Relational reasoning via probabilistic coupling

Gilles Barthe, Thomas Espitau, Benjamin Grégoire, Justin Hsu,
Léo Stefanesco, Pierre-Yves Strub

IMDEA Software, ENS Cachan, ENS Lyon, Inria, University of Pennsylvania

November 28, 2015

Relational properties

Properties about two runs of the same program

- Assume inputs are related by Ψ
- Want to prove the outputs are related by Φ

Examples

Monotonicity

- $\Psi : in_1 \leq in_2$
- $\Phi : out_1 \leq out_2$
- "Bigger inputs give bigger outputs"

Examples

Monotonicity

- $\Psi : in_1 \leq in_2$
- $\Phi : out_1 \leq out_2$
- "Bigger inputs give bigger outputs"

Non-interference

- $\Psi : low_1 = low_2$
- $\Phi : out_1 = out_2$
- "If low-security inputs are the same, then outputs are the same"

Probabilistic relational properties

Richer properties

- Differential privacy
- Cryptographic indistinguishability

Probabilistic relational properties

Richer properties

- Differential privacy
- Cryptographic indistinguishability

Verification tool: pRHL [BGZ-B]

- Imperative while language + command for random sampling
- Deterministic input, randomized output
- Hoare-style logic

Inspiration from probability theory

Probabilistic couplings

- Used by mathematicians for proving relational properties
- Applications: Markov chains, probabilistic processes

Idea

- Place two processes in the same probability space
- Coordinate the sampling

Our results

Main observation

$\boxed{\text{The logic pRHL internalizes coupling}}$

Our results

Main observation

$\boxed{\text{The logic pRHL internalizes coupling}}$

Consequences

- Constructing pRHL proof $\rightarrow$ constructing a coupling
- Can verify classic examples of couplings in mathematics with proof assistant EasyCrypt (built on pRHL)

The plan

Today

- Introducing probabilistic couplings
- Introducing the relational logic pRHL
- Example: convergence of random walks

Probabilistic couplings

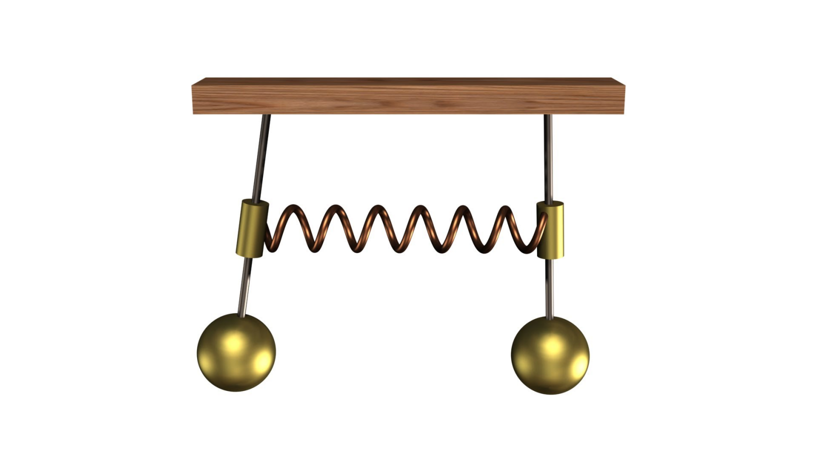

Introducing to probabilistic couplings

Basic ingredients

- Given: two distributions X_1, X_2 over set A
- Produce: joint distribution Y over $A \times A$
  - Distribution over the first component is X_1
  - Distribution over the second component is X_2

Introducing to probabilistic couplings

Basic ingredients

- Given: two distributions X_1, X_2 over set A
- Produce: joint distribution Y over $A \times A$
  - Distribution over the first component is X_1
  - Distribution over the second component is X_2

Definition

Given two distributions X_1, X_2 over a set A, a coupling Y is a distribution over $A \times A$ such that $\pi_1(Y) = X_1$ and $\pi_2(Y) = X_2$.

Example: mirrored random walks

Simple random walk on integers

- Start at position $p = 0$
- Each step, flip coin $x \xleftarrow{\$} \mathit{flip}$
- Heads: $p \leftarrow p + 1$
- Tails: $p \leftarrow p - 1$

Example: mirrored random walks

Simple random walk on integers

- Start at position $p = 0$
- Each step, flip coin $x \xleftarrow{\$} \mathit{flip}$
- Heads: $p \leftarrow p + 1$
- Tails: $p \leftarrow p - 1$

