

# Probabilistic Programming

## Lecture #6: Syntax and Operational Semantics of pGCL

Joost-Pieter Katoen

RWTH Lecture Series on Probabilistic Programming 2022-23


## Overview

- 1 Probabilistic Guarded Command Language
- 2 Operational semantics of pGCL
- 3 Expected Rewards
- 4 Recursion

## Overview

- 1 Probabilistic Guarded Command Language
- 2 Operational semantics of pGCL
- 3 Expected Rewards
- 4 Recursion

## Dijkstra's guarded command language: Syntax



- |  |                          |
|--|--------------------------|
| ▶ <code>skip</code>                    | empty statement          |
| ▶ <code>diverge</code>                 | divergence               |
| ▶ <code>x := E</code>                  | assignment               |
| ▶ <code>prog1 ; prog2</code>           | sequential composition   |
| ▶ <code>if (G) prog1 else prog2</code> | choice                   |
| ▶ <code>prog1 [] prog2</code>          | non-deterministic choice |
| ▶ <code>while (G) prog</code>          | iteration                |

## Elementary pGCL ingredients

- ▶ Program variables  $x \in \text{Vars}$  whose values are fractional numbers
- ▶ Arithmetic expressions  $E$  over the program variables
- ▶ Boolean expressions  $G$  (aka: **guards**) over the program variables
- ▶ **Distribution expressions**  $\mu : \Sigma \rightarrow \text{Dist}(\mathbb{Q})$
- ▶ **Probability expressions**  $p : \Sigma \rightarrow [0, 1] \cap \mathbb{Q}$

where  $\Sigma$  is the set of program states; made precise later

## Let's start simple

---

```
x := 0 [0.5] x := 1;
y := -1 [0.5] y := 0
```

---

This program admits four runs and yields the outcome:

$$Pr[x=0, y=0] = Pr[x=0, y=-1] = Pr[x=1, y=0] = Pr[x=1, y=-1] = 1/4$$

## Probabilistic GCL: Syntax

Dexter  
Kozen



Annabelle  
McIver



Carroll  
Morgan



- ▶ **skip** empty statement
- ▶ **diverge** divergence
- ▶  $x := E$  assignment
- ▶  $x := \text{r} \leftarrow \mu$  **random assignment** ( $x : \approx \mu$ )
- ▶  $\text{prog1} ; \text{prog2}$  sequential composition
- ▶ **if** ( $G$ )  $\text{prog1}$  **else**  $\text{prog2}$  choice
- ▶  $\text{prog1} [p] \text{prog2}$  **probabilistic choice**
- ▶ **while** ( $G$ )  $\text{prog}$  iteration

Conditioning in the form of observe-statements omitted for now.

## A loopy program

For  $0 < p < 1$  an arbitrary probability:

---

```
bool c := true;
int i := 0;
while (c) {
  i++;
  (c := false [p] c := true)
}
```

---

The loopy program models a geometric distribution with parameter  $p$ .

$$Pr[i = N] = (1-p)^{N-1} \cdot p \quad \text{for } N > 0$$

## On termination

---

```

bool c := true;
int i := 0;
while (c) {
  i++;
  (c := false [p] c := true)
}

```

---

This program does **not** always terminate. It **almost** surely terminates.

## Duelling cowboys

---

```

int cowboyDuel(float a, b) {
  int t := A [0.5] t := B;
  bool c := true;
  while (c) {
    if (t = A) {
      (c := false [a] t := B);
    } else {
      (c := false [b] t := A);
    }
  }
  return t;
}

```

---

## The good, the bad, and the ugly



## Random assignments

The **random assignment**  $x : \approx \mu$  works as follows:

1. evaluate distribution expression  $\mu$  in the current program state  $s$
2. sample from the resulting probability distribution  $\mu(s)$   
this yields the value  $v$  with probability  $\mu(s)(v)$
3. assign the value  $v$  to the variable  $x$ .

