

Probabilistic Programming

Lecture #16+#17: Expected Runtime Analysis

Joost-Pieter Katoen



RWTH Lecture Series on Probabilistic Programming 2018

Overview

- 1 Motivation
- 2 An unsound approach
- 3 The expected runtime transformer
- 4 Properties
- 5 Proof rules for runtimes of loops
- 6 Proving positive almost-sure termination
- 7 Case studies

Overview



- 1 Motivation
- 2 An unsound approach
- 3 The expected runtime transformer
- 4 Properties
- 5 Proof rules for runtimes of loops
- 6 Proving positive almost-sure termination
- 7 Case studies

The runtime of a probabilistic program

The **runtime** of a probabilistic program depends
on the **input** and
on the internal **randomness** of the program.

The runtime of a probabilistic program is random

```
int i := 0;
repeat {i++; (c := false [0.5] c := true)}
until (c)
```



The **expected runtime** is $1 + 3 \cdot 1/2 + 5 \cdot 1/4 + \dots + (2n+1) \cdot 1/2^n = 5$.

Efficiency of randomised algorithms

Quicksort:

```
QS(A) =
  if |A| <= 1 { return A; }
  i := ceil(|A|/2);
  A< := {a in A | a < A[i]};
  A> := {a in A | a > A[i]};
  return QS(A<) ++ A[i] ++ QS(A>)
```

Worst case complexity:
 $O(N^2)$ comparisons



Randomised Quicksort:

```
rQS(A) =
  if |A| <= 1 { return A; }
  i := Unif[1...|A|];
  A< := {a in A | a < A[i]};
  A> := {a in A | a > A[i]};
  return rQS(A<) ++ A[i] ++ rQS(A>)
```

Worst case complexity:
 $O(N \log N)$ expected comparisons



Expected runtimes

Expected run-time of program P on input s :

$$\sum_{i=1}^{\infty} i \cdot Pr \left(\begin{array}{l} "P \text{ terminates after} \\ i \text{ steps on input } s" \end{array} \right)$$

Coupon collector's problem

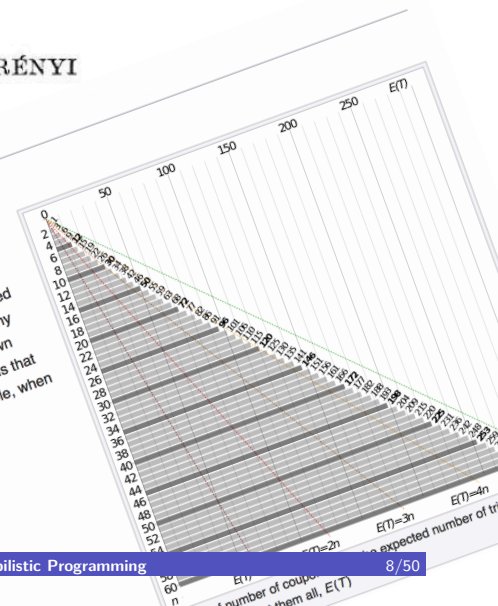
ON A CLASSICAL PROBLEM OF PROBABILITY THEORY

by
P. ERDŐS and A. RÉNYI

Coupon collector's problem

From Wikipedia, the free encyclopedia

In **probability theory**, the **coupon collector's problem** describes the "collect all coupons and win" contests. It asks the following question: Suppose that there is an urn of n different **coupons**, from which coupons are being collected, equally likely, with replacement. What is the probability that more than t sample trials are needed to collect all n coupons? An alternative statement is: Given n coupons, how many coupons do you expect you need to draw with replacement before having drawn each coupon at least once? The mathematical analysis of the problem reveals that each coupon is expected number of trials needed grows as $\Theta(n \log(n))$.^[1] For example, when $n = 50$, about 225 trials to collect all 50 coupons.



