



Model Checking

A Hands-On Introduction

A. Cimatti, M. Pistore, and M. Roveri

13th International Conference on Automated Planning & Scheduling
June 10 2003
Trento (Italy)

Model Checking: A Hands-On Introduction— June 10 2003, Trento (Italy)

– p. 1

In this course



Speakers:

- Alessandro Cimatti (ITC-irst)
- Marco Pistore (ITC-irst and University of Trento)
- Marco Roveri (ITC-irst)

The goals of the course are:

- to provide a practical introduction to symbolic model checking,
- to describe the basic features of the NuSMV symbolic model checker.

Model Checking: A Hands-On Introduction— June 10 2003, Trento (Italy)

– p. 2

The course at a glance



- Introduction
 - Introduction to Formal Methods and Model Checking
 - Models for Reactive Systems: Kripke Structures
 - Properties of Reactive Systems: CTL, LTL
- NuSMV model checker
 - The NuSMV Open Source project
 - The SMV language
 - NuSMV demo
- Advanced Topics
 - Symbolic Model Checking Techniques: BDD-based and SAT-based techniques
 - The SMV language (advanced)
- Play with NuSMV

Model Checking: A Hands-On Introduction— June 10 2003, Trento (Italy)

– p. 3



Part 1 - Course Overview

— *Model Checking* — *A Hands-On Introduction*

A. Cimatti, M. Pistore, and M. Roveri

13th International Conference on Automated Planning & Scheduling, June 10 2003, Trento (Italy)

Model Checking: A Hands-On Introduction— June 10 2003, Trento (Italy)

— p. 4

The Need for Formal Methods



- Development of Industrial Systems:
 - System reliability depends on the “correct” functioning of hardware and software.
 - Safety-critical, money-critical systems.
 - Growing complexity of environment and required functionalities.
 - Market issues: time-to-market, development costs.

Model Checking: A Hands-On Introduction— June 10 2003, Trento (Italy)

— p. 5

The Need for Formal Methods (II)



- Difficulties with traditional methodologies:
 - Ambiguities (in requirements, architecture, detailed design).
 - Errors in specification/design refinements.
 - The assurance provided by testing can be too limited.
 - Testing came too late in development.
- Consequences:
 - Expensive errors in the early design phases.
 - Infeasibility of achieving high reliability requirements.
 - Low software quality: limited modifiability.

Model Checking: A Hands-On Introduction— June 10 2003, Trento (Italy)

— p. 6

Formal Methods: Solution and Benefits

- The problem:
 - Certain (sub)systems can be too complicated/critical to design with traditional techniques.
 - Formal Methods:
 - *Formal Specification*: precise, unambiguous description.
 - *Formal Validation & Verification Tools*: exhaustive analysis of the formal specification.
 - Potential Benefits:
 - Find design bugs in early design stages.
 - Achieve higher quality standards.
 - Shorten time-to-market reducing manual validation phases.
 - Produce well documented, maintainable products.
- Highly recommended by ESA, NASA for the design of safety-critical systems.

Formal Methods: Potential Problems

- Main Issue: Effective use of Formal Methods
 - debug/verify during the design process;
 - without slowing down the design process;
 - without increasing costs too much.
- Potential Problems of Formal Methods:
 - formal methods can be too costly;
 - formal methods can be not effective;
 - training problems;
 - the verification problem can be too difficult.
- How can we get benefits and avoid these problems?
 - by adapting our technologies and tools to the specific problem at hand;
 - using “automated” verification techniques (e.g. model checking).

Formal Verification Techniques

- Model-based simulation or testing
 - method: test for ϕ by exploring possible behaviors.
 - tools: test case generators.
 - applicable if: the system defines an executable model.
- Deductive Methods
 - method: provide a formal proof that ϕ holds.
 - tool: theorem prover, proof assistant or proof checker.
 - applicable if: systems can be described as a mathematical theory.
- Model Checking
 - method: systematic check of ϕ in all states of the system.
 - tools: model checkers.
 - applicable if: system generates (finite) behavioral model.

Simulation and Testing



- Basic procedure:
 - take a model (simulation) or a realization (testing).
 - stimulate it with certain inputs, i.e. the test cases.
 - observe produce behavior and check whether this is “desired”.
- Drawback:
 - The number of possible behaviors can be too large (or even infinite).
 - Unexplored behaviors may contain the fatal bug.
- Testing and simulation can show the presence of bugs, *not their absence*.

Theorem Proving



- Basic procedure:
 - describe the system as a mathematical theory.
 - express the property in the mathematical theory.
 - prove that the property is a theorem in the mathematical theory.
- Drawback:
 - Express the system as a mathematical theory can be difficult.
 - Find a proof can require a big effort.
- Theorem proving can be used to prove *absence of bugs*.

Model Checking



- Basic procedure:
 - describe the system as Finite State Model.
 - express properties in Temporal Logic.
 - formal V&V by automatic exhaustive search over the state space.
- Drawback:
 - State space explosion.
 - Expressivity – hard to deal with parametrized systems.
- Model checking can be used to prove *absence of bugs*.

Industrial success of Model Checking



- From academics to industry in a decade.
- Easier to integrate within industrial development cycle:
 - input from practical design languages (e.g. VHDL, SDL, StateCharts);
 - expressiveness limited but often sufficient in practice.
- Does not require deep training (“push-button” technology).
 - Easy to explain as exhaustive simulation.
- Powerful debugging capabilities:
 - detect costly problems in early development stages (cfr. Pentium bug);
 - exhaustive, thus effective (often bugs are also in scaled-down problems).
 - provides counterexamples (directs the designer to the problem).

Model Checking in a nutshell



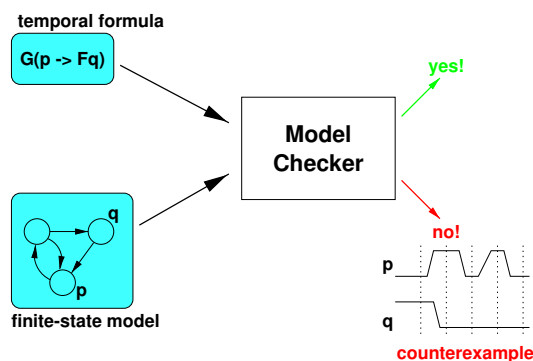
- Reactive systems represented as a finite state models (in this course, Kripke models).
- System behaviors represented as (possibly) infinite sequences of states.
- Requirements represented as formulae in temporal logics.
- “*The system satisfies the requirement*” represented as truth of the formula in the Kripke model.
- Efficient model checking algorithms based on exhaustive exploration of the Kripke model.

What is a Model Checker



- A model checker is a software tool that
- given a description of a Kripke model M ...
 - ... and a property Φ ,
 - decides whether $M \models \Phi$,
 - returns “yes” if the property is satisfied,
 - otherwise returns “no”, and provides a counterexample.

What is a Model Checker



Model Checking: A Hands-On Introduction—June 10 2003, Trento (Italy)

– p. 16

What is not covered in this course



- A deep theoretical background. We will focus on practice.
- Advanced model checking techniques:
 - abstraction;
 - compositional, assume-guarantee reasoning;
 - symmetry reduction;
 - approximation techniques (e.g. directed to bug hunting);
 - model transformation techniques (e.g. minimization wrt to bisimulation).
- Many other approaches and tools for model checking.
- Model checking for infinite-state (e.g. hybrid, timed) systems.

Model Checking: A Hands-On Introduction—June 10 2003, Trento (Italy)

– p. 17

Part 2 - Symbolic Model Checking



– *Model Checking* –
A Hands-On Introduction
A. Cimatti, M. Pistore, and M. Roveri

13th International Conference on Automated Planning & Scheduling, June 10 2003, Trento (Italy)

Model Checking: A Hands-On Introduction—June 10 2003, Trento (Italy)

– p. 18

Description languages for Kripke Models



A Kripke model is usually presented using a structured programming language.

Each component is presented by specifying

- state variables: determine the state space S and the labeling L .
- initial values for state variables: determine the set of initial states I .
- instructions: determine the transition relation R .

Components can be combined via

- synchronous composition,
- asynchronous composition.

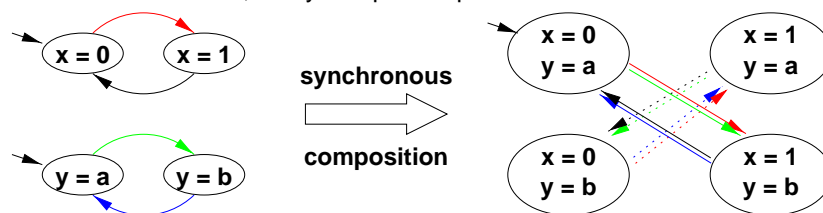
State explosion problem in model checking:

- linear in model size, but model is exponential in number of components.

Synchronous Composition



- Components evolve in parallel.
- At each time instant, every component performs a transition.

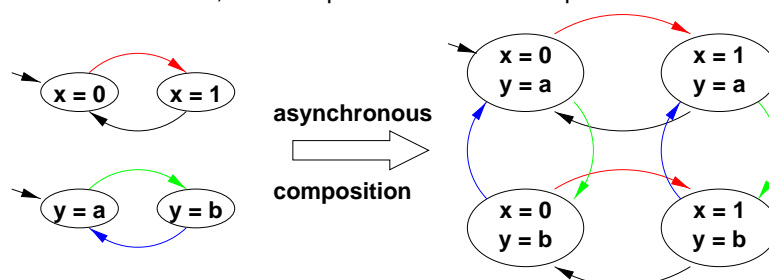


- Typical example: sequential hardware circuits.
- Synchronous composition is the default in NuSMV.

Asynchronous Composition



- Interleaving of evolution of components.
- At each time instant, one component is selected to perform a transition.



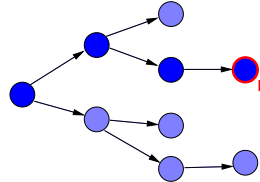
- Typical example: communication protocols.
- Asynchronous composition can be represented with NuSMV processes.

Properties of Reactive Systems (I)



Safety properties:

- nothing bad ever happens
 - deadlock: two processes waiting for input from each other, the system is unable to perform a transition.
 - no reachable state satisfies a “bad” condition, e.g. never two processes in critical section at the same time
- can be refuted by a finite behaviour
- it is never the case that p .



Model Checking: A Hands-On Introduction—June 10 2003, Trento (Italy)

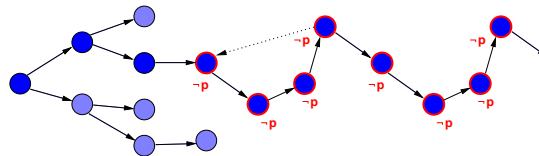
— p. 25

Properties of Reactive Systems (II)



Liveness properties:

- Something desirable will eventually happen
 - whenever a subroutine takes control, it will always return it (sooner or later)
- can be refuted by infinite behaviour
 - a subroutine takes control and never returns it



- an infinite behaviour can be presented as a loop

Model Checking: A Hands-On Introduction—June 10 2003, Trento (Italy)

— p. 26

Temporal Logics



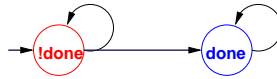
- Express properties of “Reactive Systems”
 - nonterminating behaviours,
 - without explicit reference to time.
- Linear Time Temporal Logic (LTL)
 - interpreted over each path of the Kripke structure
 - linear model of time
 - temporal operators
- Computation Tree Logic (CTL)
 - interpreted over computation tree of Kripke model
 - branching model of time
 - temporal operators plus path quantifiers

Model Checking: A Hands-On Introduction—June 10 2003, Trento (Italy)

— p. 27

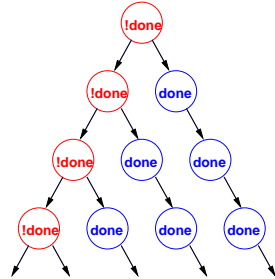
Computation tree vs. computation paths

Consider the following Kripke structure:

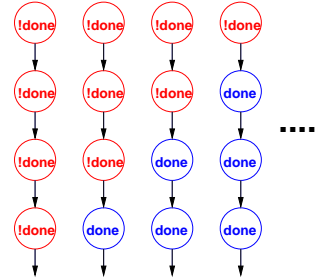


Its execution can be seen as:

• an infinite *computation tree*



• a set of infinite *computation paths*



Linear Time Temporal Logic (LTL)

LTL properties are evaluated over paths, i.e., over infinite, linear sequences of states:

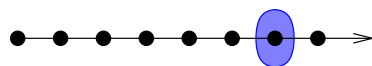
$$s[0] \rightarrow s[1] \rightarrow \dots \rightarrow s[t] \rightarrow s[t+1] \rightarrow \dots$$

LTL provides the following temporal operators:

- “Finally” (or “future”): Fp is true in $s[t]$ iff p is true in **some** $s[t']$ with $t' \geq t$
- “Globally” (or “always”): Gp is true in $s[t]$ iff p is true in **all** $s[t']$ with $t' \geq t$
- “Next”: Xp is true in $s[t]$ iff p is true in $s[t+1]$
- “Until”: pUq is true in $s[t]$ iff
 - q is true in some state $s[t']$ with $t' \geq t$
 - p is true in all states $s[t'']$ with $t \leq t'' < t'$

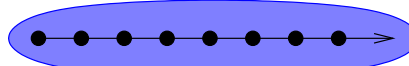
LTL

finally P



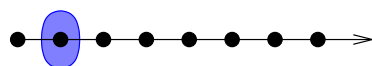
$F P$

globally P



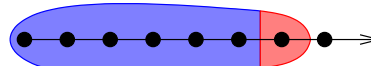
$G P$

next P



$X P$

P until q



$P U q$

LTL: Examples



- Liveness: “if input, then eventually output”

$$G(\text{input} \rightarrow F\text{output})$$

- Strong fairness: “infinitely send implies infinitely recv.”

$$GF\text{send} \rightarrow GF\text{recv}$$

- Weak until: “no output before input”

$$\neg\text{output } W \text{ input}$$

$$\text{where } p W q \leftrightarrow (p U q \vee Gp)$$

Computation Tree Logic (CTL)

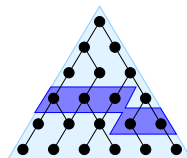


- CTL properties are evaluated over trees.
- Every temporal operator (F, G, X, U) preceded by a path quantifier (A or E).
- Universal modalities (AF, AG, AX, AU): the temporal formula is true in **all** the paths starting in the current state.
- Existential modalities (EF, EG, EX, EU): the temporal formula is true in **some** of the paths starting in the current state.

CTL

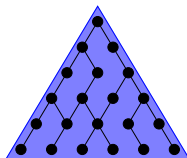


finally P



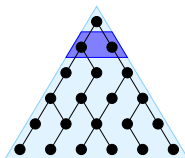
$AF P$

globally P



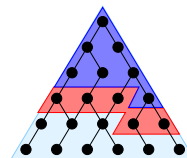
$AG P$

next P

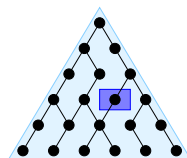


$AX P$

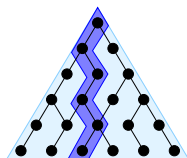
P until q



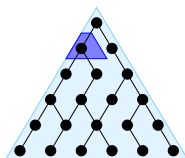
$A[P U q]$



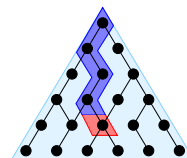
$EF P$



$EG P$



$EX P$



$E[P U q]$

CTL



- Some dualities:

$$AGp \leftrightarrow \neg EF\neg p$$

$$AFp \leftrightarrow \neg EG\neg p$$

$$AXp \leftrightarrow \neg EX\neg p$$

- Example: specifications for the mutual exclusion problem.

$$AG\neg(C_1 \wedge C_2)$$

mutual exclusion

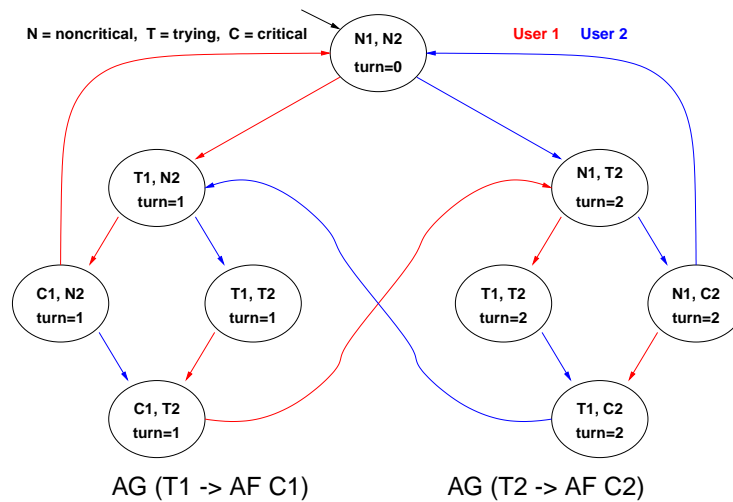
$$AG(T_1 \rightarrow AF C_1)$$

liveness

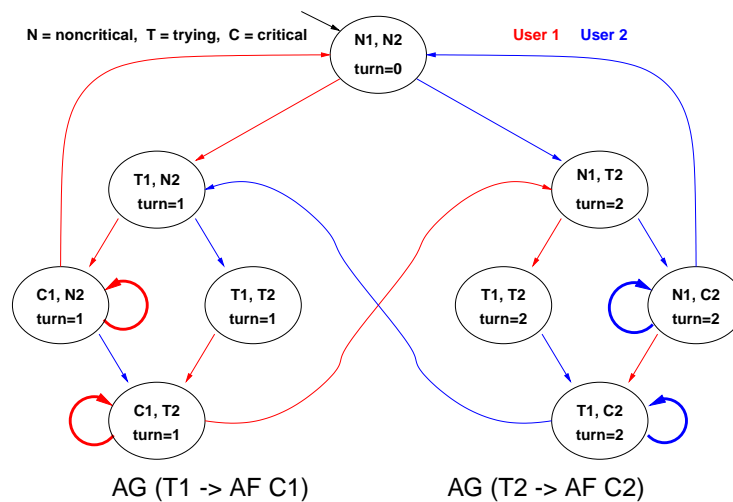
$$AG(N_1 \rightarrow EX T_1)$$

non-blocking

The need for fairness conditions



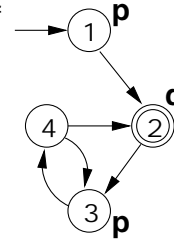
The need for fairness conditions



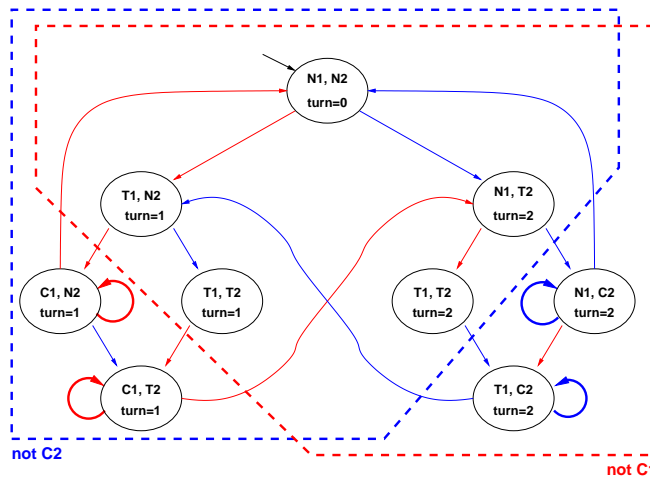
Fair Kripke models



- Intuitively, fairness conditions are used to eliminate behaviours where a condition never holds
 - e.g. once a process is in critical section, it never exits
- Formally, a Fair Kripke model (S, R, I, AP, L, F) consists of
 - a set of states S ;
 - a set of initial states $I \subseteq S$;
 - a set of transitions $R \subseteq S \times S$;
 - a set of atomic propositions AP ;
 - a labeling $L \subseteq S \times AP$;
 - \Rightarrow a set of fairness conditions $F = \{f_1, \dots, f_n\}$, with $f_i \subseteq S$.
- Fair path: at least one state for each f_i occurs in the path an infinite number of times.
- Fair state: a state from which at least one fair path originates.



