

CITS5501 Software Testing and Quality Assurance

Formal methods

Unit coordinator: Arran Stewart

May 1, 2018

Sources

- Pressman, R., *Software Engineering: A Practitioner's Approach*, McGraw-Hill, 2005
- Huth and Ryan, *Logic in Computer Science*
- Pierce et al, *Software Foundations vol 1*

formal methods

Overview

- When doing software engineering – specifying and developing software systems – the activities done can be done with varying levels of mathematical rigor.
- For instance, we could write a requirement
 - informally, just using natural language, and perhaps tables and diagrams. This is easy, but can be imprecise and ambiguous (and hard to spot when that has occurred)
 - semi-formally, perhaps using occasional mathematical formulas or bits of pseudocode to express what's required
 - mostly using mathematical notation, with a bit of English to clarify what the notation represents. This is typically a lot more work, and it can be harder to ensure the notation matches our intuitive idea of what the system should do, but has little or no vagueness or ambiguity.

Overview (2)

- Things towards the “more formal” side of this spectrum will tend to get called “lightweight formal methods” or “formal methods”.
- They are sometimes described as “mathematically based techniques for the specification, development and verification of software and hardware systems” [Wikipedia], but without more detail that’s not very illuminating.
- So we’ll start with an example formal method (verification of programs using Hoare logic), then come back to the definition, and look at other sorts of formal methods.

A typical approach

Often, we'll apply formal methods in the following way:

- We'll have some specification – some property that we want our system to have
 - e.g., that it calculates the factorial of a natural number; or never gets deadlocked; or has certain security properties.
- And we'll have something representing the system – this is called a *model*
 - This could be actual code, or it could be annotated code, or it could be some more abstract model of the system (like state machines, which we have seen earlier)
- And we will try to show that the model meets the specification.

Rationale

- Why use formal methods?
- Building reliable software is very, very hard.
 - Software systems can be hugely complex, and knowing exactly what it is doing at any point of time is likewise hard.
- So computer scientists and software engineers have come up with all sorts of techniques for improving reliability (many of which we've seen) – testing, risk management, quality controls, maths-based techniques for reasoning about the properties of software
 - And this last sort of technique is what we call formal methods.

Rationale

- By *reasoning* about the properties of software – i.e., proving things about it – we can get much greater certainty that our programs are reliable and error-free, than we can through testing
- Testing is a sort of *empirical investigation* – we go out and check whether we can find something (bugs, in this case)
- But if we don't find it, that doesn't mean that whatever we were looking for doesn't exist – we may not have looked hard enough or in the right places.
 - (People once thought it was an eternal and obvious truth that there weren't such things as black swans, but it turned out they weren't looking in the right places.)

Example

- Consider a bit of code, in some Python-like language, for multiplying i by j
 - (We will look at more complex examples later)

```
n := 0
while j > 0:
    n := n + i
    j := j - 1
return n
```

How can we show that this code is correct? (Assuming it is.)

Example (2)

- We could try it on a number of inputs. (We might try using 0, values near 0, and values not especially near zero.)
- We could inspect the code and see if it conforms to our idea of multiplication
- We could try to *prove* that, once the code has finished executing, for *any* integers i and j , n will equal $i * j$

Proofs of correctness using Hoare logic

Program verification

- *Proofs of correctness* use techniques from formal logic to prove that if the starting state (i.e., “input” variables) of a program satisfies particular properties, then the end state after executing a program (i.e., “output” variables) satisfies some other properties.
- The first lot of properties are called *preconditions* (assertions that hold prior to execution of a piece of code), and the second lot are *postconditions* (assertions that hold after execution)

Program verification (2)

- For instance, if our program P is the snippet of code from before –

```
n := 0
while j > 0:
  n := n + i
  j := j - 1
return n
```

– then our input variables are i and j , our output variable is n , and our precondition might be

- i and j are any integers

and our postcondition might be

- n equals (the original value of) i times (the original value of) j

Assertions

- Formally, these assertions are *predicates* that hold of variables. (i.e., they're things that are true or false, given variables as input).
- “is greater than 3” is a predicate.

Example assertions

- Bounds on elements of the data:

$$n \geq 0$$

- Ordering properties of the data:

$$\text{for all } j : 0 \leq j < n - 1 : a_j \leq a_{j+1}$$

- “Finding the maximum”

e.g. Asserting that p is the position of the maximum element in some array $a[0..n - 1]$

$$0 \leq p < n \vee (\text{for all } j : 0 \leq j < n : a_j \leq a_p)$$

Assertions in a program - multiplication

We'll distinguish our input variables from our working variables, for the moment. We'll make a and b input.

```
{ pre :  $\top$  } // no precons -- "top" or "true"  
i := a  
j := b  
n := 0  
while j > 0:  
  n = n + i  
  j = j - 1  
{ post:  $n = a * b$  }
```


Assertions in a program - old values

- The problem here is that we want to refer to the *old* values of i and j in our postcondition
- Systems and languages for program verification often will have a special syntax for this, to avoid having to introduce new variables
- e.g. In the Eiffel language, you can refer to the value of i in the pre-state as just `old i`

