

# LTL Model Checking

Wishnu Prasetya

[wishnu@cs.uu.nl](mailto:wishnu@cs.uu.nl)

[www.cs.uu.nl/docs/vakken/pv](http://www.cs.uu.nl/docs/vakken/pv)

# Overview

- This pack :
  - Abstract model of programs
  - Temporal properties
  - Verification (via model checking) algorithm
  - Concurrency

# Abstract Model

- Temporarily back off from concrete SPIN level, and look instead at a more abstract view on the problem.
- Model a “program” as a finite automaton.
- More interested in “run-time properties” (as opposed to e.g. pre/post conditions):
  - Whenever R receives, the value it receives is never 0.
  - If a holds, then eventually b

# To keep in mind: the term “model” is heavily overloaded...

*(but generally, a model simply means a simplified version of a real thing)*

Automaton for explaining SPIN


Automaton SPIN constructs during verification

SPIN's Promela model

UML model

Real Program

# Finite State Automaton



Many variations, depending on modeling purposes:


- Single or multiple init states
- With or without accepting states
- With or without label on arrows, or on states
- How it executes

**Execution** : a path through the automaton, starting with an initial state.

# State?

- Concrete state of a program → too verbose.
- More abstract → the values of program's variables
  - How SPIN works
- Even more abstractly, through a fixed set of propositions
  - Define your set of propositions
  - Specify which propositions hold on which states
  - How I will explain SPIN

# Example



*Convention: this implies that  $x > 0$  does **not** hold in state 0.*

With *Prop* = { isOdd  $x$ ,  $x > 0$  }

$$V(0) = \{ \text{isOdd } x \}$$

$$V(1) = \{ \text{isOdd } x, x > 0 \}$$

Abstract: we can't say anything about properties which were not taken in *Prop*.

# Kripke Structure

- A finite automaton (  $S, s_0, R, Prop, V$  )
  - $S$  : the set of possible states, with  $s_0$  the initial state.
  - $R$  : the arrows
    - $R(s)$  gives the set of possible next-state from  $s$
    - non-deterministic
  - $Prop$  : set of *atomic* propositions
  - $V$  : labeling function

$a \in V(s)$  means  $a$  holds in  $s$ , else it does *not* hold.



# Prop

- It consists of *atomic* propositions.
- We'll require them to be non-contradictive. That is, for any subset Q of Prop :

$$(\bigwedge Q) \wedge (\bigwedge \{ \neg p \mid p \notin Q \})$$


is satisfiable. Else you may get inconsistent labeling.

- Example:
  - $Prop = \{ x > 0, y > 0 \}$  is ok.
  - $Prop = \{ x > 0, x > 1 \}$  is not ok. E.g. the subset  $\{ x > 1 \}$  is inconsistent.

# Modelling execution

- Recall: an *execution* is a path through your automaton.
- Limit to infinite executions  $\rightarrow$  to simplify discussion.
- This induces what we will call an **'abstract' execution**, which is a *sequence of set of propositions* that hold along that path.
- Overloading the term "execution"...

# Example



$$Prop = \{ \text{isOdd } x, x > 0 \}$$

$$V(0) = \{ \text{isOdd } x \}$$

$$V(1) = \{ \text{isOdd } x, x > 0 \}$$

Consider execution: 0, 0, 1, 1, ...

It induces abs-exec:

$\{ \text{isOdd } x \}$  ,  $\{ \text{isOdd } x \}$ ,  $\{ \text{isOdd } x, x > 0 \}$ ,  $\{ \text{isOdd } x, x > 0 \}$  , ...

Note that it is a sequence over  $\text{power}(Prop)$ .

# Properties


- Recall that we want to express “run-time” properties → we’ll use “temporal properties” from Linear Temporal Logic (LTL)
- Originally designed by philosophers to study the way that time is used in natural language arguments 😊

Based on a number of operators to express relation over time: “next”, “always”, “eventually”


- Brought to Computer Science by Pnueli, 1977.