Fairness: $\{\{\text{not } C1\}, \{\text{not } C2\}\}$

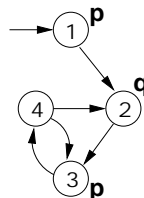


Model Checking



Model Checking is a formal verification technique where...

- ...the system is represented as Finite State Machine



- ...the properties are expressed as temporal logic formulae
 - LTL: $\mathbf{G}(p \rightarrow \mathbf{F}q)$
 - CTL: $\mathbf{AG}(p \rightarrow \mathbf{AF}q)$
- ...the model checking algorithm checks whether all the executions of the model satisfy the formula.

The Main Problem: State Space Explosion



- The bottleneck:
 - Exhaustive analysis may require to store all the states of the Kripke structure
 - The state space may be exponential in the number of components
 - State Space Explosion: too much memory required
- Symbolic Model Checking:
 - Symbolic representation
 - Different search algorithms

Symbolic Model Checking

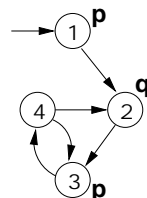


- Symbolic representation:
 - manipulation of *sets of states* (rather than single states);
 - sets of states represented by formulae in propositional logic;
 - set cardinality not directly correlated to size
 - expansion of *sets of transitions* (rather than single transitions);
 - two main symbolic techniques:
 - Binary Decision Diagrams (BDDs)
 - Propositional Satisfiability Checkers (SAT solvers)
- Different model checking algorithms:
 - Fix-point Model Checking (historically, for CTL)
 - Bounded Model Checking (historically, for LTL)
 - Invariant Checking

CTL Model Checking: Example



Consider a simple system and a specification:

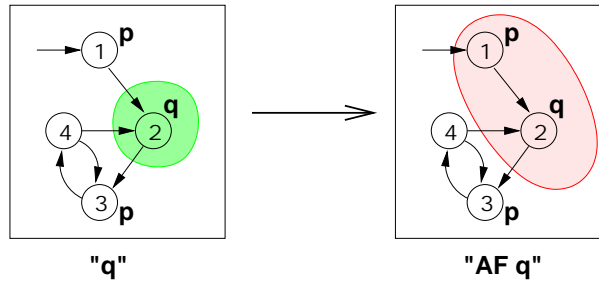


$AG(p \rightarrow AFq)$

Idea:

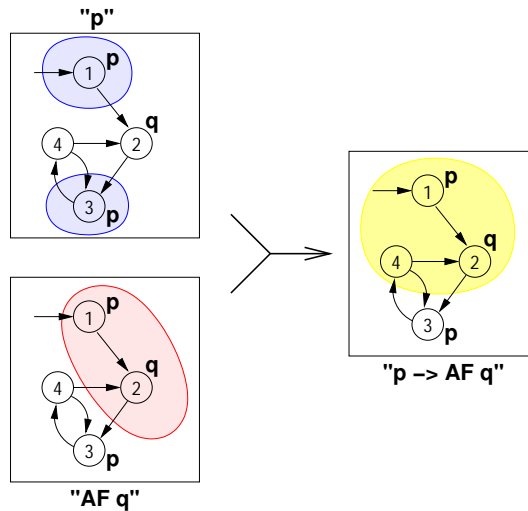
- construct the set of states where the formula holds
- proceeding “bottom-up” on the structure of the formula
- $q, AFq, p, p \rightarrow AF q, AG(p \rightarrow AF q)$

CTL Model Checking: Example

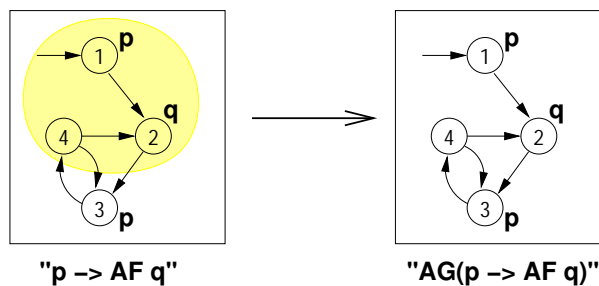


AF q is the union of **q**, **AX q**, **AX AX q**, ...

CTL Model Checking: Example



CTL Model Checking: Example



The set of states where the formula holds is empty!
Counterexample reconstruction is based on the intermediate sets.

Fix-Point Symbolic Model Checking



Model Checking Algorithm for CTL formulae based on fix-point computation:

- traverse formula structure, for each subformula build set of satisfying states; compare result with initial set of states.
- boolean connectives: apply corresponding boolean operation;
- on $AX \Phi$, apply preimage computation
 - $\forall s'. (\mathcal{T}(s, s') \rightarrow \Phi(s'))$
- on $AF \Phi$, compute least fixpoint using
 - $AF \Phi \leftrightarrow (\Phi \vee AX AF \Phi)$
- on $AG \Phi$, compute greatest fixpoint using
 - $AG \Phi \leftrightarrow (\Phi \wedge AX AG \Phi)$

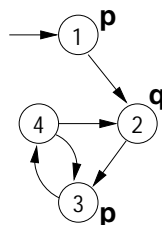
Bounded Model Checking

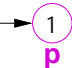


Key ideas:

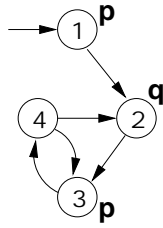
- looks for counter-example paths of increasing length k
 - oriented to finding bugs
- for each k , builds a boolean formula that is satisfiable iff there is a counter-example of length k
 - can be expressed using $k \cdot |s|$ variables
 - formula construction is not subject to state explosion
- satisfiability of the boolean formulas is checked using a *SAT procedure*
 - can manage complex formulae on several 100K variables
 - returns satisfying assignment (i.e., a counter-example)

Bounded Model Checking: Example



- Formula: $G(p \rightarrow Fq)$
- Negated Formula (violation): $F(p \ \& \ G \ ! \ q)$
- $k = 0$: 
- No counter-example found.

Bounded Model Checking: Example

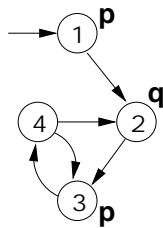


• Formula: $G(p \rightarrow Fq)$

• $k = 1$:

• No counter-example found.

Bounded Model Checking: Example

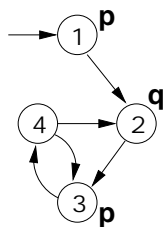


• Formula: $G(p \rightarrow Fq)$

• $k = 2$:

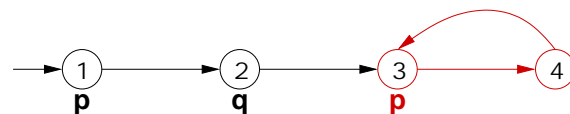
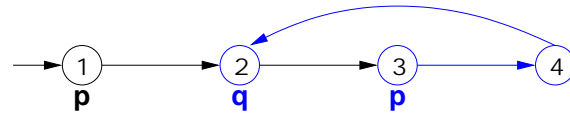
• No counter-example found.

Bounded Model Checking: Example



• Formula: $G(p \rightarrow Fq)$

• $k = 3$:



• The 2nd trace is a counter-example!

Bounded Model Checking



● **Bounded Model Checking:**

Given a FSM $\mathcal{M} = \langle \mathcal{S}, \mathcal{I}, \mathcal{T} \rangle$, an LTL property ϕ and a bound $k \geq 0$:

$$\mathcal{M} \models_k \phi$$

● This is equivalent to the satisfiability problem on formula:

$$[[\mathcal{M}, \phi]]_k \equiv [[\mathcal{M}]]_k \wedge [[\phi]]_k$$

where:

- $[[\mathcal{M}]]_k$ is a k -path compatible with \mathcal{I} and \mathcal{T} :

$$\mathcal{I}(s_0) \wedge \mathcal{T}(s_0, s_1) \wedge \dots \wedge \mathcal{T}(s_{k-1}, s_k)$$

- $[[\phi]]_k$ says that the k -path satisfies ϕ

Bounded Model Checking: Examples



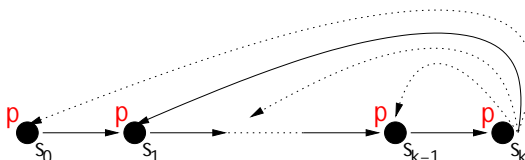
● $\phi = F p$

$$[[F p]]_k = \bigvee_{i=0}^k p(s_i)$$



● $\phi = G p$

$$[[G p]]_k = \bigvee_{i=0}^k \left(\mathcal{T}(s_k, s_i) \wedge \bigwedge_{i=0}^k p(s_i) \right)$$

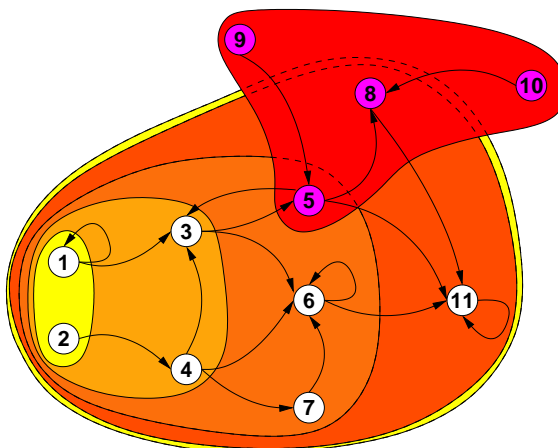


Symbolic Model Checking of Invariants



Checking invariant properties (e.g. **AG ! bad** is a reachability problem):

- is there a reachable state that is also a bad state (⊗)?

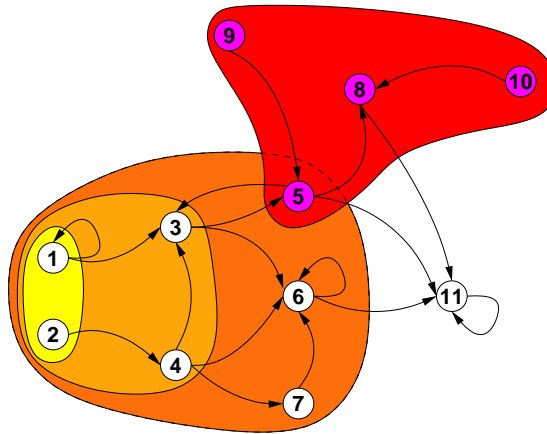


On the fly Checking of Invariants



Anticipate bug detection:

- at each layer, check if a new state is a bug



Model Checking: A Hands-On Introduction—June 10 2003, Trento (Italy)

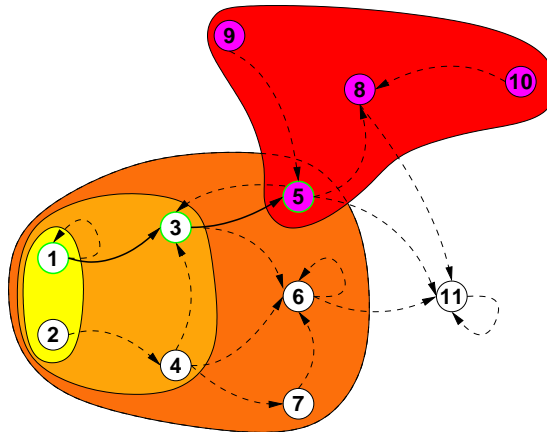
— p. 55

On the fly Checking of Invariants: Counterexamples



If a bug is found,

- a counterexample can be reconstructed proceeding backwards



Model Checking: A Hands-On Introduction—June 10 2003, Trento (Italy)

— p. 56

Inductive Reasoning on Invariants



1. If all the initial states are good,
 2. and if from any good state we only go to good states
- ⇒ then we can conclude that the system is correct for all reachable states.

Model Checking: A Hands-On Introduction—June 10 2003, Trento (Italy)

— p. 57



Part 3 - The NuSMV Model Checker

— *Model Checking* —
A Hands-On Introduction
A. Cimatti, M. Pistore, and M. Roveri

13th International Conference on Automated Planning & Scheduling, June 10 2003, Trento (Italy)

Model Checking: A Hands-On Introduction— June 10 2003, Trento (Italy)

— p. 58

Introduction



- ☞ NuSMV is a symbolic model checker developed by ITC-IRST and UniTN with the collaboration of CMU and UniGE.
- ☞ The NuSMV project aims at the development of a state-of-the-art model checker that:
 - is robust, open and customizable;
 - can be applied in technology transfer projects;
 - can be used as research tool in different domains.
- ☞ NuSMV is *OpenSource*:
 - developed by a distributed community,
 - “Free Software” license.

Model Checking: A Hands-On Introduction— June 10 2003, Trento (Italy)

— p. 59

History: NuSMV 1



- NuSMV is a reimplement and extension of SMV.
- ☞ NuSMV started in 1998 as a joint project between ITC-IRST and CMU:
 - the starting point: SMV version 2.4.4.
 - SMV is the first BDD-based symbolic model checker (McMillan, 90).
 - ☞ NuSMV version 1 has been released in July 1999.
 - limited to BDD-based model checking
 - extends and upgrades SMV along three dimensions:
 - functionalities (LTL, simulation)
 - architecture
 - implementation
 - ☞ Results:
 - used for teaching courses and as basis for several PhD theses
 - interest by industrial companies and academics

Model Checking: A Hands-On Introduction— June 10 2003, Trento (Italy)

— p. 60

History: NuSMV 2



- ☞ The NuSMV 2 project started in September 2000 with the following goals:
 - Introduction of SAT-based model checking
 - OpenSource licensing
 - Larger team (Univ. of Trento, Univ. of Genova, ...)
- ☞ NuSMV 2 has been released in November 2001.
 - first freely available model checker that combines BDD-based and SAT-based techniques
 - extended functionalities wrt NuSMV 1 (cone of influence, improved conjunctive partitioning, multiple FSM management)
- ☞ Results: in the first two months:
 - more than 60 new registrations of NuSMV users
 - more than 300 downloads

OpenSource License

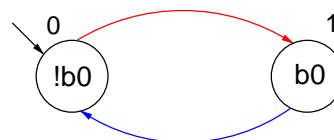


- The idea of OpenSource:
- The System is developed by a distributed community
 - Notable examples: Netscape, Apache, Linux
 - *Potential* benefits: shared development efforts, faster improvements...
- Aim: provide a *publicly available, state-of-the-art* symbolic model checker.
- *publicly available*: free usage in research and commercial applications
 - *state of the art*: improvements should be made freely available
- Distribution license for NuSMV 2: *GNU Lesser General Public License* (LGPL):
- anyone can freely download, copy, use, modify, and redistribute NuSMV 2
 - any modification and extension should be made publicly available under the terms of LGPL (“copyleft”)

The first SMV program



```
MODULE main
  VAR
    b0 : boolean;
  ASSIGN
    init(b0) := 0;
    next(b0) := !b0;
```



- An SMV program consists of:
- ☞ Declarations of the state variables (*b0* in the example); the state variables determine the state space of the model.
 - ☞ Assignments that define the valid initial states (*init(b0) := 0*).
 - ☞ Assignments that define the transition relation (*next(b0) := !b0*).

Declaring state variables



The SMV language provides booleans, enumerative and bounded integers as data types:

boolean:

```
VAR
  x : boolean;
```

enumerative:

```
VAR
  st : {ready, busy, waiting, stopped};
```

bounded integers (intervals):

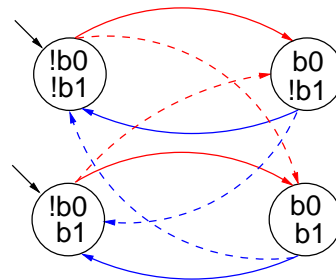
```
VAR
  n : 1..8;
```

Adding a state variable



```
MODULE main
VAR
  b0 : boolean;
  b1 : boolean;

ASSIGN
  init(b0) := 0;
  next(b0) := !b0;
```



Remarks:

- ☞ The new state space is the cartesian product of the ranges of the variables.
- ☞ Synchronous composition between the “subsystems” for b0 and b1.



Declaring the set of initial states



- ☞ For each variable, we constrain the values that it can assume in the *initial states*.

```
init(<variable>) := <simple_expression> ;
```

- ☞ <simple_expression> must evaluate to values in the domain of <variable>.
- ☞ If the initial value for a variable is not specified, then the variable can initially assume any value in its domain.

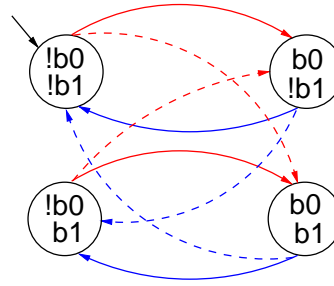
Declaring the set of initial states



```
MODULE main
  VAR
    b0 : boolean;
    b1 : boolean;

  ASSIGN
    init(b0) := 0;
    next(b0) := !b0;

    init(b1) := 0;
```



Expressions



Arithmetic operators:

+ - * / mod - (unary)

Comparison operators:

= != > < <= >=

Logic operators:

& | xor ! (not) -> <->

Conditional expression:

```
case
  c1 : e1;
  c2 : e2;          if c1 then e1 else if c2 then e2 else if ... else
  ...
  1  : en;          en
esac
```

Set operators:

{v1, v2, ..., vn} (enumeration) in (set inclusion) union (set union)

Expressions



- Expressions in SMV do not necessarily evaluate to one value. In general, they can represent a set of possible values.

```
init(var) := {a,b,c} union {x,y,z} ;
```

- The meaning of := in assignments is that the lhs can assume non-deterministically a value in the set of values represented by the rhs.
- A constant c is considered as a syntactic abbreviation for $\{c\}$ (the singleton containing c).