For denoting distribution expressions, we use the **bra-ket notation**.

$$\frac{1}{2} \cdot [a] + \frac{1}{3} \cdot [b] + \frac{1}{6} \cdot [c]$$

denotes the distribution  $\mu$  with  $\mu(a) = 1/2$ ,  $\mu(b) = 1/3$ , and  $\mu(c) = 1/6$ . The support set of  $\mu$  equals  $\{a, b, c\}$

Examples on the black board.

# Overview

1 Probabilistic Guarded Command Language

2 Operational semantics of pGCL

3 Expected Rewards

4 Recursion

## The inventors of semantics



Tony Hoare



Robert W. Floyd



Gordon Plotkin



Christopher Strachey



Dana Scott

## Why formal semantics matters

► Unambiguous meaning to all programs

- Basis for proving correctness
  - of programs
  - of program transformations
  - of program equivalence
  - of static analysis
  - of compilers
  - .....

## Approaches to semantics

- **Operational semantics:** (developed by Plotkin)
  - The meaning of a program = how it executes on an abstract machine.
  - Useful for modelling the execution behaviour of a program.
- **Axiomatic semantics:** (developed by Floyd and Hoare)
  - Provides correctness assertions for each program construct.
  - Useful for verifying that the program's computed results are correct with respect to the specification.
- **Denotational semantics:** (developed by Strachey and Scott)
  - Provides a mapping of language constructs onto mathematical objects.
  - Useful for obtaining an abstract insight into the working of a program.

Today: operational semantics of pGCL in terms of Markov chains.

Later: denotational semantics of pGCL in terms of weakest preconditions.

## Structural operational semantics: ingredients

- ▶ **Variable valuation**  $s : \text{Vars} \rightarrow \mathbb{Q}$  maps each program variable onto a value (here: rational numbers)

- ▶ Variable update: for variable  $x$  and value  $v \in \mathbb{Q}$ , let

$$s[x := v](y) = s(y) \quad \text{if } x \neq y \quad \text{and} \quad s[x := v](y) = v \text{ otherwise.}$$

- ▶ Let  $\llbracket E \rrbracket$  denote the **valuation of expression**  $E$

- ▶ **Program configuration** (aka: state)  $\langle P, s \rangle$  denotes that

- ▶ program  $P$  is next to be executed (aka: program counter), and
- ▶ the current variable valuation equals  $s$ .

- ▶ **Transition rules** for the execution of commands:  $\langle P, s \rangle \longrightarrow \langle P', s' \rangle$   
denoted as  $\frac{\text{premise}}{\text{conclusion}}$  where the premise is omitted if it equals true.

## Operational semantics

**Aim:** Model the behaviour of a program  $P$  by the MC  $\llbracket P \rrbracket$ .

**Approach:**

- ▶ Take states of the form
  - ▶  $\langle Q, s \rangle$  with program  $Q$  or  $\downarrow$ , and variable valuation  $s : \text{Var} \rightarrow \mathbb{Q}$
  - ▶  $\langle \text{sink} \rangle$  models program **termination** (successful)
- ▶ Take initial state  $\langle P, s \rangle$  where  $s$  fulfils the initial conditions
- ▶ Take transition relation  $\rightarrow$  as **smallest** relation satisfying the transition rules

## Operational semantics

**Aim:** Model the behaviour of a program  $P$  by the MC  $\llbracket P \rrbracket$ .

This MC is defined using **structured operational semantics**

## Transition rules (1)

$$\langle \text{skip}, s \rangle \rightarrow \langle \downarrow, s \rangle$$

$$\langle \downarrow, s \rangle \rightarrow \langle \text{sink} \rangle \quad \langle \text{sink} \rangle \rightarrow \langle \text{sink} \rangle$$

$$\langle x := E, s \rangle \rightarrow \langle \downarrow, s[x := s(\llbracket E \rrbracket)] \rangle$$

$$\frac{\mu(s)(v) = a > 0}{\langle x : \approx \mu, s \rangle \xrightarrow{a} \langle \downarrow, s[x := v] \rangle}$$

$$\langle P[p] Q, s \rangle \rightarrow \mu \text{ with } \mu(\langle P, s \rangle) = p \text{ and } \mu(\langle Q, s \rangle) = 1-p$$

## Random assignments

The **random assignment**  $x : \approx \mu$  works as follows:

1. evaluate distribution expression  $\mu$  in the current program state  $s$
2. sample from the resulting probability distribution  $\mu(s)$   
this yields the value  $v$  with probability  $\mu(s)(v)$
3. assign the value  $v$  to the variable  $x$ .

