

Automated Analysis of Probabilistic Programs

Joost-Pieter Katoen

Software Modeling and Verification Group
RWTH Aachen University

joint work with [Friedrich Gretz](#) and [Annabelle McIver](#)

September 24, 2013

Overview

- 1 Introduction
- 2 Probabilistic guarded command language
- 3 Operational semantics of pGCL
- 4 Denotational semantics of pGCL
- 5 Denotational vs. operational semantics of pGCL
- 6 Synthesizing loop invariants
- 7 Epilogue



Probabilistic programs

What are probabilistic programs?

Sequential, possibly non-deterministic, programs with **random assignments**.

Applications

Cryptography, privacy, quantum computing, and randomized algorithms.

The scientific challenge

- ▶ Such programs are **small**, but **hard** to understand and analyse¹.
 - ▶ Problems: infinite variable domains, (lots of) parameters, and loops.
- ⇒ **Our aim: push the limits of automated analysis**

¹Their analysis is undecidable.

Once upon a time



Duelling cowboys

```
int cowboyDuel(float a, b) { // 0 < a < 1, 0 < b < 1
    int t := A [] t := B; // decide cowboy for first shooting
    turn
    bool c := true;
    while (c) {
        if (t = A) {
            (c := false [a] t := B); // A shoots B with prob. a
        } else {
            (c := false [b] t := A); // B shoots A with prob. b
        }
    }
    return t; // the survivor
}
```

Claim:

Cowboy A wins the duel with probability at least $\frac{(1-b) \cdot a}{a+b-a \cdot b}$.

Playing with geometric distributions

- ▶ X is a random variable, geometrically distributed with parameter p
- ▶ Y is a random variable, geometrically distributed with parameter q

Q: generate a sample x , say, according to the random variable $X - Y$

```
int XminY1(float p, q){ // 0 <= p, q <= 1
    int x := 0;
    bool flip := false;
    while (not flip) { // take a sample of X to increase x
        (x += 1 [p] flip := true);
    }
    flip := false;
    while (not flip) { // take a sample of Y to decrease x
        (x -= 1 [q] flip := true);
    }
    return x; // a sample of X-Y
}
```

An alternative program

```
int XminY2(float p, q){
  int x := 0;
  bool flip := false;
  (flip := false [0.5] flip := true); // flip a fair coin
  if (not flip) {
    while (not flip) { // sample X to increase x
      (x += 1 [p] flip := true);
    }
  } else {
    flip := false; // reset flip
    while (not flip) { // sample Y to decrease x
      x -= 1;
      (skip [q] flip := true);
    }
  }
  return x; // a sample of X-Y
}
```

Program equivalence

```

int XminY1(float p, q){
  int x, f := 0, 0;
  while (f = 0) {
    (x += 1 [p] f := 1);
  }
  f := 0;
  while (f = 0) {
    (x -= 1 [q] f := 1);
  }
  return x;
}

```

```

int XminY2(float p, q){
  int x, f := 0, 0;
  (f := 0 [0.5] f := 1);
  if (f = 0) {
    while (f = 0) {
      (x += 1 [p] f := 1);
    }
  } else {
    f := 0;
    while (f = 0) {
      x -= 1;
      (skip [q] f := 1);
    }
  }
  return x;
}

```

Claim: [Kiefer et. al., 2012]

Both programs are equivalent for $(p, q) = (\frac{1}{2}, \frac{2}{3})$. **Q:** No other ones?

Correctness of probabilistic programs

Question:

How to verify the correctness of such programs? In an automated way?

Apply model checking?

- ▶ Apply **MDP model checking**. LiQuor, PRISM
 - ⇒ works for program instances, but no general solution.
- ▶ Use **abstraction-refinement** techniques. PASS, POGAR
 - ⇒ loop analysis with real variables does not work well.
- ▶ Check **language equivalence**. APEX
 - ⇒ cannot deal with parameterised probabilistic programs.
- ▶ Apply **parameterised** probabilistic model checking. PARAM
 - ⇒ deals with fixed-sized probabilistic programs.

Apply deductive verification!

[McIver & Morgan]

Duelling cowboys

```
int cowboyDuel(float a, b) { // 0 < a < 1, 0 < b < 1
  int t := A [] t := B; // decide which cowboy has first
    shooting turn
  bool c := true;
  while (c) {
    if (t = A) {
      (c := false [a] t := B); // A shoots B with prob. a
    } else {
      (c := false [b] t := A); // B shoots A with prob. b
    }
  }
  return t; // the survivor
}
```

We can infer:

Cowboy A wins the duel with probability at least $\frac{(1-b) \cdot a}{a + b - a \cdot b}$.