Declaring the transition relation



- ☞ The transition relation is specified by constraining the values that variables can assume in the *next state*.

```
next(<variable>) := <next_expression> ;
```

- ☞ <next_expression> must evaluate to values in the domain of <variable>.

- ☞ <next_expression> depends on “current” and “next” variables:

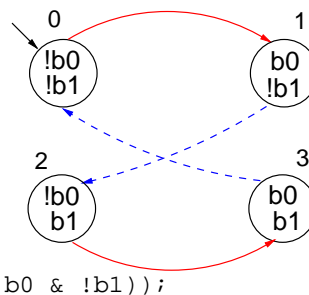
```
next(a) := { a, a+1 } ;  
next(b) := b + (next(a) - a) ;
```

- ☞ If no `next()` assignment is specified for a variable, then the variable can evolve non deterministically, i.e. it is unconstrained.
Unconstrained variables can be used to model non-deterministic *inputs* to the system.

Declaring the transition relation



```
MODULE main  
VAR  
  b0 : boolean;  
  b1 : boolean;  
  
ASSIGN  
  init(b0) := 0;  
  next(b0) := !b0;  
  
  init(b1) := 0;  
  next(b1) := ((!b0 & b1) | (b0 & !b1));
```



Specifying normal assignments



- ☞ Normal assignments constrain the *current value* of a variable to the current values of other variables.
- ☞ They can be used to model *outputs* of the system.

```
<variable> := <simple_expression> ;
```

- ☞ <simple_expression> must evaluate to values in the domain of the <variable>.

Specifying normal assignments

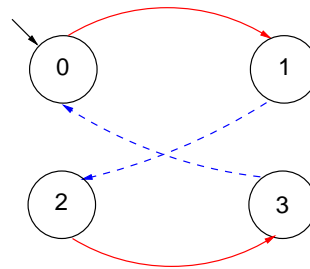


```
MODULE main
VAR
  b0 : boolean;
  b1 : boolean;
  out : 0..3;

ASSIGN
  init(b0) := 0;
  next(b0) := !b0;

  init(b1) := 0;
  next(b1) := ((!b0 & b1) | (b0 & !b1));

  out := b0 + 2*b1;
```



Restrictions on the ASSIGN



For technical reasons, the transition relation must be *total*, i.e., for every state there must be at least one successor state.

In order to guarantee that the transition relation is total, the following restrictions are applied to the SMV programs:

- ☞ Double assignments rule – Each variable may be assigned only once in the program.
- ☞ Circular dependencies rule – A variable cannot have “cycles” in its dependency graph that are not broken by delays.

If an SMV program does not respect these restrictions, an error is reported by NuSMV.

Double assignments rule



Each variable may be assigned only once in the program.

All of the following combinations of assignments are illegal:

```
init(status) := ready;
init(status) := busy;
```

```
next(status) := ready;
next(status) := busy;
```

```
status := ready;
status := busy;
```

```
init(status) := ready;
status := busy;
```

```
next(status) := ready;
status := busy;
```

Circular dependencies rule



A variable cannot have “cycles” in its dependency graph that are not broken by delays.

All the following combinations of assignments are illegal:

```
x := (x + 1) mod 2;  
  
x := (y + 1) mod 2;  
y := (x + 1) mod 2;  
  
next(x) := x & next(x);  
  
next(x) := x & next(y);  
next(y) := y & next(x);
```

The following example is *legal*, instead:

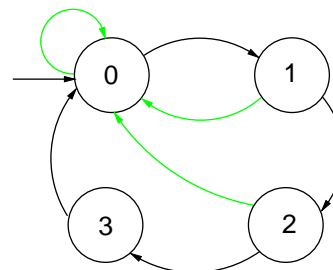
```
next(x) := x & next(y);  
next(y) := y & x;
```

The modulo 4 counter with reset



The counter can be reset by an external “uncontrollable” reset signal.

```
MODULE main  
VAR  
  b0 : boolean;  
  b1 : boolean;  
  reset : boolean;  
  out : 0..3;  
  
ASSIGN  
  init(b0) := 0;  
  next(b0) := case  
    reset = 1 : 0;  
    reset = 0 : !b0;  
  esac;  
  
  init(b1) := 0;  
  next(b1) := case  
    reset : 0;  
    1 : ((!b0 & b1) | (b0 & !b1));  
  esac;  
  
  out := b0 + 2*b1;
```

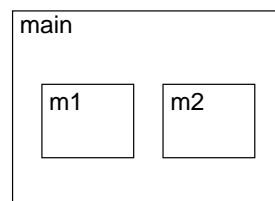


Modules



An SMV program can consist of one or more *module declarations*.

```
MODULE mod  
  VAR out: 0..9;  
  ASSIGN next(out) :=  
    (out + 1) mod 10;  
  
MODULE main  
  VAR m1 : mod;  
  VAR m2 : mod;  
  sum: 0..18;  
  ASSIGN sum := m1.out + m2.out;
```



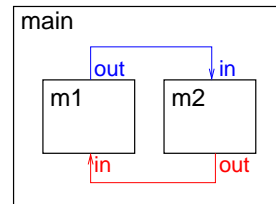
- ☞ Modules are instantiated in other modules. The instantiation is performed inside the VAR declaration of the parent module.
- ☞ In each SMV specification there must be a module `main`. It is the top-most module.
- ☞ All the variables declared in a module instance are visible in the module in which it has been instantiated via the dot notation (e.g., `m1.out`, `m2.out`).

Module parameters



Module declarations may be *parametric*.

```
MODULE mod(in)
  VAR out: 0..9;
  ...
MODULE main
  VAR m1 : mod(m2.out);
      m2 : mod(m1.out);
  ...
```



- ☞ *Formal parameters* (*in*) are substituted with the *actual parameters* (*m2.out, m1.out*) when the module is instantiated.
- ☞ Actual parameters can be any legal expression.
- ☞ Actual parameters are passed by reference.

Example: The modulo 8 counter revisited



```
MODULE counter_cell(tick)

  VAR
    value : boolean;
    done  : boolean;

  ASSIGN
    init(value) := 0;
    next(value) := case
      tick = 0 : value;
      tick = 1 : (value + 1) mod 2;
    esac;

    done := tick & (((value + 1) mod 2) = 0);
```

Remarks:

- ☞ *tick* is the formal parameter of module *counter_cell*.

Example: The modulo 8 counter revisited



```
MODULE main
  VAR
    bit0 : counter_cell(1);
    bit1 : counter_cell(bit0.done);
    bit2 : counter_cell(bit1.done);
    out  : 0..7;

  ASSIGN
    out := bit0.value + 2*bit1.value + 4*bit2.value;
```

Remarks:

- ☞ Module *counter_cell* is instantiated three times.
- ☞ In the instance *bit0*, the formal parameter *tick* is replaced with the actual parameter *1*.
- ☞ When a module is instantiated, all variables/symbols defined in it are preceded by the module instance name, so that they are unique to the instance.

Module hierarchies



A module can contain instances of others modules, that can contain instances of other modules... provided the module references are not circular.

```
MODULE counter_8 (tick)
  VAR
    bit0 : counter_cell(tick);
    bit1 : counter_cell(bit0.done);
    bit2 : counter_cell(bit1.done);
    out  : 0..7;
    done : boolean;
  ASSIGN
    out := bit0.value + 2*bit1.value + 4*bit2.value;
    done := bit2.done;

MODULE counter_512(tick) -- A counter modulo 512
  VAR
    b0 : counter_8(tick);
    b1 : counter_8(b0.done);
    b2 : counter_8(b1.done);
    out : 0..511;
  ASSIGN
    out := b0.out + 8*b1.out + 64*b2.out;
```

Model Checking: A Hands-On Introduction—June 10 2003, Trento (Italy)

– p. 82

Specifications



In the SMV language:

- ☞ Specifications can be added in any module of the program.
- ☞ Each property is verified separately.
- ☞ Different kinds of properties are allowed:
 - Properties on the reachable states
 - *invariants* (INVARSPEC)
 - Properties on the computation paths (*linear time* logics):
 - LTL (LTLSPEC)
 - qualitative characteristics of models (COMPUTE)
 - Properties on the computation tree (*branching time* logics):
 - CTL (SPEC)
 - Real-time CTL (SPEC)

Model Checking: A Hands-On Introduction—June 10 2003, Trento (Italy)

– p. 83

Invariant specifications



☞ Invariant properties are specified via the keyword INVARSPEC:

```
INVARSPEC <simple_expression>
```

☞ Example:

```
MODULE counter_cell(tick)
  ...
MODULE counter_8(tick)
  VAR
    bit0 : counter_cell(tick);
    bit1 : counter_cell(bit0.done);
    bit2 : counter_cell(bit1.done);
    out  : 0..7;
    done : boolean;
  ASSIGN
    out := bit0.value + 2*bit1.value + 4*bit2.value;
    done := bit2.done;

INVARSPEC
  done <-> (bit0.done & bit1.done & bit2.done)
```

Model Checking: A Hands-On Introduction—June 10 2003, Trento (Italy)

– p. 84

LTl specifications



- LTl properties are specified via the keyword `LTLSPEC`:

```
LTLSPEC <ltl_expression>
```

where `<ltl_expression>` can contain the following temporal operators:

```
X _ F _ G _ _ U _
```

- A state in which `out = 3` is eventually reached.

```
LTLSPEC F out = 3
```

- Condition `out = 0` holds until `reset` becomes false.

```
LTLSPEC (out = 0) U (!reset)
```

- Even time a state with `out = 2` is reached, a state with `out = 3` is reached afterwards.

```
LTLSPEC G (out = 2 -> F out = 3)
```

Quantitative characteristics computations



It is possible to compute the minimum and maximum length of the paths between two specified conditions.

- Quantitative characteristics are specified via the keyword `COMPUTE`:

```
COMPUTE  
MIN/MAX [ <simple_expression> , <simple_expression> ]
```

- For instance, the shortest path between a state in which `out = 0` and a state in which `out = 3` is computed with

```
COMPUTE  
MIN [ out = 0 , out = 3 ]
```

- The length of the longest path between a state in which `out = 0` and a state in which `out = 3`.

```
COMPUTE  
MAX [ out = 0 , out = 3 ]
```

CTL properties



- CTL properties are specified via the keyword `SPEC`:

```
SPEC <ctl_expression>
```

where `<ctl_expression>` can contain the following temporal operators:

```
AX _ AF _ AG _ A[_ U _]  
EX _ EF _ EG _ E[_ U _]
```

- It is possible to reach a state in which `out = 3`.

```
SPEC EF out = 3
```

- A state in which `out = 3` is always reached.

```
SPEC AF out = 3
```

- It is always possible to reach a state in which `out = 3`.

```
SPEC AG EF out = 3
```

- Even time a state with `out = 2` is reached, a state with `out = 3` is reached afterwards.

```
SPEC AG (out = 2 -> AF out = 3)
```

Bounded CTL specifications



NuSMV provides *bounded CTL* (or *real-time CTL*) operators.

- ☞ There is no state that is reachable in 3 steps where $\text{out} = 3$ holds.

```
SPEC
!EBF 0..3 out = 3
```

- ☞ A state in which $\text{out} = 3$ is reached in 2 steps.

```
SPEC
ABF 0..2 out = 3
```

- ☞ From any reachable state, a state in which $\text{out} = 3$ is reached in 3 steps.

```
SPEC
AG ABF 0..3 out = 3
```

Fairness Constraints



Let us consider again the counter with reset.

- ☞ The specification $\text{AF out} = 1$ is not verified.
- ☞ On the path where *reset* is always 1, then the system loops on a state where $\text{out} = 0$, since the counter is always reset:

```
reset = 1,1,1,1,1,1,1...
out = 0,0,0,0,0,0,0...
```

- ☞ Similar considerations hold for the property $\text{AF out} = 2$. For instance, the sequence:

```
reset = 0,1,0,1,0,1,0...
```

generates the loop:

```
out = 0,1,0,1,0,1,0...
```

which is a counterexample to the given formula.

Fairness Constraints



- ☞ NuSMV allows to specify *fairness* constraints.
- ☞ Fairness constraints are formulas which are assumed to be true infinitely often in all the execution paths of interest.
- ☞ During the verification of properties, NuSMV considers path quantifiers to apply only to fair paths.
- ☞ Fairness constraints are specified as follows:

```
FAIRNESS <simple_expression>
```

Fairness Constraints



☞ With the fairness constraint

```
FAIRNESS
  out = 1
```

we restrict our analysis to paths in which the property `out = 1` is true infinitely often.

- ☞ The property $AF \text{ out} = 1$ under this fairness constraint is now verified.
- ☞ The property $AF \text{ out} = 2$ is still not verified.
- ☞ Adding the fairness constraint `out = 2`, then also the property $AF \text{ out} = 2$ is verified.

The DEFINE declaration



In the following example, the values of variables `out` and `done` are defined by the values of the other variables in the model.

```
MODULE main          -- counter_8
VAR
  b0  : boolean;
  b1  : boolean;
  b2  : boolean;
  out : 0..8;
  done : boolean;

ASSIGN
  init(b0) := 0;
  init(b1) := 0;
  init(b2) := 0;

  next(b0) := !b0;
  next(b1) := (!b0 & b1) | (b0 & !b1);
  next(b2) := ((b0 & b1) & !b2) | (!(b0 & b1) & b2);

  out := b0 + 2*b1 + 4*b2;
  done := b0 & b1 & b2;
```

The DEFINE declaration



DEFINE declarations can be used to define *abbreviations*:

```
MODULE main          -- counter_8
VAR
  b0  : boolean;
  b1  : boolean;
  b2  : boolean;

ASSIGN
  init(b0) := 0;
  init(b1) := 0;
  init(b2) := 0;

  next(b0) := !b0;
  next(b1) := (!b0 & b1) | (b0 & !b1);
  next(b2) := ((b0 & b1) & !b2) | (!(b0 & b1) & b2);

DEFINE
  out := b0 + 2*b1 + 4*b2;
  done := b0 & b1 & b2;
```

The DEFINE declaration



- ☞ The syntax of `DEFINE` declarations is the following:

```
DEFINE <id> := <simple_expression> ;
```

- ☞ They are similar to macro definitions.
- ☞ No new state variable is created for defined symbols (hence, no added complexity to model checking).
- ☞ Each occurrence of a defined symbol is replaced with the body of the definition.

Arrays



The SMV language provides also the possibility to define *arrays*.

```
VAR
```

```
  x : array 0..10 of boolean;
```

```
  y : array 2..4 of 0..10;
```

```
  z : array 0..10 of array 0..5 of {red, green, orange};
```

```
ASSIGN
```

```
  init(x[5]) := 1;
```

```
  init(y[2]) := {0,2,4,6,8,10};
```

```
  init(z[3][2]) := {green, orange};
```

- ☞ Remark: Array indexes in SMV *must be constants*.

Records



Records can be defined as modules without parameters and assignments.

```
MODULE point
```

```
  VAR x: -10..10;
```

```
      y: -10..10;
```

```
MODULE circle
```

```
  VAR center: point;
```

```
      radius: 0..10;
```

```
MODULE main
```

```
  VAR c: circle;
```

```
  ASSIGN
```

```
    init(c.center.x) := 0;
```

```
    init(c.center.y) := 0;
```

```
    init(c.radius)   := 5;
```


The constraint style of model specification

The following SMV program:

```
MODULE main
VAR request : boolean;
    state    : {ready,busy};
ASSIGN
  init(state) := ready;
  next(state) := case
    state = ready & request : busy;
    1                        : {ready,busy};
  esac;
```

can be alternatively defined in a *constraint style*, as follows:

```
MODULE main
VAR request : boolean;
    state    : {ready,busy};
INIT
  state = ready
TRANS
  (state = ready & request) -> next(state) = busy
```

The constraint style of model specification

☞ The SMV language allows for specifying the model by defining constraints on:

● the *states*:

```
INVAR <simple_expression>
```

● the *initial states*:

```
INIT <simple_expression>
```

● the *transitions*:

```
TRANS <next_expression>
```

☞ There can be zero, one, or more constraints in each module, and constraints can be mixed with assignments.

☞ Any propositional formula is allowed in constraints.

☞ Very useful for writing translators from other languages to NuSMV.

☞ $\text{INVAR } p$ is equivalent to $\text{INIT } p$ and $\text{TRANS } \text{next}(p)$, but is more efficient.

☞ Risk of defining *inconsistent models* ($\text{INIT } p \ \& \ !p$).

Assignments versus constraints

☞ Any ASSIGN-based specification can be easily rewritten as an equivalent constraint-based specification:

```
ASSIGN
  init(state) := {ready,busy};   INIT state in {ready,busy}
  next(state) := ready;          TRANS next(state) = ready
  out := b0 + 2*b1;              INVAR out = b0 + 2*b1
```

☞ The converse is not true: constraint

```
TRANS
  next(b0) + 2*next(b1) + 4*next(b2) =
  (b0 + 2*b1 + 4*b2 + tick) mod 8
```

cannot be easily rewritten in terms of ASSIGNS.

Assignments versus constraints



☞ Models written in **assignment style**:

- by construction, there is always *at least one initial state*;
- by construction, all states have *at least one next state*;
- *non-determinism is apparent* (unassigned variables, set assignments...).

☞ Models written in **constraint style**:

- INIT constraints *can be inconsistent*:
 - inconsistent model: no initial state,
 - any specification (also SPEC 0) is vacuously true.
- TRANS constraints *can be inconsistent*:
 - the transition relation is not total (there are deadlock states),
 - NuSMV detects and reports this case.