# Informal meaning


$\square f$  // always  $f$



$\times f$  // next  $f$



$f U g$  //  $f$  holds  
until  $g$



# Example

```
active proctype S () {  
  do  
  :: { c!x ;  
      passed: x++ }  
  od  
}
```

```
active proctype R () {  
  do  
  :: c?y  
  od  
}
```

- [] ( $y == x \parallel y == x - 1$ )
- [] ( $\text{true } \mathbf{U} \text{ } S @ \text{passed}$  )
- Could have been expressed as a Monitor process and progress-labels.
- But e.g. []( $a \mathbf{U} q$ ) cannot be expressed with progress labels.

# Quite expressive! For example...

- $\square (p \rightarrow (\text{true} \cup q))$  // whenever p holds, eventually q will hold
- $p \cup (q \cup r)$
- $\text{true} \cup \square p$  // eventually stabilizing to p

# Let's do this more formally...

- Syntax:

$\varphi ::= p$  // atomic proposition from *Prop*

|  $\neg\varphi$  |  $\varphi \wedge \psi$  |  $X\varphi$  |  $\varphi U \psi$

- Derived operators:

- $\varphi \vee \psi = \neg(\neg\varphi \wedge \neg\psi)$

- $\varphi \rightarrow \psi = \neg\varphi \vee \psi$

- $[], \langle \rangle, W, \dots$

- Interpreted over (abstract) executions.



# Defining the meaning of temporal formulas

- First we'll define the meaning wrt to a single abstract execution. Let  $\Pi$  be such an execution:
  - $\Pi, i \models \varphi$
  - $\Pi \models \varphi = \Pi, 0 \models \varphi$
- If  $P$  is a Kripke structure,

$P \models \varphi$  means that  $\varphi$  holds on all abs. executions of  $P$

# Meaning

- Let  $\Pi$  be an (abstract) execution.

- $\Pi, i \models p \quad = \quad p \in \Pi(i) \quad // \quad p \in Prop$

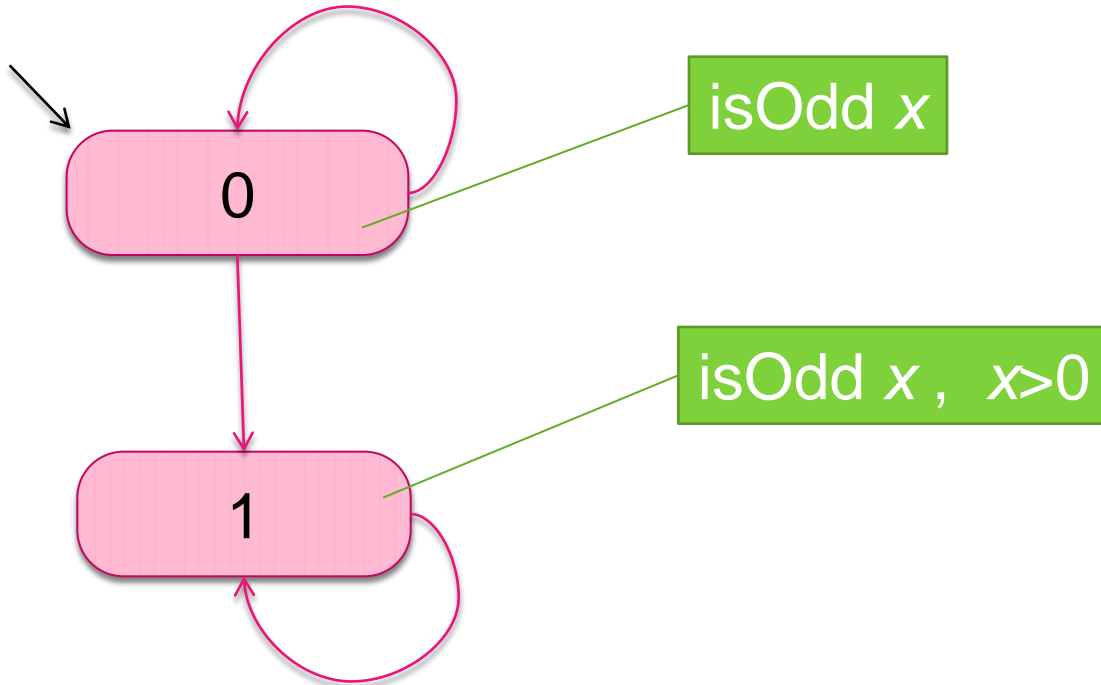
- $\Pi, i \models \neg\varphi \quad = \quad \text{not } (\Pi, i \models \varphi)$