Program equivalence

```

int XminY1(float p, q){
  int x, f := 0, 0;
  while (f = 0) {
    (x += 1 [p] f := 1);
  }
  f := 0;
  while (f = 0) {
    (x -= 1 [q] f := 1);
  }
  return x;
}

```

```


int XminY2(float p, q){
  int x, f := 0, 0;
  (f := 0 [0.5] f := 1);
  if (f = 0) {
    while (f = 0) {
      (x += 1 [p] f := 1);
    }
  } else {
    f := 0;
    while (f = 0) {
      x -= 1;
      (skip [q] f := 1);
    }
  }
  return x;
}

```

Our analysis yields:

Both programs are equivalent for any q with $q = \frac{1}{2-p}$.

Graphically this means ...



Both programs yield the same expected outcome for all points on the curve

$$q = \frac{1}{2-p}$$

Roadmap of the talk

- 1 Introduction
- 2 Probabilistic guarded command language
- 3 Operational semantics of pGCL
- 4 Denotational semantics of pGCL
- 5 Denotational vs. operational semantics of pGCL
- 6 Synthesizing loop invariants
- 7 Epilogue

Overview

- 1 Introduction
- 2 Probabilistic guarded command language**
- 3 Operational semantics of pGCL
- 4 Denotational semantics of pGCL
- 5 Denotational vs. operational semantics of pGCL
- 6 Synthesizing loop invariants
- 7 Epilogue

Dijkstra's guarded command language



- ▶ `skip` empty statement
- ▶ `abort` abortion
- ▶ `x := E` assignment
- ▶ `prog1 ; prog2` sequential composition
- ▶ `if (G) prog1 else prog2` choice
- ▶ `prog1 [] prog2` non-deterministic choice
- ▶ `while (G) prog` iteration

Probabilistic guarded command language pGCL



- ▶ `skip` empty statement
- ▶ `abort` abortion
- ▶ `x := E` assignment
- ▶ `prog1 ; prog2` sequential composition
- ▶ `if (G) prog1 else prog2` choice
- ▶ `prog1 [] prog2` non-deterministic choice
- ▶ `prog1 [p] prog2` probabilistic choice
- ▶ `while (G) prog` iteration

Overview


- 1 Introduction
- 2 Probabilistic guarded command language
- 3 Operational semantics of pGCL**
- 4 Denotational semantics of pGCL
- 5 Denotational vs. operational semantics of pGCL
- 6 Synthesizing loop invariants
- 7 Epilogue

Markov decision processes

Markov decision process

An MDP \mathcal{M} is a tuple (S, S_0, \rightarrow) where

- ▶ S is a countable set of **states** with **initial state-set** $S_0 \subseteq S$, $S_0 \neq \emptyset$
- ▶ $\rightarrow \subseteq S \times \text{Dist}(S)$ is a **transition relation**



Operational semantics of pGCL

Aim: Model the behaviour of a program $P \in \text{pGCL}$ by an MDP $\mathcal{M}[[P]]$.

Approach:


- ▶ Let η be a variable valuation of the program variables
- ▶ Use the special (semantic) construct **exit** for successful termination
- ▶ States are of the form $\langle Q, \eta \rangle$ with $Q \in \text{pGCL}$ or $Q = \text{exit}$
- ▶ Initial states are tuples $\langle P, \eta \rangle$ where η fulfils the initial conditions
- ▶ Transition relation \rightarrow is the smallest relation satisfying the inference rules

MDP of duelling cowboys

```

int cowboyDuel(float a, b) {
  int t := A [] t := B;
  bool c := true;
  while (c) {
    if (t = A) {
      (c := false [a] t := B);
    } else {
      (c := false [b] t := A);
    }
  }
  return t;
}

```



This MDP is parameterized but finite. Once we count the number of shots before one of the cowboys dies, the MDP becomes **infinite**. Our approach however allows to determine e.g., the expected number of shots before success.

