

Formal Verification of Probabilistic Programs: Two Approaches

Joe Hurd

`joe.hurd@cl.cam.ac.uk`

University of Cambridge

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

Contents

- **Introduction**
- Approach 1: Monads
 - Formalizing Probability
 - Modelling Probabilistic Programs
 - Example Verifications
- Approach 2: pGCL
 - Formalizing Probabilistic Guarded Commands
 - wlp Verification Condition Generator
 - Example Verifications
- Conclusion

Introduction

Probabilistic programs are useful for many applications:

- Symmetry breaking
 - Rabin's mutual exclusion algorithm
- Eliminating pathological cases
 - Miller-Rabin primality test
- Algorithm complexity
 - Sorting nuts and bolts
- Defeating a powerful adversary
 - Mixed strategies in game theory
- Solving a problem in an extremely simple way
 - Finding minimal cuts

Introduction

- Quicksort Algorithm (Hoare, 1962):

```
fun quicksort elements =  
  if length elements <= 1 then elements  
  else  
    let  
      val pivot          = choose_pivot elements  
      val (left, right) = partition pivot elements  
    in  
      quicksort left @ [pivot] @ quicksort right  
    end;
```

- Usually $O(n \log n)$ comparisons, unless choice of pivot interacts badly with data.

Introduction

- Example of bad behaviour when pivot is first element:

```
input:      [5, 4, 3, 2, 1]
pivot 5:    [4, 3, 2, 1]--5--[]
pivot 4:    [3, 2, 1]--4--[]
pivot 3:    [2, 1]--3--[]
pivot 2:    [1]--2--[]
output:     [1, 2, 3, 4, 5]
```

- Lists in reverse order take $O(n^2)$ comparisons.
- So do lists that are in the right order!

Introduction

- Solution: Introduce randomization into the algorithm itself.
- Pick pivots uniformly at random from the list of elements.
- Every list has exactly the same performance profile:
 - Expected number of comparisons is $O(n \log n)$.
 - Small class $C \subset S_n$ of lists with guaranteed bad performance has been replaced with a small probability $|C|/n!$ of bad performance on any input.

Introduction

- Broken procedure for choosing a pivot:

```
fun choose_pivot elements =  
  if length elements = 1 orelse coin_flip ()  
  then hd elements  
  else choose_pivot (tl elements);
```

- Not a uniform distribution when length of elements > 2 .
- Actually reinstates a bad class of input lists taking $O(n^2)$ (expected) comparisons.
- Would like to verify probabilistic programs in a theorem prover.

The HOL Theorem Prover

- Developed by Mike Gordon's Hardware Verification Group in Cambridge, first release was HOL88.
- Latest release in mid-2002 called HOL4, developed jointly by Cambridge, Utah and ANU.
- Implements classical Higher-Order Logic with Hindley-Milner polymorphism.
- Sprung from the Edinburgh LCF project, so has a small logical kernel to ensure soundness.
- Links to external proof tools, either as oracles (e.g., SAT solvers) or by translating their proofs (e.g., Gandalf).
- Comes with a large library of theorems contributed by many users over the years, including theories of lists, real analysis, groups etc.

Contents

- Introduction
- **Approach 1: Monads**
 - Formalizing Probability
 - Modelling Probabilistic Programs
 - Example Verifications
- Approach 2: pGCL
 - Formalizing Probabilistic Guarded Commands
 - wlp Verification Condition Generator
 - Example Verifications
- Conclusion

Introduction: Monads

To verify a probabilistic program in HOL:

- Must be able to formalize its probabilistic specification;

$$\mathcal{E} : \mathcal{P}(\mathcal{P}(\mathbb{B}^\infty)), \quad \mathbb{P} : \mathcal{E} \rightarrow \mathbb{R}$$

- and model the probabilistic program in the logic;

$$\text{prob_program} : \mathbb{N} \rightarrow \mathbb{B}^\infty \rightarrow \{\text{success, failure}\} \times \mathbb{B}^\infty$$

- then finally **prove** that the program satisfies its specification.

$$\vdash \forall n. \mathbb{P} \{s \mid \text{fst} (\text{prob_program } n \ s) = \text{failure}\} \leq 2^{-n}$$