Coupon collector's problem

```

cp := [0,...,0]; // no coupons yet
i := 1; // coupon to be collected next
x := 0; // number of coupons collected
while (x < N) {
  while (cp[i] != 0) {
    i := uniform(1..N) // next coupon
  }
  cp[i] := 1; // coupon i obtained
  x++; // one coupon less to go
}

```

The expected runtime of this program is in $\Theta(N \cdot \log N)$.

Randomised primality test


Problem: is N prime or not?

Basic structure of a **randomised** primality test:

1. Randomly pick a number a , say
2. Do the **primality test**: Check some equality involving a and N
3. If equality fails, N is composite (with witness a)
4. Otherwise repeat the process.

If after $K > 0$ iterations, N is not found to be composite, then N is **probably prime**.

Closest-pair problem



Closest-pair problem: find two distinct points $u, v \in \mathbb{R}^2$ among N points in the plane that minimise the Euclidean distance among all pairs of these points.

A naive deterministic approach takes $O(N^2)$. More efficient version in $O(N \cdot \log N)$.

Rabin's randomised algorithm has an **expected runtime** in $O(N)$.

Some primality tests

► **Fermat primality test**:

Select $a \in \mathbb{Z}$ relative prime to N . If $a^{N-1} \bmod N \neq 1$, then N is composite.

► **Rabin-Miller test**:

Select $0 < a < N$. Let $2^s \cdot d = N-1$ where d is odd. If $a^d \not\equiv 1 \pmod{N}$ and $a^{2^r \cdot d} \not\equiv -1 \pmod{N}$ for all $0 \leq r \leq s-1$, then N is composite.

► **Solovay and Strassen test**:

For N odd, pick $a < N$. If $a^{(N-1)/2} \not\equiv \pm 1 \pmod{N}$, then N is composite.

Adleman and Huang (1992) provided a randomised primality test that terminates with **expected polynomial runtime** and certainly provides the correct answer.¹

¹Decision problems with this characteristic constitute the complexity class ZPP (zero-error probabilistic polynomial time).

The aim of this lecture

A wp-calculus to reason about runtimes at [the source code level](#).

No “descend” into the underlying probabilistic model.

The calculus should be [compositional](#).

Hurdles in runtime analysis

1. Programs may admit [diverging runs](#) while still having a [finite expected runtime](#)

```
while (x > 0) { x-- [1/2] skip }
```

admits a diverging run but has expected runtime $O(x)$.

2. Having a finite expected time is [not compositional](#) w.r.t. sequencing
3. Expected runtimes are [extremely sensitive](#) to variations in probabilities

```
while (x > 0) { x-- [1/2+p] x++ } // 0 <= p <= 1/2
```

- ▶ For $p=0$, the expected runtime is infinite.
- ▶ For arbitrary small $p > 0$, the expected runtime is $\frac{1}{2} \cdot p \cdot x$, linear in x .

Proving [positive](#) almost-sure termination

- ▶ [What?](#) AST+[termination in finite expected time](#)
- ▶ [Generalise. How?](#)
 - ▶ Provide an weakest-precondition calculus
 - ▶ for [expected runtimes](#)
- ▶ [Why?](#)
 - ▶ Reason about the efficiency of randomised algorithms
 - ▶ Reason about simulation efficiency of Bayesian networks
 - ▶ Is compositional and reasons at the program's code

Overview

- 1 Motivation
- 2 **An unsound approach**
- 3 The expected runtime transformer
- 4 Properties
- 5 Proof rules for runtimes of loops
- 6 **Proving positive almost-sure termination**
- 7 Case studies

Re-use weakest preconditions?

Idea: equip the program with a counter `rc`
and use standard wp-reasoning to determine its expected value.

Determine $wp(P, rc)$ for program P .



Dexter Kozen
A probabilistic PDL
1983

An example

Consider the program Q :

```
x := 1;
while (x > 0) { x := 0 [1/2] while(true) { skip } }
```

Equipping Q with a runtime counter yields Q_{rc} :

```
x := 1; rc := 0;
while (x > 0) {
  rc++;
  (x := 0 [1/2] while(true) { rc++ ; skip })
}
```

As $wp(\text{inner loop}, f) = 0$ for every f , it follows $\Phi_{Q_{rc}} \leq \Phi_{P_{rc}}$.