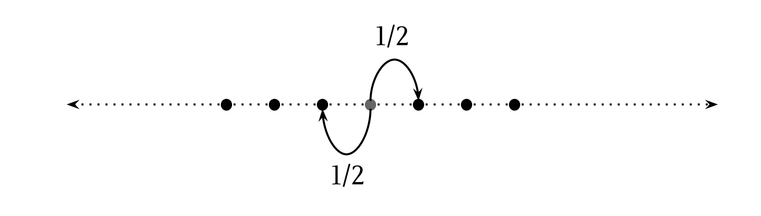

Figure: Simple random walk

Coupling the walks to meet

Case $p_1 = p_2$: Walks have met

- Arrange samplings $x_1 = x_2$
- Continue to have $p_1 = p_2$

Coupling the walks to meet

Case $p_1 = p_2$: Walks have met

- Arrange samplings $x_1 = x_2$
- Continue to have $p_1 = p_2$

Case $p_1 \neq p_2$: Walks have not met

- Arrange samplings $x_1 = \neg x_2$
- Walks make mirror moves

Coupling the walks to meet

Case $p_1 = p_2$: Walks have met

- Arrange samplings $x_1 = x_2$
- Continue to have $p_1 = p_2$

Case $p_1 \neq p_2$: Walks have not met

- Arrange samplings $x_1 = \neg x_2$
- Walks make mirror moves

Under coupling, if walks meet, they move together

Why is this interesting?

Goal: memorylessness

- Start two random walks at w and $w + 2k$
- To show: position distributions converge as we take more steps

Why is this interesting?

Goal: memorylessness

- Start two random walks at w and $w + 2k$
- To show: position distributions converge as we take more steps

Coupling bounds distance between distributions

- Once walks meet, they stay equal
- Distance is at most probability walks don't meet

Why is this interesting?

Goal: memorylessness

- Start two random walks at w and $w + 2k$
- To show: position distributions converge as we take more steps

Coupling bounds distance between distributions

- Once walks meet, they stay equal
- Distance is at most probability walks don't meet

Theorem

If Y is a coupling of two distributions (X_1, X_2), then

$$\|X_1 - X_2\|_{TV} \triangleq \sum_{a \in A} |X_1(a) - X_2(a)| \leq \Pr_{(y_1, y_2) \sim Y}[y_1 \neq y_2].$$

The logic pRHL

The program logic pRHL

Probabilistic Relational Hoare Logic

- Hoare-style logic for probabilistic relational properties
- Proposed by Barthe, Grégoire, Zanella-Béguelin
- Implemented in the EasyCrypt proof assistant for crypto proofs

Language and judgments

The pWhile imperative language

$$c ::= x \leftarrow e \mid x \xleftarrow{\$} d \mid \textsf{if } e \textsf{ then } c \textsf{ else } c \mid \textsf{while } e \textsf{ do } c \mid \textsf{skip} \mid c;\ c$$

Language and judgments

The pWhile imperative language

$$c ::= x \leftarrow e \mid \boxed{x \xleftarrow{\$} d} \mid \text{if } e \text{ then } c \text{ else } c \mid \text{while } e \text{ do } c \mid \text{skip} \mid c;\ c$$

Language and judgments

The pWhile imperative language

$$c ::= x \leftarrow e \mid \boxed{x \xleftarrow{\$} d} \mid \text{if } e \text{ then } c \text{ else } c \mid \text{while } e \text{ do } c \mid \text{skip} \mid c;\ c$$

Basic pRHL judgments

$$\vDash c_1 \sim c_2 : \Psi \Rightarrow \Phi$$

- Ψ and Φ are formulas over labeled program variables x_1, x_2
- Ψ is precondition, Φ is postcondition

Interpreting the judgment

$$\vDash c_1 \sim c_2 : \Psi \Rightarrow \Phi$$

Interpreting the judgment

$$\models c_1 \sim c_2 : \Psi \Rightarrow \Phi$$

Interpreting pre- and post-conditions

- Ψ interpreted as a relation on two memories
- Φ interpreted as a relation $\Phi^\dagger$ on distributions over memories

Interpreting the judgment

$$\models c_1 \sim c_2 : \Psi \Rightarrow \Phi$$

Interpreting pre- and post-conditions

- Ψ interpreted as a relation on two memories
- Φ interpreted as a relation $\Phi^\dagger$ on distributions over memories

Definition (Couplings in disguise!)

If Φ is a relation on A, the lifted relation $\Phi^\dagger$ is a relation on $\mathbf{Distr}(A)$ where $\mu_1 \, \Phi^\dagger \mu_2$ if there exists $\mu \in \mathbf{Distr}(A \times A)$ with

- $\mathrm{supp}(\mu) \subseteq \Phi$; and
- $\pi_1(\mu) = \mu_1$ and $\pi_2(\mu) = \mu_2$.

