

Model Checking and Linear Temporal Logic

Jonathan Aldrich
Carnegie Mellon University

Based on slides developed by Natasha Sharygina

17-654: Analysis of Software Artifacts

Formal Verification by Model Checking

Domain: *Continuously operating concurrent systems* (e.g. operating systems, hardware controllers and network protocols)

- Ongoing, reactive semantics
 - Non-terminating, infinite computations
 - Manifest non-determinism

Instrument: Temporal logic [Pnueli 77] is a formalism for reasoning about behavior of reactive systems

Motivation: What can be Verified?


- Architecture
 - Will these two components interact properly?
 - Allen and Garlan: Wright system checks architectures for deadlock
- Code
 - Am I using this component correctly?
 - Microsoft's Static Driver Verifier ensures complex device driver rules are followed
 - Substantially reduced Windows blue screens
 - Is my code safe
 - Will it avoid error conditions?
 - Will it be responsive, eventually accepting the next input?
- Security
 - Is the protocol I'm using secure
 - Model checking has found defects in security protocols

Temporal Logic Model Checking

[Clarke, Emerson 81][Queille, Sifakis 82]

- Systems are modeled by **finite state machines**
- **Properties** are written in **propositional temporal logic**
- Verification procedure is an **exhaustive search of the state space** of the design
- **Diagnostic counterexamples**

Temporal Logic Model Checking



What is Model Checking?

Does model M satisfy a property P ?
(written $M \models P$)

What is “ M ”?

What is “ P ”?

What is “satisfy”?

What is “M”?


Example Program:

```

precondition: numTickets > 0
reserved = false;
while (true) {
    getQuery();
    if (numTickets > 0 && !reserved)
        reserved = true;
    if (numTickets > 0 && reserved) {
        reserved = false;
        numTickets--;
    }
}

```

State Transition Diagram



What is “M”?


Example Program:

```

precondition: numTickets > 0
reserved = false;
while (true) {
    getQuery();
    if (numTickets > 0 && !reserved)
        reserved = true;
    if (numTickets > 0 && reserved) {
        reserved = false;
        numTickets--;
    }
}

```

State Transition Diagram



What is “M”?

Example Program:

```

precondition: numTickets > 0
reserved = false;
while (true) {
  getQuery();
  if (numTickets > 0 && !reserved)
    reserved = true;
  if (numTickets > 0 && reserved) {
    reserved = false;
    numTickets--;
  }
}


```

What is interesting about this?

Are tickets available?

Is a ticket reserved?

a
r



What is “M”?

Example Program:

```

precondition: numTickets > 0
reserved = false;
while (true) {
  getQuery();
  if (numTickets > 0 && !reserved)
    reserved = true;
  if (numTickets > 0 && reserved) {
    reserved = false;
    numTickets--;
  }
}


```

What is interesting about this?

Are tickets available?

Is a ticket reserved?

a
r



What is “M”?

Abstracted Program: fewer states

```
precondition: available == true
reserved = false;
while (true) {
    getQuery();
    if (available && !reserved)
        reserved = true;
    if (available && reserved) {
        reserved = false;
        available = ?;
    }
}
```

What is “M”?

Abstracted Program: fewer states

```

precondition: available == true
reserved = false;
while (true) {
  getQuery();
  if (available && !reserved)
    reserved = true;
  if (available && reserved) {
    reserved = false;
    available = ?;
  }
}

```



State Transition Graph or Kripke Model


What is “M”?

Abstracted Program: fewer states

```

precondition: available == true
reserved = false;
while (true) {
  getQuery();
  if (available && !reserved)
    reserved = true;
  if (available && reserved) {
    reserved = false;
    available = ?;
  }
}

```



State Transition Graph or Kripke Model

What is “M”?

State: valuations to all variables

concrete state: (numTickets=5, reserved=false)

abstract state: (a=true, r=false)


Initial states: subset of states

Arcs: transitions between states

Atomic Propositions:

a: numTickets > 0

r: reserved = true



State Transition Graph or Kripke Model


What is “M”?

$$M = \langle S, S_0, R, L \rangle$$

Kripke structure:



What is “M”?

$$M = \langle S, S_0, R, L \rangle$$




Kripke structure:

S – finite set of states



What is “M”?



$$M = \langle S, S_0, R, L \rangle$$



Kripke structure:

S – finite set of states

$S_0 \subseteq S$ – set of initial states



What is “M”?


$$M = \langle S, S_0, R, L \rangle$$

Kripke structure:

S – finite set of states

$S_0 \subseteq S$ – set of initial states

$R \subseteq S \times S$ – set of arcs



What is “M”?

$$M = \langle S, S_0, R, L \rangle$$


Kripke structure:

S – finite set of states


$S_0 \subseteq S$ – set of initial states

$R \subseteq S \times S$ – set of arcs

$L: S \rightarrow 2^{\text{AP}}$ – mapping from states to a set of atomic propositions