- *non-determinism is hidden* in the constraints:

```
TRANS (state = ready & request) -> next(state) = busy
```

Synchronous composition

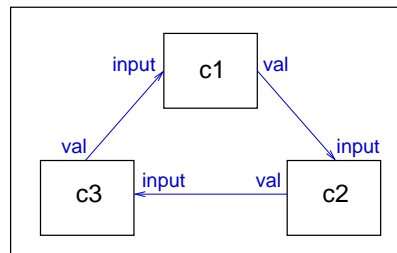


☞ By default, composition of modules is **synchronous**:

all modules move at each step.

```
MODULE cell(input)
  VAR
    val : {red, green, blue};
  ASSIGN
    next(val) := {val, input};

MODULE main
  VAR
    c1 : cell(c3.val);
    c2 : cell(c1.val);
    c3 : cell(c2.val);
```



Synchronous composition



A possible execution:

step	c1.val	c2.val	c3.val
0	red	green	blue
1	red	red	green
2	green	red	green
3	green	red	green
4	green	red	red
5	red	green	red
6	red	red	red
7	red	red	red
8	red	red	red
9	red	red	red
10	red	red	red

Asynchronous composition



- ☞ **Asynchronous** composition can be obtained using keyword `process`.
- ☞ In asynchronous composition *one process moves at each step*.
- ☞ Boolean variable `running` is defined in each process:
 - it is true when that process is selected;
 - it can be used to guarantee a fair scheduling of processes.

```
MODULE cell(input)
  VAR
    val : {red, green, blue};
  ASSIGN
    next(val) := {val, input};
  FAIRNESS
    running

MODULE main
  VAR
    c1 : process cell(c3.val);
    c2 : process cell(c1.val);
    c3 : process cell(c2.val);
```

Model Checking: A Hands-On Introduction— June 10 2003, Trento (Italy)

– p. 103

Asynchronous composition



A possible execution:

<i>step</i>	<i>running</i>	<i>c1.val</i>	<i>c2.val</i>	<i>c3.val</i>
0	-	red	green	blue
1	c2	red	red	blue
2	c1	blue	red	blue
3	c1	blue	red	blue
4	c2	blue	red	blue
5	c3	blue	red	red
6	c2	blue	blue	red
7	c1	blue	blue	red
8	c1	red	blue	red
9	c3	red	blue	blue
10	c3	red	blue	blue

Model Checking: A Hands-On Introduction— June 10 2003, Trento (Italy)

– p. 104

NuSMV resources



- ☞ NuSMV home page:
 - <http://nusmv.irst.itc.it/>
- ☞ Mailing lists:
 - nusmv-users@irst.itc.it (public discussions)
 - nusmv-announce@irst.itc.it (announces of new releases)
 - nusmv@irst.itc.it (the development team)
 - to subscribe: <http://nusmv.irst.itc.it/mail.html>
- ☞ Course notes and slides:
 - <http://nusmv.irst.itc.it/courses/>

Model Checking: A Hands-On Introduction— June 10 2003, Trento (Italy)

– p. 105

NuSMV 2.1 User Manual

**Roberto Cavada, Alessandro Cimatti,
Emanuele Olivetti, Marco Pistore,
and Marco Roveri**

IRST - Via Sommarive 18, 38055 Povo (Trento) – Italy

Email: nusmv@irst.itc.it

This document is part of the distribution package of the NuSMV model checker, available at <http://nusmv.irst.itc.it>.

Parts of this documents have been taken from "The SMV System - Draft", by K. McMillan, available at <http://www.cs.cmu.edu/~modelcheck/smv/smvmanual.r2.2.ps>.

Copyright © 1998-2002 by CMU and ITC-irst.

Table of Contents

1	Introduction	1
2	Tutorial	2
2.1	Examples	2
2.2	Simulation	6
2.3	CTL model checking	10
2.4	LTL model checking	13
2.5	Bounded Model Checking	14
3	Syntax	21
3.1	Expressions	21
3.1.1	Simple Expressions	21
3.1.1.1	Case Expressions	22
3.1.1.2	Set Expressions	22
3.1.2	Next Expressions	22
3.2	Definition of the FSM	23
3.2.1	State Variables	23
3.2.1.1	Type Specifiers	23
3.2.2	Input Variables	24
3.2.3	ASSIGN declarations	24
3.2.4	TRANS declarations	25
3.2.5	INIT declarations	25
3.2.6	INVAR declarations	25
3.2.7	DEFINE declarations	25
3.2.8	ISA declarations	26
3.2.9	MODULE declarations	26
3.2.10	Identifiers	27
3.2.11	The <code>main</code> module	28
3.2.12	Processes	28
3.2.13	FAIRNESS declarations	29
3.3	Specifications	29
3.3.1	CTL Specifications	29
3.3.2	LTL Specifications	30
3.3.3	Real Time CTL Specifications and Computations	31
4	Running NuSMV interactively	32
4.1	Model Reading and Building	33
4.2	Commands for Checking Specifications	36
4.3	Commands for Bounded Model Checking	39
4.4	Simulation Commands	46
4.5	Traces Inspection Commands	48
4.6	Interface to the DD Package	48
4.7	Administration Commands	51
4.8	Other Environment Variables	57
5	Running NuSMV batch	58
	Bibliography	60

Appendix A	Compatibility with CMU SMV	61
Command Index		63
Variable Index		64
Index		65

1 Introduction

NuSMV is a symbolic model checker originated from the reengineering, reimplementing and extension of CMU SMV, the original BDD-based model checker developed at CMU [McMil93]. The NuSMV project aims at the development of a state-of-the-art symbolic model checker, designed to be applicable in technology transfer projects: it is a well structured, open, flexible and documented platform for model checking, and is robust and close to industrial systems standards [CCGR00].

Version 1 of NuSMV basically implements BDD-based symbolic model checking. Version 2 of NuSMV (NuSMV2 in the following) inherits all the functionalities of the previous version, and extends them in several directions [CCG+02]. The main novelty in NuSMV2 is the integration of model checking techniques based on propositional satisfiability (SAT) [BCCZ99]. SAT-based model checking is currently enjoying a substantial success in several industrial fields, and opens up new research directions. BDD-based and SAT-based model checking are often able to solve different classes of problems, and can therefore be seen as complementary techniques.

Starting from NuSMV2, we are also adopting a new development and license model. NuSMV2 is distributed with an OpenSource license (see <http://www.opensource.org>), that allows anyone interested to freely use the tool and to participate in its development. The aim of the NuSMV OpenSource project is to provide to the model checking community a common platform for the research, the implementation, and the comparison of new symbolic model checking techniques. Since the release of NuSMV2, the NuSMV team has received code contributions for different parts of the system. Several research institutes and commercial companies have expressed interest in collaborating to the development of NuSMV.

The main features of NuSMV are the following:

- **Functionalities.** NuSMV allows for the representation of synchronous and asynchronous finite state systems, and for the analysis of specifications expressed in Computation Tree Logic (CTL) and Linear Temporal Logic (LTL), using BDD-based and SAT-based model checking techniques. Heuristics are available for achieving efficiency and partially controlling the state explosion. The interaction with the user can be carried on with a textual interface, as well as in batch mode.
- **Architecture.** A software architecture has been defined. The different components and functionalities of NuSMV have been isolated and separated in modules. Interfaces between modules have been provided. This reduces the effort needed to modify and extend NuSMV.
- **Quality of the implementation.** NuSMV is written in ANSI C, is POSIX compliant, and has been debugged with Purify in order to detect memory leaks. Furthermore, the system code is thoroughly commented. NuSMV uses the state of the art BDD package developed at Colorado University, and provides a general interface for linking with state-of-the-art SAT solvers. This makes NuSMV very robust, portable, efficient, and easy to understand by other people than the developers.

This document is structured as follows.

- In Chapter 2 [Tutorial], page 2 we give an introduction to the usage of the main functionalities of NuSMV.
- In Chapter 3 [Syntax], page 21 we define the syntax of the input language of NuSMV.
- In Chapter 4 [Running NuSMV interactively], page 32 the commands of the interaction shell are described.
- In Chapter 5 [Running NuSMV batch], page 58 we define the batch mode of NuSMV.

NuSMV is available at <http://nusmv.irst.itc.it>.

2 Tutorial

In this chapter we give a short introduction to the usage of the main functionalities of NuSMV. In Section 2.1 [Examples], page 2 we describe the input language of NuSMV by presenting some examples of NuSMV models. Section 2.2 [Simulation], page 6 shows how the user can get familiar with the behavior of a NuSMV model by exploring its possible executions. Section 2.3 [CTL model checking], page 10 and Section 2.4 [LTL model checking], page 13 give an overview of BDD-based model checking, while Section 2.5 [Bounded Model Checking], page 14 presents SAT-based model checking in NuSMV.

2.1 Examples

In this section we describe the input language of NuSMV by presenting some examples of NuSMV models. A complete description of the NuSMV language can be found in Chapter 3 [Syntax], page 21.

The input language of NuSMV is designed to allow for the description of Finite State Machines (FSM from now on) which range from completely synchronous to completely asynchronous, and from the detailed to the abstract. One can specify a system as a synchronous Mealy machine, or as an asynchronous network of nondeterministic processes. The language provides for modular hierarchical descriptions, and for the definition of reusable components. Since it is intended to describe finite state machines, the only data types in the language are finite ones – booleans, scalars and fixed arrays. Static data types can also be constructed.

The primary purpose of the NuSMV input is to describe the transition relation of the FSM; this relation describes the valid evolutions of the state of the FSM. In general, any propositional expression in the propositional calculus can be used to define the transition relation. This provides a great deal of flexibility, and at the same time a certain danger of inconsistency. For example, the presence of a logical contradiction can result in a deadlock – a state or states with no successor. This can make some specifications vacuously true, and makes the description unimplementable. While the model checking process can be used to check for deadlocks, it is best to avoid the problem when possible by using a restricted description style. The NuSMV system supports this by providing a parallel-assignment syntax. The semantics of assignment in NuSMV is similar to that of single assignment data flow language. By checking programs for multiple parallel assignments to the same variable, circular assignments, and type errors, the interpreter insures that a program using only the assignment mechanism is implementable. Consequently, this fragment of the language can be viewed as a description language, or a programming language.

Consider the following simple program in the NuSMV language.

```

MODULE main
VAR
  request : boolean;
  state   : {ready, busy};
ASSIGN
  init(state) := ready;
  next(state) := case
    state = ready & request = 1 : busy;
    1                             : {ready, busy};
  esac;

```

The space of states of the FSM is determined by the declarations of the state variables (in the above example `request` and `state`). The variable `request` is declared to be of (predefined) type `boolean`. This means that it can assume the (integer) values 0 and 1. The variable `state` is a scalar variable, which can take the symbolic values `ready` or `busy`. The following assignment sets the initial value of the variable `state` to `ready`. The initial value of `request` is completely unspecified, i.e. it can be either 0 or 1. The transition relation of the FSM is expressed by

defining the value of variables in the next state (i.e. after each transition), given the value of variables in the current states (i.e. before the transition). The `case` segment sets the next value of the variable `state` to the value `busy` (after the column) if its current value is `ready` and `request` is 1 (i.e. true). Otherwise (the 1 before the column) the next value for `state` can be any in the set `{ready, busy}`. The variable `request` is not assigned. This means that there are no constraints on its values, and thus it can assume any value. `request` is thus an unconstrained input to the system.

The following program illustrates the definition of reusable modules and expressions. It is a model of a three bit binary counter circuit. The order of module definitions in the input file is not relevant.

```

MODULE counter_cell(carry_in)
  VAR
    value : boolean;
  ASSIGN
    init(value) := 0;
    next(value) := (value + carry_in) mod 2;
  DEFINE
    carry_out := value & carry_in;

MODULE main
  VAR
    bit0 : counter_cell(1);
    bit1 : counter_cell(bit0.carry_out);
    bit2 : counter_cell(bit1.carry_out);

```

The FSM is defined by instantiating three times the module type `counter_cell` in the module `main`, with the names `bit0`, `bit1` and `bit2` respectively. The `counter_cell` module has one formal parameter `carry_in`. In the instance `bit0`, this parameter is given the actual value 1. In the instance `bit1`, `carry_in` is given the value of the expression `bit0.carry_out`. This expression is evaluated in the context of the `main` module. However, an expression of the form `'a.b'` denotes component `'b'` of module `'a'`, just as if the module `'a'` were a data structure in a standard programming language. Hence, the `carry_in` of module `bit1` is the `carry_out` of module `bit0`.

The keyword `'DEFINE'` is used to assign the expression `value & carry_in` to the symbol `carry_out`. A definition can be thought of as a variable with value (functionally) depending on the current values of other variables. The same effect could have been obtained as follows (notice that the *current* value of the variable is assigned, rather than the *next* value.):

```

  VAR
    carry_out : boolean;
  ASSIGN
    carry_out := value & carry_in;

```

Defined symbols do not require introducing a new variable, and hence do not increase the state space of the FSM. On the other hand, it is not possible to assign to a defined symbol a value non-deterministically. Another difference between defined symbols and variables is that while the type of variables is declared a priori, for definitions this is not the case.

The previous examples describe synchronous systems, where the assignments statements are taken into account in parallel and simultaneously. NuSMV allows to model asynchronous systems. It is possible to define a collection of parallel processes, whose actions are interleaved, following an asynchronous model of concurrency. This is useful for describing communication protocols, or asynchronous circuits, or other systems whose actions are not synchronized (including synchronous circuits with more than one clock region).

The following program represents a ring of three asynchronous inverting gates.

```

MODULE inverter(input)

```

```

VAR
  output : boolean;
ASSIGN
  init(output) := 0;
  next(output) := !input;

MODULE main
  VAR
    gate1 : process inverter(gate3.output);
    gate2 : process inverter(gate1.output);
    gate3 : process inverter(gate2.output);

```

Among all the modules instantiated with the `process` keyword, one is nondeterministically chosen, and the assignment statements declared in that process are executed in parallel. It is implicit that if a given variable is not assigned by the process, then its value remains unchanged. Because the choice of the next process to execute is non-deterministic, this program models the ring of inverters independently of the speed of the gates.

We remark that the system is not forced to eventually choose a given process to execute. As a consequence the output of a given gate may remain constant, regardless of its input. In order to force a given process to execute infinitely often, we can use a *fairness constraint*. A fairness constraint restricts the attention of the model checker to only those execution paths along which a given formula is true infinitely often. Each process has a special variable called `running` which is 1 if and only if that process is currently executing.

By adding the declaration:

```

FAIRNESS
  running

```

to the module `inverter`, we can effectively force every instance of `inverter` to execute infinitely often.

An alternative to using processes to model an asynchronous circuit is to allow all gates to execute simultaneously, but to allow each gate to choose non-deterministically to re-evaluate its output or to keep the same output value. Such a model of the inverter ring would look like the following:

```

MODULE inverter(input)
  VAR
    output : boolean;
  ASSIGN
    init(output) := 0;
    next(output) := (!input) union output;

MODULE main
  VAR
    gate1 : inverter(gate3.output);
    gate2 : inverter(gate1.output);
    gate3 : inverter(gate2.output);

```

The `union` operator (set union) coerces its arguments to singleton sets as necessary. Thus, the next `output` of each gate can be either its current `output`, or the negation of its current `input` – each gate can choose non-deterministically whether to delay or not. As a result, the number of possible transitions from a given state can be as 2^n , where n is the number of gates. This sometimes (but not always) makes it more expensive to represent the transition relation.

We remark that in this case we cannot force the inverters to be effectively active infinitely often using a fairness declaration. In fact, a valid scenario for the synchronous model is the one where all the inverters are idle and assign to the next `output` the current value of `output`.

In the declaration of the `main` module it is possible to mix `process` instances and normal instances of modules. Let us consider the following example:

```
MODULE main
  VAR p0: m;
      p1: m;
      p2: process m;
      p3: process m;

MODULE m
  VAR x: 0..3;
  ASSIGN
    init(x) := 0;
    next(x) := (x+1) mod 4;
```

In the example, four copies of module `m` are instantiated. Two of them (`p0` and `p1`) are asynchronous processes, while the other two (`p2` and `p3`) are normal, synchronous modules. The non-process modules instantiated inside `main` are considered part of a special "top-level" process that runs asynchronously with respect to the other processes. That is, during each transition of the model, either process `p0` is active (in which case the value of `p0.x` is updated), or process `p1` is active (in which case the value of `p1.x` is updated), or the top-level process is active (in which case the values of `p2.x` and `p3.x` are updated).

The following program is another example of asynchronous model. It uses a variable semaphore to implement mutual exclusion between two asynchronous processes. Each process has four states: `idle`, `entering`, `critical` and `exiting`. The `entering` state indicates that the process wants to enter its critical region. If the variable `semaphore` is 0, it goes to the `critical` state, and sets `semaphore` to 1. On exiting its critical region, the process sets `semaphore` to 0 again.

```
MODULE main
  VAR
    semaphore : boolean;
    proc1      : process user(semaphore);
    proc2      : process user(semaphore);
  ASSIGN
    init(semaphore) := 0;

MODULE user(semaphore)
  VAR
    state : {idle, entering, critical, exiting};
  ASSIGN
    init(state) := idle;
    next(state) :=
      case
        state = idle           : {idle, entering};
        state = entering & !semaphore : critical;
        state = critical       : {critical, exiting};
        state = exiting        : idle;
        1                       : state;
      esac;
    next(semaphore) :=
      case
        state = entering : 1;
        state = exiting  : 0;
        1                 : semaphore;
```

```

    esac;
FAIRNESS
    running

```

NuSMV allows to specify the FSM directly in terms of propositional formulas. The set of possible initial states is specified as a formula in the current state variables. A state is initial if it satisfies the formula. The transition relation is directly specified as a propositional formula in terms of the *current* and *next* values of the state variables. Any current state/next state pair is in the transition relation if and only if it satisfies the formula.

These two functions are accomplished by the ‘INIT’ and ‘TRANS’ keywords. As an example, here is a description of the three inverter ring using only TRANS and INIT:

```

MODULE main
VAR
    gate1 : inverter(gate3.output);
    gate2 : inverter(gate1.output);
    gate3 : inverter(gate2.output);

MODULE inverter(input)
VAR
    output : boolean;
INIT
    output = 0
TRANS
    next(output) = !input | next(output) = output

```

According to the TRANS declaration, for each inverter, the next value of the `output` is equal either to the negation of the `input`, or to the current value of the `output`. Thus, in effect, each gate can choose non-deterministically whether or not to delay.

Using TRANS and INIT it is possible to specify inadmissible FSMs, where the set of initial states is empty or the transition relation is not total. This may result in logical absurdities.

2.2 Simulation

Simulation offers to the user the possibility of exploring the possible executions (*traces* from now on) of a NuSMV model. In this way, the user can get familiar with a model and can acquire confidence with its correctness before the actual verification of properties. This section describes the basic features of simulation in NuSMV. Further details on the simulation commands can be found in Section 4.4 [Simulation Commands], page 46.

In order to achieve maximum flexibility and degrees of freedom in a simulation session, NuSMV permits three different trace generation strategies: deterministic, random and interactive. Each of them corresponds to a different way a state is picked from a set of possible choices. In deterministic simulation mode the first state of a set (whatever it is) is chosen, while in the random one the choice is performed nondeterministically. In these two first modes traces are automatically generated by NuSMV: the user obtains the whole of the trace in a time without control over the generation itself (except for the simulation mode and the number of states entered via command line).