Contents

- Introduction
- Approach 1: Monads
 - **Formalizing Probability**
 - Modelling Probabilistic Programs
 - Example Verifications
- Approach 2: pGCL
 - Formalizing Probabilistic Guarded Commands
 - wlp Verification Condition Generator
 - Example Verifications
- Conclusion

Formalizing Probability

- Need to construct a probability space of Bernoulli($\frac{1}{2}$) sequences, to give meaning to specifications like

$$\mathbb{P} \{s \mid \text{fst} (\text{prob_program } n \ s) = \text{failure}\}$$

- To ensure soundness, would like it to be a purely definitional extension of HOL (no axioms).
- Use measure theory, and end up with a set \mathcal{E} of events and a probability function \mathbb{P} :

$$\mathcal{E} = \{S \subset \mathbb{B}^\infty \mid S \text{ is a measurable set}\}$$

$$\mathbb{P}(S) = \text{the probability measure of } S \text{ (for } S \in \mathcal{E}\text{)}$$

Formalizing Probability

- Formalized some general measure theory in HOL, including Carathéodory's extension theorem.
- Next defined the measure of prefix sets (or cylinders):

$$\forall l. \mu \{s_0 s_1 s_2 \cdots \mid [s_0, \dots, s_{n-1}] = l\} = 2^{-(\text{length } l)}$$

- Finally extended this measure to a σ -algebra:

$$\mathcal{E} = \sigma(\text{prefix sets})$$

$$\mathbb{P} = \text{Carathéodory extension of } \mu \text{ to } \mathcal{E}$$

- Similar to the definition of Lebesgue measure.

Contents

- Introduction
- Approach 1: Monads
 - Formalizing Probability
 - **Modelling Probabilistic Programs**
 - Example Verifications
- Approach 2: pGCL
 - Formalizing Probabilistic Guarded Commands
 - wlp Verification Condition Generator
 - Example Verifications
- Conclusion

Modelling Probabilistic Programs

- Given a probabilistic ‘function’:

$$\hat{f} : \alpha \rightarrow \beta$$

- Model \hat{f} with a higher-order logic function

$$f : \alpha \rightarrow \mathbb{B}^\infty \rightarrow \beta \times \mathbb{B}^\infty$$

that passes around ‘an infinite sequence of coin-flips.’

- The probability that $\hat{f}(a)$ meets a specification $B : \beta \rightarrow \mathbb{B}$ can then be formally defined as

$$\mathbb{P} \{s \mid B(\text{fst } (f \ a \ s))\}$$

Modelling Probabilistic Programs

- Can use state-transformer monadic notation to express HOL models of probabilistic programs:

$$\text{unit } a = \lambda s. (a, s)$$

$$\text{bind } f \ g = \lambda s. \text{let } (x, s') \leftarrow f(s) \text{ in } g \ x \ s'$$

$$\text{coin_flip } f \ g = \lambda s. (\text{if shd } s \text{ then } f \text{ else } g, \text{stl } s)$$

- For example, if `dice` is a program that generates a dice throw from a sequence of coin flips, then

$$\text{two_dice} = \text{bind } \text{dice} \ (\lambda x. \text{bind } \text{dice} \ (\lambda y. \text{unit } (x + y)))$$

generates the sum of two dice.

Example: The Binomial($n, \frac{1}{2}$) Distribution

- Definition of a sampling algorithm for the Binomial($n, \frac{1}{2}$) distribution:

$\vdash \text{bit} = \text{coin_flip} (\text{unit } 1) (\text{unit } 0)$

$\vdash \text{binomial } 0 = \text{unit } 0 \wedge$

$\forall n.$

$\text{binomial} (\text{suc } n) =$


$\text{bind bit } (\lambda x. \text{bind} (\text{binomial } n) (\lambda y. \text{unit } (x + y)))$

- Correctness theorem:

$$\vdash \forall n, r. \mathbb{P} \{s \mid \text{fst} (\text{binomial } n \ s) = r\} = \binom{n}{r} \left(\frac{1}{2}\right)^n$$

Probabilistic Termination

- The Binomial($n, \frac{1}{2}$) sampling algorithm is **guaranteed to terminate** within n coin-flips.
- The following algorithm generates dice throws from coin-flips (Knuth and Yao, 1976):



- The backward loops introduce the possibility of looping forever.
- But the probability of this happening is 0.
- **Probabilistic termination:** the program terminates with probability 1.

