

Probabilistic Guarded Commands Mechanized in HOL

Joe Hurd

`joe.hurd@comlab.ox.ac.uk`

Oxford University

Joint work with Annabelle McIver (Macquarie University) and
Carroll Morgan (University of New South Wales)

Contents

- **Introduction**
- Formalizing Probabilistic Guarded Commands
- wlp Verification Condition Generator
- Example: Rabin's Mutual Exclusion Algorithm
- Conclusion

Introduction: pGCL

- pGCL stands for probabilistic Guarded Command Language.
- It's Dijkstra's GCL extended with probabilistic choice

$$c_1 \text{ } p \oplus \text{ } c_2$$

- Like GCL, the semantics is based on weakest preconditions.
- **Important:** retains demonic choice

$$c_1 \sqcap c_2$$

- Developed by Morgan et al. in the Programming Research Group, Oxford, 1994–

The HOL Theorem Prover

- Developed by Mike Gordon's Hardware Verification Group in Cambridge, first release was HOL88.
- Latest release called HOL4, developed jointly by Cambridge, Utah and ANU.
- Implements classical Higher-Order Logic: essentially first-order logic with quantification over functions.
- Sprung from the Edinburgh LCF project, so has a small logical kernel to ensure soundness.

Contents

- Introduction
- **Formalizing Probabilistic Guarded Commands**
- wlp Verification Condition Generator
- Example: Rabin's Mutual Exclusion Algorithm
- Conclusion

pGCL Semantics

- Given a standard GCL program C and a postcondition Q , let P be the weakest precondition that satisfies

$$[P]C[Q]$$

- Precondition P is weaker than P' if $P' \Rightarrow P$.
- Think of C as a function that transforms postconditions into weakest preconditions.
- pGCL generalizes this to probabilistic programs:
 - Conditions $\alpha \rightarrow \mathbb{B}$ become *expectations* $\alpha \rightarrow [0, +\infty]$.
 - Expectation P is weaker than P' if $P' \sqsubseteq P$.
 - Think of programs as *expectation transformers*.

Expectations

- Expectations are reward functions, from states to expected rewards.
- Modelled in HOL as functions $\alpha \rightarrow [0, +\infty]$.
- Define the following operations on expectations:
 - $\text{Min } e_1 \ e_2 \equiv \lambda s. \min (e_1 \ s) (e_2 \ s)$
 - $e_1 \sqsubseteq e_2 \equiv \forall s. e_1 \ s \leq e_2 \ s$
 - $\text{Cond } b \ e_1 \ e_2 \equiv \lambda s. \text{if } b \ s \text{ then } e_1 \ s \text{ else } e_2 \ s$
 - $\text{Lin } p \ e_1 \ e_2 \equiv \lambda s. [p \ s]^{\leq 1} \times e_1 \ s + (1 - [p \ s]^{\leq 1}) \times e_2 \ s$

States

- Fix states to be mappings from variable names to integers:

$$\text{state} \equiv \text{string} \rightarrow \mathbb{Z}$$

- For convenience, define a state update function:

$$\text{assign } v \ f \ s \equiv \lambda w. \text{ if } v = w \text{ then } f \ s \text{ else } s \ w$$

pGCL Commands

Model pGCL commands with a HOL datatype:

command \equiv Abort
| Skip
| Assign of string \times (state $\rightarrow \mathbb{Z}$)
| Seq of command \times command
| Demon of command \times command
| Prob of (state \rightarrow posreal) \times command \times command
| While of (state $\rightarrow \mathbb{B}$) \times command

Note: the probability in Prob can depend on the state.