In the third simulation mode, the user has a complete control over traces generation by interactively building the trace. During an interactive simulation session, the system stops at every step, showing a list of possible future states: the user is requested to choose one of the items. This feature is particularly useful when one wants to inspect some particular reactions of the model to be checked. When the number of possible future states exceeds an internal limit, rather than "confusing" the user with a choice from a high number of possible evolutions, the system asks the user to "guide" the simulation via the insertion of some further constraints over the possible future states. The system will continue to ask for constraints insertion until the

number of future states will be under the predefined threshold. The constraints entered during this phase are accumulated (in a logical product) in a single big constraint. This constraint is used only for the current step of the simulation and is discarded before the next step. The system checks the expressions entered by the user and does not accept them whenever an inconsistency arises. Cases of inconsistency (i.e. empty set of states) may be caused by:

- the entered expressions (i.e. $a \ \& \ \sim a$);
- the result of the entered expressions conjoined with previous accumulated ones;
- the result of accumulated constraints conjoined with the set of possible future states.

A typical execution sequence of a simulation session could be as follows. Suppose to use the model described below.

```
MODULE main
VAR
  request : boolean;
  state : {ready,busy};
ASSIGN
  init(state) := ready;
  next(state) := case
    state = ready & request : busy;
    1 : {ready,busy};
  esac;
```

As a preliminary step, this model has to read into the NuSMV system. This can be obtained by executing the following commands (we assume that the model is saved in file `short.smv`):¹

```
system_prompt> NuSMV -int short.smv
NuSMV > go
NuSMV >
```

In order to start the simulation, an initial state has to be chosen. This can be done in three ways:

- by default, the simulator uses the *current state* as a starting point of every new simulation; this behavior is possible only if a current state is defined (e.g., if we are exploring a trace, as described in Section 4.5 [Traces Inspection Commands], page 48);
- if command `goto_state` is used, the user can select any state of an already existent trace as the *current state* (see Section 4.5 [Traces Inspection Commands], page 48);
- if `pick_state` is used, then the user can choose the starting state of the simulation among the initial states of the model; this command has to be used when a *current state* does not exist yet (that is when the system has being reset).

At this point of the example *current state* does not exist, and there is no trace currently stored in the system. Therefore, an item from the set of initial states has to be picked using command `pick_state`. A simulation session can be started now, using the `simulate` command. Consider for instance the following simulation session:

```
system_prompt> NuSMV -int short.smv
NuSMV > go
NuSMV > pick_state -r
NuSMV > print_current_state -v
Current state is 1.1
request = 0
state = ready
NuSMV > simulate -r 3
***** Starting Simulation From State 1.1 *****
```

¹ We assume that every NuSMV command is followed by a `(RET)` keystroke. In the following examples, NuSMV commands are written in a **bold face** to distinguish them from system output messages.

```

NuSMV > show_traces -t
There is 1 trace currently available.
NuSMV > show_traces -v
##### Trace number: 1 #####
-> State 1.1 <-
    request = 0
    state = ready
-> State 1.2 <-
    request = 1
    state = busy
-> State 1.3 <-
    request = 1
    state = ready
-> State 1.4 <-
    request = 1
    state = busy

```

Command **pick.state -r** requires to pick the starting state of the simulation *randomly* from the set of initial states of the model. Command **simulate -r 3** asks to build a three-steps simulation by picking randomly the next states of the steps. As shown by command **show_traces -v**, the resulting trace contains 4 states (the initial one, and the three ones that have been added by the random simulation). We remark that the generated traces are numbered: every trace is identified by an integer number, while every state belonging to a trace is identified by a *dot notation*: for example state *1.3* is the third state of the first generated trace.

Now the user can start a new simulation by choosing a new starting state. In the next example, for instance, the user extends trace 1 by first choosing state 1.4 as the *current state* and by then running a random simulation of length 3.

```

NuSMV > goto_state 1.4
The starting state for new trace is:
-> State 2.4 <-
    request = 1
    state = busy
NuSMV > simulate -r 3
***** Simulation Starting From State 2.4 *****
NuSMV > show_traces 2
##### Trace number: 2 #####
-> State 2.1 <-
    request = 1
    state = ready
-> State 2.2 <-
    state = busy
-> State 2.3 <-
    request = 0
-> State 2.4 <-
    request = 1
-> State 2.5 <-
    request = 0
-> State 2.6 <-
    state = ready
-> State 2.7 <-
NuSMV >

```

As the reader can see from the previous example, the new trace is stored as trace 2.

The user is also able to interactively choose the states of the trace he wants to build: an example of an interactive simulation is shown below:

```
NuSMV > pick_state -i
***** AVAILABLE FUTURE STATES *****

0) -----
   request = 1
   state = ready

1) -----
   request = 0
```

Choose a state from the above (0-1): 1

Chosen state is: 1

```
NuSMV > simulate -i 1
***** Simulation Starting From State 3.1 *****
***** AVAILABLE FUTURE STATES *****

0) -----
   request = 1
   state = ready

1) -----
   state = busy

2) -----
   request = 0
   state = ready

3) -----
   state = busy
```

Choose a state from the above (0-3): 0

Chosen state is:
request = 1
state = ready

```
NuSMV > show_traces 3
##### Trace number: 3 #####
-> State 3.1 <-
   request = 0
   state = ready
-> State 3.2 <-
   request = 1
```

The user can also specify some constraints to restrict the set of states from which the simulator will pick out. Constraints can be set for both the `pick_state` command and the `simulate` command using option `-c`. For example the following command picks an initial state by defining a simple constraint:

```
NuSMV > pick_state -c "request = 1" -i
***** AVAILABLE FUTURE STATES *****

0) -----
```



```
request = 1
state = ready
```

There's only one future state. Press Return to Proceed. `(RET)`

Chosen state is: 0

```
NuSMV > quit
system_prompt>
```

Note how the set of possible states to choose has being restricted (in this case there is only one future state, so the system will automatically pick it, waiting for the user to press the `(RET)` key).

We remark that, in the case of command `simulate`, the constraints defined using option `-c` are "global" for the actual trace to be generated, in the sense that they are always included in every step of the simulation. They are hence complementary to the constraints entered with the `pick_state` command, or during an interactive simulation session when the number of future states to be displayed is too high, since these are "local" only to a single simulation step and are "forgotten" in the next one.

2.3 CTL model checking

The main purpose of a model checker is to verify that a model satisfies a set of desired properties specified by the user. In NuSMV, the specifications to be checked can be expressed in two different temporal logics: the Computation Tree Logic CTL, and the Linear Temporal Logic LTL extended with Past Operators. CTL and LTL specifications are evaluated by NuSMV in order to determine their truth or falsity in the FSM. When a specification is discovered to be false, NuSMV constructs and prints a counterexample, i.e. a trace of the FSM that falsifies the property. In this section we will describe model checking of specifications expressed in CTL, while the next section we consider the case of LTL specifications.

CTL is a *branching-time* logics: its formulas allow for specifying properties that take into account the non-deterministic, branching evolution of a FSM. More precisely, the evolution of a FSM from a given state can be described as an infinite tree, where the nodes are the states of the FSM and the branching in due to the non-determinism in the transition relation. The paths in the tree that start in a given state are the possible alternative evolutions of the FSM from that state. In CTL one can express properties that should hold for *all the paths* that start in a state, as well as for properties that should hold just for *some of the paths*.

Consider for instance CTL formula $AF\ p$. It expresses the condition that, for *all* the paths (A) starting from a state, *eventually in the future* (F) condition p must hold. That is, all the possible evolutions of the system will eventually reach a state satisfying condition p . CTL formula $EF\ p$, on the other hand, requires that there *exists* some path (E) that eventually in the future satisfies p .

Similarly, formula $AG\ p$ requires that condition p is always, or *globally*, true in all the states of all the possible paths, while formula $EG\ p$ requires that there is some path along which condition p is continuously true.

Other CTL operators are:

- $A\ [p\ U\ q]$ and $E\ [p\ U\ q]$, requiring condition p to be true *until* a state is reached that satisfies condition q ;
- $AX\ p$ and $EX\ p$, requiring that condition p is true in all or in some of the next states reachable from the current state.

CTL operators can be nested in an arbitrary way and can be combined using logic operators ($!$, $\&$, $|$, \rightarrow , \leftarrow ...). Typical examples of CTL formulas are $AG\ !\ p$ ("condition p is absent in all the evolutions"), $AG\ EF\ p$ ("it is always possible to reach a state where p holds"), and $AG\ (p\ \rightarrow\ AF\ q)$ ("each occurrence of condition p is followed by an occurrence of condition q ").

In NuSMV a CTL specification is given as CTL formula introduced by the keyword ‘SPEC’ (see Section 3.3.1 [CTL Specifications], page 29). Whenever a CTL specification is processed, NuSMV checks whether the CTL formula is true in all the initial states of the model. If this is not a case, then NuSMV generates a counter-example, that is, a (finite or infinite) trace that exhibits a valid behavior of the model that does not satisfy the specification. Traces are very useful for identifying the error in the specification that leads to the wrong behavior. We remark that the generation of a counter-example trace is not always possible for CTL specifications. Temporal operators corresponding to existential path quantifiers cannot be proved false by a showing of a single execution path. Similarly, sub-formulas preceded by universal path quantifier cannot be proved true by a showing of a single execution path.

Consider the case of the semaphore program described in Section 2.1 [Examples], page 2. A desired property for this program is that it should never be the case that the two processes `proc1` and `proc2` are at the same time in the `critical` state (this is an example of a "safety" property). This property can be expressed by the following CTL formula:

```
AG ! (proc1.state = critical & proc2.state = critical)
```

Another desired property is that, if `proc1` wants to enter its critical state, it eventually does (this is an example of a "liveness" property). This property can be expressed by the following CTL formula:

```
AG (proc1.state = entering -> AF proc1.state = critical)
```

In order to verify the two formulas on the semaphore model, we add the two corresponding CTL specification to the program, as follows:

```
MODULE main
  VAR
    semaphore : boolean;
    proc1     : process user(semaphore);
    proc2     : process user(semaphore);
  ASSIGN
    init(semaphore) := 0;
  SPEC AG ! (proc1.state = critical & proc2.state = critical)
  SPEC AG (proc1.state = entering -> AF proc1.state = critical)

MODULE user(semaphore)
  VAR
    state : {idle, entering, critical, exiting};
  ASSIGN
    init(state) := idle;
    next(state) :=
      case
        state = idle           : {idle, entering};
        state = entering & !semaphore : critical;
        state = critical       : {critical, exiting};
        state = exiting        : idle;
        1                       : state;
      esac;
    next(semaphore) :=
      case
        state = entering : 1;
        state = exiting  : 0;
        1                 : semaphore;
      esac;
  FAIRNESS
    running
```

By running NuSMV with the command

```
system_prompt> NuSMV semaphore.smv
```

we obtain the following output:

```
-- specification AG (!(proc1.state = critical & proc2.state = critical))
-- is true

-- specification AG (proc1.state = entering -> AF proc1.state = critical)
-- is false
-- as demonstrated by the following execution sequence

-> State 1.1 <-
  semaphore = 0
  proc1.state = idle
  proc2.state = idle
  [executing process proc2]

-> State 1.2 <-
  [executing process proc1]

-- loop starts here --
-> State 1.3 <-
  proc1.state = entering
  [executing process proc2]

-> State 1.4 <-
  proc2.state = entering
  [executing process proc2]

-> State 1.5 <-
  semaphore = 1
  proc2.state = critical
  [executing process proc1]

-> State 1.6 <-
  [executing process proc2]

-> State 1.7 <-
  proc2.state = exiting
  [executing process proc2]

-> State 1.8 <-
  semaphore = 0
  proc2.state = idle
  [executing process proc2]
```

NuSMV tells us that the first CTL specification is true: it is never the case that the two processes will be at the same time in the critical region. On the other hand, the second specification is false. NuSMV produces a counter-example path where initially `proc1` goes to state `entering` (state 1.3), and then a loop starts in which `proc2` repeatedly enters its critical region (state 1.5) and then returns to its `idle` state (state 1.8); in the loop, `proc1` is activated only when `proc2` is in the critical region, and is therefore not able to enter its critical region (state 1.6). This path not only shows that the specification is false, it also points out why can it happen that `proc1` never enters its critical region.

Note that in the printout of a cyclic, infinite counter-example the starting point of the loop is marked by `-- loop starts here --`. Moreover, in order to make it easier to follow the action in systems with a large number of variables, only the values of variables that have changed in the last step are printed in the states of the trace.

2.4 LTL model checking

NuSMV allows for specifications expressed in LTL. Intuitively, while CTL specifications express properties over the computation tree of the FSM (branching-time approach), LTL characterizes each linear path induced by the FSM (linear-time approach). The two logics have in general different expressive power, but also share a significant intersection that includes most of the common properties used in practice. Typical LTL operators are:

- **F p** (read "in the future p"), stating that a certain condition p holds in one of the future time instants;
- **G p** (read "globally p"), stating that a certain condition p holds in all future time instants;
- **p U q** (read "p until q"), stating that condition p holds until a state is reached where condition q holds;
- **X p** (read "next p"), stating that condition p is true in the next state.

We remark that, differently from CTL, LTL temporal operators do not have path quantifiers. In fact, LTL formulas are evaluated on linear paths, and a formula is considered true in a given state if it is true for all the paths starting in that state.

Consider the case of the semaphore program and of the safety and liveness properties already described in Section 2.3 [CTL model checking], page 10. These properties correspond to LTL formulas

```
G ! (proc1.state = critical & proc2.state = critical)
```

expressing that the two processes cannot be in the critical region at the same time, and

```
G (proc1.state = entering -> F proc1.state = critical)
```

expressing that whenever a process wants to enter its critical session, it eventually does.

If we add the two corresponding LTL specification to the program, as follows:²

```
MODULE main
  VAR
    semaphore : boolean;
    proc1      : process user(semaphore);
    proc2      : process user(semaphore);
  ASSIGN
    init(semaphore) := 0;
  LTLSPEC G ! (proc1.state = critical & proc2.state = critical)
  LTLSPEC G (proc1.state = entering -> F proc1.state = critical)

MODULE user(semaphore)
  VAR
    state : {idle, entering, critical, exiting};
  ASSIGN
    init(state) := idle;
    next(state) :=
      case
        state = idle           : {idle, entering};
        state = entering & !semaphore : critical;
```

² In NuSMV a LTL specification are introduced by the keyword 'LTLSPEC' (see Section 3.3.2 [LTL Specifications], page 30).

```

        state = critical           : {critical, exiting};
        state = exiting           : idle;
        1                         : state;
    esac;
next(semaphore) :=
    case
        state = entering : 1;
        state = exiting  : 0;
        1                 : semaphore;
    esac;
FAIRNESS
    running

```

NuSMV produces the following output:

```

-- specification G (!(proc1.state = critical & proc2.state = critical))
-- is true
-- specification G (proc1.state = entering -> F proc1.state = critical)
-- is false
-- as demonstrated by the following execution sequence
-> State 1.1 <-
    semaphore = 0
    proc1.state = idle
    proc2.state = idle
    [executing process proc2]

-> State 1.2 <-
[...]
```

That is, the first specification is true, while the second is false and a counter-example path is generated.

In NuSMV, LTL properties can also include *past* temporal operators. Differently from standard temporal operators, that allow to express properties over the future evolution of the FSM, past temporal operators allow to characterize properties of the path that lead to the current situation. The typical past operators are:

- $\bigcirc p$ (read "once p "), stating that a certain condition p holds in one of the past time instants;
- $H p$ (read "historically p "), stating that a certain condition p holds in all previous time instants;
- $p S q$ (read " p since q "), stating that condition p holds since a previous state where condition q holds;
- $Y p$ (read "yesterday p "), stating that condition p holds in the previous time instant.

Past temporal operators can be combined with future temporal operators, and allow for the compact characterization of complex properties.

A detailed description of the syntax of LTL formulas can be found in Section 3.3.2 [LTL Specifications], page 30.

2.5 Bounded Model Checking

In this section we give a short introduction to the use of Bounded Model Checking (BMC) in NuSMV. Further details on BMC can be found in Section 4.3 [Commands for Bounded Model Checking], page 39. For a more in-depth introduction to the theory underlying BMC please refer to [BCCZ99].

Consider the following model, representing a simple, deterministic counter modulo 8 (we assume that the following specification is contained in file `modulo8.smv`):

```

MODULE main
VAR
  y : 0..15;

ASSIGN
  init(y) := 0;

TRANS
  case
    y = 7 : next(y) = 0;
    1      : next(y) = ((y + 1) mod 16);
  esac

```

This slightly artificial model has only the state variable y , ranging from 0 to 15. The values of y are limited by the transition relation to the $[0, 7]$ interval. The counter starts from 0, deterministically increments by one the value of y at each transition up to 7, and then restarts from zero.

We would like to check with BMC the LTL specification $G (y=4 \rightarrow X y=6)$ expressing that "each time the counter value is 4, the next counter value will be 6". This specification is obviously false, and our first step is to use NuSMV BMC to demonstrate its falsity. To this purpose, we add the following specification to file `modulo8.smv`:

```
LTLSPEC G ( y=4 -> X y=6 )
```

and we instruct NuSMV to run in BMC by using command-line option `-bmc`:

```

system_prompt> NuSMV -bmc modulo8.smv
-- no counterexample found with bound 0 for specification
  G(y = 4 -> X y = 6)
-- no counterexample found with bound 1 for ...
-- no counterexample found with bound 2 for ...
-- no counterexample found with bound 3 for ...
-- no counterexample found with bound 4 for ...
-- specification G (y = 4 -> X y = 6) is false
-- as demonstrated by the following execution sequence
State 1.1: y = 0
State 1.2: y = 1
State 1.3: y = 2
State 1.4: y = 3
State 1.5: y = 4
State 1.6: y = 5
system_prompt>

```

NuSMV has found that the specification is false, and is showing us a counterexample, i.e. a trace where the value of y becomes 4 (at time 4) and at the next step is not 6.

```

bound:   0   1   2   3   4   5
         o--->o--->o--->o--->o--->o
state:  y=0 y=1 y=2 y=3 y=4 y=5

```

The output produced by NuSMV shows that, before the counterexample of length 5 is found, NuSMV also tried to find counterexamples of lengths 0 to 4. However, there are no such counterexamples. For instance, in the case of bound 4, the traces of the model have the following form:

```

bound:   0   1   2   3   4
         o--->o--->o--->o--->o
state:  y=0 y=1 y=2 y=3 y=4

```

In this situation, y gets the value 4, but it is impossible for NuSMV to say anything about the following state.