Thus, $\Phi_{Q_{rc}}(I) \leq \Phi_{P_{rc}}(I) \leq I$ for $I = rc + [x > 0] \cdot 2$.

This contradicts the fact that the true expected runtime of Q is ∞ .

An example

Consider the program P :

```
x := 1;
while (x > 0) { x := 0 [1/2] skip }
```

Equipping P with a runtime counter yields P_{rc} :

```
x := 1; rc := 0;
while (x > 0) { rc++; (x := 0 [1/2] skip) }
```

It follows $\Phi(I) \leq I$ for $I = rc + [x > 0] \cdot 2$.

In total, we thus obtain $wp(P_{rc}, rc) = 2$.


Overview

- 1 Motivation
- 2 An unsound approach
- 3 The expected runtime transformer
- 4 Properties
- 5 Proof rules for runtimes of loops
- 6 Proving positive almost-sure termination
- 7 Case studies

The basic idea

Let $\text{ert}() : \text{pGCL} \rightarrow (\mathbb{T} \rightarrow \mathbb{T})$ where:

- ▶ $\text{ert}(P, t)(s)$ is the expected runtime of P on input state s if t captures the runtime of the computation following P .
- ▶ $\text{ert}(P, 0)(s)$ is the expected runtime of P on input state s .



The runtime model

We assume the following runtimes:

- ▶ Executing a **skip**-statement takes a single time unit
- ▶ Executing an (ordinary or random) assignment takes a single time unit
- ▶ Evaluating a guard takes a single time unit
- ▶ Flipping a coin in a probabilistic choice takes a single time unit
- ▶ Sequential composition does not take time

The ert-calculus can be easily adapted to other runtime models.

Runtimes

Expectations

A **expectation** $f : \mathbb{S} \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$.

Let \mathbb{E} be the set of all expectations and let \sqsubseteq be defined for $f, g \in \mathbb{E}$ by:

$$f \sqsubseteq g \quad \text{if and only if} \quad f(s) \leq g(s) \quad \text{for all } s \in \mathbb{S}.$$

Runtimes

A **runtime** $t : \mathbb{S} \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$.

Let \mathbb{T} denote the set of all runtimes and let \leq be defined for $t, u \in \mathbb{T}$ by:

$$t \leq u \quad \text{if and only if} \quad t(s) \leq u(s) \quad \text{for all } s \in \mathbb{S}.$$

A runtime **transformer** is defined in a similar way as an expectation transformer

Expected runtime transformer for pGCL

Syntax

- ▶ **skip**
- ▶ **diverge**
- ▶ $x := E$
- ▶ $x := \text{mu}$
- ▶ $P_1 ; P_2$
- ▶ **if** (G) P_1 **else** P_2
- ▶ $P_1 \text{ [p] } P_2$
- ▶ **while**(G) P

Expected runtime $\text{ert}(P, t)$

- ▶ $1 + t$
- ▶ ∞
- ▶ $1 + t[x := E]$
- ▶ $1 + \lambda s. \int_{\mathbb{Q}} (\lambda v. t(s[x := v])) d\mu_s$
- ▶ $\text{ert}(P_1, \text{ert}(P_2, t))$
- ▶ $1 + [G] \cdot \text{ert}(P_1, t) + [\neg G] \cdot \text{ert}(P_2, t)$
- ▶ $1 + p \cdot \text{ert}(P_1, t) + (1-p) \cdot \text{ert}(P_2, t)$
- ▶ $\text{lfp } X. (1 + [G] \cdot \text{ert}(P, X) + [\neg G] \cdot t)$