Model of Computation

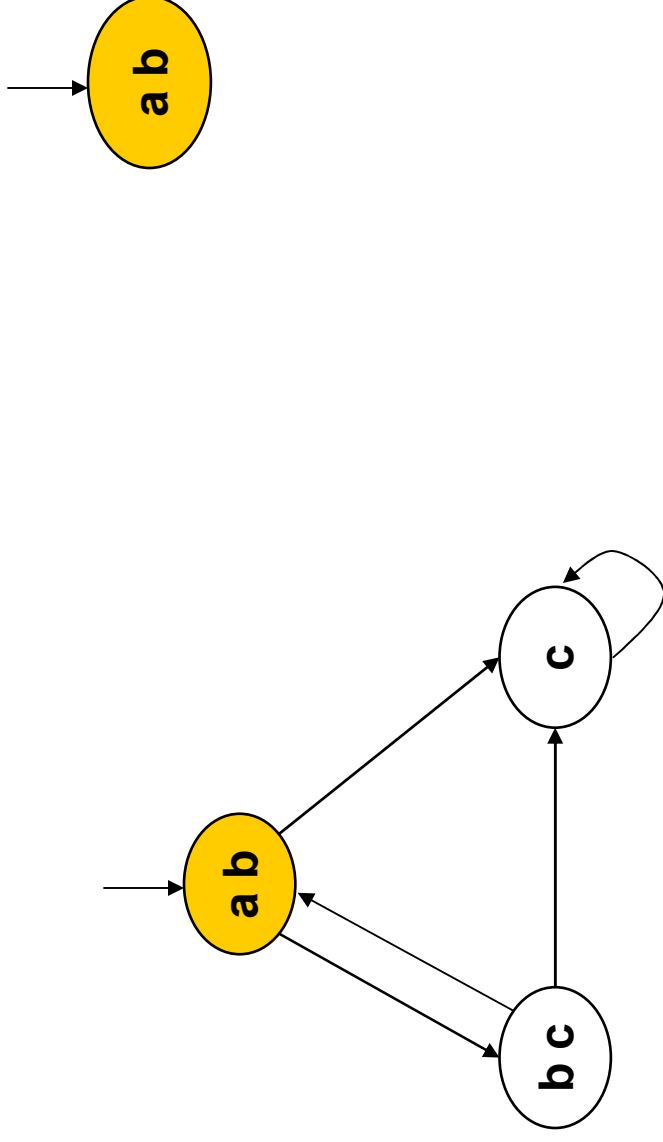


State Transition Graph

Computation Traces

Unwind State Graph to obtain traces. A *trace* is an infinite sequence of states. The *semantics* of a FSM is a set of traces.

Model of Computation





State Transition Graph

Computation Traces

Unwind State Graph to obtain traces. A *trace* is an infinite sequence of states. The *semantics* of a FSM is a set of traces.

Model of Computation




State Transition Graph

Computation Traces

Unwind State Graph to obtain traces. A *trace* is an infinite sequence of states. The *semantics* of a FSM is a set of traces.

Model of Computation




State Transition Graph

Computation Traces

Unwind State Graph to obtain traces. A *trace* is an infinite sequence of states. The *semantics* of a FSM is a set of traces.

Model of Computation




State Transition Graph

Unwind State Graph to obtain traces. A *trace* is an infinite sequence of states. The *semantics* of a FSM is a set of traces.

Computation Traces

Model of Computation




State Transition Graph


Computation Traces

Unwind State Graph to obtain traces. A *trace* is an infinite sequence of states. The *semantics* of a FSM is a set of traces.

Model of Computation



State Transition Graph



Infinite Computation Tree

Represent all traces with an infinite computation tree

What is “P”?

Different kinds of temporal logics

Syntax: What are the formulas in the logic?

Semantics: What does it mean for model **M** to satisfy formula **P**?

Formulas:

- Atomic propositions: properties of states
- Temporal Logic Specifications: properties of traces.

Computation Tree Logics

Examples:

Safety (mutual exclusion): no two processes can be at a critical section at the same time

Liveness (absence of starvation): every request will be eventually granted

Temporal logics differ according to how they handle branching in the underlying computation tree.

Our Focus




In a **linear temporal logic (LTL)**, operators are provided for describing system behavior along a single computation path.

In a **branching-time logic (CTL)**, the temporal operators quantify over the paths that are possible from a given state.


The Logic LTL

Linear Time Logic (LTL) [Pnueli 77]: logic of temporal sequences.


- α : α holds in the current state




- $X\alpha$: α holds in the next state




- $F\gamma$: γ holds eventually



- $G\lambda$: λ holds from now on



- $(\alpha U \beta)$: α holds until β holds



Typical LTL Formulas

- **G** (*Req* \Rightarrow **F** *Ack*): whenever *Request* occurs, it will be eventually *Acknowledged*.
- **G** (*DeviceEnabled*): *DeviceEnabled* always holds on every computation path.
- **G** (**F** *Restart*): Fairness: from any state one will eventually get to a *Restart* state. I.e. *Restart* states occur infinitely often.
- **G** (*Reset* \Rightarrow **F** *Restart*): whenever the reset button is pressed one will eventually get to the *Restart* state.