Proof rules

The key rule: Sampling

$$\textsc{Sample}\ \frac{f \in T \xrightarrow{1-1} T \qquad \forall v \in T.\ d_1(v) = d_2(f\ v)}{\vDash x_1 \xleftarrow{\$} d_1 \sim x_2 \xleftarrow{\$} d_2 : \forall v,\ \Phi[v/x_1, f(v)/x_2] \Rightarrow \Phi}$$

Notes

Proof rules

The key rule: Sampling

$$\textsc{Sample}\ \frac{\boxed{f \in T \xrightarrow{1-1} T} \qquad \forall v \in T.\ d_1(v) = d_2(f\ v)}{\models x_1 \xleftarrow{\$} d_1 \sim x_2 \xleftarrow{\$} d_2 : \forall v,\ \Phi[v/x_1, f(v)/x_2] \Rightarrow \Phi}$$

Notes

Proof rules

The key rule: Sampling

$$\textsc{Sample}\ \frac{\boxed{f \in T \xrightarrow{1-1} T} \qquad \forall v \in T.\ d_1(v) = d_2(f\ v)}{\models x_1 \xleftarrow{\$} d_1 \sim x_2 \xleftarrow{\$} d_2 : \forall v,\ \Phi[v/x_1, f(v)/x_2] \ \Rightarrow \Phi}$$

Notes

- Bijection f: specifies how to coordinate the samples

Proof rules

The key rule: Sampling

$$\textsc{Sample}\ \frac{f \in T \xrightarrow{1-1} T \qquad \boxed{\forall v \in T.\ d_1(v) = d_2(f\ v)}}{\vDash x_1 \overset{\$}{\leftarrow} d_1 \sim x_2 \overset{\$}{\leftarrow} d_2 : \forall v,\ \Phi[v/x_1, f(v)/x_2] \ \Rightarrow \Phi}$$

Notes

- Bijection f: specifies how to coordinate the samples

Proof rules

The key rule: Sampling

$$\textsc{Sample}\ \frac{f \in T \xrightarrow{1-1} T \qquad \boxed{\forall v \in T.\ d_1(v) = d_2(f\ v)}}{\models x_1 \xleftarrow{\$} d_1 \sim x_2 \xleftarrow{\$} d_2 : \forall v,\ \Phi[v/x_1, f(v)/x_2] \ \Rightarrow \Phi}$$

Notes

- Bijection f: specifies how to coordinate the samples
- Side condition: marginals are preserved under f

Proof rules

The key rule: Sampling

$$\textsc{Sample}\ \frac{f \in T \xrightarrow{1-1} T \qquad \forall v \in T.\ d_1(v) = d_2(f\ v)}{\vDash x_1 \xleftarrow{\$} d_1 \sim x_2 \xleftarrow{\$} d_2 : \forall v, \boxed{\Phi[v/x_1, f(v)/x_2]} \Rightarrow \Phi}$$

Notes

- Bijection f: specifies how to coordinate the samples
- Side condition: marginals are preserved under f

Proof rules

The key rule: Sampling

$$\textsc{Sample}\ \frac{f \in T \xrightarrow{1-1} T \qquad \forall v \in T.\ d_1(v) = d_2(f\ v)}{\models x_1 \xleftarrow{\$} d_1 \sim x_2 \xleftarrow{\$} d_2 : \forall v, \boxed{\Phi[v/x_1, f(v)/x_2]} \Rightarrow \Phi}$$

Notes

- Bijection f: specifies how to coordinate the samples
- Side condition: marginals are preserved under f
- Assume: samples coupled when proving postcondition Φ

Examples

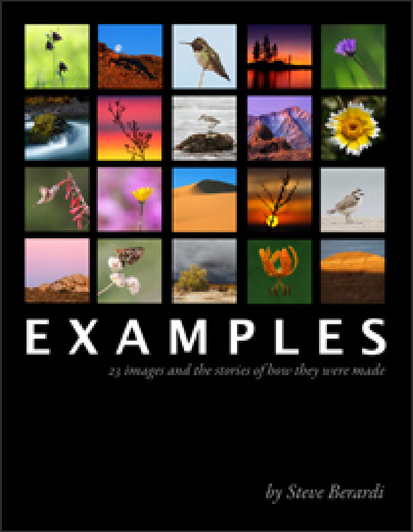
EXAMPLES
23 images and the stories of how they were made
by Steve Berardi

Example: mirroring random walks in pRHL

The code

```
pos ← start;          // Start position
i ← 0;
H ← [];               // Ghost code
while i < N do
  b ←$ flip;
  H ← b :: H;         // Ghost code
  if b then
    pos ← pos + 1;
  else
    pos ← pos - 1;
  fi
i ← i + 1;
end
return pos           // Final position
```

