This induction principle is abstractly captured by:

$$\frac{\epsilon > 0 \quad k > 1 \quad \forall \epsilon'. (\downarrow(k \cdot \epsilon') \multimap P) * \downarrow(\epsilon') \vdash P}{\downarrow(\epsilon) \vdash P}$$

Error Credit Induction Rule

This induction principle is abstractly captured by:

$$\frac{\epsilon > 0 \quad k > 1 \quad \forall \epsilon'. (\zeta(k \cdot \epsilon') \multimap P) * \zeta(\epsilon') \vdash P}{\zeta(\epsilon) \vdash P}$$

To use: own $\zeta(\epsilon)$ and pick amplification factor $k > 1$.

Get an arbitrary $\zeta(\epsilon')$ and an IH usable after paying $\zeta(k \cdot \epsilon')$.

Error Credit Induction Rule

This induction principle is abstractly captured by:

$$\frac{\epsilon > 0 \quad k > 1 \quad \forall \epsilon'. (\not\downarrow(k \cdot \epsilon') \multimap P) * \not\downarrow(\epsilon') \vdash P}{\not\downarrow(\epsilon) \vdash P}$$

To use: own $\not\downarrow(\epsilon)$ and pick amplification factor $k > 1$.

Get an arbitrary $\not\downarrow(\epsilon')$ and an IH usable after paying $\not\downarrow(k \cdot \epsilon')$.

Soundness: induct on $\lceil -\log_k(\epsilon) \rceil$, the number of scalings until reaching $\not\downarrow(1)$.

Eris_t is relatively **complete** for allocation-free programs.

Meaning: if e is an expression that:

1. almost-surely terminates
2. does not get stuck
3. does not allocate

then you can derive a triple in Eris_t for e .

Only bounding error/termination might seem **limited**. But can characterize arbitrary distributions this way.

Distribution Correctness

Only bounding error/termination might seem **limited**. But can characterize arbitrary distributions this way.

Suppose f is an implementation of a sampler for a distribution D and we want to prove f is correct.

Marionneau et al. extend Eris_t soundness proof to show that it suffices to establish

$$\frac{\mathbb{E}_{x \sim D}[\epsilon(x)] = \epsilon_0}{[\not\epsilon(\epsilon_0)] f () [x. \not\epsilon(\epsilon(x))]}$$

Termination:

- Higher-order rejection sampler
- Incremental randomized SAT solving
- 1-D symmetric random walk

Sampler Correctness:

- Geometric
- Binomial
- Negative Binomial
- Beta Binomial
- ...



Rocq Exercises