LTL Conventions

- G is sometimes written \square
- F is sometimes written \diamond

Notation

- A path π in M is an infinite sequence of states s_0, s_1, \dots such that for every $i \geq 0$, $(s_i, s_{i+1}) \in R$
- π^i denotes the suffix of π starting at s_i
- $M, \pi \models f$ means that f holds along path π in the Kripke structure M

Semantics of LTL Formulas



Semantics of LTL Formulas

$$\begin{array}{c}
 \bullet \xrightarrow{\quad} \circ \xrightarrow{\quad} \circ \xrightarrow{\quad} \circ \xrightarrow{\quad} \circ \\
 \mathbf{p}
 \end{array}
 \quad M, \pi \models \mathbf{p}
 \quad \Leftrightarrow \quad \pi \models \mathbf{S} \dots \wedge \mathbf{p} \in L(\mathbf{s})$$

$$M, \pi \models \neg \mathbf{g} \quad \Leftrightarrow \quad M, \pi \not\models \mathbf{g}$$

$$M, \pi \models \mathbf{g}_1 \wedge \mathbf{g}_2 \quad \Leftrightarrow \quad M, \pi \models \mathbf{g}_1 \wedge M, \pi \models \mathbf{g}_2$$

$$M, \pi \models \mathbf{g}_1 \vee \mathbf{g}_2 \quad \Leftrightarrow \quad M, \pi \models \mathbf{g}_1 \vee M, \pi \models \mathbf{g}_2$$

Semantics of LTL Formulas

$$\begin{array}{c} \bullet \longrightarrow \circ \longrightarrow \circ \longrightarrow \circ \longrightarrow \circ \longrightarrow \circ \longrightarrow \circ \end{array} \quad M, \pi \models p \quad \Leftrightarrow \quad \pi \models s \dots \wedge p \in L(s)$$

$$M, \pi \models \neg g \quad \Leftrightarrow \quad M, \pi \not\models g$$

$$M, \pi \models g_1 \wedge g_2 \quad \Leftrightarrow \quad M, \pi \models g_1 \wedge M, \pi \models g_2$$

$$M, \pi \models g_1 \vee g_2 \quad \Leftrightarrow \quad M, \pi \models g_1 \vee M, \pi \models g_2$$

$$\begin{array}{c} \bullet \longrightarrow \circ \longrightarrow \circ \longrightarrow \circ \longrightarrow \circ \longrightarrow \circ \longrightarrow \circ \end{array} \quad M, \pi \models \mathbf{X} g \quad \Leftrightarrow \quad M, \pi^1 \models g$$

Semantics of LTL Formulas



$$M, \pi \models p \iff \pi = s \dots \wedge p \in L(s)$$

$$M, \pi \models \neg g \iff M, \pi \not\models g$$

$$M, \pi \models g_1 \wedge g_2 \iff M, \pi \models g_1 \wedge M, \pi \models g_2$$

$$M, \pi \models g_1 \vee g_2 \iff M, \pi \models g_1 \vee M, \pi \models g_2$$



$$M, \pi \models \mathbf{X} g \iff M, \pi^1 \models g$$



$$M, \pi \models \mathbf{F} g \iff \exists k \geq 0 \mid M, \pi^k \models g$$

Semantics of LTL Formulas

$$\begin{array}{c} \bullet \rightarrow \bullet \rightarrow \bullet \rightarrow \bullet \rightarrow \bullet \end{array} \quad M, \pi \models p \quad \Leftrightarrow \quad \pi \models s \dots \wedge p \in L(s)$$

$$M, \pi \models \neg g \quad \Leftrightarrow \quad M, \pi \not\models g$$

$$M, \pi \models g_1 \wedge g_2 \quad \Leftrightarrow \quad M, \pi \models g_1 \wedge M, \pi \models g_2$$

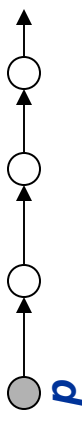
$$M, \pi \models g_1 \vee g_2 \quad \Leftrightarrow \quad M, \pi \models g_1 \vee M, \pi \models g_2$$

$$\begin{array}{c} \bullet \rightarrow \bullet \rightarrow \bullet \rightarrow \bullet \rightarrow \bullet \\ \mathbf{X} \end{array} g \quad \Leftrightarrow \quad M, \pi^1 \models g$$

$$\begin{array}{c} \bullet \rightarrow \bullet \rightarrow \bullet \rightarrow \bullet \rightarrow \bullet \\ \mathbf{F} \end{array} g \quad \Leftrightarrow \quad \exists k \geq 0 \mid M, \pi^k \models g$$

$$\begin{array}{c} \bullet \rightarrow \bullet \rightarrow \bullet \rightarrow \bullet \rightarrow \bullet \\ \mathbf{G} \end{array} g \quad \Leftrightarrow \quad \forall k \geq 0 \mid M, \pi^k \models g$$

Semantics of LTL Formulas



$$M, \pi \models p \iff \pi = s \dots \wedge p \in L(s)$$

$$M, \pi \models \neg g \iff M, \pi \not\models g$$

$$M, \pi \models g_1 \wedge g_2 \iff M, \pi \models g_1 \wedge M, \pi \models g_2$$

$$M, \pi \models g_1 \vee g_2 \iff M, \pi \models g_1 \vee M, \pi \models g_2$$



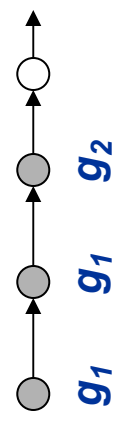
$$M, \pi \models \mathbf{X} g \iff M, \pi^1 \models g$$



$$M, \pi \models \mathbf{F} g \iff \exists k \geq 0 \mid M, \pi^k \models g$$



$$M, \pi \models \mathbf{G} g \iff \forall k \geq 0 \mid M, \pi^k \models g$$



$$M, \pi \models g_1 \mathbf{U} g_2 \iff \exists k \geq 0 \mid M, \pi^k \models g_2 \wedge \forall 0 \leq j < k \mid M, \pi^j \models g_1$$