In general, in BMC mode NuSMV tries to find a counterexample of increasing length, and immediately stops when it succeeds, declaring that the formula is false. The maximum number of iterations can be controlled by using command-line option `-bmc_length`. The default value is 10. If the maximum number of iterations is reached and no counter-example is found, then NuSMV exits, and the truth of the formula is not decided. We remark that in this case we cannot conclude that the formula is true, but only that any counter-example should be longer than the maximum length.

```
system_prompt> NuSMV -bmc -bmc_length 4 modulo8.smv
-- no counterexample found with bound 0 for ...
-- no counterexample found with bound 1 for ...
-- no counterexample found with bound 2 for ...
-- no counterexample found with bound 3 for ...
-- no counterexample found with bound 4 for ...
system_prompt>
```

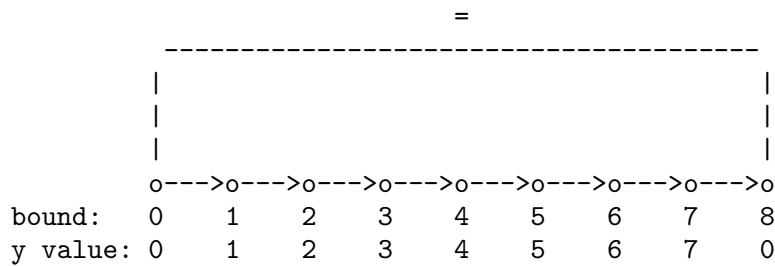
Let us consider now another property, $!G F (y = 2)$, stating that y gets the value 2 only a finite number of times. Again, this is a false property due to the cyclic nature of the model. Let us modify the specification of file `model8.smv` as follows:

```
LTLSPEC !G F (y = 2)
```

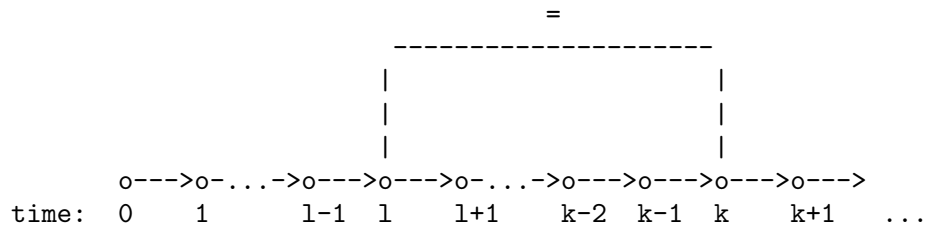
and let us run NuSMV in BMC mode:

```
system_prompt> NuSMV -bmc modulo8.smv
-- no counterexample found with bound 0 for specification ! G F y = 2
-- no counterexample found with bound 1 for ...
-- no counterexample found with bound 2 for ...
-- no counterexample found with bound 3 for ...
-- no counterexample found with bound 4 for ...
-- no counterexample found with bound 5 for ...
-- no counterexample found with bound 6 for ...
-- no counterexample found with bound 7 for ...
-- specification ! G F y = 2 is false
-- as demonstrated by the following execution sequence
-- loop starts here --
State 2.1: y = 0
State 2.2: y = 1
State 2.3: y = 2
State 2.4: y = 3
State 2.5: y = 4
State 2.6: y = 5
State 2.7: y = 6
State 2.8: y = 7
State 2.9: y = 0
system_prompt>
```

In this example NuSMV has increased the problem bound until a cyclic behavior of length 8 is found that contains a state where y value is 2. Since the behavior is cyclic, state 2.3 is entered infinitely often and the property is false.



In general, BMC can find two kinds of counterexamples, depending on the property being analyzed. For safety properties (e.g. like the first one used in this tutorial), a counterexample is a finite sequence of transitions through different states. For liveness properties, counterexamples are infinite but periodic sequences, and can be represented in a bounded setting as a finite prefix followed by a loop, i.e. a finite sequence of states ending with a loop back to some previous state. So a counterexample which demonstrates the falsity of a liveness property as " $\neg \text{G F } p$ " cannot be a finite sequence of transitions. It must contain a loop which makes the infinite sequence of transitions as well as we expected.



Consider the above figure. It represents an examples of a generic infinite counterexample, with its two parts: the prefix part (times from 0 to $l-1$), and the loop part (indefinitely from l to $k-1$). Because the loop always jumps to a previous time it is called 'loopback'. The loopback condition requires that state k is identical to state l . As a consequence, state $k+1$ is forced to be equal to state $l+1$, state $k+2$ to be equal to state $l+2$, and so on.

A fine-grained control of the length and of the loopback condition for the counter-example can be specified by using command `check_ltlspec_bmc_onepb` in interactive mode (see Section 4.3 [Commands for Bounded Model Checking], page 39). This command accepts options `-k`, that specifies the length of the counter-example we are looking for, and `-l`, that defines the loopback condition. Consider the following interactive session:

```

system_prompt> NuSMV -int modulo8.smv
NuSMV > go_bmc
NuSMV > check_ltlspec_bmc_onepb -k 9 -l 0
-- no counterexample found with bound 9 and loop at 0 for specification
! G F y = 2
NuSMV > check_ltlspec_bmc_onepb -k 8 -l 1
-- no counterexample found with bound 8 and loop at 1 for specification
! G F y = 2
NuSMV > check_ltlspec_bmc_onepb -k 9 -l 1
-- specification ! G F y = 2 is false
-- as demonstrated by the following execution sequence
State 1.1: y = 0
-- loop starts here --
State 1.2: y = 1
State 1.3: y = 2
State 1.4: y = 3
State 1.5: y = 4
State 1.6: y = 5

```



```

State 1.7: y = 6
State 1.8: y = 7
State 1.9: y = 0
State 1.10: y = 1
NuSMV > quit
system_prompt>

```

NuSMV did not find a counterexample for cases (k=9, l=0) and (k=8, l=1). The following figures show that these case look for counterexamples that do not match with the model of the counter, so it is not possible for NuSMV to satisfy them.

k = 9, l = 0:

```

=
-----
|                                     |
|                                     |
|                                     |
o--->o--->o--->o--->o--->o--->o--->o--->o--->o
bound:  0   1   2   3   4   5   6   7   8   9
y value: 0   1   2   3   4   5   6   7   0   1

```

k = 8, l = 1:

```

=
-----
|                                     |
|                                     |
|                                     |
o--->o--->o--->o--->o--->o--->o--->o--->o
bound:  0   1   2   3   4   5   6   7   8
y value: 0   1   2   3   4   5   6   7   0

```

Case (k=9, l=1), instead allows for a counter-example:

k = 9, l = 1:

```

=
-----
|                                     |
|                                     |
|                                     |
o--->o--->o--->o--->o--->o--->o--->o--->o
bound:  0   1   2   3   4   5   6   7   8   9
y value: 0   1   2   3   4   5   6   7   0   1

```

In NuSMV it is possible to specify the loopback condition in four different ways:

- **The loop as a precise time-point.** Use a natural number as the argument of option -l.
- **The loop length.** Use a negative number as the argument of option -l. The negative number is the loop length, and you can also imagine it as a precise time-point loop relative to the path bound.
- **No loopback.** Use symbol 'X' as the argument of option -l. In this case NuSMV will not find infinite counterexamples.
- **All possible loops.** Use symbol '*' as the argument of option -l. In this case NuSMV will search counterexamples for paths with any possible loopback structure. A counterexample with no loop will be also searched. This is the default value for option -l.

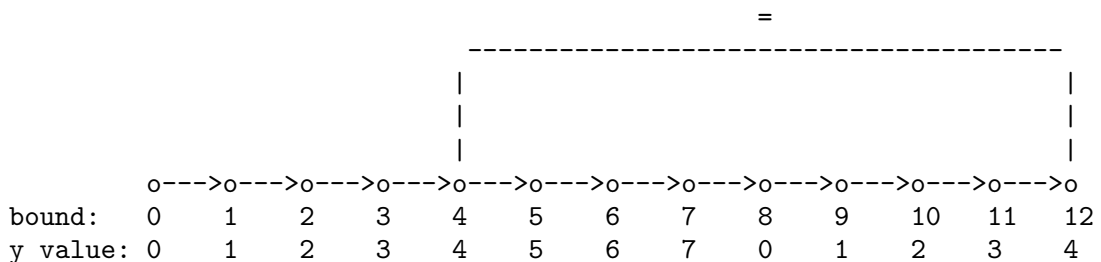
In the following example we look for a counter-example of length 12 with a loop of length 8:

```

system_prompt> NuSMV -int modulo8.smv
NuSMV > go_bmc
NuSMV > check_ltlspec_bmc_onepb -k 12 -l -7
-- specification ! G F y = 2 is false
-- as demonstrated by the following execution sequence
State 1.1: y = 0
State 1.2: y = 1
State 1.3: y = 2
State 1.4: y = 3
-- loop starts here --
State 1.5: y = 4
State 1.6: y = 5
State 1.7: y = 6
State 1.8: y = 7
State 1.9: y = 0
State 1.10: y = 1
State 1.11: y = 2
State 1.12: y = 3
State 1.13: y = 4
NuSMV >

```

This picture illustrates the produced counterexample in a more effective way:



If no loopback is specified, NuSMV is not able to find a counterexample for the given liveness property:

```

system_prompt> NuSMV -int modulo8.smv
NuSMV > go_bmc
NuSMV > check_ltlspec_bmc_onepb -k 12 -l X
-- no counterexample found with bound 12 and no loop for ...
NuSMV >

```

Bounded Model Checking in NuSMV can be used not only for checking LTL specification, but also for checking invariants. An invariant is a propositional property which must always hold. BMC tries to prove the truth of invariants via an inductive reasoning, by checking if (i) the property holds in every initial state, and (ii) if it holds in any state reachable from any state where it holds.

Let us modify file `modulo8.smv` by replacing the LTL specification with the following invariant specification:

```
INVARSPEC y in (0..12)
```

and let us run NuSMV in BMC mode:

```

system_prompt> NuSMV -bmc modelo8.smv
-- cannot prove the invariant y in (0 .. 12) : the induction fails
-- as demonstrated by the following execution sequence

```

```

State 1.1: y = 12
State 1.2: y = 13
system_prompt>

```

NuSMV reports that the given invariant cannot be proved, and it shows a state satisfying "y in (0..12)" that has a successor state not satisfying "y in (0..12)". This two-steps sequence of assignments shows why the induction fails. Note that NuSMV does not say the given formula is really false, but only that it cannot be proven to be true using the inductive reasoning described previously.

If we try to prove the stronger invariant `y in (0..7)` we obtain:

```

system_prompt> NuSMV -bmc modelo8.smv
-- invariant y in (0 .. 7) is true
system_prompt>

```

In this case NuSMV is able to prove that `y in (0..7)` is true. As a consequence, also the weaker invariant `y in (0..12)` is true, even if NuSMV is not able to prove it in BMC mode. On the other hand, the returned counter-example can be used to *strengthen* the invariant, until NuSMV is able to prove it.

Now we check the false invariant `y in (0..6)`:

```

-- cannot prove the invariant y in (0 .. 6) : the induction fails
-- as demonstrated by the following execution sequence
State 1.1: y = 6
State 1.2: y = 7
NuSMV >

```

As for property `y in (0..12)`, NuSMV returns a two steps sequence showing that the induction fails. The difference is that, in the former case state 'y=12' is NOT reachable, while in the latter case the state 'y=6' can be reached. Unfortunately enough, the BMC-based invariant checker is not able to distinguish these two cases.

3 Syntax

We present now the complete syntax of the input language of NuSMV. In the following, an *atom* may be any sequence of characters starting with a character in the set {A-Za-z_} and followed by a possibly empty sequence of characters belonging to the set {A-Za-z0-9_ \ \$ # -}. A *number* is any sequence of digits. A digit belongs to the set {0-9}.

All characters and case in a name are significant. Whitespace characters are space (SPACE), tab (TAB) and newline (RET). Any string starting with two dashes ('--') and ending with a newline is a comment. Any other tokens recognized by the parser are enclosed in quotes in the syntax expressions below. Grammar productions enclosed in square brackets ('[]') are optional.

3.1 Expressions

Expressions are constructed from variables, constants, and a collection of operators, including boolean connectives, integer arithmetic operators, case expressions and set expressions.

3.1.1 Simple Expressions

Simple expressions are expressions built only from current state variables. Simple expressions can be used to specify sets of states, e.g. the initial set of states. The syntax of simple expressions is as follows:

```

simple_expr ::
    atom                ;; a symbolic constant
  | number              ;; a numeric constant
  | "TRUE"              ;; The boolean constant 1
  | "FALSE"             ;; The boolean constant 0
  | var_id              ;; a variable identifier
  | "(" simple_expr ")"
  | "!" simple_expr     ;; logical not
  | simple_expr "&" simple_expr ;; logical and
  | simple_expr "|" simple_expr ;; logical or
  | simple_expr "xor" simple_expr ;; logical exclusive or
  | simple_expr "->" simple_expr ;; logical implication
  | simple_expr "<->" simple_expr ;; logical equivalence
  | simple_expr "=" simple_expr ;; equality
  | simple_expr "!=" simple_expr ;; inequality
  | simple_expr "<" simple_expr ;; less than
  | simple_expr ">" simple_expr ;; greater than
  | simple_expr "<=" simple_expr ;; less than or equal
  | simple_expr ">=" simple_expr ;; greater than or equal
  | simple_expr "+" simple_expr ;; integer addition
  | simple_expr "-" simple_expr ;; integer subtraction
  | simple_expr "*" simple_expr ;; integer multiplication
  | simple_expr "/" simple_expr ;; integer division
  | simple_expr "mod" simple_expr ;; integer remainder
  | set_simple_expr     ;; a set simple_expression
  | case_simple_expr    ;; a case expression

```

A *var_id*, (see Section 3.2.10 [Identifiers], page 27) or identifier, is a symbol or expression which identifies an object, such as a variable or a defined symbol. Since a *var_id* can be an *atom*, there is a possible ambiguity if a variable or defined symbol has the same name as a symbolic constant. Such an ambiguity is flagged by the interpreter as an error.

The order of parsing precedence for operators from high to low is:

```

*,/
+,-
mod
=,!=,<,>,<=,>=
!
&
|,xor
<->
->

```

Operators of equal precedence associate to the left, except `->` that associates to the right. Parentheses may be used to group expressions.

3.1.1.1 Case Expressions

A case expression has the following syntax:

```

case_simple_expr ::
    "case"
        simple_expr ":" simple_expr ";"
        simple_expr ":" simple_expr ";"
        ...
        simple_expr ":" simple_expr ";"
    "esac"

```

A `case_simple_expr` returns the value of the first expression on the right hand side of `:`, such that the corresponding condition on the left hand side evaluates to 1. Thus, if `simple_expr` on the left side is true, then the result is the corresponding `simple_expr` on the right side. If none of the expressions on the left hand side evaluates to 1, the result of the `case_expression` is the numeric value 1. It is an error for any expression on the left hand side to return a value other than the truth values 0 or 1.

3.1.1.2 Set Expressions

A set expression has the following syntax:

```

set_expr ::
    "{" set_elem "," ... "," set_elem "}" ;; set definition
    | simple_expr "in" simple_expr           ;; set inclusion test
    | simple_expr "union" simple_expr        ;; set union
set_elem :: simple_expr

```

A set can be defined by enumerating its elements inside curly braces `{...}`. The inclusion operator `in` tests a value for membership in a set. The union operator `union` takes the union of two sets. If either argument is a number or a symbolic value instead of a set, it is coerced to a singleton set.

3.1.2 Next Expressions

While simple expressions can represent sets of states, next expressions relate current and next state variables to express transitions in the FSM. The structure of next expressions is similar to the structure of simple expressions (See Section 3.1.1 [Simple Expressions], page 21). The difference is that next expression allow to refer to next state variables. The grammar is depicted below.

```

next_expr ::
    atom                ;; a symbolic constant
    | number            ;; a numeric constant
    | "TRUE"           ;; The boolean constant 1

```

```

| "FALSE"                ;; The boolean constant 0
| var_id                 ;; a variable identifier
| "(" next_expr ")"
| "next" "(" simple_expr ")" ;; next value of an "expression"
| "!" next_expr          ;; logical not
| next_expr "&" next_expr  ;; logical and
| next_expr "|" next_expr ;; logical or
| next_expr "xor" next_expr ;; logical exclusive or
| next_expr "->" next_expr ;; logical implication
| next_expr "<->" next_expr ;; logical equivalence
| next_expr "=" next_expr ;; equality
| next_expr "!=" next_expr ;; inequality
| next_expr "<" next_expr  ;; less than
| next_expr ">" next_expr  ;; greater than
| next_expr "<=" next_expr ;; less than or equal
| next_expr ">=" next_expr ;; greater than or equal
| next_expr "+" next_expr  ;; integer addition
| next_expr "-" next_expr  ;; integer subtraction
| next_expr "*" next_expr  ;; integer multiplication
| next_expr "/" next_expr  ;; integer division
| next_expr "mod" next_expr ;; integer remainder
| set_next_expr          ;; a set next_expression
| case_next_expr         ;; a case expression

```

`set_next_expr` and `case_next_expr` are the same as `set_simple_expr` (see Section 3.1.1.2 [Set Expressions], page 22) and `case_simple_expr` (see Section 3.1.1.1 [Case Expressions], page 22) respectively, with the replacement of "simple" with "next". The only additional production is "next" "(" simple_expr ")", which allows to "shift" all the variables in `simple_expr` to the *next* state. The `next` operator distributes on every operator. For instance, the formula `next((A & B) | C)` is a shorthand for the formula `(next(A) & next(B)) | next(C)`. It is an error if in the scope of the `next` operator occurs another `next` operator.

3.2 Definition of the FSM

3.2.1 State Variables

A state of the model is an assignment of values to a set of state variables. These variables (and also instances of modules) are declared by the notation:

```

var_declaration :: "VAR"
                atom ":" type ";"
                atom ":" type ";"
                ...

```

The type associated with a variable declaration can be either a boolean, a scalar, a user defined module, or an array of any of these (including arrays of arrays).

3.2.1.1 Type Specifiers

A type specifier has the syntax:

```

type :: boolean
      | "{" val "," val "," ... val "}"
      | number ".." number
      | "array" number ".." number "of" type
      | atom [ "(" simple_expr "," simple_expr "," ... ")" ]

```

```

    | "process" atom [ "(" simple_expr "," ... "," simple_expr ")" ]

val  :: atom
    | number

```

A variable of type `boolean` can take on the numerical values 0 and 1 (representing false and true, respectively). In the case of a list of values enclosed in quotes (where atoms are taken to be symbolic constants), the variable is a scalar which take any of these values. In the case of an `array` declaration, the first `simple_expr` is the lower bound on the subscript and the second `simple_expr` is the upper bound. Both of these expressions must evaluate to integer constants. Finally, an atom optionally followed by a list of expressions in parentheses indicates an instance of module atom (See Section 3.2.9 [MODULE declarations], page 26). The keyword causes the module to be instantiated as an asynchronous process (See Section 3.2.12 [Processes], page 28).

