

CSE503:  
SOFTWARE ENGINEERING  
ELECTRIFYING SPECIFICATIONS

David Notkin  
Spring 2011

## Formal methods

- We've covered some of the basics, although there is much more out there – we'll get to some of that
- One thing is clear so far: programs *execute* – they are kinetic – while formalisms just sit there – they represent potential
- This property makes many formalisms less attractive to many people, as the benefits are harder to see

“It is easier to change the specification to fit the program than vice versa.” –Perlis

503 11sp © UW CSE • D. Notkin

## Electrifying formalisms

- Daniel Jackson (and others) have worked on addressing this concern by “electrifying” formalisms – that is, making them “executable” in some sense, or at least providing useful feedback to a developer quickly
- Alloy is Jackson’s core approach to this, but it’s not the only one out there

503 11sp © UW CSE • D. Notkin

## Executable specifications

- One way to electrify a formalism is to execute it – many formalisms represent high-level programs
  - Google Scholar found ~95K entries to “executable specifications”
  - Many such executable specifications look a lot like (various kinds of) logic programs or functional programs; much of this work is related to *automatic programming*
- The execution gives insight into what the specification means
- Performance of these “programs” is usually poor
- And automated refinement techniques to evolve from an executable specification to an efficient program seem to be limited
  - This work goes back to at least 1976 with Darlington, Burstall, Manna and others

503 11sp © UW CSE • D. Notkin

## Type checking

- Type checking is a good example of an electrified formalism – based on the type system and proof rules, the compiler reports immediately on situations where the program may be misusing typed values
- This is also a good example of an alternative style of electrification
- Rather than executing to “see what it does,” it rather compares two different views of the computation – what the types constrain and whether the program satisfies those constraints

503 11sp © UW CSE • D. Notkin

## Comparing

- I believe that this comparative approach is extremely powerful – more powerful than the “let’s make higher-level specifications that we can execute” in the long-run
- In some sense, that was the idea of proving desirable properties from axioms – the axioms described relationships of an ADT, and the posited properties were alternative views of what should be true for that ADT
  - But this wasn’t electrified in any sense

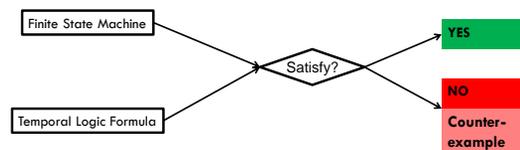
503 11sp © UW CSE • D. Notkin

## Model checking

- Model checking is one of the bases for electrifying comparisons of program views
- We’ll look at this first from the high-level notion of model checking, then look at Alloy as an instance of electrified (bounded) model checking, and then look back at more conventional model checking

503 11sp © UW CSE • D. Notkin

## Model checking



- What can the finite state machines and temporal logic formulae represent?
- What does “satisfy” mean? How does “satisfy” work?
- Why should we care?
- What is a counterexample?
- What does “Yes” actually mean?

503 11sp © UW CSE • D. Notkin

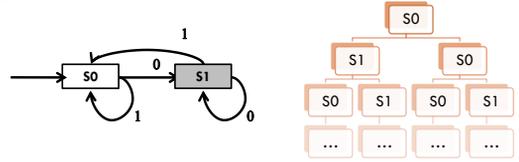
## ACM 2007 Turing Award Citation

9

In 1981, Edmund M. Clarke and E. Allen Emerson, working in the USA, and Joseph Sifakis working independently in France, authored seminal papers that founded what has become the highly successful field of Model Checking. This verification technology provides an algorithmic means of determining whether an abstract model--representing, for example, a hardware or software design--satisfies a formal specification expressed as a temporal logic formula. Moreover, if the property does not hold, the method identifies a counterexample execution that shows the source of the problem. The progression of Model Checking to the point where it can be successfully used for complex systems has required the development of sophisticated means of coping with what is known as the state explosion problem. Great strides have been made on this problem over the past 27 years by what is now a very large international research community. As a result many major hardware and software companies are now using Model Checking in practice. Examples of its use include the verification of VLSI circuits, communication protocols, software device drivers, real-time embedded systems, and security algorithms. ...

503 11sp © UW CSE • D. Notkin

## A computation tree over an FSM



- Represent all possible FSM paths with a computation tree – even when infinite, the structure is constrained because of finite states
  - The computation tree is generated from the state machine
  - The temporal logic formula queries the computation tree
- Model checking answers questions about this tree structure – kinds of queries include
  - Does every accepting input include a 0? A 1?
  - Does any accepting input include a 0? A 1?
  - Does every accepting input that has a 1 have a 1 in the remaining input?
- More generally, safety and liveness properties of many kinds

503 11sp © UW CSE • D. Notkin

## FAQ

11

- What can the finite state machines and temporal logic formulae represent?
  - See the Turing Award citation (and more...)
- What does "satisfy" mean? How does "satisfy" work?
  - Satisfy means that the temporal logic formula is guaranteed to hold over the computation tree defined by the FSM.
- Why should we care?
  - Guarantees can be a good thing
- What is a counterexample?
  - A path through the computation tree that contradicts the temporal logic formula
  - There is a mismatch between the two descriptions – but one cannot tell which is "wrong" without further work
- What does "Yes" actually mean?
  - It is a guarantee that the property holds, but it provides no guarantee that the property or the FSM are what the developers think they are