Derived Commands

Define the following *derived commands* as syntactic sugar:

$$v := f \quad \equiv \quad \text{Assign } v \ f$$

$$c_1 ; c_2 \quad \equiv \quad \text{Seq } c_1 \ c_2$$

$$c_1 \sqcap c_2 \quad \equiv \quad \text{Demon } c_1 \ c_2$$

$$c_1 \ p \oplus \ c_2 \quad \equiv \quad \text{Prob } (\lambda s. p) \ c_1 \ c_2$$

$$\text{Cond } b \ c_1 \ c_2 \quad \equiv \quad \text{Prob } (\lambda s. \text{if } b \ s \ \text{then } 1 \ \text{else } 0) \ c_1 \ c_2$$

$$v := \{e_1, \dots, e_n\} \quad \equiv \quad v := e_1 \ \sqcap \ \dots \ \sqcap \ v := e_n$$

$$v := \langle e_1, \dots, e_n \rangle \quad \equiv \quad v := e_1 \ 1/n \oplus \ v := \langle e_2, \dots, e_n \rangle$$

$$b_1 \rightarrow c_1 \mid \dots \mid b_n \rightarrow c_n \quad \equiv$$

$$\left\{ \begin{array}{l} \text{Abort} \quad \text{if none of the } b_i \text{ hold on the current state} \\ \prod_{i \in I} c_i \quad \text{where } I = \{i \mid 1 \leq i \leq n \wedge b_i \text{ holds}\} \end{array} \right.$$

Weakest Preconditions

Define weakest preconditions (wp) directly on commands:

$$\vdash (\text{wp } \text{Abort} = \lambda e. \text{Zero})$$

$$\wedge (\text{wp } \text{Skip} = \lambda e. e)$$

$$\wedge (\text{wp } (\text{Assign } v \ f) = \lambda e, s. e \ (\text{assign } v \ f \ s))$$

$$\wedge (\text{wp } (\text{Seq } c_1 \ c_2) = \lambda e. \text{wp } c_1 \ (\text{wp } c_2 \ e))$$

$$\wedge (\text{wp } (\text{Demon } c_1 \ c_2) = \lambda e. \text{Min } (\text{wp } c_1 \ e) \ (\text{wp } c_2 \ e))$$

$$\wedge (\text{wp } (\text{Prob } p \ c_1 \ c_2) = \lambda e. \text{Lin } p \ (\text{wp } c_1 \ e) \ (\text{wp } c_2 \ e))$$

$$\wedge (\text{wp } (\text{While } b \ c) = \lambda e. \text{expect_lfp } (\lambda e'. \text{Cond } b \ (\text{wp } c \ e') \ e))$$

Weakest Preconditions: Example

- The goal is to end up with variables i and j containing the same value:

$$post \equiv \text{if } i = j \text{ then } 1 \text{ else } 0.$$

- First program:

$$\begin{aligned} pd &\equiv i := \langle 0, 1 \rangle ; j := \{0, 1\} \\ &\vdash wp \text{ } pd \text{ } post = \text{Zero} \end{aligned}$$

- Second program:

$$\begin{aligned} dp &\equiv j := \{0, 1\} ; i := \langle 0, 1 \rangle \\ &\vdash wp \text{ } dp \text{ } post = \lambda s. 1/2. \end{aligned}$$

Example: Monty Hall

contestant *switch* \equiv

$pc := \{1, 2, 3\} ;$

$cc := \langle 1, 2, 3 \rangle ;$

$pc \neq 1 \wedge cc \neq 1 \rightarrow ac := 1$

| $pc \neq 2 \wedge cc \neq 2 \rightarrow ac := 2$

| $pc \neq 3 \wedge cc \neq 3 \rightarrow ac := 3 ;$

if \neg *switch* then Skip else

$cc :=$ (if $cc \neq 1 \wedge ac \neq 1$ then 1

else if $cc \neq 2 \wedge ac \neq 2$ then 2 else 3)

The postcondition is simply the desired goal of the contestant, i.e.,

$win \equiv$ if $cc = pc$ then 1 else 0.

Example: Monty Hall

- Verification proceeds by:
 1. Rewriting away all the syntactic sugar.
 2. Expanding the definition of wp .
 3. Carrying out the numerical calculations.
- After 22 seconds and 250536 primitive inferences in the logical kernel:
$$\vdash wp(\text{contestant } \mathit{switch}) \text{ win} = \lambda s. \text{ if } \mathit{switch} \text{ then } 2/3 \text{ else } 1/3$$
- In other words, by switching the contestant is twice as likely to win the prize.
- Not trivial to do by hand, because the intermediate expectations get rather large.