- $\Pi, i \models \varphi \wedge \psi \quad = \quad \Pi, i \models \varphi$   
and  
 $\Pi, i \models \psi$

# Meaning

- $\Pi, i \models \neg \varphi = \Pi, i+1 \models \varphi$
- $\Pi, i \models \varphi \cup \psi =$  there is a  $j \geq i$  such that  $\Pi, j \models \psi$   
and  
for all  $h, i \leq h < j$ , we have  $\Pi, h \models \varphi$ .

# Example




Consider abs-exec  $\Pi$  :  $\{isOdd\ x\}$  ,  $\{isOdd\ x\}$  ,  $\{isOdd\ x, x > 0\}$  ,  $\{isOdd\ x, x > 0\}$  , ...

$\Pi \models isOdd\ x \cup x > 0$


*However, this is not a valid property of the program.*

# Some derived operators


$\langle \rangle f$  // eventually  $f$



$f W g$  // weak until



$g R f$  // releases



# Derived operators

$$\langle \rangle \varphi = \text{true} \text{ U } \varphi$$

$$\llbracket \rrbracket \varphi = \neg \langle \rangle \neg \varphi$$

$$\varphi \text{ W } \psi = \llbracket \rrbracket \varphi \vee (\varphi \text{ U } \psi)$$

$$\varphi \text{ R } \psi = \psi \text{ W } (\varphi \wedge \psi)$$

# Past operators

- Useful, e.g. to say: if  $P$  is doing something with  $x$ , then it must have asked a permission to do so.
- “previous”  
 $\Pi, i \models \mathbf{Y} \varphi$       =  $\varphi$  holds in the previous state
- “since”  
 $\Pi, i \models \varphi \mathbf{S} \psi$       =  $\psi$  held in the past, and since that to now  $\varphi$  holds
- Unfortunately, not supported by SPIN.

# Ok, so how can I verify $P \models \varphi$ ?

- We can't directly check all executions  $\rightarrow$  infinite (in two dimensions).
- Try to prove it directly using the definitions?
- We'll take a look another strategy...
- First, let's view abstract executions as *sentences*.

View  $P$  as a sentence-generator. Define:

$$L(P) = \{ \Pi \mid \Pi \text{ is an abs-exec of } P \}$$

*these are sentences over  $\text{pow}(erProp)$*



# Representing $\varphi$ as an automaton ...

- Let  $\varphi$  be the temporal formula we want to verify.
- Suppose we can construct automaton  $A_\varphi$  that ‘accepts’ exactly those infinite sentences over power( $Prop$ ) for which  $\varphi$  holds.
- So  $A_\varphi$  is such that :