For denoting distribution expressions, we use the **bra-ket notation**.

$$\frac{1}{2} \cdot [a] + \frac{1}{3} \cdot [b] + \frac{1}{6} \cdot [c]$$

denotes the distribution  $\mu$  with  $\mu(a) = 1/2$ ,  $\mu(b) = 1/3$ , and  $\mu(c) = 1/6$ . The support set of  $\mu$  equals  $\{a, b, c\}$

Examples on the black board.

## Example

---

```
x := 0 [1/3] x := 1;
y := unif[1..5]
```

---



---

```
x := unif[1..5];
if (x >= 2) { y := unif[1..x] } else { y := x }
```

---



---

```
x += 1 [1/(abs(x)+1)] x -= 1
```

---

## Transition rules (2)

$$\frac{\langle P, s \rangle \rightarrow \mu}{\langle P; Q, s \rangle \rightarrow \nu} \text{ with } \nu(\langle P'; Q', s' \rangle) = \mu(\langle P', s' \rangle) \text{ where } \downarrow; Q = Q$$

$$\frac{s \models G}{\langle \text{if } (G)\{P\} \text{ else } \{Q\}, s \rangle \rightarrow \langle P, s \rangle} \quad \frac{s \not\models G}{\langle \text{if } (G)\{P\} \text{ else } \{Q\}, s \rangle \rightarrow \langle Q, s \rangle}$$


$$\frac{s \models G}{\langle \text{while}(G)\{P\}, s \rangle \rightarrow \langle P; \text{while}(G)\{P\}, s \rangle} \quad \frac{s \not\models G}{\langle \text{while}(G)\{P\}, s \rangle \rightarrow \langle \downarrow, s \rangle}$$

## Duelling cowboys

---

```
int cowboyDuel(float a, b) {
  int t := A [0.5] t := B;
  bool c := true;
  while (c) {
    if (t = A) {
      (c := false [a] t := B);
    } else {
      (c := false [b] t := A);
    }
  }
  return t;
}
```

---



This (parametric) MC is finite. Once we count the number of shots before one of the cowboys dies, the MC becomes **countably infinite**.