Probabilistic Termination

- Probabilistic termination is more expressive than guaranteed termination.
- No coin-flip algorithm that is guaranteed to terminate can sample from the following distributions:
 - Uniform(3): choosing one of 0, 1, 2 each with probability $\frac{1}{3}$.
 - Geometric($\frac{1}{2}$): choosing $n \in \mathbb{N}$ with probability $(\frac{1}{2})^{n+1}$.
The index of the first head in a sequence of coin-flips.
- We model probabilistic termination in HOL using a probabilistic while loop:

$\vdash \forall c, b, a.$

$\text{while } c \ b \ a = \text{if } c(a) \text{ then bind } (b \ a) \ (\text{while } c \ b) \ \text{else unit } a$

Contents

- Introduction
- Approach 1: Monads
 - Formalizing Probability
 - Modelling Probabilistic Programs
 - **Example Verifications**
- Approach 2: pGCL
 - Formalizing Probabilistic Guarded Commands
 - wlp Verification Condition Generator
 - Example Verifications
- Conclusion

Example: The Uniform(3) Distribution

- First make a raw definition of `unif3`:

\vdash `unif3 =`
`while ($\lambda n. n = 3$)`
`(coin_flip (coin_flip (unit 0) (unit 1)) (coin_flip (unit 2) (unit 3))) 3`

- Next prove `unif3` satisfies probabilistic termination.
- This allows us to derive a recursive definition of `unif3`:


\vdash `unif3 = coin_flip (coin_flip (unit 0) (unit 1)) (coin_flip (unit 2) unif3)`

- The correctness theorem also follows:

$\vdash \forall n. \mathbb{P} \{s \mid \text{fst} (\text{unif3 } s) = n\} = \text{if } n < 3 \text{ then } \frac{1}{3} \text{ else } 0$

Example: Optimal Dice

A probabilistic finite state automaton:



```
dice =  
coin_flip  
(prob_repeat  
  (coin_flip  
    (coin_flip  
      (unit none)  
      (unit (some 1)))  
    (mmap some  
      (coin_flip  
        (unit 2)  
        (unit 3))))))  
(prob_repeat  
  (coin_flip  
    (mmap some  
      (coin_flip  
        (unit 4)  
        (unit 5)))  
    (coin_flip  
      (unit (some 6))  
      (unit none))))))
```

Example: Optimal Dice

- Correctness theorem:

$$\vdash \forall n. \mathbb{P} \{s \mid \text{fst}(\text{dice } s) = n\} = \text{if } 1 \leq n \wedge n \leq 6 \text{ then } \frac{1}{6} \text{ else } 0$$


- The dice program takes $3\frac{2}{3}$ coin flips (on average) to output a dice throw.
- Knuth and Yao (1976) show this to be optimal.
- To generate the sum of two dice throws, is it possible to do better than $7\frac{1}{3}$ coin flips?

Example: Optimal Dice

On average, this program takes $4\frac{7}{18}$ coin flips to produce a result, and this is also optimal.


⊢ $\forall n.$

$\mathbb{P}\{s \mid \text{fst}(\text{two_dice } s) = n\} =$
if $n = 2 \vee n = 12$ then $\frac{1}{36}$
else if $n = 3 \vee n = 11$ then $\frac{2}{36}$
else if $n = 4 \vee n = 10$ then $\frac{3}{36}$
else if $n = 5 \vee n = 9$ then $\frac{4}{36}$
else if $n = 6 \vee n = 8$ then $\frac{5}{36}$
else if $n = 7$ then $\frac{6}{36}$
else 0



Example: Random Walk

- A drunk exits a pub at point n , and lurches left and right with equal probability until he hits home at point 0.



- Will the drunk always get home?