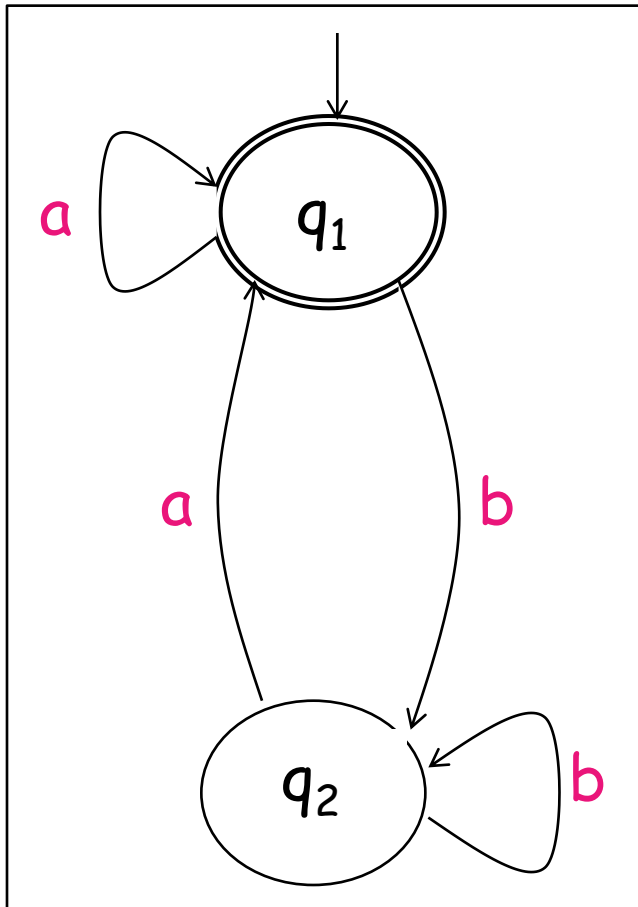
$$L(A_\varphi) = \{ \Pi \mid \Pi \models \varphi \}$$

# Re-express as a language problem

- Well,  $P \models \varphi$  iff
  - There is no  $\Pi \in L(P)$  which will violate  $\varphi$ .
  - In other words, there is no  $\Pi \in L(P)$  that will be accepted by  $L(A_{\neg\varphi})$ .
- So:

$$P \models \varphi \quad \text{iff} \quad L(P) \cap L(A_{\neg\varphi}) = \emptyset$$

# Automaton for accepting sentences




- Add *acceptance states*.
- Accepted sentence:  

“*aba*” and “*aa*” is accepted  
“*bb*” is not accepted.
- But this is for finite sentences.  

For infinite ...?

# Buchi Automaton



- Just different criterion for acceptance
- Examples

“*abab*” → not an infinite

“*ababab...*” → accepted


“*abbbb...*” → not accepted!

# Expressing temporal formulas as Buchi

Use  $\text{power}(\text{Prop})$  as the alphabet  $\Sigma$  of arrow-labels.

Example:  $\neg Xp$  ( $= X\neg p$ )

We'll take  $\text{Prop} = \{p\}$




*Indirectly saying that  $p$  is false.*

*We can drop this, since we only need to (fully) cover accepted sentences.*


# To make the drawing less verbose...

$\neg Xp$ , using  $Prop = \{p\}$



So we have 4 subsets.

$\neg Xp$ , using  $Prop = \{p, q\}$



$p \notin$


Stands for all subsets of  $Prop$  that do not contain  $p$ ; thus implying “ $p$  does not hold”.

$p \in$


Stands for all subsets of  $Prop$  that contain  $p$ ; thus implying “ $p$  holds”.

# Always and Eventually


$\Box p$



$\Diamond p$




$\Diamond \Box p$




# Until

$p \mathbf{U} q$  :



$p \mathbf{U} \mathbf{X} q$  :





# Not Until

Formula:  $\neg (p \mathbf{U} q)$

We'll use the help of these properties:


$$\neg(p \mathbf{U} q) = p \wedge \neg q \mathbf{W} \neg p \wedge \neg q$$

$$\neg(p \mathbf{W} q) = p \wedge \neg q \mathbf{U} \neg p \wedge \neg q$$

$$= \neg q \mathbf{W} \neg p \wedge \neg q$$

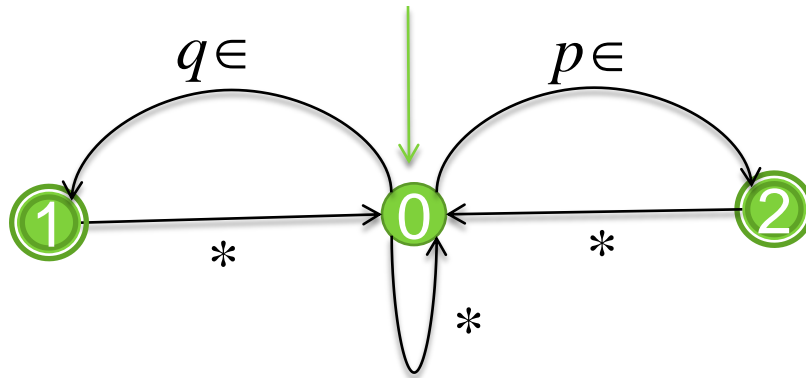
$$= \neg q \mathbf{U} \neg p \wedge \neg q$$

(also generally when  $p, q$  are LTL formulas)



# Generalized Buchi Automaton

$$\Box \langle \rangle p \quad \wedge \quad \Box \langle \rangle q$$



Sets of accepting states:  $\mathbf{F} = \{ \{1\}, \{2\} \}$

which is different than just  $F = \{ 1, 2 \}$  in an ordinary Buchi.

Every GBA can be converted to BA.

# Difficult cases

- How about nested formulas like:

$$\begin{aligned} &(\mathbf{X}p) \mathbf{U} q \\ &(p \mathbf{U} q) \mathbf{U} r \end{aligned}$$

Their Buchi is not trivial to construct.


- Still, any LTL formula can be converted to a Buchi. SPIN implements an automated conversion algorithm; unfortunately it is quite complicated.

# Check list



$$P \models \varphi \quad \text{iff} \quad L(P) \cap L(A_{\neg\varphi}) = \emptyset$$

1. How to construct  $A_{\neg\varphi}$ ?  $\rightarrow$  Buchi ✓
2. We still have a mismatch, because  $P$  is a Kripke structure!
  - Fortunately, we can easily convert it to a Buchi.
3. We still have to construct the intersection.
4. We still to figure out a way to check emptiness.