Overview

- 1 Introduction
- 2 Probabilistic guarded command language
- 3 Operational semantics of pGCL
- 4 Denotational semantics of pGCL**
- 5 Denotational vs. operational semantics of pGCL
- 6 Synthesizing loop invariants
- 7 Epilogue

Weakest preconditions

Weakest precondition

[Dijkstra 1975]

A **predicate transformer** is a total function between two predicates on the state of a program.

The predicate transformer $wp(P, F)$ for program P and postcondition F yields the **"weakest" precondition** E on the initial state of P ensuring that the execution of P terminates in a final state satisfying F .

Hoare triple $\{E\} P \{F\}$ holds for **total** correctness iff $E \Rightarrow wp(P, F)$.

Predicate transformer semantics of Dijkstra's GCL

Syntax

- ▶ skip
- ▶ abort
- ▶ $x := E$
- ▶ $P_1 ; P_2$
- ▶ if (G) P_1 else P_2
- ▶ $P_1 [] P_2$
- ▶ while (G) P

Semantics $wp(P, F)$

- ▶ F
- ▶ false
- ▶ $F[x := E]$
- ▶ $wp(P_1, wp(P_2, F))$
- ▶ $(G \Rightarrow wp(P_1, F)) \wedge (\neg G \Rightarrow wp(P_2, F))$
- ▶ $wp(P_1, F) \wedge wp(P_2, F)$
- ▶ $\mu X. ((G \Rightarrow wp(P, X)) \wedge (\neg G \Rightarrow F))$

μ is the least fixed point operator wrt. the ordering \Rightarrow on predicates.

Expectations

Weakest pre-expectation

[McIver & Morgan 2004]

An **expectation** maps program states onto non-negative reals. It's the quantitative analogue of a predicate.

An **expectation transformer** is a total function between two **expectations** on the state of a program.

The transformer $wp(P, f)$ for program P and post-expectation f yields the **least expectation** e on P 's initial state ensuring that P 's execution terminates with an expectation f .

Annotation $\{e\} P \{f\}$ holds for **total** correctness iff $e \leq wp(P, f)$, where \leq is to be interpreted in a point-wise manner.

Expectation transformer semantics of pGCL

Syntax

- ▶ skip
- ▶ abort
- ▶ $x := E$
- ▶ $P_1 ; P_2$
- ▶ if (G) P_1 else P_2
- ▶ $P_1 [] P_2$
- ▶ $P_1 [p] P_2$
- ▶ while (G) P

Semantics $wp(P, f)$

- ▶ f
- ▶ 0
- ▶ $f[x := E]$
- ▶ $wp(P_1, wp(P_2, f))$
- ▶ $[G] \cdot wp(P_1, f) + [\neg G] \cdot wp(P_2, f)$
- ▶ $\min(wp(P_1, f), wp(P_2, f))$
- ▶ $p \cdot wp(P_1, f) + (1-p) \cdot wp(P_2, f)$
- ▶ $\mu X. ([G] \cdot wp(P, X) + [\neg G] \cdot f)$

μ is the least fixed point operator wrt. the ordering \leq on expectations.

A simple slot machine

```

void flip {
  d1 := ♥ [1/2] ♦;
  d2 := ♥ [1/2] ♦;
  d3 := ♥ [1/2] ♦;
}

```

Example weakest pre-expectations

Let $all(x) \equiv (x = d_1 = d_2 = d_3)$.

- ▶ If $f = [all(\heartsuit)]$, then $wlp(flip, f) = \frac{1}{8}$.
- ▶ If $g = 10 \cdot [all(\heartsuit)] + 5 \cdot [all(\diamondsuit)]$, then:

$$wlp(flip, g) = 6 \cdot \frac{1}{8} \cdot 0 + 1 \cdot \frac{1}{8} \cdot 10 + 1 \cdot \frac{1}{8} \cdot 5 = \frac{15}{8}$$

So the least fraction of the jackpot the gamer can expect to win is $\frac{15}{8}$.

Overview

- 1 Introduction
- 2 Probabilistic guarded command language
- 3 Operational semantics of pGCL
- 4 Denotational semantics of pGCL
- 5 Denotational vs. operational semantics of pGCL**
- 6 Synthesizing loop invariants
- 7 Epilogue

MDPs with rewards

To compare the operational and wp- and wlp-semantics, we use [rewards](#).

MDP with rewards

An MDP with [rewards](#) is a pair (\mathcal{M}, r) with \mathcal{M} an MDP with state space S and $r : S \rightarrow \mathbb{R}$ a function assigning a real [reward](#) to each state.