lfp is the least fixed point operator wrt. the ordering \leq on runtimes

Examples

Elementary properties

- ▶ **Continuity:** $ert(P, t)$ is continuous on (\mathbb{T}, \leq)
- ▶ **Monotonicity:** $t \leq t'$ implies $ert(P, t) \leq ert(P, t')$
- ▶ **Constant propagation:** $ert(P, k + t) = k + ert(P, t)$
- ▶ **Preservation of ∞ :** $ert(P, \infty) = \infty$
- ▶ **Connection to wp:** $ert(P, t) = ert(P, 0) + wp(P, t)$
- ▶ **Affinity:** $ert(P, a \cdot t + t') = ert(P, 0) + r \cdot ert(P, t) + ert(P, t')$

Overview

- 1 Motivation
- 2 An unsound approach
- 3 The expected runtime transformer
- 4 Properties**
- 5 Proof rules for runtimes of loops
- 6 Proving positive almost-sure termination
- 7 Case studies

(Positive) almost-sure termination

For every pGCL program P and input state s :

$$\underbrace{ert(P, 0)(s) < \infty}_{\text{positive a.s-termination on } s} \quad \text{implies} \quad \underbrace{wp(P, 1)(s) = 1}_{\text{almost-sure termination on } s}$$

Moreover:

$$\underbrace{ert(P, 0) \leq \infty}_{\text{universal positive a.s-termination}} \quad \text{implies} \quad \underbrace{wp(P, 1) = 1}_{\text{universal almost-sure termination}}$$

A Markov chain perspective on runtimes

- ▶ Consider $ert(P, t)$ for pCGL program P
- ▶ Consider the Markov chain $\llbracket P \rrbracket$ of program P
- ▶ Attach rewards to each Markov chain state in $\llbracket P \rrbracket$:
 - ▶ State $\langle \downarrow, s \rangle$ gets reward $t(s)$
 - ▶ State $\langle \text{skip}, s \rangle$ gets reward one
 - ▶ State $\langle \text{diverge}, s \rangle$ gets reward ∞
 - ▶ State $\langle x := E, s \rangle$ gets reward one
 - ▶ State $\langle x \approx \mu, s \rangle$ gets reward one
 - ▶ State $\langle \text{if } G \dots, s \rangle$ gets reward one
 - ▶ State $\langle P[p]Q, s \rangle$ gets reward one
 - ▶ State $\langle \text{while}(G)P' \dots, s \rangle$ gets reward one
 - ▶ All other states get reward zero

Correspondence between $ert()$ and Markov chains

Compatibility theorem

For every pGCL program P and input s :

$$ert(P, \mathbf{0})(s) = ER^{\llbracket P \rrbracket}(s, \diamond \text{sink})$$

In words: the $ert(P, \mathbf{0})$ for input s equals the expected reward to reach final state *sink* in MC $\llbracket P \rrbracket$ where reward function r in $\llbracket P \rrbracket$ is defined as defined on the previous slide.

Example

Backward compatibility

Deterministic programs

For any GCL program P , $ert(P, \mathbf{0})$ equals the number of executed computational steps² of P until P terminates.

²This equals the number of skip statements, guard evaluations and assignments.

Overview

- 1 Motivation
- 2 An unsound approach
- 3 The expected runtime transformer
- 4 Properties
- 5 Proof rules for runtimes of loops
- 6 Proving positive almost-sure termination
- 7 Case studies

Runtime invariants

Runtime invariants

Let Φ_t be the wp-characteristic function of $P' = \text{while}(G)\{P\}$ with respect to post-runtime $t \in \mathbb{T}$ and let $I \in \mathbb{T}$. Then:

1. I is a **runtime-superinvariant** of P' w.r.t. t iff $\Phi_t(I) \leq I$.
2. I is a **runtime-subinvariant** of P' w.r.t. t iff $I \leq \Phi_t(I)$.

If I is a **runtime-superinvariant** of $\text{while}(G)\{P\}$ with respect to $t \in \mathbb{T}$, then:

$$\text{ert}(\text{while}(G)\{P\}, t) \leq I$$

Loops

Reasoning about loops requires — like for wp — invariants.