# Label on state or label on arrow...




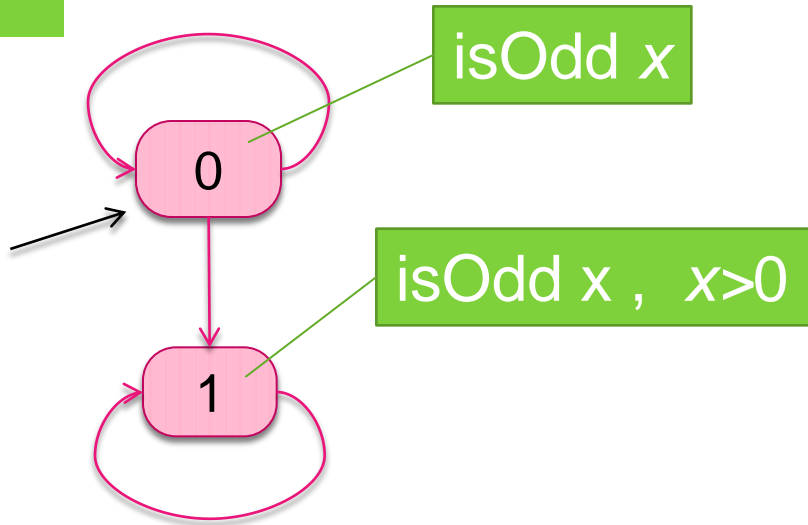
generate the same sentences



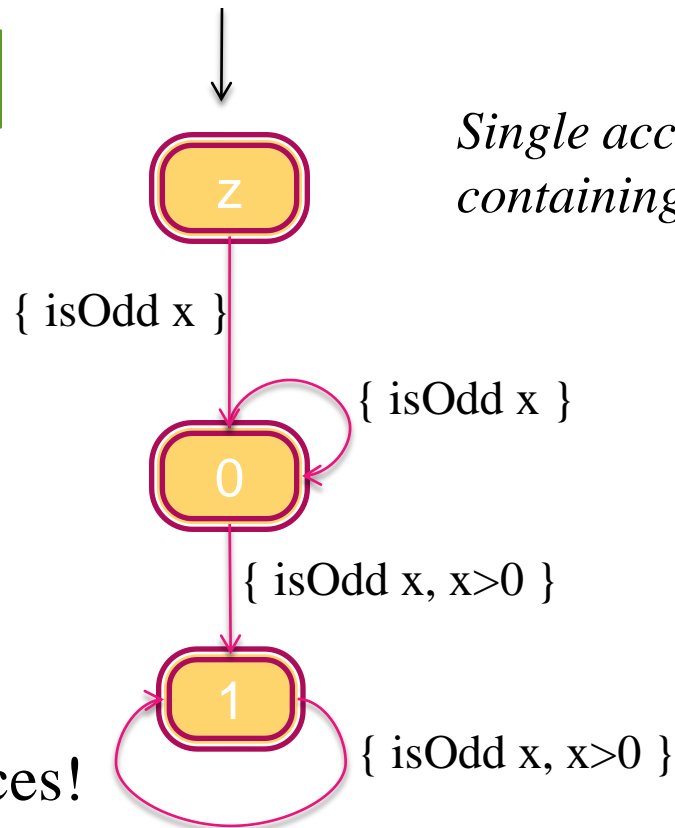
generate the same sentences



# Converting Kripke to Buchi



Generate the same inf. sentences!



*Single accepting set  $F$ , containing all states.*

# Computing intersection

- Rather than directly checking:

*The Buchi version of Kripke P*  
😊

$$L(A_P) \cap L(A_{\neg\varphi}) = \emptyset$$

We check:

$$L(A_P \cap A_{\neg\varphi}) = \emptyset$$

*So we need to figure out how to construct this intersection of two Buchis. Execution over this intersection is also called a “lock-step” execution.*


# Intersection

- Two buchi automata  $A$  and  $B$  over the same alphabet
  - The set of states are respectively  $\Sigma_A$  and  $\Sigma_B$ .
  - starting at respectively  $s_A$  and  $s_B$ .
  - Single accepting set, respectively  $F_A$  and  $F_B$ .
  - $F_A$  is assumed to be  $\Sigma_A$
- $A \cap B$  can be thought as defining lock-step execution of both:
  - The states :  $\Sigma_A \times \Sigma_B$ , starting at  $(s_A, s_B)$
  - Can make a transition only if  $A$  and  $B$  can *both* make the corresponding transition.
  - A single acceptance set  $F$ ;  $(s, t)$  is accepting if  $t \in F_B$ .




# Constructing Intersection, example

$A_p$ :



$p$ : isOdd  $x$   
 $q$ :  $x > 0$

$A_{\neg \langle \rangle q}$ :



$A_p \cap A_{\neg \langle \rangle q}$ :



# Verification

- Sufficient to have an algorithm to check if  $L(C) = \emptyset$ , for the intersection-automaton  $C$ .

$L(C) \neq \emptyset$  iff there is a finite path from  $C$ 's initial state to an accepting state  $f$ , followed by a cycle back to  $f$ .

- So, it comes down to a cycle finding in a finite graph! Solvable.
- The path leading to such a cycle also acts as your counter example!

# Approaches

- View  $C = A_p \cap A_{\neg\varphi}$  as a directed graph.  
Approach 1 :
  1. Calculate all strongly connected component (SCCs) of  $C$  (e.g. with Tarjan) .
  2. Check if there is an SCC containing an accepting state, reachable from  $C$ 's initial state.
- Approach 2, based on Depth First Search (DFS); can be made *lazy* :
  - the full graph is constructed as-we-go, as you search for a cycle.

