

Hoare logic and Model checking

Part II: Model checking

Lecture 7: Introduction to model checking

Jean Pichon-Pharabod

University of Cambridge

CST Part II – 2019/20

Warning about terminology and notation

The field of model checking is a terminological and notational nightmare.

In violation of the Countryside Code, we leave the field in a state worse than the one we found it in.

We also write meta-level and object-level constructors with the same symbols when not ambiguous.

Acknowledgements

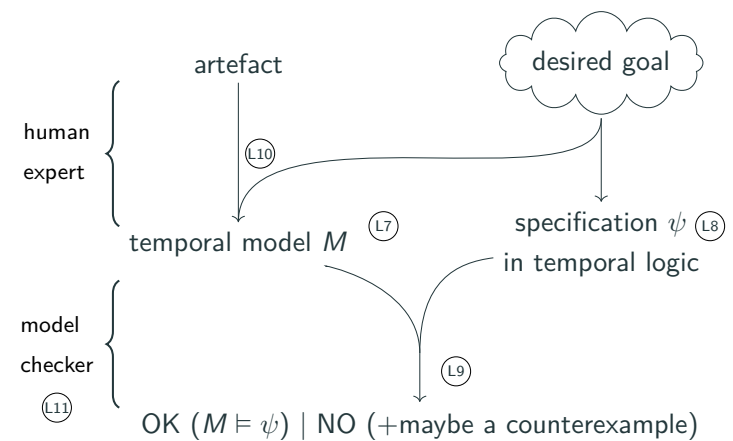
These slides are heavily based on previous versions by Mike Gordon, Dominic Mulligan, Alan Mycroft, and Conrad Watt.

They are also inspired by slides and lecture notes by John Gallagher, Gourinath Banda, and Pierre Ganty, by Paul Gastin, by Orna Grumberg, by Arie Gurfinkel, by Daniel Kroening, by Antoine Miné, by Julien Schmalz, by David A. Schmidt, and by Carsten Sinz and Tomáš Balyo.

Thanks to Neel Krishnaswami, David Kaloper Meršinjak, and Ben Simner for remarks and reporting mistakes.

1

Model checking



This diagram gives a very static, top-down picture, but it is the feedback that provides the value.

2

3

Example


Suppose we are given an algorithm that is supposed to transfer, from one bank of the Cam to the other, using only a punt with seat for one, a wolf, a goat, and a cabbage¹.

The success criteria are

- safety: the cabbage and the goat, and the wolf and the goat, cannot be left alone on a bank;
- liveness: all three items are moved to the other bank.


¹it is a large cabbage, so it takes up the whole seat

How to find good models



A pretty good model of the solar system

Encyclopédie, Diderot, d'Alembert, *et al.*



The need to go beyond excessively simple models

Flammarion engraving, anonymous

“All models are wrong, some are useful” applies. The designer must ensure the model captures the significant aspects of the real system. Achieving it is a special skill, the acquisition of which requires thoughtful practice

— How Amazon Web Services Uses Formal Methods

4

Example

How to model the problem?

- Option 1:

$$\mathcal{L} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + i\bar{\psi}\not{D}\psi + \text{h.c.} + \bar{\psi}_i y_{ij} \psi_j \phi + \text{h.c.} + |D_\mu \phi|^2 - V(\phi) + ???$$

- Option 2: (G. Doré, anonymous (Wellcome coll.), G. Waddington)



- Option 3: (apologies to the Phaistos cat)

$$\text{Side} ::= \curvearrowright \mid \curvearrowleft \quad \text{Item} ::= \text{flower} \mid \text{cat} \mid \text{goat} \mid \text{cabbage}$$

$$\text{State} \stackrel{\text{def}}{=} \text{Item} \rightarrow \text{Side}$$

...

5

Temporal models

A **temporal model** over **atomic propositions** AP is a left-total transition system where states are labelled with some of AP , and where some states are distinguished as initial:

$$M, \dots \in \text{TModel} \stackrel{\text{def}}{=} (S \in \text{Set}) \times \text{states}$$

$$(S_0 \in \text{sub } S) \times \text{initial states}$$

$$(\textcircled{1} T \textcircled{2} \in \text{relation } S S) \times \text{transition}$$

$$(\ell \in S \rightarrow \text{sub } AP) \times \text{state labelling}$$

$$(\forall s \in S. \exists s' \in S. s T s') \text{ left-total}$$

Elements of S are denoted s .