Semantics of LTL Formulas

$$\begin{array}{c} \bullet \longrightarrow \bullet \longrightarrow \bullet \longrightarrow \bullet \longrightarrow \bullet \longrightarrow \bullet \\ p \end{array} \quad M, \pi \models p \quad \Leftrightarrow \quad \pi = s \dots \wedge p \in L(s)$$

$$M, \pi \models \neg g \quad \Leftrightarrow \quad M, \pi \not\models g$$

$$M, \pi \models g_1 \wedge g_2 \quad \Leftrightarrow \quad M, \pi \models g_1 \wedge M, \pi \models g_2$$

$$M, \pi \models g_1 \vee g_2 \quad \Leftrightarrow \quad M, \pi \models g_1 \vee M, \pi \models g_2$$

$$\begin{array}{c} \bullet \longrightarrow \bullet \longrightarrow \bullet \longrightarrow \bullet \longrightarrow \bullet \longrightarrow \bullet \\ g \end{array} \quad M, \pi \models \mathbf{X} g \quad \Leftrightarrow \quad M, \pi^1 \models g$$

$$\begin{array}{c} \bullet \longrightarrow \bullet \longrightarrow \bullet \longrightarrow \bullet \longrightarrow \bullet \longrightarrow \bullet \\ g \end{array} \quad M, \pi \models \mathbf{F} g \quad \Leftrightarrow \quad \exists k \geq 0 \mid M, \pi^k \models g$$

$$\begin{array}{c} \bullet \longrightarrow \bullet \longrightarrow \bullet \longrightarrow \bullet \longrightarrow \bullet \longrightarrow \bullet \\ g \quad g \quad g \end{array} \quad M, \pi \models \mathbf{G} g \quad \Leftrightarrow \quad \forall k \geq 0 \mid M, \pi^k \models g$$

$$\begin{array}{c} \bullet \longrightarrow \bullet \longrightarrow \bullet \longrightarrow \bullet \longrightarrow \bullet \longrightarrow \bullet \\ g_1 \quad g_1 \quad g_2 \end{array} \quad M, \pi \models g_1 \mathbf{U} g_2 \quad \Leftrightarrow \quad \exists k \geq 0 \mid M, \pi^k \models g_2 \quad \wedge \quad \forall 0 \leq j < k \quad M, \pi^j \models g_1$$

semantics of “until” in English are potentially unclear—
that’s why we have a formal definition

Practice Writing Properties


- If the door is locked, it will not open until someone unlocks it
 - assume atomic predicates locked, unlocked, open
- If you press ctrl-C, you will get a command line prompt
- The saw will not run unless the safety guard is engaged

Practice Writing Properties

- If the door is locked, it will not open until someone unlocks it
 - assume atomic predicates locked, unlocked, open
 - $G(\text{locked} \Rightarrow (\neg \text{open} \cup \text{unlocked}))$
- If you press ctrl-C, you will get a command line prompt
 - $G(\text{ctrlC} \Rightarrow \text{F prompt})$
- The saw will not run unless the safety guard is engaged
 - $G(\neg \text{safety} \Rightarrow \neg \text{running})$

LTL Model Checking Example


- Pressing Start will eventually result in heat



LTL Model Checking Example

- Pressing Start will eventually result in heat

$G(\text{Start} \Rightarrow \mathbf{F} \text{Heat})$




LTL Model Checking

- f (primitive formula)
 - Just check the properties of the current state
- $X f$
 - Verify f holds in all successors of the current state
- $G f$
 - Find all reachable states from the current state, and ensure f holds in all of them
 - use depth-first or breadth-first search
- $f U g$
 - Do a depth-first search from the current state. Stop when you get to a g or you loop back on an already visited state. Signal an error if you hit a state where f is false before you stop.
- $F f$
 - Harder. Intuition: look for a path from the current state that loops back on itself, such that f is false on every state in the path. If no such path is found, the formula is true.
 - Reality: use Büchi automata

LTL Model Checking Example


- Pressing Start will eventually result in heat

$G(\text{Start} \Rightarrow \mathbf{F} \text{Heat})$



LTL Model Checking Example

- The oven doesn't heat up until the door is closed.




LTL Model Checking Example

- The oven doesn't heat up until the door is closed.

$(\neg \text{Heat}) \mathbf{U} \text{Close}$

$(\neg \text{Heat}) \mathbf{W} \text{Close}$

$\mathbf{G} (\text{not Closed} \Rightarrow \text{not Heat})$



Efficient Algorithms for LTL Model Checking


- Use Büchi automata
 - Beyond the scope of this course
- Canonical reference on Model Checking:
 - Edmund Clarke, Orna Grumberg, and Doron A. Peled. Model Checking. MIT Press, 1999.