The reward $r(s)$ stands for the reward earned on entering state s .

Cumulative reward for reachability

Let $\pi = s_0 \xrightarrow{\mu_0} s_1 \xrightarrow{\mu_1} \dots$ be an infinite path in (\mathcal{M}, r) and $T \subseteq S$ a set of [target](#) states such that $\pi \models \diamond T$. The [cumulative reward](#) along π before reaching T is defined by:

$$r_T(\pi) = r(s_0) + \dots + r(s_k) \text{ where } s_i \notin T \text{ for all } i < k \text{ and } s_k \in T.$$

If $\pi \not\models \diamond T$, then $r_T(\pi) = 0$.

Reward-bounded reachability

Expected reward for reachability

The **minimal expected reward** until reaching $T \subseteq S$ from $s \in S$ is:

$$ERew(s \models \diamond T) = \min_{\mathfrak{P}} \int_0^\infty c \cdot Pr^{\mathfrak{P}} \{ \pi \in Paths^{\mathfrak{P}}(s, \diamond T) \mid r_T(\pi) = c \} dc$$

A **demonic** positional policy corresponds to a **weakest** pre-expectation.

Relating operational and wp-semantics of pGCL

Weakest pre-expectations vs. expected reachability rewards

For pGCL-program P , variable valuation η , and post-expectation f :

$$wp(P, f)(\eta) = ERew^{\mathcal{M}[[P]]}(\langle P, \eta \rangle \models \diamond P \checkmark)$$

where rewards in MDP $\mathcal{M}[[P]]$ are: $r(\langle \text{exit}, \eta' \rangle) = f(\eta')$ and 0 otherwise.

Thus, $wp(P, f)$ evaluated at η is the minimal expected value of f over any of the resulting distributions of P . The weakest liberal pre-expectation $wp(P, f)$ is similar under the condition that the program terminates.

Overview

- 1 Introduction
- 2 Probabilistic guarded command language
- 3 Operational semantics of pGCL
- 4 Denotational semantics of pGCL
- 5 Denotational vs. operational semantics of pGCL
- 6 Synthesizing loop invariants**
- 7 Epilogue

Qualitative loop invariants

Recall that for while-loops we have:

$$wp(\text{while}(G)\{P\}, F) = \mu X. (G \Rightarrow wp(P, X) \wedge \neg G \Rightarrow F)$$

To determine this wp , one exploits an “invariant” I such that $\neg G \wedge I \Rightarrow F$.

Loop invariant

Predicate I is a **loop invariant** if it is preserved by loop iterations:

$$G \wedge I \Rightarrow wp(P, I) \quad (\text{consecution condition})$$

Then: $\{I\} \text{while}(G)\{P\} \{F\}$ is a correct program annotation.

Linear invariant generation [Colón et al., 2003]

Linear programs

A program is **linear** program whenever all guards are linear constraints, and updates are linear expressions (in the real program variables).

Approach by Colón et al.

1. Speculatively annotate a program with **linear** boolean expressions:

$$\alpha_1 \cdot x_1 + \dots + \alpha_n \cdot x_n + \alpha_{n+1} \leq 0$$

where α_i is a parameter and x_i a program variable.

2. Express verification conditions as **inequality constraints** over α_i, x_i .
3. Transform these inequality constraints into **polynomial constraints** (e.g., using Farkas lemma).
4. Use off-the-shelf constraint-solvers to solve them (e.g., REDLOG).
5. Exploit resulting assertions to infer program correctness.

Quantitative loop invariants

Recall that for while-loops we have:

$$wp(\text{while}(G)\{P\}, f) = \mu X. ([G] \cdot wp(P, X) + [\neg G] \cdot f)$$

To determine this wp , we use an “invariant” I such that $[\neg G] \cdot I \leq f$.

Quantitative loop invariant

Expectation I is a **quantitative loop invariant** if —by consecution—

- ▶ it is preserved by loop iterations: $[G] \cdot I \leq wp(P, I)$.

To guarantee soundness, I has to fulfill either:

1. I is bounded from below and by above by some constants, or
2. on each iteration there is a probability $\epsilon > 0$ to exit the loop

Then: $\{I\} \text{while}(G)\{P\} \{f\}$ is a correct program annotation.