Example

A **w**rong proof rule for lower bonds

Probabilistic programs do **not** satisfy:

if $I \leq \Phi_t(I)$ then $I \leq \text{ert}(\text{while}(G)P, t)$.

These “metering” functions I do work for ordinary programs

[Frohn *et al.*, IJCAR 2016]

Runtime ω -invariants

Runtime ω -invariants

Let $n \in \mathbb{N}$, $t \in \mathbb{T}$ and Φ_t the ert-characteristic function of $\text{while}(G)\{P\}$.

The monotonically increasing³ sequence $(I)_{n \in \mathbb{N}}$ is a **runtime- ω -subinvariant** of the loop w.r.t. runtime t iff

$$I_0 \leq \Phi_t(0) \quad \text{and} \quad I_{n+1} \leq \Phi_t(I_n) \quad \text{for all } n.$$

In a similar way, **runtime ω -superinvariants** can be defined, but we will not use them here.

³But not necessarily strictly increasing.

A counterexample

```
while (true) { skip [1/2] x++ }
```

- ▶ Characteristic functional $F(X) = 1 + 1/2(1 + 1 + X[x/x+1])$
- ▶ Least fixed point is **4** as $F(4) = 2 + 1/2 \cdot 4 = 4$
- ▶ $4 + 2^i$ is a fixed point of F too:

$$F(4 + 2^i) = 2 + \frac{1}{2}(4 + 2^{i+1}) = 4 + 2^i$$

- ▶ Thus: $4 + 2^i \leq F(4 + 2^i)$ but $4 + 2^i \not\leq 4 = \text{lfp } F$
- ▶ In fact, $4 + 2^{i+c}$ is a fixed point of F for any c :

$$F(4 + 2^{i+c}) = 2 + \frac{1}{2}(4 + 2^{i+c+1}) = 4 + 2^{i+c}$$

Lower bounds

Runtime lower bounds

If I_n is a runtime ω -subinvariant of $\text{while}(G)\{P\}$ with respect to t , then:

$$\sup_n I_n \leq \text{ert}(\text{while}(G)P, t)$$

Example

Consider the same program as for proving an upper bound on the expected runtime.

Overview

- 1 Motivation
- 2 An unsound approach
- 3 The expected runtime transformer
- 4 Properties
- 5 Proof rules for runtimes of loops
- 6 Proving positive almost-sure termination
- 7 Case studies

Proving that PAST is not compositional (1)

```
while (x > 0) { x := x-1 }
```

It is easy to check that a **lower ω -invariant** is:

$$J_n = 1 + \underbrace{[0 < x < n] \cdot 2x}_{\text{on iteration}} + \underbrace{[x \geq n] \cdot (2n-1)}_{\text{on termination}}$$

Thus we obtain that:

$$\lim_{n \rightarrow \infty} (1 + [0 < x < n] \cdot 2x + [x \geq n] \cdot (2n-1)) = 1 + [x > 0] \cdot 2x$$

is a **lower bound** on the runtime of the above program.

PAST is not compositional

Consider the two probabilistic programs:

```
int x := 1;
bool c := true;
while (c) {
  c := false [0.5] c := true;
  x := 2*x
}
```

Finite expected termination time

```
while (x > 0) {
  x--
}
```

Finite termination time

Running the right after the left program
yields an **infinite** expected termination time

Proving that PAST is not compositional (2)

```
while (c) { {c := false [0.5] c := true}; x := 2*x; }
while (x > 0) { x := x-1 }
```

Template for a lower ω -invariant of composed program:

$$I_n = 1 + \underbrace{[c \neq 1] \cdot (1 + [x > 0] \cdot 2x)}_{\text{on termination}} + \underbrace{[c = 1] \cdot (a_n + b_n \cdot [x > 0] \cdot 2x)}_{\text{on iteration}}$$

The constraints on being a lower ω -invariant yield:

$$a_0 \leq 2 \quad \text{and} \quad a_{n+1} \leq 7/2 + 1/2 \cdot a_n \quad \text{and} \quad b_0 \leq 0 \quad \text{and} \quad b_{n+1} \leq 1 + b_n$$

This admits the solution $a_n = 7 - 5/2^n$ and $b_n = n$. Then: $\lim_{n \rightarrow \infty} I_n = \infty$.