503 11sp © UW CSE • D. Notkin

## Two Approaches to Model Checking

12

- Explicit – represent all states
  - Use conventional state-space search
  - Reduce state space by folding equivalent states together
- Symbolic – represent sets of states using boolean formulae
  - Reduce huge state spaces by considering large sets of states simultaneously – to the first order, this is the meeting of BDDs (binary decision diagrams) and model checking (more later)
  - Convert state machines, logic formulae, etc. to boolean representations
  - Perform state space exploration using boolean operators to perform set operations
  - SAT solvers are often at the base of symbolic model checking

503 11sp © UW CSE • D. Notkin

## Bounded model checking

13

- Restricting to finite abstractions work in many situations – protocols, etc.
- But programs themselves are not generally restricted to finite abstractions – so how could model checking apply?
- Simply, one can constrain the search space: exhaustively search within a finite space with the expectation (hope?) that the finite space is a good approximation to the infinite space
- D. Jackson calls this the “small scope hypothesis”

503 11sp © UW CSE • D. Notkin

## Basic claim and idea

14

- Most counterexamples can be identified with a failing trace that is short, by some definition
  - The number of operations
  - The complexity of the input structures (e.g., how many nodes)
  - Size of a device
  - Length of a pathname
  - ...
- If no counterexample is returned when checking the truncated space, one builds confidence in the model but cannot be certain of the properties
- Eliminating counterexamples likely helps through the infinite space

503 11sp © UW CSE • D. Notkin

## A little pre-Alloy history

15

- Model-based descriptions
- The Z specification language

503 11sp © UW CSE • D. Notkin

## Model-oriented

16

- Model a system by describing its state together with operations over that state
  - An operation is a function that maps a value of the state together with values of parameters to the operation onto a new state value
- A model-oriented language typically describes mathematical objects (e.g. data structures or functions) that are structurally similar to the required computer software

503 11sp © UW CSE • D. Notkin

## Z ("zed")

- Perhaps the most widely known and used model-based specification language
- Good for describing state-based abstract descriptions roughly in the abstract data type style
- Based on typed set theory and predicate logic
- A few commercial and honorary successes
  - I'll come back to one reengineering story afterwards

503 11sp © UW CSE • D. Notkin

## Basics

- Static schemas
  - States a system can occupy
  - Invariants that must be maintained in every system state
- Dynamic schemas
  - Operations that are permitted
  - Relationship between inputs and outputs of those operations
  - Changes of states

503 11sp © UW CSE • D. Notkin

## The classic example

- A "birthday book" that tracks people's birthdays and can issue reminders of those birthdays
  - There are tons of web-based versions of these now
- There are two basic types of atomic elements in this example
  - [NAME,DATE]
  - An inherent degree of abstraction: nothing about formats, possible values, etc.

503 11sp © UW CSE • D. Notkin

## Taken directly from Spivey's ZRM

```

BirthdayBook
known : P NAME
birthday : NAME ↔ DATE
known = dom birthday

known = { John, Mike, Susan }
birthday = { John ↔ 25 Mar,
            Mike ↔ 20 Dec,
            Susan ↔ 20 Dec }.

AddBirthday
Δ BirthdayBook
name? : NAME
date? : DATE
name? ≠ known
birthday = birthday ∪ { name? ↔ date? }

known' = known ∪ { name? }.
In fact we can prove this from the specification of AddBirthday, using the invariants on the state before and after the operation:
known' = dom birthday'
= dom (birthday ∪ { name? ↔ date? })
= dom birthday ∪ dom { name? ↔ date? }
= dom birthday ∪ { name? }
= known ∪ { name? }
            
```

← static schema

← legal system state

← dynamic schema

← Proof of expected property

503 11sp © UW CSE • D. Notkin

## Some more schemas

21

```

FindBirthday
  ⊆ BirthdayBook
  name? : NAME
  date! : DATE
  name? ∈ known
  date! = birthday(name?)
        
```

← schema w/ unchanged inputs

```

Remind
  ⊆ BirthdayBook
  today? : DATE
  card! : P NAME
  card! = [ n : known | birthday(n) = today? ]
        
```

← schema w/ unchanged inputs

```

InitBirthdayBook
  BirthdayBook
  known = ∅
        
```

← initializing schema

503 11sp © UW CSE • D. Notkin

## What about errors?

22

- Entering a birthday for someone already in the book?
- Looking up a birthday for an unknown person?
- What is the meaning of the specification in the face of errors like these – unsatisfied pre-conditions, in most situations?
- Rewrite the whole specification?