SPIN: The Promela Language

- PROcess MEta Language
- Asynchronous composition of independent processes
- Communication using channels and global variables
- Non-deterministic choices and interleavings

Mutual Exclusion

- Peterson's solution to the mutual exclusion problem

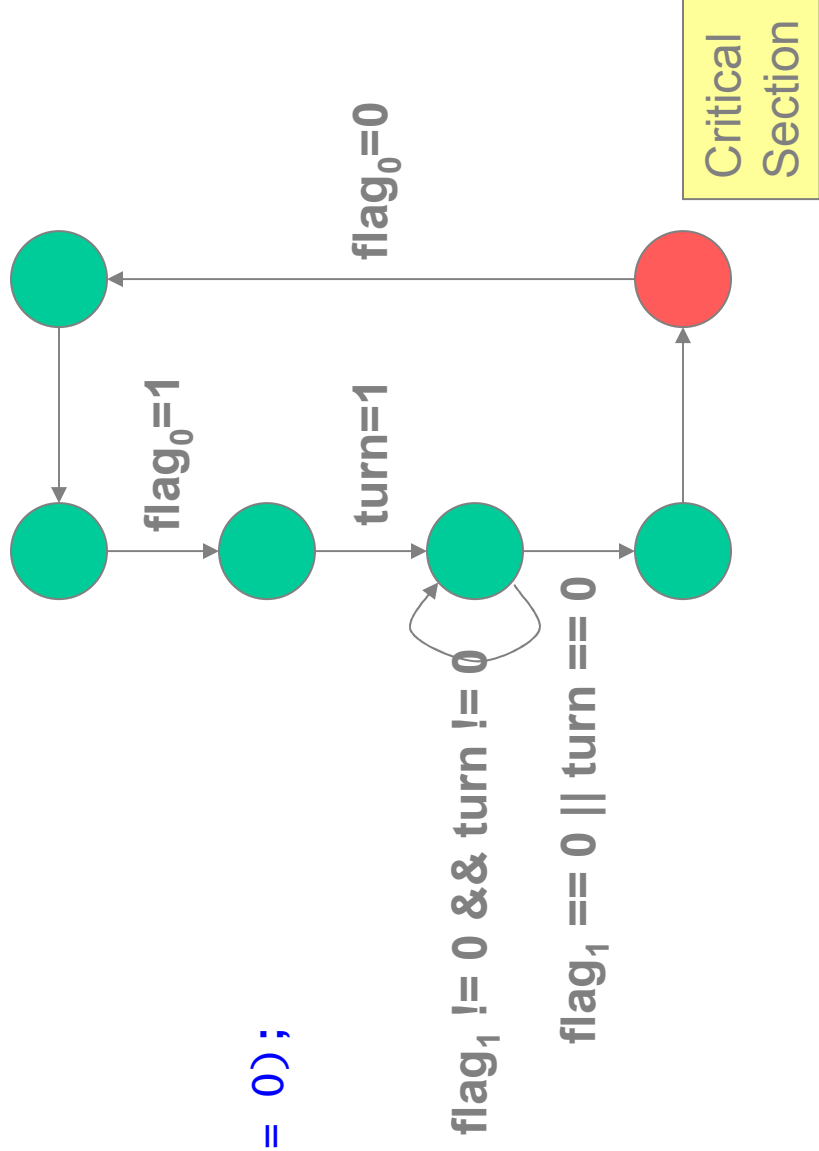


Mutual Exclusion in SPIN

```

bool turn;
bool flag[2];
proctype mutex0() {
again:
    flag[0] = 1;
    turn = 1;
    (flag[1] == 0 || turn == 0);
    /* critical section */
    flag[0] = 0;
    goto again;
}

```




Mutual Exclusion in SPIN

```

bool turn;
bool flag[2];
proctype mutex0() {
again:
    flag[0] = 1;
    turn = 1;
    (flag[1] == 0 || turn == 0);
    /* critical section */
    flag[0] = 0;
    goto again;
}

```

guard:
Cannot go past this point
until the condition is true



Mutual Exclusion in SPIN

```
bool turn, flag[2];

active [2] proctype user0
{
  assert(!_pid == 0 || __pid == 1);
again:
  flag[_pid] = 1;
  turn = 1 - _pid;
  (flag[1 - _pid] == 0 || turn == _pid);

  /* critical section */

  flag[_pid] = 0;
  goto again;
}
```

Mutual Exclusion in SPIN

Active process:
automatically creates instances of processes

```
bool turn, flag[2];

active [2] proctype user()
{
    assert(!_pid == 0 || !_pid == 1);
again:
    flag[_pid] = 1;
    turn = 1 - _pid;
    (flag[1 - _pid] == 0 || turn == _pid);

    /* critical section */

    flag[_pid] = 0;
    goto again;
}
```

Mutual Exclusion in SPIN

```

bool turn, flag[2];

active [2] proctype user0
{
  assert(!_pid == 0 || !_pid == 1);
again:
  flag[_pid] = 1;
  turn = 1 - _pid;
  (flag[1 - _pid] == 0 || turn == _pid);

  /* critical section */

  flag[_pid] = 0;
  goto again;
}

```

Mutual Exclusion in SPIN

```
bool turn, flag[2];

active [2] proctype user0
{
  assert(_pid == 0 || __pid == 1);
  again:
  flag[_pid] = 1;
  turn = 1 - _pid;
  (flag[1 - _pid] == 0 || turn == _pid);
  /* critical section */

  flag[_pid] = 0;
  goto again;
}
```

assert:
Checks that there are only
at most two instances with
identifiers 0 and 1