3.2.2 Input Variables

A state of the model is an assignment of values to a set of state variables. These variables (and also instances of modules) are declared by the notation:

```

ivar_declaration :: "IVAR"
                atom ":" type ";"
                atom ":" type ";"
                ...

```

The type associated with a variable declaration can be either a boolean, a scalar, a user defined module, or an array of any of these (including arrays of arrays) (See Section 3.2.1 [State Variables], page 23).

3.2.3 ASSIGN declarations

An assignment has the form:

```

assign_declaration :: "ASSIGN"
                  assign_body ";"
                  assign_body ";"
                  ...

```

```

assign_body ::
    atom           "!=" simple_expr      ;; normal assignment
  | "init" "(" atom ")" "!=" simple_expr  ;; init assignment
  | "next" "(" atom ")" "!=" next_expr    ;; next assignment

```

On the left hand side of the assignment, `atom` denotes the current value of a variable, `'init(atom)'` denotes its initial value, and `'next(atom)'` denotes its value in the next state. If the expression on the right hand side evaluates to an integer or symbolic constant, the assignment simply means that the left hand side is equal to the right hand side. On the other hand, if the expression evaluates to a set, then the assignment means that the left hand side is contained in that set. It is an error if the value of the expression is not contained in the range of the variable on the left hand side.

In order for a program to be implementable, there must be some order in which the assignments can be executed such that no variable is assigned after its value is referenced. This is not the case if there is a circular dependency among the assignments in any given process. Hence, such a condition is an error. It is also an error for a variable to be assigned more than once at any given time. More precisely, it is an error if:

1. the next or current value of a variable is assigned more than once in a given process, or
2. the initial value of a variable is assigned more than once in the program, or
3. the current value and the initial value of a variable are both assigned in the program, or
4. the current value and the next value of a variable are both assigned in the program.

3.2.4 TRANS declarations

The transition relation R of the model is a set of current state/next state pairs. Whether or not a given pair is in this set is determined by a boolean valued expression T , introduced by the ‘TRANS’ keyword. The syntax of a TRANS declaration is:

```
trans_declaration :: "TRANS" trans_expr [";"]

trans_expr      :: next_expr
```

It is an error for the expression to yield any value other than 0 or 1. If there is more than one TRANS declaration, the transition relation is the conjunction of all of TRANS declarations.

3.2.5 INIT declarations

The set of initial states of the model is determined by a boolean expression under the ‘INIT’ keyword. The syntax of a INIT declaration is:

```
init_declaration :: "INIT" init_expr [";"]

init_expr       :: simple_expr
```

It is an error for the expression to contain the `next()` operator, or to yield any value other than 0 or 1. If there is more than one INIT declaration, the initial set is the conjunction of all of the INIT declarations.

3.2.6 INVAR declarations

The set of invariant states (i.e. the analogous of normal assignments, as described in Section 3.2.3 [ASSIGN declarations], page 24) can be specified using a boolean expression under the ‘INVAR’ keyword. The syntax of a INVAR declaration is:

```
invar_declaration      :: "INVAR" invar_expr [";"]

invar_expr            :: simple_expr
```

It is an error for the expression to contain the `next()` operator, or to yield any value other than 0 or 1. If there is more than one INVAR declaration, the invariant set is the conjunction of all of the INVAR declarations.

3.2.7 DEFINE declarations

In order to make descriptions more concise, a symbol can be associated with a commonly expression. The syntax for this kind of declaration is:

```
define_declaration :: "DEFINE"
                    atom ":@" simple_expr ";"
                    atom ":@" simple_expr ";"
                    ...
                    atom ":@" simple_expr ";"
```

Whenever an identifier referring to the symbol on the left hand side of the ‘:=’ in a DEFINE occurs in an expression, it is replaced by the expression on the right hand side. The expression on the right hand side is always evaluated in its context, however (see Section 3.2.9 [MODULE declarations], page 26 for an explanation of contexts). Forward references to defined symbols are allowed, but circular definitions are not allowed, and result in an error.

It is not possible to assign values to defined symbols non-deterministically. Another difference between defined symbols and variables is that while variables are statically typed, definitions are not.

3.2.8 ISA declarations

There are cases in which some parts of a module could be shared among different modules, or could be used as a module themselves. In NuSMV it is possible to declare the common parts as separate modules, and then use the ISA declaration to import the common parts inside a module declaration.

The syntax of an ISA declaration is as follows:

```
isa_declaration :: "ISA" atom
```

where *atom* must be the name of a declared module. The ISA declaration can be thought as a simple macro expansion command, because the body of the module referenced by an ISA command is replaced to the ISA declaration.

3.2.9 MODULE declarations

A module is an encapsulated collection of declarations. Once defined, a module can be reused as many times as necessary. Modules can also be so that each instance of a module can refer to different data values. A module can contain instances of other modules, allowing a structural hierarchy to be built.

The syntax of a module declaration is as follows.

```
module ::
  "MODULE" atom [ "(" atom "," atom "," ... atom ")" ]
  [ var_declaration      ]
  [ ivar_declaration     ]
  [ assign_declaration   ]
  [ trans_declaration    ]
  [ init_declaration     ]
  [ invar_declaration    ]
  [ spec_declaration     ]
  [ checkinvar_declaration ]
  [ ltlspec_declaration  ]
  [ compute_declaration  ]
  [ fairness_declaration  ]
  [ define_declaration   ]
  [ isa_declaration      ]
```

The *atom* immediately following the keyword "MODULE" is the name associated with the module. Module names are drawn from a separate name space from other names in the program, and hence may clash with names of variables and definitions. The optional list of atoms in parentheses are the formal parameters of the module. Whenever these parameters occur in expressions within the module, they are replaced by the actual parameters which are supplied when the module is instantiated (see below).

An *instance* of a module is created using the VAR declaration (see Section 3.2.1 [State Variables], page 23). This declaration supplies a name for the instance, and also a list of actual parameters, which are assigned to the formal parameters in the module definition. An actual parameter can be any legal expression. It is an error if the number of actual parameters is different from the number of formal parameters. The semantic of module instantiation is similar to call-by-reference. For example, in the following program fragment:

```
MODULE main
...
VAR
  a : boolean;
  b : foo(a);
...
MODULE foo(x)
```

```

ASSIGN
  x := 1;

```

the variable `a` is assigned the value 1. This distinguishes the call-by-reference mechanism from a call-by-value scheme.

Now consider the following program:

```

MODULE main
...
DEFINE
  a := 0;
VAR
  b : bar(a);
...
MODULE bar(x)
DEFINE
  a := 1;
  y := x;

```

In this program, the value of `y` is 0. On the other hand, using a call-by-name mechanism, the value of `y` would be 1, since `a` would be substituted as an expression for `x`.

Forward references to module names are allowed, but circular references are not, and result in an error.

3.2.10 Identifiers

An *id*, or identifier, is an expression which references an object. Objects are instances of modules, variables, and defined symbols. The syntax of an identifier is as follows.

```

id ::
    atom
    | "self"
    | id "." atom
    | id "[" simple_expr "]"

```

An *atom* identifies the object of that name as defined in a `VAR` or `DEFINE` declaration. If `a` identifies an instance of a module, then the expression `a.b` identifies the component object named `b` of instance `a`. This is precisely analogous to accessing a component of a structured data type. Note that an actual parameter of module `a` can identify another module instance `b`, allowing `a` to access components of `b`, as in the following example:

```

MODULE main
... VAR
  a : foo(b);
  b : bar(a);
...

MODULE foo(x)
DEFINE
  c := x.p | x.q;

MODULE bar(x)
VAR
  p : boolean;
  q : boolean;

```

Here, the value of `c` is the logical or of `p` and `q`.

If ‘a’ identifies an array, the expression ‘a[b]’ identifies element ‘b’ of array ‘a’. It is an error for the expression ‘b’ to evaluate to a number outside the subscript bounds of array ‘a’, or to a symbolic value.

It is possible to refer the name the current module has been instantiated to by using the `self` builtin identifier.

```

MODULE element(above, below, token)
  VAR
    Token : boolean;

  ASSIGN
    init(Token) := token;
    next(Token) := token-in;

  DEFINE
    above.token-in := Token;
    grant-out := below.grant-out;

MODULE cell
  VAR
    e2 : element(self, e1, 0);
    e1 : element(e1, self, 1);

  DEFINE
    e1.token-in := token-in;
    grant-out := grant-in & !e1.grant-out;

MODULE main
  VAR c1 : cell;

```

In this example the name the `cell` module has been instantiated to is passed to the sub-module `element`. In the main module, declaring `c1` to be an instance of module `cell` and defining `above.token-in` in module `e2`, really amounts to defining the symbol `c1.token-in`. When you, in the `cell` module, declare `e1` to be an instance of module `element`, and you define `grant-out` in module `e1` to be `below.grant-out`, you are really defining it to be the symbol `c1.grant-out`.

3.2.11 The main module

The syntax of a NuSMV program is:

```

program ::
  module_1
  module_2
  ...
  module_n

```

There must be one module with the name `main` and no formal parameters. The module `main` is the one evaluated by the interpreter.

3.2.12 Processes

Processes are used to model interleaving concurrency. A *process* is a module which is instantiated using the keyword ‘`process`’ (see Section 3.2.1 [State Variables], page 23). The program executes a step by non-deterministically choosing a process, then executing all of the assignment statements in that process in parallel. It is implicit that if a given variable is not assigned by the process, then its value remains unchanged. Each instance of a process has a

special boolean variable associated with it called `running`. The value of this variable is 1 if and only if the process instance is currently selected for execution. A process may run only when its parent is running. In addition no two processes with the same parents may be running at the same time.

3.2.13 FAIRNESS declarations

A *fairness constraint* restricts the attention only to *fair execution paths*. When evaluating specifications, the model checker considers path quantifiers to apply only to fair paths.

NuSMV supports two types of fairness constraints, namely justice constraints and compassion constraints. A *justice constraint* consists of a formula `f` which is assumed to be true infinitely often in all the fair paths. In NuSMV justice constraints are identified by keywords `JUSTICE` and, for backward compatibility, `FAIRNESS`. A *compassion constraint* consists of a pair of formulas `(p,q)`; if property `p` is true infinitely often in a fair path, then also formula `q` has to be true infinitely often in the fair path. In NuSMV compassion constraints are identified by keyword `COMPASSION`.¹

Fairness constraints are declared using the following syntax:

```

fairness_declaration ::
    "FAIRNESS" simple_expr [";"]
  | "JUSTICE" simple_expr [";"]
  | "COMPASSION" "(" simple_expr "," simple_expr ")" [";"]

```

A path is considered fair if and only if it satisfies all the constraints declared in this manner.

3.3 Specifications

The specifications to be checked on the FSM can be expressed in two different temporal logics: the Computation Tree Logic CTL, and the Linear Temporal Logic LTL extended with Past Operators. It is also possible to analyze quantitative characteristics of the FSM by specifying real-time CTL specifications. Specifications can be positioned within modules, in which case they are preprocessed to rename the variables according to the containing context.

CTL and LTL specifications are evaluated by NuSMV in order to determine their truth or falsity in the FSM. When a specification is discovered to be false, NuSMV constructs and prints a counterexample, i.e. a trace of the FSM that falsifies the property.

3.3.1 CTL Specifications

A CTL specification is given as a formula in the temporal logic CTL, introduced by the keyword `'SPEC'`. The syntax of this declaration is:

```

spec_declaration :: "SPEC" spec_expr [";"]

```

```

spec_expr      :: ctl_expr

```

The syntax of CTL formulas recognized by the NuSMV parser is as follows:

```

ctl_expr ::
    simple_expr                ;; a simple boolean expression
  | "(" ctl_expr ")"
  | "!" ctl_expr              ;; logical not
  | ctl_expr "&" ctl_expr      ;; logical and
  | ctl_expr "|" ctl_expr      ;; logical or
  | ctl_expr "xor" ctl_expr    ;; logical exclusive or

```

¹ In the current version of NuSMV, compassion constraints are supported only for BDD-based LTL model checking. We plan to add support for compassion constraints also for CTL specifications and in Bounded Model Checking in the next releases of NuSMV.

```

| ctl_expr "->" ctl_expr      ;; logical implies
| ctl_expr "<->" ctl_expr     ;; logical equivalence
| "EG" ctl_expr              ;; exists globally
| "EX" ctl_expr              ;; exists next state
| "EF" ctl_expr              ;; exists finally
| "AG" ctl_expr              ;; forall globally
| "AX" ctl_expr              ;; forall next state
| "AF" ctl_expr              ;; forall finally
| "E" "[" ctl_expr "U" ctl_expr "]" ;; exists until
| "A" "[" ctl_expr "U" ctl_expr "]" ;; forall until

```

It is an error for an expressions in a CTL formula to contain a ‘next()’ operator, or to have non-boolean components, i.e. subformulas which evaluate to a value other than 0 or 1.

It is also possible to specify invariants, i.e. propositional formulas which must hold invariantly in the model. The corresponding command is ‘INVARSPEC’, with syntax:

```
checkinvar_declaration :: "INVARSPEC" simple_expr ";"
```

This statement corresponds to

```
SPEC AG simple_expr ";"
```

but can be checked by a specialized algorithm during reachability analysis.

3.3.2 LTL Specifications

LTL specifications are introduced by the keyword ‘LTLSPEC’. The syntax of this declaration is:

```
ltlspec_declaration :: "LTLSPEC" ltl_expr [";"]
```

where

```

ltl_expr ::
  simple_expr                ;; a simple boolean expression
| "(" ltl_expr ")"
| "!" ltl_expr              ;; logical not
| ltl_expr "&" ltl_expr      ;; logical and
| ltl_expr "|" ltl_expr     ;; logical or
| ltl_expr "xor" ltl_expr   ;; logical exclusive or
| ltl_expr "->" ltl_expr    ;; logical implies
| ltl_expr "<->" ltl_expr    ;; logical equivalence
;; FUTURE
| "X" ltl_expr              ;; next state
| "G" ltl_expr              ;; globally
| "F" ltl_expr              ;; finally
| ltl_expr "U" ltl_expr     ;; until
| ltl_expr "V" ltl_expr     ;; releases
;; PAST
| "Y" ltl_expr              ;; previous state
| "Z" ltl_expr              ;; not previous state not
| "H" ltl_expr              ;; historically
| "O" ltl_expr              ;; once
| ltl_expr "S" ltl_expr     ;; since
| ltl_expr "T" ltl_expr    ;; triggered

```

In NuSMV, LTL specifications can be analyzed both by means of BDD-based reasoning, or by means of SAT-based bounded model checking. In the first case, NuSMV proceeds along the lines described in [CGH97]. For each LTL specification, a tableau able to recognize the behaviors falsifying the property is constructed, and then synchronously composed with the model. With

respect to [CGH97], the approach is fully integrated within NuSMV, and allows for full treatment of past temporal operators. In the case of BDD-based reasoning, the counterexample generated to show the falsity of a LTL specification may contain state variables which have been introduced by the tableau construction procedure.

In the second case, a similar tableau construction is carried out to encode the existence of a path of limited length violating the property. NuSMV generates a propositional satisfiability problem, that is then tackled by means of an efficient SAT solver [BCCZ99].

In both cases, the tableau constructions are completely transparent to the user.

3.3.3 Real Time CTL Specifications and Computations

NuSMV allows for Real Time CTL specifications [EMSS90]. NuSMV assumes that each transition takes unit time for execution. RTCTL extends the syntax of CTL path expressions with the following bounded modalities:

```
rtctl_expr ::
    ctl_expr
  | "EBF" range rtctl_expr
  | "ABF" range rtctl_expr
  | "EBG" range rtctl_expr
  | "ABG" range rtctl_expr
  | "A" "[" rtctl_expr "BU" range rtctl_expr "]"
  | "E" "[" rtctl_expr "BU" range rtctl_expr "]"

range    :: number ".." number
```

Intuitively, in the formula $E [a \text{ BU } m..n b]$ m (n , resp.) represents the minimum (maximum) number of permitted transition along a path of a structure before the eventuality holds.

Real time CTL specifications can be defined with the following syntax, which extends the syntax for CTL specifications.

```
spec_declaration :: "SPEC" rtctl_expr [";"]
```

With the 'COMPUTE' statement, it is also possible to compute quantitative information on the FSM. In particular, it is possible to compute the exact bound on the delay between two specified events, expressed as CTL formulas. The syntax is the following:

```
compute_declaration :: "COMPUTE" compute_expr [";"]
where
compute_expr ::
    "MIN" "[" rtctl_expr "," rtctl_expr "]"
  | "MAX" "[" rtctl_expr "," rtctl_expr "]"
```

MIN [*start* , *final*] computes the set of states reachable from *start*. If at any point, we encounter a state satisfying *final*, we return the number of steps taken to reach the state. If a fixed point is reached and no states intersect *final* then *infinity* is returned.