Contents

- Introduction
- Formalizing Probabilistic Guarded Commands
- wlp **Verification Condition Generator**
- Example: Rabin's Mutual Exclusion Algorithm
- Conclusion

Weakest Liberal Preconditions

Weakest liberal conditions (wlp) model partial correctness.

- ⊢ (wlp Abort = $\lambda e.$ **Infty**)
- ∧ (wlp Skip = $\lambda e.$ e)
- ∧ (wlp (Assign $v f$) = $\lambda e, s.$ e (assign $v f s$))
- ∧ (wlp (Seq $c_1 c_2$) = $\lambda e.$ wlp c_1 (wlp $c_2 e$))
- ∧ (wlp (Demon $c_1 c_2$) = $\lambda e.$ Min (wlp $c_1 e$) (wlp $c_2 e$))
- ∧ (wlp (Prob $p c_1 c_2$) = $\lambda e.$ Lin p (wlp $c_1 e$) (wlp $c_2 e$))
- ∧ (wlp (While $b c$) = $\lambda e.$ **expect_gfp** ($\lambda e'.$ Cond b (wlp $c e'$) e))

Weakest Liberal Preconditions: Example

- We illustrate the difference between wp and wlp on the simplest infinite loop:

$$\text{loop} \equiv \text{While } (\lambda s. \top) \text{ Skip}$$

- For any postcondition post , we have

$$\vdash \text{wp loop } \text{post} = \text{Zero} \wedge \text{wlp loop } \text{post} = \text{Infty}$$

- These correspond to the Hoare triples

$$[\perp] \text{ loop } [\text{post}] \quad \{\top\} \text{ loop } \{\text{post}\}$$

as we would expect from an infinite loop.

Calculating wlp Lower Bounds

- Suppose we have a pGCL command c and a postcondition q .
- We wish to derive a lower bound on the weakest liberal precondition.
 - In general, programs are shown to have desirable properties by proving *lower bounds*.
 - Example: $\vdash (\lambda s. 0.95) \sqsubseteq \text{wp Prog (if ok then 1 else 0)}$
- Can think of this as the query $P \sqsubseteq \text{wlp } c \ q$.
- **Idea:** use a Prolog interpreter to solve for the variable P .

Calculating wlp: Rules

Simple rules:

- $\text{Infty} \sqsubseteq \text{wlp Abort } Q$
- $Q \sqsubseteq \text{wlp Skip } Q$
- $R \sqsubseteq \text{wlp } C_2 Q \wedge P \sqsubseteq \text{wlp } C_1 R$
 \Rightarrow
 $P \sqsubseteq \text{wlp (Seq } C_1 C_2) Q$

Note: the Prolog interpreter automatically calculates the ‘middle condition’ in a Seq command.

Calculating wlp: While Loops

- Define an assertion command: $\text{Assert } p \ c \equiv c$.
- Provide a while rule that requires an assertion:
 - $R \sqsubseteq \text{wlp } C \ P \ \wedge \ P \sqsubseteq \text{wlp_cond } b \ R \ Q$
 \Rightarrow
 $P \sqsubseteq \text{wlp } (\text{Assert } P \ (\text{While } b \ c)) \ Q$
- The second premise generates a *verification condition* as an extra subgoal.
- It is left to the user to provide a useful loop invariant in the Assert around the while loop.

Contents

- Introduction
- Formalizing Probabilistic Guarded Commands
- wlp Verification Condition Generator
- **Example: Rabin's Mutual Exclusion Algorithm**
- Conclusion

Rabin's Mutual Exclusion Algorithm

- Suppose N processors are executing concurrently, and from time to time some of them need to enter a critical section of code.
- The mutual exclusion algorithm of Rabin (1982, 1992) works by electing a leader who is permitted to enter the critical section:
 1. Each of the waiting processors repeatedly tosses a fair coin until a head is shown
 2. The processor that required the largest number of tosses wins the election.
 3. If there is a tie, then have another election.
- Could implement the coin tossing using
$$n := 0 ; b := 0 ; \text{While } (b = 0) (n := n + 1 ; b := \langle 0, 1 \rangle)$$