Mutual Exclusion in SPIN

```
bool turn, flag[2];
byte ncrit;
active [2] proctype user0
```

ncrit:
Counts the number of
Process in the critical section

```
{
  assert(!_pid == 0 || __pid == 1);
again:
  flag[_pid] = 1;
  turn = 1 - _pid;
  (flag[1 - _pid] == 0 || turn == _pid);

  ncrit++;
  assert(ncrit == 1); /* critical section */
  ncrit--;

  flag[_pid] = 0;
  goto again;
}
```

Mutual Exclusion in SPIN

```
bool turn, flag[2];
byte ncrit;
active [2] proctype user0
{
    assert(_pid == 0 || __pid == 1);
again:
    flag[_pid] = 1;
    turn = 1 - _pid;
    (flag[1 - _pid] == 0 || turn == _pid);

    ncrit++;
    assert(ncrit == 1); /* critical section */
    ncrit--;

    flag[_pid] = 0;
    goto again;
}
```

assert:
Checks that there are always
at most one process in the
critical section

Mutual Exclusion in SPIN

```

bool turn, flag[2];
bool critical[2];
#define critical1 critical[1]

active [2] proctype user0
{
    assert(!_pid == 0 || __pid == 1);
again:
    flag[_pid] = 1;
    turn = 1 - _pid;
    (flag[1 - _pid] == 0 || turn == _pid);

    critical[_pid] = 1;
    /* critical section */
    critical[_pid] = 0;

    flag[_pid] = 0;
    goto again;
}

```

LTL Properties:

The processes are never both in the critical section

No matter what happens, a process will eventually get to a critical section

Once process 0 raises its flag, it will eventually get into the critical section

Mutual Exclusion in SPIN

```

bool turn, flag[2];
bool critical[2];
#define critical1 critical[1]

active [2] proctype user0
{
    assert(!_pid == 0 || __pid == 1);
again:
    flag[_pid] = 1;
    turn = 1 - _pid;
    (flag[1 - _pid] == 0 || turn == _pid);

    critical[_pid] = 1;
    /* critical section */
    critical[_pid] = 0;

    flag[_pid] = 0;
    goto again;
}

```

LTL Properties:

The processes are never both in the critical section

$G (\text{critical}[0] == 0 \parallel \text{critical}[1] == 0)$
 $[\] (\text{!critical}0 \parallel \text{!critical}1)$

No matter what happens, a process will eventually get to a critical section

$F \text{critical}[0]$
 $\langle \rangle \text{critical}0$

Once process 0 raises its flag, it will eventually get into the critical section

$G (\text{turn} \rightarrow F \text{critical}0)$
 $[\] (\text{turn} \rightarrow \langle \rangle \text{critical}0)$

Mutual Exclusion in SPIN

```

bool turn, flag[2];
bool critical[2];
#define critical1 critical[1]

active [2] proctype user0
{
    assert(!_pid == 0 || __pid == 1);
again:
    flag[_pid] = 1;
    turn = 1 - _pid;
    (flag[1 - _pid] == 0 || turn == _pid);

    critical[_pid] = 1;
    /* critical section */
    critical[_pid] = 0;

    flag[_pid] = 0;
    goto again;
}

```

LTL Properties:

The processes are never both in the critical section

```

G (critical[0] == 0 || critical[1] == 0)
[] (!critical0 || !critical1)

```

* caveat: can't use array indexes in SPIN LTL properties--have to duplicate code


** caveat: can't use next (X) due to SPIN limitations

*** caveat: SPIN accepts a "never" property rather than an "always" property, so LTL formulas should be negated before passing them to SPIN.

State Space Explosion

Problem:

Size of the state graph can be exponential in size of the program (both in the number of the program *variables* and the number of program *components*)



$$M = M_1 \parallel \dots \parallel M_n$$

If each M_i has just 2 local states, potentially 2^n global states

Research Directions: State space reduction

Model Checking Performance

- Model Checkers today can routinely handle systems with between 100 and 300 state variables.
- Systems with 10^{120} reachable states have been checked.
- By using appropriate abstraction techniques, systems with an essentially **unlimited number of states** can be checked.

Notable Examples

- **IEEE Scalable Coherent Interface** – In 1992 Dill’s group at Stanford used **Murphi** to find several errors, ranging from uninitialized variables to subtle logical errors
- **IEEE Futurebus** – In 1992 Clarke’s group at CMU found previously undetected design errors
- **PowerScale multiprocessor** (processor, memory controller, and bus arbiter) was verified by Verimag researchers using CAESAR toolbox
- **Lucent telecom. protocols** were verified by FormalCheck – errors leading to lost transitions were identified
- **PowerPC 620 Microprocessor** was verified by Motorola’s Verdict model checker.

The Grand Challenge: Model Check Software

Extract finite state machines from programs written in conventional programming languages

Use a finite state programming language:

- executable design specifications (Statecharts, xUML, etc.).

Unroll the state machine obtained from the executable of the program.

The Grand Challenge: Model Check Software

Use a combination of the state space reduction techniques to avoid generating too many states.

- **Verisoft** (Bell Labs)
- **FormalCheck/xUML** (UT Austin, Bell Labs)
- **ComFoRT** (CMU/SEI)

Use static analysis to extract a finite state skeleton from a program.

Model check the result.