⚠ We interpret not being labelled as a lack of information, not as a negation.

6

7

Corner case: the initial temporal model

$$\mathbb{0} \in \text{TModel}$$


$$\mathbb{0} \stackrel{\text{def}}{=} \left\langle \begin{array}{l} \emptyset, \\ \emptyset, \\ \emptyset, \\ s \mapsto \emptyset, \\ \dots \end{array} \right\rangle \quad (\text{it is empty})$$

12

Corner case: the terminal temporal model

$$\mathbb{1} \in \text{TModel}$$

$$\mathbb{1} \stackrel{\text{def}}{=} \left\langle \begin{array}{l} AP \rightarrow \mathbb{B}, \\ \{s \mid \top\}, \\ \{s_0, s_1 \mid \top\}, \\ s \mapsto \{p \mid s p\}, \\ \dots \end{array} \right\rangle$$



13

Temporal model of a terrible punter

A punter with no concern for goat welfare or cabbage welfare:

$$AP = \text{State}$$

$$M = \left\langle \begin{array}{l} \text{State}, \\ \{s \mid \forall i. s i = \cdot\}, \\ \left\{ s, s' \mid \left(s \overset{\text{flip}}{\curvearrowright} = \text{flip}(s' \overset{\text{flip}}{\curvearrowright}) \wedge \right. \right. \\ \left. \left. (\text{moveone } s s' \vee \text{movezero } s s') \right) \right\}, \\ s \mapsto \{i \mid s = i\}, \\ \dots \end{array} \right\rangle$$

14

Temporal model of a terrible punter

$$\text{flip} \cdot \stackrel{\text{def}}{=} \cdot \text{flip} \cdot \stackrel{\text{def}}{=} \cdot$$


$$\text{moveone } s s' \stackrel{\text{def}}{=} \left(\begin{array}{l} \text{aux } s s' \text{ 🌸 🐐 🐐 } \vee \\ \text{aux } s s' \text{ 🐐 🌸 🐐 } \vee \\ \text{aux } s s' \text{ 🐐 🌸 🐐 } \end{array} \right)$$

$$\text{aux } s s' a b c \stackrel{\text{def}}{=} \left(\begin{array}{l} s a = s \overset{\text{flip}}{\curvearrowright} \wedge s' a = s' \overset{\text{flip}}{\curvearrowright} \wedge \\ s' b = s b \wedge s' c = s c \end{array} \right)$$

$$\text{movezero } s s' \stackrel{\text{def}}{=} s \text{ 🌸 } = s' \text{ 🌸 } \wedge s \text{ 🐐 } = s' \text{ 🐐 } \wedge s \text{ 🐐 } = s' \text{ 🐐 }$$

15

Temporal model of a terrible punter




Safety: we never go through a red state.

Liveness: we eventually reach the blue state.

Both are pretty clearly false! :-)

16

Informal temporal model of an elevator



Let us try to describe how an elevator for a building with 3 levels works:

- it starts at the ground floor, with the door closed, and goes back there when it is not called;
- if going through a level where it is called, it stops there and opens its door;
- ...

Textual descriptions do not scale very well.

17

Temporal model of an elevator: statics and specification

Direction ::= home | up | down

Level ::= 0 | 1 | 2

Location ::= 0 | 1/2 | 1 | 3/2 | 2

Called $\stackrel{def}{=} Level \rightarrow \mathbb{B}$

DoorStatus ::= open | closed

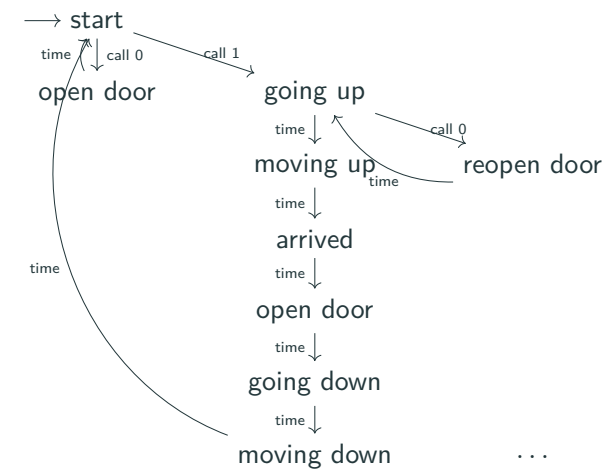
ElevatorStatus $\stackrel{def}{=} Direction \times Location \times Called \times DoorStatus$