Assertions in a program - square root example

```
a := 0
b := n + 1
{ pre :  $n \geq 0$  }
while a + 1 != b:
  d := (a + b) div 2
  if d * d <= n:
    a := d
  else:
    b := d
{ post:  $a^2 \leq n < (a+1)^2$  }
```

- Where we have a sequence [*preconditions*, *code fragment*, *postconditions*], we call this a **Hoare triple** (after logician and computer scientist Tony Hoare of Oxford, who also invented the Quicksort algorithm, amongst other things)

Invariants

Where we have an assertion which should hold before and after *every* statement, we call this an *invariant*

```
a := 0
b := n + 1
{ pre :  $n \geq 0$  }
{ inv P:  $a < b \leq n + 1 \wedge a^2 \leq n < b^2$  }
  // ^ an *invariant*, which *always* holds
while a + 1 != b:
  d := (a + b) div 2
  if d * d <= n:
    a := d
  else:
    b := d
{ post:  $a^2 \leq n < (a+1)^2$  }
```

Termination

- What if we have a loop that doesn't terminate? What can we conclude then?
- There are two ways of handling this:
 - *partial correctness* doesn't require a program to terminate
 - *total correctness* does

Partial correctness

- if we have a triple $\{ \phi \} P \{ \psi \}$, then we say that it is “satisfied under partial correctness” if
 - for all states which satisfy the preconditions, ϕ
 - the state resulting from executing P satisfies the postcondition, ψ , *if P actually terminates.*
- (If the pre- or post-conditions contain multiple assertions, we can always “and” them all together into just one assertion.)

Partial correctness

- So if our precondition is just \top , and our postcondition is $n = i * j$, then here is a program that has partial correctness:

```
while true:
    pass
```

- It *cannot* terminate, so partial correctness is satisfied.

Partial correctness

- So what use is that?
- Well, often it's handy to tackle a proof of correctness in two stages:
 1. Prove that *if* the program terminates, *then* it produces the results we want
 2. Prove that the program terminates
- And then we have *total* correctness.

Total correctness

- if we have a triple $\{ \phi \} P \{ \psi \}$, then we say that it is “satisfied under total correctness” if
 - for all states which satisfy the preconditions, ϕ
 - the state resulting from executing P satisfies the postcondition, ψ , and P is guaranteed to terminate.

Proving things

- So how do we actually build up our proof?
 - We need rules about how to combine triples together.
 - We cover these rules now

Composition

- The composition rule says:

If we have $\{ a \} P_1 \{ b \}$ and $\{ b \} P_2 \{ c \}$

then we can derive $\{ a \} P_1; P_2 \{ c \}$

Assignment

- The assignment rule states that, after an assignment, any predicate that was previously true for the right-hand side of the assignment now holds for the variable on the left-hand side.
- Formally:

$$A[E/x] \quad x := E \quad \{ A \}$$

Here, A represents some (probably and'ed together) assertion about variables;

$A[E/x]$ means, “ A , but wherever x appears, substitute for it the expression E instead”

- Which gives us the meaning we want: whatever *was* true of E is now true of x .

Conditional

- We represent “if” using the conditional rule
- The rule is:

If we have

$$\{ a \wedge C \} P_1 \{ b \}, \text{ and we have } \{ a \wedge \neg C \} P_2 \{ b \}$$

then we can derive

$$\{ a \} \text{ if } C \text{ then } P_1 \text{ else } P_2 \{ b \}$$

- (i.e., If we know that we end up with b holding regardless of whether C is true or not, then we know it must be true after the `if` statement)

“Partial while”

- The rule for partial correctness of while loops is:

If we have a triple $\{ A \wedge B \} P \{ A \}$

then we can derive:

$$\{ A \} \text{ while } B \text{ do } P \text{ done } \{ \neg B \wedge A \}$$

- Here, A is what is called the **loop invariant** – it is something that should be *preserved by* (i.e. is true before and true after) the loop body.
- After the loop is done, A should still hold, no matter how many times the loop executed.
- *if* we got out of the loop, it must be the case afterwards that $\neg B$ (since otherwise, “while B ” would’ve continued)

“Total while”

- If we can prove partial correctness for a while loop; and we can prove the loop terminates; then we've proved total correctness.
 - (There is syntax for this in Hoare logic, but we won't go into it.)

Worked example

- program fragment:

```
{ pre :  $b \geq 0$  }  
i := a  
j := b  
n := 0  
while j > 0:  
  n = n + i  
  j = j - 1  
{ post:  $n = a * b$  }
```

Worked example

- program fragment:

```
{ pre :  $b \geq 0$  }  
i := a  
j := b  
n := 0  
while j > 0:  
  n = n + i  
  j = j - 1  
{ post:  $n = a * b$  }
```

- ... we try to work out a loop invariant ...