Our approach

Main steps

1. Speculatively annotate a program with **linear** expressions:

$$[\alpha_1 \cdot x_1 + \dots + \alpha_n \cdot x_n + \alpha_{n+1} \ll 0] \cdot (\beta_1 \cdot x_1 + \dots + \beta_n \cdot x_n + \beta_{n+1})$$

with real parameters α_i, β_i , program variable x_i , and $\ll \in \{<, \leq\}$.

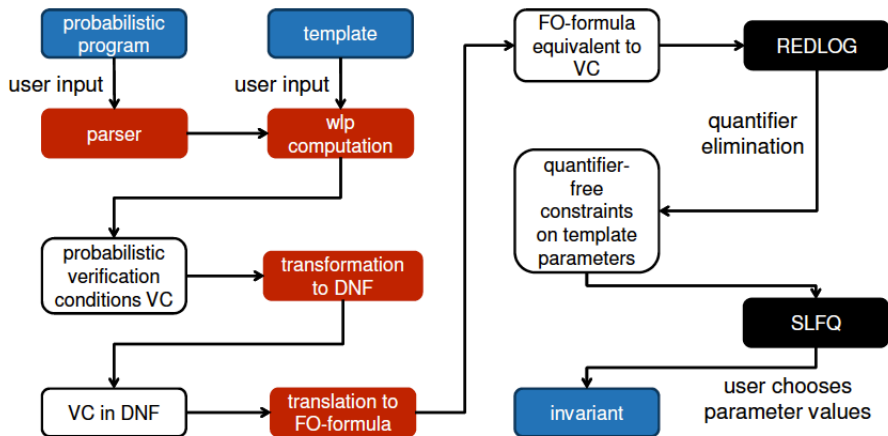
2. Transform these numerical constraints into Boolean predicates.
3. Transform these predicates into non-linear FO formulas
4. Use constraint-solvers for quantifier elimination (e.g., REDLOG).
5. Simplify the resulting formulas (e.g., using SLFQ and SMT solving).
6. Exploit resulting assertions to infer program correctness.

Soundness and completeness

Theorem

For any linear pGCL program annotated with propositionally linear expressions, our method will find all parameter solutions that make the annotation valid, and no others.

PRINSYS Tool: Synthesis of Probabilistic Invariants



download from moves.rwth-aachen.de/prinsys

Duelling cowboys: when does A win?

```

int cbDuel(float a, b) {
  int t := A;
  int c := 1;
  while (c = 1) {
    if (t = A) {
      (c := 0 [a] t := B);
    } else {
      (c := 0 [b] t := A);
    }
  }
  return t;
}

```

Aim: find expectation \mathcal{T}

Satisfying $\mathcal{T} \leq [t = A]$ upon termination.

Observation

On entering the loop, $c = 1$ and either $t = A$ or $t = B$.

Template suggestion

$$\begin{aligned}
 \mathcal{T} = & \underbrace{[t = A \wedge c = 0]}_{\substack{A \text{ wins duel} \\ A\text{'s turn}}} \cdot 1 \\
 & + \underbrace{[t = A \wedge c = 1]}_{\substack{A\text{'s turn} \\ B\text{'s turn}}} \cdot \alpha \\
 & + \underbrace{[t = B \wedge c = 1]}_{\substack{A\text{'s turn} \\ B\text{'s turn}}} \cdot \beta
 \end{aligned}$$

Duelling cowboys: when does A win?

Invariant template

$$\mathcal{T} = [t = A \wedge c = 0] \cdot 1 + [t = A \wedge c = 1] \cdot \alpha + [t = B \wedge c = 1] \cdot \beta$$

Initially, $t = A \wedge c = 1$ and thus $\alpha = Pr\{A \text{ wins duel}\}$.

Running PRINSYS yields

$$a \cdot \beta - a + \alpha - \beta \leq 0 \quad \wedge \quad b \cdot \alpha - \alpha + \beta \leq 0$$

Simplification yields

$$\beta \leq (1 - b) \cdot \alpha \quad \text{and} \quad \alpha \leq \frac{a}{a + b - a \cdot b}$$

As we want to maximise the probability to win

$$\beta = (1 - b) \cdot \alpha \quad \text{and} \quad \alpha = \frac{a}{a + b - a \cdot b}$$

It follows that cowboy A wins the duel with probability $\frac{a}{a + b - a \cdot b}$.