- **Bandera** – Kansas State
- **Java PathFinder** – NASA Ames
- **SLAM/Bebop** - Microsoft

Bonus: Computation Tree Logic (CTL)

Computation Tree Logics

Formulas are constructed from *path quantifiers* and *temporal operators*:

1. Path Quantifiers:

- **A** – “for every path”
- **E** – “there exists a path”


2. Temporal Operator:

- **X** α - α holds **next** time
- **F** α - α holds sometime in the **future**
- **G** α - α holds **globally** in the **future**
- α **U** β - α holds **until** β holds


LTL: start with an A and then use only Temporal Operators

The Logic CTL


In a branching-time logic (CTL), the temporal operators quantify over the paths that are possible from a given state (s_0). Requires each temporal operator (**X**, **F**, **G**, and **U**) to be preceded by a path quantifier (**A** or **E**).




$$\mathbf{M}, s_0 \models \mathbf{AG} \mathbf{c}$$



$$\mathbf{M}, s_0 \models \mathbf{AF} \mathbf{c}$$



$$\mathbf{M}, s_0 \models \mathbf{EF} \mathbf{c}$$




$$\mathbf{M}, s_0 \models \mathbf{EG} \mathbf{c}$$

Typical CTL Formulas

- **EF** ($Started \wedge \neg Ready$): it is possible to get to a state where *Started* holds but *Ready* does not hold.
- **AG** ($Req \Rightarrow \mathbf{AF} Ack$): whenever *Request* occurs, it will be eventually *Acknowledged*.
- **AG** (*DeviceEnabled*): *DeviceEnabled* always holds on every computation path.
- **AG (EF Restart)**: from any state it is possible to get to the *Restart* state.


Trivia

- **AG(EF p)** cannot be expressed in LTL
 - Reset property: from every state it is possible to get to p
 - But there might be paths where you never get to p
 - Different from **A(GF p)**
 - Along each possible path, for each state in the path, there is a future state where p holds
 - Counterexample: ababab...



Trivia

- **A(FG p)** cannot be expressed in CTL
 - Along all paths, one eventually reaches a point where p always holds from then on
 - But at some points in some paths where p always holds, there might be a diverging path where p does not hold
 - Different from **AF(AG p)**
 - Along each possible path there exists a state such that p always holds from then on
 - Counterexample: the path that stays in s_0



Linear vs. branching-time logics

some advantages of LTL

- LTL properties are preserved under “abstraction”: i.e., if \mathcal{M} “approximates” a more complex model \mathcal{M}' , by introducing more paths, then

$$\mathcal{M} \models \psi \Rightarrow \mathcal{M}' \models \psi$$
- “counterexamples” for LTL are simpler: consisting of single executions (rather than trees).
- The automata-theoretic approach to LTL model checking is simpler (no tree automata involved).
- anecdotally, it seems most properties people are interested in are linear-time properties.

some advantages of BT logics

- BT allows expression of some useful properties like ‘**reset**’.
- CTL, a limited fragment of the more complete BT logic CTL*, can be model checked in time linear in the formula size (as well as in the transition system). But formulas are usually far smaller than system models, so this isn’t as important as it may first seem.
- Some BT logics, like μ -calculus and CTL, are well-suited for the kind of fixed-point computation scheme used in symbolic model checking.

Formulas over States and Paths

- State formulas
 - Describe a property of a state in a model M
 - If $p \in AP$, then p is a state formula
 - If f and g are state formulas, then $\neg f$, $f \wedge g$ and $f \vee g$ are state formulas
 - If f is a path formula, then $\mathbf{E} f$ and $\mathbf{A} f$ are state formulas
- Path formulas
 - Describe a property of an infinite path through a model M
 - If f is a state formula, then f is also a path formula
 - If f and g are path formulas, then $\neg f$, $f \wedge g$, $f \vee g$, $\mathbf{X} f$, $\mathbf{F} f$, $\mathbf{G} f$, and $\mathbf{U} g$ are path formulas

Notation

- A path π in M is an infinite sequence of states s_0, s_1, \dots such that for every $i \geq 0$, $(s_i, s_{i+1}) \in R$
- π^i denotes the suffix of π starting at s_i
- If f is a state formula, $M, s \models f$ means that f holds at state s in the Kripke structure M
- If f is a path formula, $M, \pi \models f$ means that f holds along path π in the Kripke structure M


Semantics of Formulas

$M, s \models p$	$\Leftrightarrow p \in L(s)$	$M, \pi \models f$	$\Leftrightarrow \pi = s \dots \wedge M, s \models f$
$M, s \models \neg f$	$\Leftrightarrow M, s \not\models f$	$M, \pi \models \neg g$	$\Leftrightarrow M, \pi \not\models g$
$M, s \models f_1 \wedge f_2$	$\Leftrightarrow M, s \models f_1 \wedge M, s \models f_2$	$M, \pi \models g_1 \wedge g_2$	$\Leftrightarrow M, \pi \models g_1 \wedge M, \pi \models g_2$
$M, s \models f_1 \vee f_2$	$\Leftrightarrow M, s \models f_1 \vee M, s \models f_2$	$M, \pi \models g_1 \vee g_2$	$\Leftrightarrow M, \pi \models g_1 \vee M, \pi \models g_2$
$M, s \models \mathbf{E} g_1$	$\Leftrightarrow \exists \pi = s \dots \mid M, \pi \models g_1$	$M, \pi \models \mathbf{X} g$	$\Leftrightarrow M, \pi^1 \models g$
$M, s \models \mathbf{A} g_1$	$\Leftrightarrow \forall \pi = s \dots \mid M, \pi \models g_1$	$M, \pi \models \mathbf{F} g$	$\Leftrightarrow \exists k \geq 0 \mid M, \pi^k \models g$
		$M, \pi \models \mathbf{G} g$	$\Leftrightarrow \forall k \geq 0 \mid M, \pi^k \models g$
		$M, \pi \models g_1 \mathbf{U} g_2$	$\Leftrightarrow \exists k \geq 0 \mid M, \pi^k \models g_2$ $\wedge \forall 0 \leq j < k \mid M, \pi^j \not\models g_1$