Rabin's Mutual Exclusion Algorithm

For our verification, we do not model i processors concurrently executing the above voting scheme, but rather the following data refinement of that system:

1. Initialize i with the number of processors waiting to enter the critical section who have just picked a number.
2. Initialize n with 1, the lowest number not yet considered.
3. If $i = 1$ then we have a unique winner: return SUCCESS.
4. If $i = 0$ then the election has failed: return FAILURE.
5. Reduce i by eliminating all the processors who picked the lowest number n (since certainly none of them won the election).
6. Increment n by 1, and jump to Step 3.

Rabin's Mutual Exclusion Algorithm

The following pGCL program implements this data refinement:

$$\begin{aligned} \text{rabin} \quad \equiv \quad & \text{While } (1 < i) (\\ & \quad n := i ; \\ & \quad \text{While } (0 < n) \\ & \quad \quad (d := \langle 0, 1 \rangle ; i := i - d ; n := n - 1) \\ &) \end{aligned}$$

The desired postcondition representing a unique winner of the election is

$$\text{post} \equiv \text{if } i = 1 \text{ then } 1 \text{ else } 0$$

Rabin's Mutual Exclusion Algorithm

- The precondition that we aim to show is

$$pre \equiv \text{if } i = 1 \text{ then } 1 \text{ else if } 1 < i \text{ then } 2/3 \text{ else } 0$$

“For any positive number of processors wanting to enter the critical section, the probability that the voting scheme will produce a unique winner is 2/3, except for the trivial case of one processor when it will always succeed.”

- **Surprising:** The probability of success is independent of the number of processors.
- We formally verify the following statement of partial correctness:

$$pre \sqsubseteq \text{wlp rabin } post$$

Rabin's Mutual Exclusion Algorithm

- Need to annotate the While loops with invariants.
- The invariant for the outer loop is simply pre .
- For the inner loop we used

if $0 \leq n \leq i$ then $2/3 \times \text{invar1 } i \ n + \text{invar2 } i \ n$ else 0

where

$\text{invar1 } i \ n \equiv$

$1 - (\text{if } i = n \text{ then } (n + 1)/2^n \text{ else if } i = n + 1 \text{ then } 1/2^n \text{ else } 0)$

$\text{invar2 } i \ n \equiv \text{if } i = n \text{ then } n/2^n \text{ else if } i = n + 1 \text{ then } 1/2^n \text{ else } 0$

- Coming up with these was the hardest part of the verification.

Rabin's Mutual Exclusion Algorithm

The verification proceeded as follows:

1. Create the annotated program `annotated_rabin`.
2. Prove `rabin = annotated_rabin`
3. Use this to reduce the goal to

$$pre \sqsubseteq \text{wlp } \text{annotated_rabin } post$$

4. This is now in the correct form to apply the VC generator.
5. Finish off the VCs with 58 lines of HOL-4 proof script.

```
| - Leq (\s. if s"i" = 1 then 1
           else if 1 < s"i" then 2/3 else 0)
      (wlp rabin (\s. if s"i" = 1 then 1 else 0))
```

Contents

- Introduction
- Formalizing Probabilistic Guarded Commands
- wlp Verification Condition Generator
- Example: Rabin's Mutual Exclusion Algorithm
- **Conclusion**

Conclusion

- Formalized the theory of pGCL in higher-order logic.
 - Definitional theory, so high assurance of consistency.
- Created an automatic tool for deriving sufficient conditions for partial correctness.
 - Useful product of mechanizing a program semantics.
- HOL-4 well suited to this task.
 - Hard VCs can be passed to the user as subgoals.

Related Work

- Formal methods for probabilistic programs:
 - Christine Paulin's work in Coq, 2002.
 - Prism model checker, Kwiatkowska et. al., 2000–
- Mechanized program semantics:
 - Formalizing Dijkstra, Harrison, 1998.
 - Mechanizing program logics in higher order logic, Gordon, 1989.