Desired goals:

- the door is not open at half-levels;
- if the elevator is called to a level, then it eventually gets there;
- the elevator does not lock people in;
- the path of the elevator is not entirely idiotic.


18

Temporal model of an elevator: partial dynamics



19

Temporal model of an elevator: complete (?) dynamics



How to have any confidence that this is correct? house by Petr Olsšák

20

Definitions

(Infinite) Paths

$$\text{stream} \in \text{Set} \rightarrow \text{Set}$$

$$\pi, \dots \in \text{stream } A \stackrel{\text{def}}{=} \mathbb{N} \rightarrow A$$

$$\text{IsPath} \in (M \in \text{TModel}) \rightarrow \text{stream } M.S \rightarrow \text{Prop}$$

$$\text{IsPath } M \pi \stackrel{\text{def}}{=} \forall n \in \mathbb{N}. (\pi \ n) \ M.T \ (\pi \ (n+1))$$

$$\text{Path} \in \text{TModel} \rightarrow \text{Set}$$

$$\text{Path } M \stackrel{\text{def}}{=} (\pi \in \text{stream } M.S) \times \text{IsPath } M \pi$$

Reachable states & the tail operation

Because the transition relation is left-total, these infinite paths are “complete”, in the sense that they capture reachability:

$$\text{Reachable} \in (M \in \text{TModel}) \rightarrow M.S \rightarrow \text{Prop}$$

$$\text{Reachable } M \ s \stackrel{\text{def}}{=} \exists \pi \in \text{stream } M.S, n \in \mathbb{N}.$$

$$\text{IsPath } M \ \pi \wedge M.S_0 \ (\pi \ 0) \wedge s = \pi \ n$$

$$\text{tailn} \in (A \in \text{Set}) \rightarrow \mathbb{N} \rightarrow \text{stream } A \rightarrow \text{stream } A$$

$$\text{tailn } A \ n \ \pi \stackrel{\text{def}}{=} i \mapsto \pi \ (i+n)$$

21

22

Stuttering

A temporal model is **stuttering** when all states loop back to themselves:

$$\begin{aligned} \text{stuttering} &\in \text{TModel} \rightarrow \text{Prop} \\ \text{stuttering } M &\stackrel{\text{def}}{=} \forall s \in M.S. s M.T s \end{aligned}$$

⚠ If the temporal model is not stuttering, then we can count transitions. This is only sound if they exactly match those of the system being analysed.


See “What good is temporal logic” §2.3, by Leslie Lamport
<https://lamport.azurewebsites.net/pubs/what-good.pdf>

23

Applications of model checking

More abstract temporal model of traffic lights

Somewhat unusually, we do not interpret a state not being labelled with a given label as that state being labelled with the negation of that label.



24

Applications of model checking

- Hardware:
 - circuits (with memory) directly translate to temporal models
 - lots of protocols
 - cache protocols
 - bus protocols
 - ...their specification involves lots of temporal “liveness” (“eventually something good”) properties
- Software: often not finite a priori, but “proper modelling”, or bounded model-checking
- Security protocols
- Distributed systems
- ...

The common denominator of many of these is the “killer app” of model checking: concurrency.

25

Examples

In the rest of this lecture, we will sketch how some of these are approached.

The point is not the details of any individual temporal model, but the overall approach.