Worked example

- program fragment:

```
{ pre :  $b \geq 0$  }  
i := a  
j := b  
n := 0  
while j > 0:  
  n = n + i  
  j = j - 1  
{ post:  $n = a * b$  }
```

- ... we try to work out a loop invariant ...
- invariant: $(j * i) + n = a * b$

Worked example

- let our loop invariant L be $(j * i) + n = a * b$

```

{ pre :  $b \geq 0$  }
i := a
j := b
n := 0
{  $j > 0 \wedge L$  }
while j > 0:
  n = n + i
  j = j - 1
{  $j \leq 0 \wedge L$  }
{ post:  $n = a * b$  }

```

- The “partial while” rule lets us put assertions round the while loop.
- So one thing we know now is that *if* the loop finishes, then L will still hold.
- If j ended up being 0, then that would be handy, because

$$j = 0 \wedge (j * i) + n = a * b$$

implies

$$n = a * b$$

which is what we would like to prove.

Worked example

- So if we can prove that
 - at the end of the loop, $j = 0$,
 - and that the loop terminates,
 - and that the initialization statements and the program preconditions satisfy the loop precondition
- ... then we've proved that the program does what we want.
(Assuming $b \geq 0$.)

Worked example

- We won't work through the proof in detail, but hopefully you have an idea of how it will go.
- Since j is some finite number, and we subtract one off it each time, we can use some math and logic to prove that the loop does terminate, and that j ends up equaling zero.

Worked example

- Coming up with the loop invariant is usually the hard part; the other rules can be applied in a more automatic kind of way.
- Edsger Dijkstra said that to properly understand a while statement is to understand its invariant.

Sorts of formal methods

Back to formal methods

- So here, our *specifications* were assertions about variable values before and after the program executed, written as mathematical formulas.
- We used a method that was largely *manual* – putting assertions around fragments of code.
- Some bits of that could be partly automated – the rules for composition and assignment could be done by machine
- The loop invariant, however, requires ingenuity to come up with
- Our *model* of the system was, in fact, the code itself.

Categorizing formal methods

- We can categorize formal methods in various ways ...

Categorizing formal methods

Degree of formality:

- how formal are the specifications and the system description?
- in natural language (informal), or something more mathematical?

Categorizing formal methods

Degree of automation:

- the extremes are fully automatic and fully manual
- most computer-aided methods are somewhere in the middle

Categorizing formal methods

Full or partial verification of properties in the specification

- What is being verified about the system? Just one property? (e.g., that it does not deadlock, say – common for concurrent systems)
- Or many/all properties?
 - (This is usually very expensive, in terms of effort)

Categorizing formal methods

Intended domain of application:

- e.g. hardware vs software;
- reactive vs terminating;
 - reactive systems run a theoretically endless “loop” and aren't intended to terminate – they just keep *reacting* to an environment
 - e.g. operating systems, embedded hardware (modelled with state machines, often)
 - terminating systems terminate, usually with some sort of *result*
- sequential vs concurrent

Categorizing formal methods

pre- vs post-development:

- Is verification done early in development, vs later or afterwards?
- Earlier is obviously better, since things are much more expensive to fix if early, if it turns out our system *doesn't* meet the specs

Categorizing formal methods

- But sometimes the system comes first, then the verification
- Often true for programming languages . . .
 - e.g. Java was released in 1995, and in 1997, a machine-checked proof of “type soundness” of a subset of Java was proved.¹
 - But: later versions of Java (from 5 onwards) turned out to have *unsound* type systems in various ways. Oops.
 - The interaction of sub-typing and inheritance turned out to make the early OO language Eiffel unsound. Also oops.²

¹Syme. “Proving Java Type Soundness”. 1997 [[pdf](#)]

²William R. Cook. A proposal for making Eiffel type-safe. The Computer Journal, 32(4):305–311, August 1989.

Categorizing formal methods

Are we trying to prove properties of an individual program? Or about *all* programs written in a particular language?

- An example of the first one is proving that a sorting function does what we want it to, or that a compiler implementation obeys some particular formal specification
- An example of the latter is proving results about the *type system* for a language, which lets us show that *all* programs in the language will have some sort of guarantees of good behaviour
 - e.g. Proving that well-typed Java programs cannot be subverted (assuming the JVM and compiler are implemented correctly) – it should be impossible to get a reference which doesn't point to a valid area of memory, for instance.

Aside – type systems

- We often don't think of type systems as being a “formal method”, but some type systems are very expressive, and allow us to prove quite strong results about our programs
- We can use them to prove that (for instance) unsanitized user data never gets output to a web page

Type systems

- An example of a type system many of us will have used in high school – consistency of SI units

Categorizing formal methods

Model-based vs proof-based approaches:

- We've seen one example of a *proof* based approach, Hoare logic.
 - Your specification is some formula in some suitable logic
 - In Hoare logic, our specification is what we want the program to *do* – it's expressed as assertions (postconditions which should hold after the program executes, if the preconditions held)
 - You try and *prove* that the system (or some abstraction of it) satisfies the specification.
- Usually requires guidance and expertise from the user

Categorizing formal methods

Model-based approaches:

- Again, our specification is some sort of formula
- This time, our system description is some mathematical structure, a **model**, \mathcal{M}
- We check whether the model \mathcal{M} **satisfies** the specification (i.e. has the properties we want)
- In many cases, this can be done automatically.
- We will see examples of model-based approaches next lecture.