Quantitative loop invariant


Annotated program for post-expectation $[t = A]$

```

1 int cowboyDuel(a, b) {
2    $\langle \frac{(1-b)a}{a+b-ab} \rangle$ 
3    $\langle \min\{\frac{a}{a+b-ab}, \frac{(1-b)a}{a+b-ab}\} \rangle$ 
4   (t := A [] t := B);
5    $\langle [t = A] \cdot \frac{a}{a+b-ab} + [t = B] \cdot \frac{(1-b)a}{a+b-ab} \rangle$ 
6   c := 1;
7    $\langle [t = A \wedge c = 0] \cdot 1 + [t = A \wedge c = 1] \cdot \frac{a}{a+b-ab} + [t = B \wedge c = 1] \cdot \frac{(1-b)a}{a+b-ab} \rangle$ 
8   while (c = 1) {
9      $\langle [t = A \wedge c = 1] \cdot \frac{a}{a+b-ab} + [t = B \wedge c = 1] \cdot \frac{(1-b)a}{a+b-ab} \rangle$ 
10     $\langle [t = A \wedge c \neq 1] \cdot a + [t = A \wedge c = 1] \cdot \frac{a}{a+b-ab}$ 
         $+ [t = B \wedge c = 0] \cdot (1-b) + [t = B \wedge c = 1] \cdot \frac{(1-b)a}{a+b-ab} \rangle$ 
11    if (t = A) {
12      (c := 0 [a] t := B);
13    } else {
14      (c := 0 [b] t := A);
15    }
16     $\langle [t = A \wedge c = 0] \cdot 1 + [t = A \wedge c = 1] \cdot \frac{a}{a+b-ab} + [t = B \wedge c = 1] \cdot \frac{(1-b)a}{a+b-ab} \rangle$ 
17  }
18   $\langle [c \neq 1] \cdot \left( [t = A \wedge c = 0] \cdot 1 + [t = A \wedge c = 1] \cdot \frac{a}{a+b-ab} \right.$ 
         $\left. + [t = B \wedge c = 1] \cdot \frac{(1-b)a}{a+b-ab} \right) \rangle$ 
19   $\langle [t = A] \rangle$ 

```


When one starts nondeterministically



Cowboy A wins the duel with probability at least $\frac{(1-b) \cdot a}{a + b - a \cdot b}$.

Program equivalence

```

int XminY1(float p, q){
  int x, f := 0, 0;
  while (f = 0) {
    (x += 1 [p] f := 1);
  }
  f := 0;
  while (f = 0) {
    (x -= 1 [q] f := 1);
  }
  return x;
}

```

```

int XminY2(float p, q){
  int x, f := 0, 0;
  (f := 0 [0.5] f := 1);
  if (f = 0) {
    while (f = 0) {
      (x += 1 [p] f := 1);
    }
  } else {
    f := 0;
    while (f = 0) {
      x -= 1;
      (skip [q] f := 1);
    }
  }
  return x;
}

```

Using template $\mathcal{T} = x + [f = 0] \cdot \alpha$ we find the invariants :

$$\alpha_{11} = \frac{p}{1-p}, \alpha_{12} = -\frac{q}{1-q}, \alpha_{21} = \alpha_{11} \text{ and } \alpha_{22} = -\frac{1}{1-q}.$$

Overview

- 1 Introduction
- 2 Probabilistic guarded command language
- 3 Operational semantics of pGCL
- 4 Denotational semantics of pGCL
- 5 Denotational vs. operational semantics of pGCL
- 6 Synthesizing loop invariants
- 7 Epilogue**

Recursive probabilistic programs

Probabilistic pushdown automata

[Esparza *et al.*, 2004]

Are a natural model for recursive probabilistic programs. Checking whether they simulate (or are simulated by) a finite Markov chain is EXPTIME-complete.

Overview of complexities

[Fu & Katoen, 2011]

	(coupled) bisimilarity	(coupled) similarity
PDA vs. finite TS	PSPACE-complete	EXPTIME-complete
pPDA vs. finite pTS	EXPTIME-complete	EXPTIME-complete

Epilogue

Take-home message

- ▶ Connection between wp-semantics and operational semantics.
- ▶ Synthesizing probabilistic loop invariants using constraint solving.
- ⇒ Large potential for automated probabilistic program analysis.
- ▶ Initial prototypical tool-support PRINSYS.

Future work

- ▶ Further development of PRINSYS.
- ▶ Non-linear probabilistic programs.
- ▶ Average time-complexity analysis.

Fin.