503 11sp © UW CSE • D. Notkin

## Or use the schema calculus!

23

```

REPORT ::= ok | already_known | not_known.
        
```

← add an enumerated type

```

Success
  result! : REPORT
  result! = ok
        
```

← schema

AddBirthday ∧ Success.

```

AlreadyKnown
  ⊆ BirthdayBook
  name? : NAME
  result! : REPORT
  name? ∈ known
  result! = already_known
        
```

← schema

RAddBirthday ⊕ (AddBirthday ∧ Success) ∨ AlreadyKnown.

← schema calculus

503 11sp © UW CSE • D. Notkin

## Schema calculus: sweet!

24

- The schema calculus allows us to combine specifications using logical operators (e.g.,  $\wedge$ ,  $\vee$ ,  $\Rightarrow$ ,  $\neg$ )
  - This allows us to define the common and error cases separately, for example, and then just  $\wedge$ -ing them together
- In some sense, it allows us to get a cleaner, smaller specification

503 11sp © UW CSE • D. Notkin

## But don't try this on programs!

25

- Wouldn't it be fantastic if we had the equivalent of the schema calculus on programs?
  - Write your error cases separately and then just  $\wedge$  them together
  - Write a text editor and a spell checker and integrate them by  $\wedge$ -ing them together
  - So you want to build a program that doesn't blow up a nuclear power plant?
    - Just build one that does, and then negate it ☺!
- Programs are not logic
  - Some classes of programming languages – largely logic and functional languages – come closer than imperative and OO languages, but they are not logic

503 11sp © UW CSE • D. Notkin

## Points about Z

26

- This isn't proving correctness between a specification and a program
  - There isn't a program!
- Even the specification without the implementation has (some) value
  - The most obvious example is when a theorem is posited and then is proven from the rest of the specification – albeit by hand
- The actual notation seems more effective than some others
  - Z is intended to be in bite-sized chunks (schema), interspersed with natural language explanations

503 11sp © UW CSE • D. Notkin

## Z used to improve CICS/ESA\_V3.1

27

- A broadly used IBM transaction processing system
- Integrated into IBM's existing and well-established development process
- Many measurements of the process indicated that they were able to reduce their costs for the development by almost five and a half million dollars
- Early results from customers also indicated significantly fewer problems, and those that have been detected are less severe than would be expected otherwise

503 11sp © UW CSE • D. Notkin

## Queen's Award for Technological Achievement

28

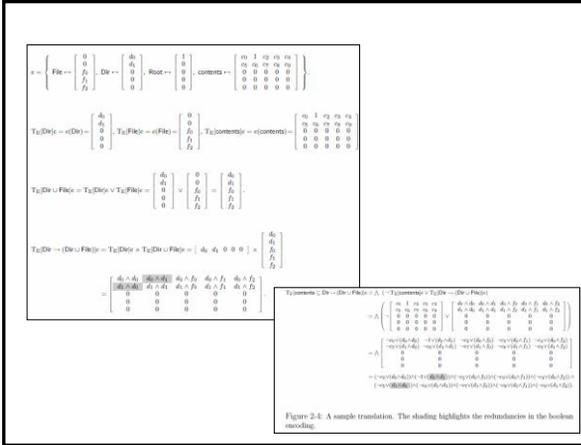
- "Her Majesty the Queen has been graciously pleased to approve the Prime Minister's recommendation that The Queen's Award for Technological Achievement should be conferred this year upon Oxford University Computing Laboratory.
- "Oxford University Computing Laboratory gains the Award jointly with IBM United Kingdom Laboratories Limited for the development of a programming method based on elementary set theory and logic known as the Z notation, and its application in the IBM Customer Information Control System (CICS) product. ...
- "The use of Z reduced development costs significantly and improved reliability and quality. Precision is achieved by basing the notation on mathematics, abstraction through data refinement, re-use through modularity and accuracy through the techniques of proof and derivation.
- "CICS is used worldwide by banks, insurance companies, finance houses and airlines etc. who rely on the integrity of the system for their day-to-day business."



Sir Charles Antony  
Richard Hoare

503 11sp © UW CSE • D. Notkin





## Facts vs. assertions

- Alloy *facts* define truths of the model
- Alloy *assertions* define truths you'd like check to make sure that they hold about the model
- Closely related to Jackson père's discussion of moods
  - Indicative mood: describes how things in the world are regardless of the behavior of the system: "Each seat is located in one and only one theater."
  - Optative mood: describes what you want the system to achieve: "Better seats should be allocated before worse seats at the same price."

503 11sp © UW CSE • D. Notkin

## And to designations and definitions

Also Jackson père

- *Designations* are atomic phenomena
  - e.g., genetic mother
- *Definitions* define terms using designations and other definitions
  - e.g., genetic child of
- Refutable descriptions can in principle be disproven
  - $\forall m, x \bullet \text{Mother}(m, x) \rightarrow \neg \text{Mother}(x, m)$
  - Can't do this with definitions
- The precise distinction between these is, to me, far more important than the formalization

503 11sp © UW CSE • D. Notkin

## Alloy ... a brief demo

503 11sp © UW CSE • D. Notkin