(so you don't immediately need the full graph)

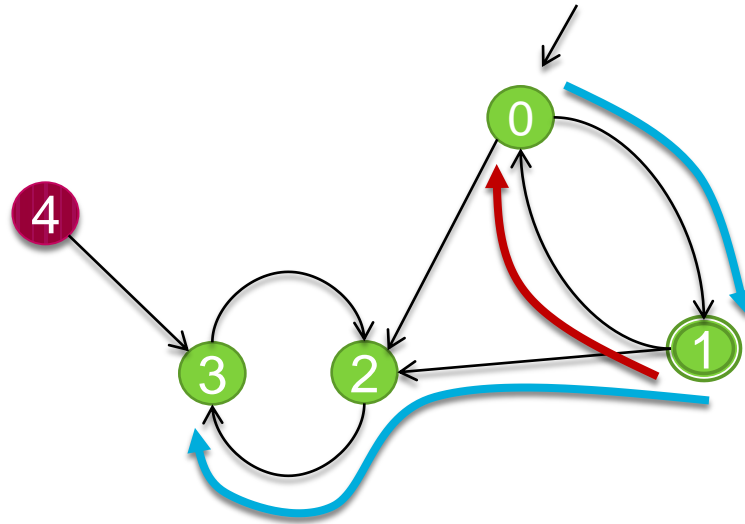
# DFS-approach (SPIN)

- DFS is a way to traverse a graph :

```
DFS(u) {  
    if (u ∈ visited) return ;  
    visited.add(u) ;  
    for (s ∈ next(u)) DFS(s) ;  
}
```

- This will visit **all** reachable nodes. You can already use this to check assertions.

# Example



# Adding cycle detection

```
DFS(u) {  
  if (u ∈ visited) return ;  
  visited.add(u) ;  
  for each (s ∈ next(u)) {  
    if (u ∈ accept) {  
      visited2 = ∅ ;  
      checkCycle (u,s) ;  
    }  
    DFS(s) ;  
  }  
}
```


# checkCycle is another DFS

```
checkCycle(root,u) {  
  
    if (u = root) throw CycleFound ;  
  
    if ( u ∈ visited2 ) return ;  
    visited2.add(u) ;  
    for each (s ∈ next(u))  
        checkCycle(root, s) ;  
}
```

*Can be extended to keep track of the path leading to the cycle → counter example.*

*See Lecture Notes.*

# Example



checkCycle(1,2)

root

*the node currently being processed*



# Lazy (on-the-fly) construction

- Remember that automaton to explore is  $C = A_P \cap A_{\neg\phi}$
- $A_P$  and  $A_{\neg\phi}$  are not literally expressed as an automata; they are Promela models. In particular  $A_P$ , when it is “expanded” to an automaton, it is usually *huge!*

Can we avoid the construction of  $A_P$ ?

- Can we avoid the construction of  $C$ ?

We can at least do lazy construction (SPIN does so).

Benefit : if a cycle is found (verification fails), effort is not wasted to first construct the full  $C$ .

# Lazy construction, representing states

- For now assume  $P$  is just a single process (no concurrency). If we would construct  $A_P$ , each of its state  $u$  is a pair  $(pc, \underline{vars})$ 
  - $pc$  is the “program counter” of  $P$ , to keep track where  $P$  is during an execution.
  - $\underline{vars}$  is a vector of the values of all variables at that state.
- $pc$  is associated to location in the Promela code; it is straight forward to locate and check all “next” statements/actions  $\alpha$  which are possible to execute.

$$\alpha \in \mathbf{enabledOn}(pc, \underline{vars})$$

# Lazy construction

- If  $\alpha$  is an action that is syntactically possible on program counter  $pc$ , let :

**exec**  $\alpha$  ( $pc, \underline{vars}$ )

denote the execution of  $\alpha$  at the state  $(pc, \underline{vars})$ , and this result a new state  $(pc', \underline{vars}')$ .

- We now modify this quantification in the DFS algorithm:

**for each** ( $s \in next(u)$ ) ....

# Lazy construction

- To (still incorrect):

```
for each (  $\alpha \in \text{enabledOn}(pc, \underline{vars})$  ) {  
    s = exec  $\alpha$  (pc, vars)  
    ...  
}
```

- $\alpha \in \text{possible}(pc, \underline{vars})$  means that  $\alpha$  is syntactically a possible next action at  $pc$ , and can execute on state  $vars$ .
- But, we also have to deal with the intersection.