Proving PAST

The ert-transformer enables to prove
that a program is positively almost-surely terminating
in a [compositional manner](#),
although PAST itself is not a compositional property.

Coupon collector's problem

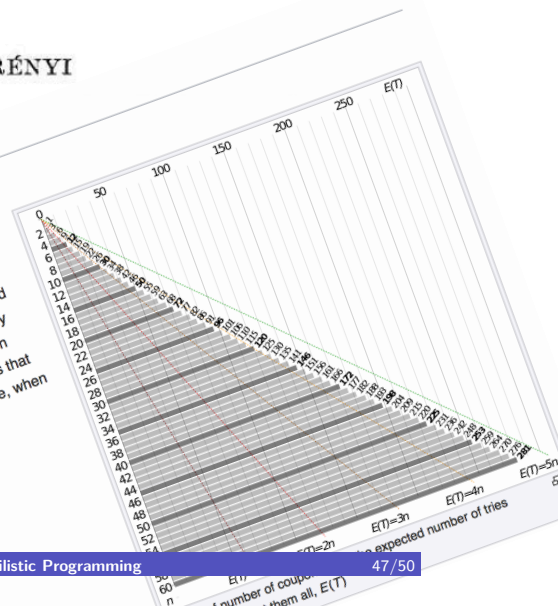
ON A CLASSICAL PROBLEM OF PROBABILITY THEORY

by
P. ERDŐS and A. RÉNYI

Coupon collector's problem

From Wikipedia, the free encyclopedia

In [probability theory](#), the **coupon collector's problem** describes the "collect all coupons and win" contests. It asks the following question: Suppose that there is an urn of n different [coupons](#), from which coupons are being collected, equally likely, with replacement. What is the probability statement is: Given n coupons, how many to collect all n coupons do you expect you need to draw with replacement before having drawn each coupon at least once? The mathematical analysis of the problem reveals that the [expected number](#) of trials needed grows as $\Theta(n \log(n))$.^[1] For example, when $n = 50$, about 225^[2] trials to collect all 50 coupons.



Overview

- 1 Motivation
- 2 An unsound approach
- 3 The expected runtime transformer
- 4 Properties
- 5 Proof rules for runtimes of loops
- 6 Proving positive almost-sure termination
- 7 Case studies

Coupon collector's problem

```
cp := [0,...,0]; i := 1; x := 0; // no coupons yet
while (x < N) {
  while (cp[i] != 0) {
    i := uniform(1..N) // next coupon
  }
  cp[i] := 1; // coupon i obtained
  x++; // one coupon less to go
}
```


Using the ert-calculus one can prove that:

$$\text{ert}(\text{cpcl}, 0) = 4 + [N > 0] \cdot 2N \cdot (2 + H_{N-1}) \in \Theta(N \cdot \log N)$$

As Harmonic number $H_{N-1} \in \Theta(\log N)$.

By systematic program verification. Machine checkable.

Random walk



Using the ert-calculus one can prove that its expected runtime is ∞ .

By systematic formal verification. Machine checkable.

Randomised binary search

```

proc BinSearch {
  mid := Unif(left, right); // pick mid uniformly
  if (left < right) {
    if (A[mid] < val) {
      left := min(mid+1, right);
      call BinSearch
    } else {
      if (A[mid] > val) {
        right := max(mid-1, left);
        call BinSearch
      } else { skip }
    } else { skip }
  }
}

```

Using the ert-calculus one can prove that its expected runtime is $\Theta(\log N)$.

By systematic formal verification. Machine checkable.