Bonus: Example

An Example

```
mtype = { NONCRITICAL, TRYING, CRITICAL };
show mtype state[2];
proctype process(int id) {
  beginning;
  noncritical:
    state[id] = NONCRITICAL;
  if
  :: goto noncritical;
  :: true;
  fi;
  trying:
    state[id] = TRYING;
  if
  :: goto trying;
  :: true;
  fi;
  critical:
    state[id] = CRITICAL;
  if
  :: goto critical;
  :: true;
  fi;
  goto beginning;}
init { run process(0); run process(1); }
```



An Example

```
mtype = { NONCRITICAL, TRYING, CRITICAL };
show mtype state[2];
proctype process(int id) {
beginning;
noncritical;
    state[id] = NONCRITICAL;
if
:: goto noncritical;
:: true;
fi;
trying;
    state[id] = TRYING;
if
:: goto trying;
:: true;
fi;
critical;
    state[id] = CRITICAL;
if
:: goto critical;
:: true;
fi;
goto beginning;}
init { run process(0); run process(1); }
```

An Example

```
mtype = { NONCRITICAL, TRYING, CRITICAL };
show mtype state[2];
proctype process(int id) {
beginning;
noncritical;
    state[id] = NONCRITICAL;
if
:: goto noncritical;
:: true;
fi;
trying;
    state[id] = TRYING;
if
:: goto trying;
:: true;
fi;
critical;
    state[id] = CRITICAL;
if
:: goto critical;
:: true;
fi;
goto beginning;}
init { run process(0); run process(1); }
```


An Example

```
mtype = { NONCRITICAL, TRYING, CRITICAL };
show mtype state[2];
proctype process(int id) {
beginning;
noncritical;
    state[id] = NONCRITICAL;
if
:: goto noncritical;
:: true;
fi;
trying;
    state[id] = TRYING;
if
:: goto trying;
:: true;
fi;
critical;
    state[id] = CRITICAL;
if
:: goto critical;
:: true;
fi;
goto beginning;}
init { run process(0); run process(1); }
```

An Example


```
mtype = { NONCRITICAL, TRYING, CRITICAL };
show mtype state[2];
proctype process(int id) {
beginning;
noncritical;
    state[id] = NONCRITICAL;
if
:: goto noncritical;
:: true;
fi;
trying;
    state[id] = TRYING;
if
:: goto trying;
:: true;
fi;
critical;
    state[id] = CRITICAL;
if
:: goto critical;
:: true;
fi;
goto beginning;}
init { run process(0); run process(1); }
```

An Example

```

mtype = { NONCRITICAL, TRYING, CRITICAL };
show mtype state[2];
proctype process(int id) {
  beginning;
  noncritical:
    state[id] = NONCRITICAL;
  if
    :: goto noncritical;
    :: true;
  fi;
  trying:
    state[id] = TRYING;
  if
    :: goto trying;
    :: true;
  fi;
  critical:
    state[id] = CRITICAL;
  if
    :: goto critical;
    :: true;
  fi;
  goto beginning;}
init { run process(0); run process(1); }

```



Enabled Statements

- A statement needs to be enabled for the process to be scheduled.

```
bool a, b;
proctype p1()
{
    a = true;
    a & b;
    a = false;
}
proctype p2()
{
    b = false;
    a & b;
    b = true;
}
init { a = false; b = false; run p1(); run p2(); }
```

Enabled Statements

- A statement needs to be enabled for the process to be scheduled.

These statements are enabled only if both **a** and **b** are true.

```
bool a, b;
proctype p1()
{
    a = true;
    a & b;
    a = false;
}
proctype p2()
{
    b = false;
    a & b;
    b = true;
}
init { a = false; b = false; run p1(); run p2(); }
```

Enabled Statements

- A statement needs to be enabled for the process to be scheduled.

```
bool a, b;
proctype p1()
{
    a = true;
    a & b;
    a = false;
}
proctype p2()
{
    b = false;
    a & b;
    b = true;
}
init { a = false; b = false; run p1(); run p2(); }
```

These statements are enabled only if both **a** and **b** are true.

In this case **b** is always false and therefore there is a deadlock.

Other constructs

- Do loops

```
do
  :: count = count + 1;
  :: count = count - 1;
  :: (count == 0) -> break
od
```

Other constructs

- Do loops
- Communication over channels

```
proctype sender(chan out)
{
  int x;

  if
  ::x=0;
  ::x=1;
  fi

  out ! x;
}
```


Other constructs

- Do loops
- Communication over channels
- Assertions

```
proctype receiver(chan in)
{
    int value;
    in ? value;
    assert(value == 0 || value == 1)
}
```

Other constructs

- Do loops
- Communication over channels
- Assertions
- Atomic Steps

```
int value;
proctype increment()
{ atomic {
    x = value;
    x = x + 1;
    value = x;
} }
```