## Duelling cowboys

```

int cowboyDuel(float a, b) { // 0 < a < 1, 0 < b < 1
  int t := A [0.5] t := B; // decide who shoots first
  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  $\frac{(1-b) \cdot \frac{1}{2} a}{a+b-a \cdot b}$ .

## The outcome of a pGCL program

Unlike a deterministic program, a pGCL program  $P$  has not a deterministic output for a given input. Instead, it yields a unique probability distribution over its final states.

In fact, this is a **sub**-distribution (probability mass at most one), as with a (possibly positive) probability,  $P$  may diverge.

Let  $P$  be a pGCL program and  $s$  an input state. Then the distribution over final states obtained by running  $P$  starting in  $s$  is given by  $Pr(s \models \Diamond \langle \downarrow, \cdot \rangle)$ .

If  $P$  is a program whose MC is finite-state, then  $Pr(s \models \Diamond \langle \downarrow, \cdot \rangle)$  can be determined by solving a linear equation system.

## Proof

## Overview

- 1 Probabilistic Guarded Command Language
- 2 Operational semantics of pGCL
- 3 Expected Rewards
- 4 Recursion

## Rewards

To reason about resource usage in MCs: use **rewards**.

A **reward** MC is a pair  $(D, r)$  with  $D$  an MC with state space  $\Sigma$  and  $r : \Sigma \rightarrow \mathbb{R}$  a function assigning a real **reward** to each state.

The reward  $r(\sigma)$  stands for the reward earned on **leaving** state  $\sigma$ .

Let  $\pi = \sigma_0 \dots \sigma_n$  be a finite path in  $(D, r)$  and  $G \subseteq \Sigma$  a set of **target** states with  $\pi \in \Diamond G$ . The **cumulative reward** along  $\pi$  until reaching  $G$  is:

$$r_G(\pi) = r(\sigma_0) + \dots + r(\sigma_{k-1}) \text{ where } \sigma_i \notin G \text{ for all } i < k \text{ and } \sigma_k \in G.$$

If  $\pi \notin \Diamond G$ , then  $r_G(\pi) = \infty$ .

## On computing expected rewards

Expected rewards in **finite** Markov chains can be computed in polynomial time by solving a system of linear equations.  
(details on the black board.)

## Expected reward for reachability

Let  $\sigma$  be such that  $Pr(\sigma \models \Diamond G) = 1$ .

Then: the **expected reward** until reaching  $G \subseteq \Sigma$  from  $\sigma \in \Sigma$  is:

$$ER(\sigma, \Diamond G) = \sum_{\hat{\pi}} Pr(\hat{\pi}) \cdot r_G(\hat{\pi})$$

where  $\hat{\pi} = \sigma_0 \dots \sigma_k$  is such that  $\sigma_k \in G$ ,  $\sigma_0 = \sigma$  and  $\sigma_i \notin G$  for all  $i < k$ .

If  $Pr(\sigma \models \Diamond G) < 1$ , then let  $ER(\sigma, \Diamond G) = \infty$ .

## Equation system for expected rewards



## Overview

1 Probabilistic Guarded Command Language

2 Operational semantics of pGCL

3 Expected Rewards

4 Recursion

## Pushdown Markov chains

A **pushdown Markov chain** consists of:


- ▶  $\Sigma = (\Sigma_{push}, \Sigma_{pop}, \Sigma_{int})$ , a finite set of states
- ▶  $\Gamma$  is a finite stack alphabet
- ▶  $(\sigma_I, Z_0) \in \Sigma \times \Gamma$ , the initial configuration
- ▶ probabilistic transition functions:
  - ▶  $P_{push} : \Sigma_{push} \rightarrow \text{Dist}(\Sigma \times \Gamma)$  push transitions
  - ▶  $P_{pop} : \Sigma_{pop} \times \Gamma \rightarrow \text{Dist}(\Sigma)$  pop transitions
  - ▶  $P_{int} : \Sigma_{int} \rightarrow \text{Dist}(\Sigma)$  internal transitions

## Probabilistic GCL with recursion: Syntax


- ▶ **skip** empty statement
- ▶ **x := E** assignment
- ▶ **x :=r= mu** random assignment ( $x : \approx \mu$ )
- ▶ **prog1 ; prog2** sequential composition
- ▶ **if (G) prog1 else prog2** choice
- ▶ **prog1 [p] prog2** probabilistic choice
- ▶ **while (G) prog** iteration
- ▶ **proc P = prog** **process definition**
- ▶ **call P** **process invocation**

Recursion does not increase the expressive power, but is often convenient.


## Pushdown MC = Markov chain + stack



## Configuration graph




## From programs to pushdown MCs



$$\{\text{skip}^1\} [1/2]^2 \{\text{call } P^3; \text{call } P^4; \text{call } P^5\}$$

## Runs of a pushdown MC



$q_0$	$q_1$	$q_2$	$q_3$	$q_4$	$q_5$	$q_6$	$q_7$	$q_8$	$q_9$	$q_{10}$	...
{A}	{A}	{A,B}	$\emptyset$	{A}	{A}	{B}	{B}	{B}	{A,B}	$\emptyset$	...

Assuming states are labeled with sets of atomic propositions

## Take-home messages

- ▶ pGCL is a “base” imperative probabilistic programming language
- ▶ Key ingredients: probabilistic choice and random assignments
- ▶ A pGCL program corresponds to a (countably infinite) Markov chain
- ▶ Computing expected rewards in finite MCs = solving linear equations
- ▶ Recursion can be added to pGCL and yields pushdown MCs

## Next lecture

Tuesday Nov 8, 16:30

No lecture on Nov 3;  
next exercise class Nov 4