26

Temporal model from operational semantics

An initial configuration for a small-step operational semantics naturally leads to a temporal model: take

- configurations as states,
- the initial configuration as the (only) initial state,
- steps as transitions, and
- some interesting properties as atomic propositions, for example


$$\begin{aligned}
 X, Y, Z, \dots &\in \text{Var} \\
 v &\in \mathbb{Z} \\
 AP &::= X \dot{=} v \mid X \dot{=} Y \mid X \dot{<} Y \mid \\
 &\quad X \dot{+} Y \dot{<} Z \mid X \dot{\times} Y \dot{<} Z \mid \\
 &\quad \dots
 \end{aligned}$$

27

Temporal model from operational semantics

Temporal model from operational semantics

For example, for a language with a concurrent composition with interleaving dynamics (as in lecture 6):



28

Dealing with the size of temporal model from operational semantics

These temporal models are very often infinite or intractably large!


Many approaches:

- bounded model checking (see lecture 9):
 - assume (and possibly check whether) loops execute no more than n times
 - consider executions of length smaller than n
 - ...
- use a model checking DSL to write an idealised version of the program (see lecture 9)
- use abstraction (see lecture 10)

29

Example circuit

Synchronous (the clock is left implicit) counter that goes 0, 1, 2, 0, 1, 2, ... (assuming all registers are initially 0):




Registers make the circuit not be a simple function, which motivates using a temporal model.

30

Temporal model from circuits

Example circuit temporal model

The states of the temporal model are the state of the registers, and the labels are which registers are set to 1:




Safety: The state $\{r_0, r_1\}$ should never be reached.

Liveness: all other states should be visited infinitely often.

31

Difference circuit

Given two circuits $C_1, C_2 \in \text{SCircuit } i \ 1$, we can define their difference circuit $C_1 \ominus C_2$:



If the answer is always 0, then they are equivalent (see lecture 10). The typical use case is to have a simple, clearly correct C_1 , and a complex C_2 to verify.

32

Temporal models of distributed algorithms

Temporal models of distributed algorithms

110 6. DISTRIBUTED CONSENSUS WITH PROCESS FAILURES

other processes; this can help to make the algorithm descriptions more uniform. These messages are technically not permitted in the model, but there is no harm in allowing them because the fictional transmissions could just be simulated by local computation.

BigStep algorithm

For every string x that occurs as a label of a node of T , each process has a variable $val(x)$. Variable $val(x)$ is used to hold the value with which the process decorates the node labelled x . Initially, each process i decorates the root of its tree with its own initial value, that is, it sets its $val(\lambda)$ to its initial value.

Round 1: Process i broadcasts $val(\lambda)$ to all processes, including i itself. Then process i records the incoming information:

1. If a message with value $v \in V$ arrives at i from j , then i sets its $val(j)$ to v .
2. If no message with a value in V arrives at i from j , then i sets $val(j)$ to $null$.

Round k , $2 \leq k \leq f+1$: Process i broadcasts all pairs (x, v) , where x is a level $k-1$ label in T that does not contain index i , $v \in V$, and $v = val(x)$.¹ Then process i records the incoming information:

1. If xj is a level k node label in T , where x is a string of process indices and j is a single index, and a message saying that $val(x) = v \in V$ arrives at i from j , then i sets $val(xj)$ to v .
2. If xj is a level k node label and no message with a value in V for $val(x)$ arrives at i from j , then i sets $val(xj)$ to $null$.

At the end of $f+1$ rounds, process i applies a decision rule. Namely, let W be the set of non- $null$ $vals$ that decorate nodes of i 's tree. If W is a singleton set, then i decides on the unique element of W ; otherwise, i decides on \perp .

¹It should not be hard to see that the trees get decorated with the values we indicated earlier. That is, process i 's root gets decorated with i 's input value. Also, if process i 's node labelled by the string i_1, \dots, i_k , $1 \leq k \leq f+1$, is decorated by a value $v \in V$, then it must be that v_i has told i at round k that v_{i_1, \dots, i_k} has told

²In order to fit our formal model, in which only one message can be sent from i to each other process at each round, all the messages with the same destination are packaged together into one large message.

Nodes in distributed algorithms are often specified in terms of interacting automata; the temporal model directly results from their interaction.

See IB Concurrent and Distributed Systems

Distributed Algorithms, by Nancy Lynch.

Models of cache algorithms

33

Summary

Temporal models make it possible to describe systems that evolve in time.

Temporal models can be extracted directly, for example from circuits, or hand-crafted to capture the relevant parts of an artefact.

In the next lecture, we will see how to use temporal logic(s) to specify the behaviour of temporal models.