# Lazy construction + intersection

```
DFSlazy(path,  $\langle u, v \rangle$ ) {  
  if ( $\langle u, v \rangle \in \text{visited}$ ) return ;  
  visited.add( $\langle u, v \rangle$ ) ;  
  
  for each ( $\alpha \in \text{enabledOn}(u)$ ,  $\beta \in \text{outArrow}(v)$ ,  
            such that label( $\beta$ ) holds on  $u$ )  
     $s = \text{exec}(\alpha, u)$   
     $t = \text{destination}(\beta)$   
    if ( $t \in \text{accept}$ ) {  
      visited2 =  $\emptyset$  ;  
      checkCycle ( ..... ) ;  
    }  
    else DFSlazy( ... ,  $\langle s, t \rangle$  ) ;  
}
```

```
DFS1(path,  $u$ ) {  
  if ( $u \in \text{visited}$ ) return ;  
  visited.add( $u$ ) ;  
  for each ( $s \in \text{next}(u)$ )  
    if ( $s \in \text{accept}$ ) {  
      visited2 =  $\emptyset$  ;  
      checkCycle ( path++[ $u$ ] ,  $u$ ,  $s$  ) ;  
    }  
    else DFS1( path++[ $u$ ] ,  $s$  ) ;  
}
```

# Concurrency

- So far, we assumed  $P$  is just a single process. What if we have a system  $S$  of  $N$  concurrent processes  $A_1 \dots A_N$ .
- Recall : interleaving model of execution.
- Solution: construct an automaton that equivalently models  $S$ , by taking the product automaton:

$$A_1 \times \dots \times A_N$$

- But... again you may prefer to construct this product lazily as well.

# Interleaved execution in SPIN

- The state of  $A_1 \parallel \dots \parallel A_N$  is represented by a vector

$$u = \langle pc_1, \dots, pc_N, \underline{vars} \rangle$$

- vars represent the vector of all variables (of all processes).
- We just need to redefine  $\alpha \in \mathbf{enabledOn}(u)$ :
  - $\alpha$  belongs to **some process**  $A_i$ , and syntactically can execute at  $pc_i$ .
  - $\alpha$  is enabled on vars

# Fairness

Consider this concurrent system :

$$P \{ \underline{\text{do}} :: x = (x+1) \% N \underline{\text{od}} \} \quad || \quad Q \{ (x==0) ; \text{print } x \}$$

Is it possible that print  $x$  is ignored forever?

- The runtime system determines which fairness assumption is reasonable :
  - No fairness
  - **Weak fairness** : any action that is persistently enabled will eventually be executed.
  - **Strong fairness** : any action that is kept recurrently enabled (but not necessarily persistently enabled) will eventually be executed.
  - There are other variations...
- A **fair execution** : an execution respecting the assumed fairness condition.



# Fairness in SPIN

```
active proctype P(){  
  do  
  :: x = (x+1) % N  
  od  
}  
active proctype Q() {  
  (x==0) ;  
  lab1 : print x }
```

*SPIN only impose “process level weak fairness” : when a process is continually enabled (it has at least one runnable action), an action of the process will eventually be executed.*

*More elaborate fairness assumptions can be encoded as LTL formulas (but gives additional verification overhead).*

- $WFair = \Box (\Box (x==0) \rightarrow \langle \rangle Q@lab1)$
- $SFair = \Box (\Box \langle \rangle (x==0) \rightarrow \langle \rangle Q@lab1)$
- To verify, e.g.  $SFair \rightarrow \langle \rangle x \text{ is printed}$

# Closing remarks

- In principle the use of LTL model checking technique is not limited to SPIN.
- Model checking real programs (as in Java Pathfinder)
  - You need a way to fully control thread scheduling
  - You have to constraint values range to make them finite state.
  - You may also need to limit the depth/length of executions
- Testing concurrent programs
  - Chose a selected set of inputs  $P(x)$ . These are your test-cases. For each test-case, use model checking to verify all possible scheduling of the threads.