Example: Random Walk

- Perhaps surprisingly, the drunk **does** always get home.
 - We formalize the proof of this in HOL.
 - Thus the formalized random walk satisfies probabilistic termination.
- This allows us to derive a natural definition of walk:

$\vdash \forall n, k.$

$\text{walk } n \ k =$

$\text{if } n = 0 \text{ then unit } k \text{ else}$

$\text{coin_flip } (\text{walk } (n+1) \ (k+1)) \ (\text{walk } (n-1) \ (k+1))$

- And prove some neat properties:

$\vdash \forall n, k. \forall^* s. \text{even } (\text{fst } (\text{walk } n \ k \ s)) = \text{even } (n + k)$

Example: Random Walk

- Can extract walk to ML and simulate it.
 - Use high-quality random bits from `/dev/random`.
- A typical sequence of results from random walks starting at level 1:

57, 1, 7, 173, 5, 49, 1, 3, 1, 11, 9, 9, 1, 1, 1547, 27, 3, 1, 1, 1, ...

- Record breakers:
 - 34th simulation yields a walk with 2645 steps
 - 135th simulation yields a walk with 603787 steps
 - 664th simulation yields a walk with 1605511 steps
- Expected number of steps to get home is infinite!

Example: Miller-Rabin Primality Test

The Miller-Rabin algorithm is a **probabilistic primality test**, used by commercial software such as Mathematica.

We formalize the test as a HOL function `miller`, and prove:

$$\vdash \forall n, t, s. \text{prime } n \Rightarrow \text{fst } (\text{miller } n \ t \ s) = \top$$

$$\vdash \forall n, t. \neg \text{prime } n \Rightarrow 1 - 2^{-t} \leq \mathbb{P} \{s \mid \text{fst } (\text{miller } n \ t \ s) = \perp\}$$

Here n is the number to test for primality, and t is the maximum number of iterations allowed.

Example: Miller-Rabin Primality Test

- Can define a pseudo-random number generator in HOL, and interpret miller in the logic to prove numbers composite:

$$\vdash \neg\text{prime}(2^{2^6} + 1) \wedge \neg\text{prime}(2^{2^7} + 1) \wedge \neg\text{prime}(2^{2^8} + 1)$$

- Or can manually extract miller to ML, and execute it using `/dev/random` and calls to GMP:

bits	$\mathbb{E}_{l,n}$	MR	Gen time	MR ₁ time
500	99424	99458	0.0443	0.2498
1000	99712	99716	0.0881	0.7284
2000	99856	99852	0.3999	4.2910

Contents

- Introduction
- Approach 1: Monads
 - Formalizing Probability
 - Modelling Probabilistic Programs
 - Example Verifications
- **Approach 2: pGCL**
 - Formalizing Probabilistic Guarded Commands
 - wlp Verification Condition Generator
 - Example Verifications
- 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–

Contents

- Introduction
- Approach 1: Monads
 - Formalizing Probability
 - Modelling Probabilistic Programs
 - Example Verifications
- Approach 2: pGCL
 - **Formalizing Probabilistic Guarded Commands**
 - wlp Verification Condition Generator
 - Example Verifications
- Conclusion

pGCL Semantics

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

$$[P]C[Q]$$

- Precondition P is weaker than P' if $P' \Rightarrow P$.
- Such a P will always exist and be unique, so 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 $\alpha \rightarrow \text{posreal}$.
 - Expectation P is weaker than P' if $P' \sqsubseteq P$.
 - Think of programs as *expectation transformers*.

pGCL Commands

Model pGCL commands with a HOL datatype:

command \equiv Assert of (state \rightarrow posreal) \times command
| 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:

$$\begin{aligned}v := e &\equiv \text{Assign } v \ e \\c_1 ; c_2 &\equiv \text{Seq } c_1 \ c_2 \\c_1 \sqcap c_2 &\equiv \text{Demon } c_1 \ c_2 \\c_1 \ p \oplus \ c_2 &\equiv \text{Prob } (\lambda s. p) \ c_1 \ c_2 \\ \text{Cond } b \ c_1 \ c_2 &\equiv \text{Prob } (\lambda s. \text{if } b \ s \ \text{then } 1 \ \text{else } 0) \ c_1 \ c_2 \\v := \{e_1, \dots, e_n\} &\equiv v := e_1 \ \sqcap \ \dots \ \sqcap \ v := e_n \\v := \langle e_1, \dots, e_n \rangle &\equiv v := e_1 \ 1/n \oplus \ v := \langle e_2, \dots, e_n \rangle \\p_1 \rightarrow c_1 \mid \dots \mid p_n \rightarrow c_n &\equiv \\ \left\{ \begin{array}{l} \text{Abort} \\ \prod_{i \in I} c_i \end{array} \right. &\quad \begin{array}{l} \text{if none of the } p_i \text{ hold on the current state} \\ \text{where } I = \{i \mid 1 \leq i \leq n \wedge p_i \text{ holds}\} \end{array}\end{aligned}$$

In addition, we write $v := n + 1$ instead of “ $v := \lambda s. s \text{ “}n\text{”} + 1$.”

Weakest Preconditions

Define weakest preconditions (wp) directly on commands:

$$\vdash (\text{wp } (\text{Assert } p \ c) = \text{wp } c)$$

$$\wedge (\text{wp } \text{Abort} = \lambda r. \text{Zero})$$

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

$$\wedge (\text{wp } (\text{Assign } v \ e) = \lambda r, s. r (\lambda w. \text{if } w = v \text{ then } e \ s \ \text{else } s \ w))$$

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

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

$$\wedge (\text{wp } (\text{Prob } p \ c_1 \ c_2) =$$

$$\lambda r, s. \text{let } x \leftarrow [p \ s]_{\leq 1} \text{ in } x(\text{wp } c_1 \ r \ s) + (1 - x)(\text{wp } c_2 \ r \ s))$$

$$\wedge (\text{wp } (\text{While } b \ c) =$$

$$\lambda r. \text{expect_lfp } (\lambda e, s. \text{if } b \ s \ \text{then } \text{wp } c \ e \ s \ \text{else } r \ s))$$

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:

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

Contents

- Introduction
- Approach 1: Monads
 - Formalizing Probability
 - Modelling Probabilistic Programs
 - Example Verifications
- Approach 2: pGCL
 - Formalizing Probabilistic Guarded Commands
 - wlp **Verification Condition Generator**
 - Example Verifications
- Conclusion

Weakest Liberal Preconditions

Weakest liberal conditions (wlp) model partial correctness.

- ⊢ (wlp (Assert p c) = wlp c)
- ∧ (wlp Abort = $\lambda r. \text{Magic}$)
- ∧ (wlp Skip = $\lambda r. r$)
- ∧ (wlp (Assign v e) = $\lambda r, s. r$ ($\lambda w. \text{if } w = v \text{ then } e \text{ s else } s \ w$))
- ∧ (wlp (Seq c_1 c_2) = $\lambda r. \text{wlp } c_1$ (wlp c_2 r))
- ∧ (wlp (Demon c_1 c_2) = $\lambda r. \text{Min}$ (wlp c_1 r) (wlp c_2 r))
- ∧ (wlp (Prob p c_1 c_2) =
 $\lambda r, s. \text{let } x \leftarrow [p \ s]_{\leq 1} \text{ in } x(\text{wlp } c_1 \ r \ s) + (1 - x)(\text{wlp } c_2 \ r \ s)$)
- ∧ (wlp (While b c) =
 $\lambda r. \text{expect_gfp}$ ($\lambda e, s. \text{if } b \ s \ \text{then } \text{wlp } c \ e \ s \ \text{else } r \ s$))

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

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

- These correspond to the Hoare triples

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

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.
- Can think of this as the first-order query $P \sqsubseteq \text{wlp } c \ q$.
- **Idea:** use a Prolog interpreter to solve for the variable P .