Example: mirroring random walks in pRHL

The code

```
pos ← start;          // Start position
i ← 0;
H ← [];               // Ghost code
while i < N do
  b $← flip;
  H ← b :: H;         // Ghost code
  if b then
    pos ← pos + 1;
  else
    pos ← pos - 1;
  fi
i ← i + 1;
end
return pos           // Final position
```

Goal: couple two walks via mirroring

Record the history

$\texttt{H}$ stores history of flips

- $\Sigma(\texttt{H})$ is the net distance that the first process moves to the right
- $\mathit{Meet}(\texttt{H})$ if there is prefix $\texttt{H'}$ of $\texttt{H}$ with $\Sigma(\texttt{H'}) = k$

Specify the coupling

Sampling rule

$$\textsc{Sample}\ \frac{f \in T \xrightarrow{1-1} T \qquad \forall v \in T.\ d_1(v) = d_2(f\ v)}{\models x_1 \xleftarrow{\$} d_1 \sim x_2 \xleftarrow{\$} d_2 : \forall v, \Phi[v/x_1, f(v)/x_2] \Rightarrow \Phi}$$

Specify the coupling

Sampling rule

$$\textsc{Sample}\ \frac{f \in T \xrightarrow{1-1} T \qquad \forall v \in T.\ d_1(v) = d_2(f\ v)}{\models x_1 \overset{\$}{\leftarrow} d_1 \sim x_2 \overset{\$}{\leftarrow} d_2 : \forall v, \Phi[v/x_1, f(v)/x_2] \Rightarrow \Phi}$$

Case on *Meet*($\mathtt{H}_1$)

- True: take bijection f to be *id*
- False: take bijection f to be negation $\neg$

Final judgment

$$\models c \sim c : \ \texttt{start}_1 + 2k = \texttt{start}_2 \ \Rightarrow \ \left(\mathit{Meet}(\texttt{H}_1) \rightarrow \texttt{pos}_1 = \texttt{pos}_2\right)$$

How to read

Final judgment

$$\vDash c \sim c : \boxed{\mathtt{start}_1 + 2k = \mathtt{start}_2} \Rightarrow (\mathit{Meet}(\mathtt{H}_1) \to \mathtt{pos}_1 = \mathtt{pos}_2)$$

How to read

Final judgment

$$\models c \sim c : \boxed{\mathtt{start}_1 + 2k = \mathtt{start}_2} \Rightarrow \left(\mathit{Meet}(\mathtt{H}_1) \rightarrow \mathtt{pos}_1 = \mathtt{pos}_2 \right)$$

How to read

- The two walks start $2k$ apart

Final judgment

$$\vDash c \sim c : \texttt{start}_1 + 2k = \texttt{start}_2 \Rightarrow (\mathit{Meet}(\texttt{H}_1) \to \texttt{pos}_1 = \texttt{pos}_2)$$

How to read

- The two walks start $2k$ apart

Final judgment

$$\vDash c \sim c : \ \mathtt{start}_1 + 2k = \mathtt{start}_2 \ \Rightarrow \boxed{(\mathit{Meet}(\mathtt{H}_1) \rightarrow \mathtt{pos}_1 = \mathtt{pos}_2)}$$

How to read

- The two walks start $2k$ apart
- If walks have met before, their positions are equal

Further examples

Lazy random walk on torus

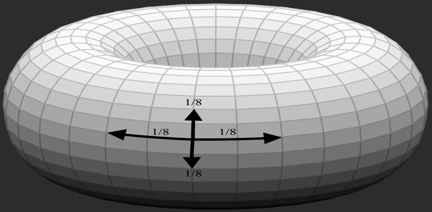

Figure: Lazy random walk on a two dimensional torus

Further examples

Lazy random walk on torus

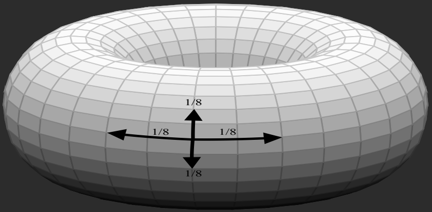

Figure: Lazy random walk on a two dimensional torus

Stochastic domination

- Notion of ordering for probabilistic processes
- Proved via couplings

Wrapping up

Open problems

Handling more advanced couplings

- Shift couplings, path couplings, etc.
- Hard example: constructive Lovász Local Lemma by Moser

Quantitative bounds

- How long does it take for the mirrored walks to meet?
- Non-relational reasoning

Borrow more ideas from the coupling literature

- Couplings from mathematics may suggest natural rules to add

Relational reasoning via probabilistic coupling

Gilles Barthe, Thomas Espitau, Benjamin Grégoire, Justin Hsu,
Léo Stefanesco, Pierre-Yves Strub

IMDEA Software, ENS Cachan, ENS Lyon, Inria, University of Pennsylvania

November 28, 2015