Calculating wlp: Rules

Example Rules:

- $\text{Magic} \sqsubseteq \text{wlp Abort } Q$
- $Q \sqsubseteq \text{wlp Skip } Q$
- $R \sqsubseteq \text{wlp } C_2 Q \wedge P \sqsubseteq \text{wlp } C_1 R \Rightarrow P \sqsubseteq \text{wlp (Seq } C_1 C_2) Q$
- $P_1 \sqsubseteq \text{wlp } C_1 Q \wedge P_2 \sqsubseteq \text{wlp } C_2 Q \Rightarrow \text{Min } P_1 P_2 \sqsubseteq \text{wlp (Demon } C_1 C_2) Q$

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

Calculating wlp: While Loops

- We use the following theorem about While loops:

$$\vdash \forall P, Q, b, c.$$

$$P \sqsubseteq \text{If } b \text{ (wlp } c \text{ } P) \text{ } Q \Rightarrow P \sqsubseteq \text{wlp (While } b \text{ } c) \text{ } Q$$

- Cannot use in this form, because of the repeated occurrence of P in the premise.
- Instead, provide a rule that requires an assertion:
 - $R \sqsubseteq \text{wlp } C \text{ } P \wedge P \sqsubseteq \text{If } b \text{ } R \text{ } Q \Rightarrow P \sqsubseteq \text{wlp (Assert } P \text{ (While } b \text{ } c)) } Q$
- 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
- Approach 1: Monads
 - Formalizing Probability
 - Modelling Probabilistic Programs
 - Example Verifications
- Approach 2: pGCL
 - Formalizing Probabilistic Guarded Commands
 - wlp Verification Condition Generator
 - **Example Verifications**
- Conclusion

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.

Example: Rabin Mutual Exclusion

- 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
$$n := 0 ; b := 0 ; \text{While } (b = 0) (n := n + 1 ; b := \langle 0, 1 \rangle)$$

Example: Rabin Mutual Exclusion

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.

Example: Rabin Mutual Exclusion

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

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

Example: Rabin Mutual Exclusion

- 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:

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

Example: Rabin Mutual Exclusion

- 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 $(2/3) * \text{invar1 } i \ n + \text{invar2 } i \ n$ 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.

Example: Rabin Mutual Exclusion

The verification proceeded as follows:

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

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

4. This is 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
- Approach 1: Monads
 - Formalizing Probability
 - Modelling Probabilistic Programs
 - Example Verifications
- Approach 2: pGCL
 - Formalizing Probabilistic Guarded Commands
 - wlp Verification Condition Generator
 - Example Verifications
- **Conclusion**

Conclusion

Advantages of Monad Approach

- Grounded in measure theory.
 - Probabilities more than real numbers.
- More suitable for verifying functional programs.
 - Simple to lift verified HOL functions to ML.
- Can reason about the distinction between probabilistic and guaranteed termination.
 - Practical difference: operating systems typically provide a source of random bits.

Conclusion

Advantages of pGCL Approach

- Supports the demonic choice programming construct.
 - Can be used to verify distributed algorithms.
- Verification easier to carry out than monad approach.
 - Modelling programs with expectation transformers is a useful abstraction.
- Deep embedding: can quantify over all programs.
 - May be useful for modelling a ‘spy’ in a security protocol verification.

Future Work: combine these approaches to get the best of both worlds.

Related Work

- Formal methods for probabilistic programs:
 - Hurd's thesis, 2002.
 - Probabilistic invariants for probabilistic machines, Hoang et. al., 2003.
 - Christine Paulin's work in Coq, 2002.
 - Prism model checker, Kwiatkowska et. al., 2000–
- Mechanized program semantics:
 - Formalizing Dijkstra, Harrison, 1998.
 - Hoare Logics in Isabelle/HOL, Nipkow, 2001.
 - Mechanizing program logics in higher order logic, Gordon, 1989.
 - A mechanically verified verification condition generator, Homeier and Martin, 1995.

Related Work

- Semantics of Probabilistic Programs:
 - *Semantics of Probabilistic Programs*, Kozen, 1979.
 - *Termination of Probabilistic Concurrent Processes*, Hart, Sharir and Pnueli, 1983.
 - *Probabilistic Non-Determinism*, Jones, 1990.
 - Probabilistic predicate transformers, Morgan, McIver, Seidel and Sanders, 1994–
 - *Notes on the Random Walk: an Example of Probabilistic Temporal Reasoning*, 1996
 - *Proof Rules for Probabilistic Loops*, Morgan, 1996