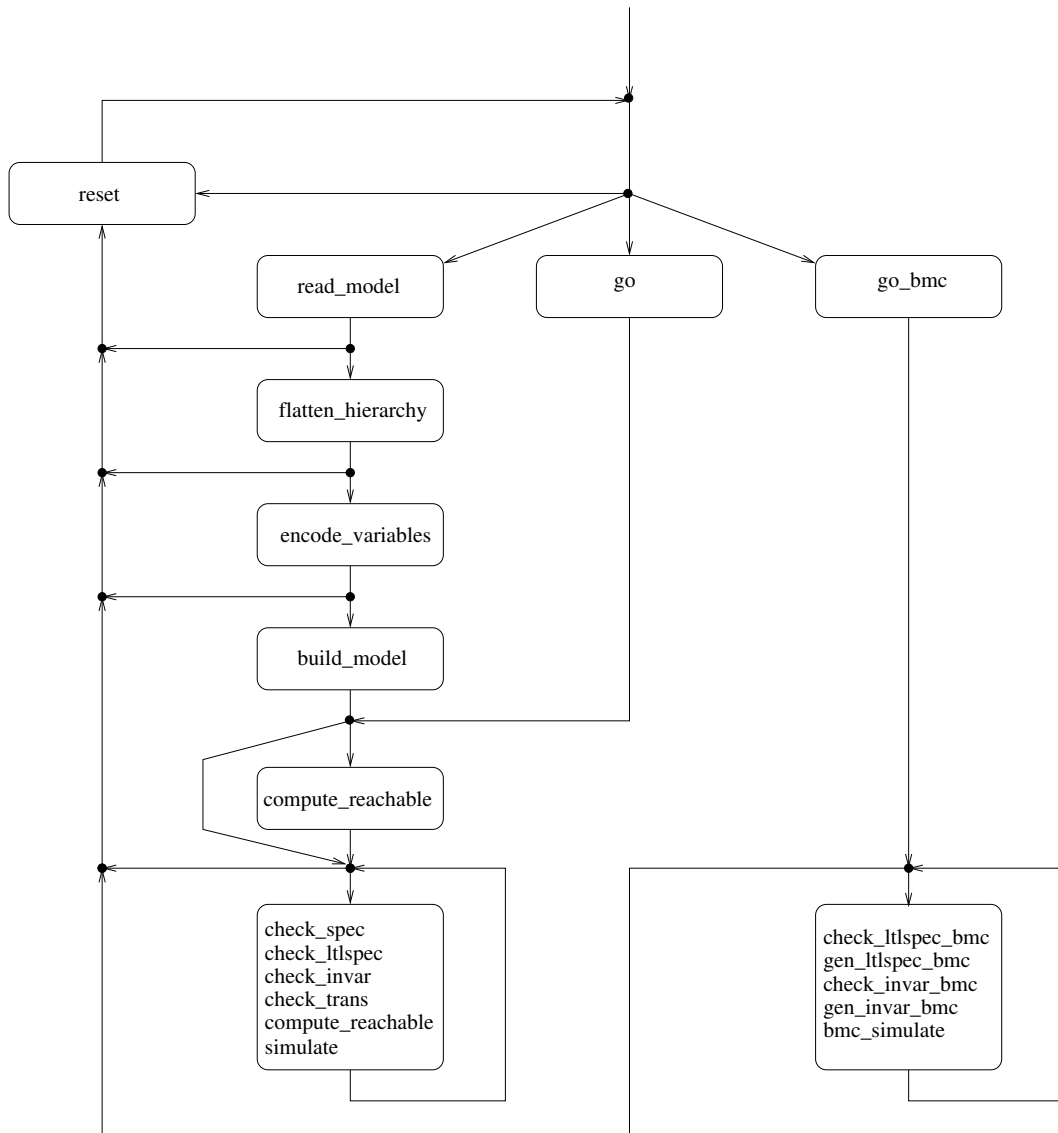
MAX [*start* , *final*] returns the length of the longest path from a state in *start* to a state in *final*. If there exists an infinite path beginning in a state in *start* that never reaches a state in *final*, then *infinity* is returned.

4 Running NuSMV interactively

The main interaction mode of NuSMV is through an interactive shell. In this mode NuSMV enters a read-eval-print loop. The user can activate the various NuSMV computation steps as system commands with different options. These steps can therefore be invoked separately, possibly undone or repeated under different modalities. These steps include the construction of the model under different partitioning techniques, model checking of specifications, and the configuration of the BDD package. The interactive shell of NuSMV is activated from the system prompt as follows ('NuSMV>' is the default NuSMV shell prompt):

```
system_prompt> NuSMV -int (RET)
NuSMV>
```

A NuSMV command is a sequence of words. The first word specifies the command to be executed. The remaining words are arguments to the invoked command. Commands separated by a ';' are executed sequentially; the NuSMV shell waits for each command to terminate in turn. The behavior of commands can depend on environment variables, similar to "csh" environment variables.



In the following we present the possible commands followed by the related environment variables, classified in different categories. Every command answers to the option `-h` by printing

out the command usage. When output is paged for some commands (option `-m`), it is piped through the program specified by the UNIX `PAGER` shell variable, if defined, or through UNIX command "more". Environment variables can be assigned a value with the `set` command.

Command sequences to NuSMV must obey the (partial) order specified in the figure depicted in the previous page. For instance, it is not possible to evaluate CTL expressions before the model is built.

The verbosity of NuSMV is controlled by the following environment variable.

verbose_level Environment Variable
 Controls the verbosity of the system. Possible values are integers from 0 (no messages) to 4 (full messages). The default value is 0.

4.1 Model Reading and Building

The following commands allow for the parsing and the compilation of the model into BDD.

read_model - *Reads a NuSMV file into NuSMV.* Command

`read_model [-h] [-i model-file]`

Reads a NuSMV file. If the `-i` option is not specified, it reads from the file specified in the environment variable `input_file`.

Command options:

`-i model-file`

Sets the environment variable `input_file` to `model-file`, and reads the model from the specified file.

input_file Environment Variable

Stores the name of the input file containing the model. It can be set by the `'set'` command or by the command line option `'-i'`. There is no default value.

flatten_hierarchy - *Flattens the hierarchy of modules* Command

`flatten_hierarchy [-h]`

This command is responsible of the instantiation of modules and processes. The instantiation is performed by substituting the actual parameters for the formal parameters, and then by prefixing the result via the instance name.

show_vars - *Shows model's symbolic variables and their values* Command

`show_vars [-h] [-s] [-i] [-m | -o output-file]`

Prints symbolic input and state variables of the model with their range of values (as defined in the input file).

Command Options:

`-s`

Prints only state variables.

`-i`

Prints only input variables.

`-m`

Pipes the output to the program specified by the `PAGER` shell variable if defined, else through the UNIX command "more".

`-o output-file`

Writes the output generated by the command to `output-file`

encode_variables - *Builds the BDD variables necessary to compile the model into BDD.* Command

`encode_variables [-h] [-i order-file]`

Generates the boolean BDD variables and the ADD needed to encode propositionally the (symbolic) variables declared in the model.

The variables are created as default in the order in which they appear in a depth first traversal of the hierarchy.

The input order file can be partial and can contain variables not declared in the model. Variables not declared in the model are simply discarded. Variables declared in the model which are not listed in the ordering input file will be created and appended at the end of the given ordering list, according to the default ordering.

Command options:

-i order-file

Sets the environment variable `input_order_file` to `order-file`, and reads the variable ordering to be used from file `order-file`. This can be combined with the `write_order` command. The variable ordering is written to a file, which can be inspected and reordered by the user, and then read back in.

input_order_file Environment Variable

Indicates the file name containing the variable ordering to be used in building the model by the 'encode_variables' command. There is no default value.

write_order - *Writes variable order to file.* Command

`write_order [-h] [(-o | -f) order-file]`

Writes the current order of BDD variables in the file specified via the `-o` option. If no option is specified the environment variable `output_order_file` will be considered. If the variable `output_order_file` is unset (or set to an empty value) then standard output will be used.

Command options:

-o order-file

Sets the environment variable `output_order_file` to `order-file` and then dumps the ordering list into that file.

-f order-file

Alias for `-o` option. Supplied for backward compatibility.

output_order_file Environment Variable

The file where the current variable ordering has to be written. The default value is 'temp.ord'.

build_model - *Compiles the flattened hierarchy into BDD* Command

`build_model [-h] [-f] [-m Method]`

Compiles the flattened hierarchy into BDD (initial states, invariants, and transition relation) using the method specified in the environment variable `partition_method` for building the transition relation.

Command options:

-m Method

Sets the environment variable `partition_method` to the value `Method`, and then builds the transition relation. Available methods are `Monolithic`, `Threshold` and `Iwls95CP`.

-f

Forces model construction. By default, only one partition method is allowed. This option allows to overcome this default, and to build the transition relation with different partitioning methods.

partition_method

Environment Variable

The method to be used in building the transition relation, and to compute images and preimages. Possible values are:

- **Monolithic.** No partitioning at all.
- **Threshold.** Conjunctive partitioning, with a simple threshold heuristic. Assignments are collected in a single cluster until its size grows over the value specified in the variable `conj_part_threshold`. It is possible (default) to use affinity clustering to improve model checking performance. See `affinity` variable.
- **Iwls95CP.** Conjunctive partitioning, with clusters generated and ordered according to the heuristic described in [RAP+95]. Works in conjunction with the variables `image_cluster_size`, `image_W1`, `image_W2`, `image_W3`, `image_W4`. It is possible (default) to use affinity clustering to improve model checking performance. See `affinity` variable. It is also possible to avoid (default) preordering of clusters (see [RAP+95]) using `iwls95preorder` variable.

conj_part_threshold

Environment Variable

The limit of the size of clusters in conjunctive partitioning. The default value is 0 BDD nodes.

affinity

Environment Variable

Enables affinity clustering heuristic described in [MOON00], possible values are 0 or 1. The default value is 1.

image_cluster_size, image_W{1,2,3,4}

Environment Variables

The parameters to configure the behavior of the *Iwls95CP* partitioning algorithm. `image_cluster_size` is used as threshold value for the clusters. The default value is 1000 BDD nodes. The other parameters attribute different weights to the different factors in the algorithm. The default values are 6, 1, 1, 2 respectively. (For a detailed description, please refer to [RAP+95].)

iwls95preorder

Environment Variable

Enables cluster preordering following heuristic described in [RAP+95], possible values are 0 or 1. The default value is 0. Preordering can be very slow.

image_verbosity

Environment Variable

Sets the verbosity for the image method *Iwls95CP*, possible values are 0 or 1. The default value is 0.

print_iwls95options - *Prints the Iwls95 Options.*

Command

`print_iwls95options [-h]`

This command prints out the configuration parameters of the IWLS95 clustering algorithm, i.e. `image_verbosity`, `image_cluster_size` and `image_W1,2,3,4`.

go - *Initializes the system for the verification.*

Command

`go [-h]`

This command initializes the system for verification. It is equivalent to the command sequence `read_model`, `flatten_hierarchy`, `encode_variables`, `build_model`, `build_flat_model`, `build_boolean_model`. If some commands have already been executed, then only the remaining ones will be invoked.

Command options:

`-h`

Prints the command usage.

process_model - *Performs the batch steps and then returns control to the interactive shell.* Command

`process_model [-h] [-i model-file] [-m Method]`

Reads the model, compiles it into BDD and performs the model checking of all the specification contained in it. If the environment variable `forward_search` has been set before, then the set of reachable states is computed. If the environment variables `enable_reorder` and `reorder_method` are set, then the reordering of variables is performed accordingly. This command simulates the batch behavior of NuSMV and then returns the control to the interactive shell.

Command options:

`-i model-file`

Sets the environment variable `input_file` to file `model-file`, and reads the model from file `model-file`.

`-m Method`

Sets the environment variable `partition_method` to `Method` and uses it as partitioning method.

4.2 Commands for Checking Specifications

The following commands allow for the BDD-based model checking of a NuSMV model.

compute_reachable - *Computes the set of reachable states* Command
`compute_reachable [-h]`

Computes the set of reachable states. The result is then used to simplify image and preimage computations. This can result in improved performances for models with sparse state spaces. Sometimes this option may slow down the performances because the computation of reachable states may be very expensive. The environment variable `forward_search` is set during the execution of this command.

print_reachable_states - *Prints out the number of reachable states. In verbose mode, prints also the list of reachable states.* Command

`print_reachable_states [-h] [-v]`

Prints the number of reachable states of the given model. In verbose mode, prints also the list of all reachable states. The reachable states are computed if needed.

check_trans - *Checks the transition relation for totality.* Command

`check_trans [-h] [-m | -o output-file]`

Checks if the transition relation is total. If the transition relation is not total then a potential deadlock state is shown out.

Command options:

`-m`

Pipes the output generated by the command to the program specified by the `PAGER` shell variable if defined, else through the UNIX command "more".

`-o output-file`

Writes the output generated by the command to the file `output-file`.

At the beginning reachable states are computed in order to guarantee that deadlock states are actually reachable.

check_trans

Environment Variable

Controls the activation of the totality check of the transition relation during the `process_model` call. Possible values are 0 or 1. Default value is 0.

check_spec - *Performs fair CTL model checking.*

Command

```
check_spec [-h] [-m | -o output-file] [-n number | -p "ctl-expr [IN context]"]
```

Performs fair CTL model checking.

A `ctl-expr` to be checked can be specified at command line using option `-p`. Alternatively, option `-n` can be used for checking a particular formula in the property database. If neither `-n` nor `-p` are used, all the SPEC formulas in the database are checked.

Command options:

`-m`

Pipes the output generated by the command in processing SPECS to the program specified by the `PAGER` shell variable if defined, else through the UNIX command "more".

`-o output-file`

Writes the output generated by the command in processing SPECS to the file `output-file`.

`-p "ctl-expr [IN context]"`

A CTL formula to be checked. `context` is the module instance name which the variables in `ctl-expr` must be evaluated in.

`-n number`

Checks the CTL property with index `number` in the property database.

If the `ag_only_search` environment variable has been set, and the set of reachable states has been already computed, then a specialized algorithm to check AG formulas is used instead of the standard model checking algorithms.

ag_only_search

Environment Variable

Enables the use of an ad hoc algorithm for checking AG formulas. The algorithm, given a formula of the form *AG alpha*, computes the set of states satisfying *alpha*, and checks whether it contains the reachable states is the empty set. If this is the case, then the property is verified, else a counterexample is printed.

forward_search

Environment Variable

Enables the computation of the reachable states during the `process_model` command and when used in conjunction with the `ag_only_search` environment variable enables the use of an ad hoc algorithm to verify invariants.

check_invar - *Performs model checking of invariants*

Command

```
check_invar [-h] [-m | -o output-file] [-n number | -p "invar-expr [IN context]"]
```

Performs invariant checking on the given model. An invariant is a set of states. Checking the invariant is the process of determining that all states reachable from the initial states lie in the invariant. Invariants to be verified can be provided as simple formulas (without any temporal operators) in the input file via the `INVARSPEC` keyword or directly at command line, using the option `-p`.

Option `-n` can be used for checking a particular invariant of the model. If neither `-n` nor `-p` are used, all the invariants are checked.

During checking of invariant all the fairness conditions associated with the model are ignored.

If an invariant does not hold, a proof of failure is demonstrated. This consists of a path starting from an initial state to a state lying outside the invariant. This path has the property that it is the shortest path leading to a state outside the invariant.

Command options:

- m**
Pipes the output generated by the program in processing INVARSPECs to the program specified by the `PAGER` shell variable if defined, else through the UNIX command "more".
- o output-file**
Writes the output generated by the command in processing INVARSPECs to the file `output-file`.
- p "invar-expr [IN context]"**
The command line specified invariant formula to be verified. `context` is the module instance name which the variables in `invar-expr` must be evaluated in.

check_ltlspec - *Performs LTL model checking* Command

```
check_ltlspec [-h] [-m | -o output-file] [-n number | -p "ltl-expr [IN context]"]
```

Performs model checking of LTL formulas. LTL model checking is reduced to CTL model checking as described in the paper by [CGH97].

A `ltl-expr` to be checked can be specified at command line using option `-p`. Alternatively, option `-n` can be used for checking a particular formula in the property database. If neither `-n` nor `-p` are used, all the LTLSPEC formulas in the database are checked.

Command options:

- m**
Pipes the output generated by the command in processing LTLSPECs to the program specified by the `PAGER` shell variable if defined, else through the Unix command "more".
- o output-file**
Writes the output generated by the command in processing LTLSPECs to the file `output-file`.
- p "ltl-expr [IN context]"**
An LTL formula to be checked. `context` is the module instance name which the variables in `ltl_expr` must be evaluated in.
- n number**
Checks the LTL property with index `number` in the property database.

compute - *Performs computation of quantitative characteristics* Command

```
compute [-h] [-m | -o output-file] [-n number | -p "compute-expr [IN context]"]
```

This command deals with the computation of quantitative characteristics of real time systems. It is able to compute the length of the shortest (longest) path from two given set of states.

MAX [alpha , beta]

MIN [alpha , beta]

Properties of the above form can be specified in the input file via the keyword `COMPUTE` or directly at command line, using option `-p`.

Option `-n` can be used for computing a particular expression in the model. If neither `-n` nor `-p` are used, all the COMPUTE specifications are computed.

Command options:

- `-m`
Pipes the output generated by the command in processing COMPUTEs to the program specified by the `PAGER` shell variable if defined, else through the UNIX command "more".
- `-o output-file`
Writes the output generated by the command in processing COMPUTEs to the file `output-file`.
- `-p "compute-expr [IN context]"`
A COMPUTE formula to be checked. `context` is the module instance name which the variables in `compute-expr` must be evaluated in.
- `-n number`
Computes only the property with index `number`

add_property - *Adds a property to the list of properties* Command
`add_property [-h] [(-c | -l | -i | -q) -p "formula [IN context]"`

Adds a property in the list of properties. It is possible to insert LTL, CTL, INVAR and quantitative (COMPUTE) properties. Every newly inserted property is initialized to unchecked. A type option must be given to properly execute the command.

Command options:

- `-c`
Adds a CTL property.
- `-l`
Adds an LTL property.
- `-i`
Adds an INVAR property.
- `-q`
Adds a quantitative (COMPUTE) property.
- `-p "formula [IN context]"`
Adds the `formula` specified on the command-line.
`context` is the module instance name which the variables in `formula` must be evaluated in.

4.3 Commands for Bounded Model Checking

In this section we describe in detail the commands for doing and controlling Bounded Model Checking in NuSMV.

Bounded Model Checking is based on the reduction of the bounded model checking problem to a propositional satisfiability problem. After the problem is generated, NuSMV internally calls a propositional SAT solver in order to find an assignment which satisfies the problem. Currently NuSMV supplies two SAT solvers: SIM and Zchaff. Notice that Zchaff is for non-commercial purposes only, and is therefore not included in the source code distribution, as well as in some of the binary distributions of NuSMV. It is also possible to generate the satisfiability problem without calling the SAT solver. Each generated problem is dumped in DIMACS format into a file. DIMACS is the standard format used as input by most external SAT solver, so it is possible to use NuSMV with an external SAT solver separately.

bmc_setup - *Builds the model in a Boolean Expression format.* Command

`bmc_setup [-h]`

You must call this command before use any other bmc-related command. Only one call per session is required.

go_bmc - *Initializes the system for the BMC verification.* Command

`go_bmc [-h]`

This command initializes the system for verification. It is equivalent to the command sequence `read_model`, `flatten_hierarchy`, `encode_variables`, `build_boolean_model`, `bmc_setup`. If some commands have already been executed, then only the remaining ones will be invoked.

Command options:

`-h`

Prints the command usage.

check_ltlspec_bmc - *Checks the given LTL specification, or all LTL* Command

specifications if no formula is given. Checking parameters are the maximum length and the loopback values

`check_ltlspec_bmc [-h | -n idx | -p "formula" [IN context]] [-k max_length] [-l loopback] [-o filename]`

This command generates one or more problems, and calls SAT solver for each one. Each problem is related to a specific problem bound, which increases from zero (0) to the given maximum problem length. Here "*length*" is the bound of the problem that system is going to generate and/or solve.

In this context the maximum problem bound is represented by the `-k` command parameter, or by its default value stored in the environment variable `bmc.length`.

The single generated problem also depends on the "*loopback*" parameter you can explicitly specify by the `-l` option, or by its default value stored in the environment variable `bmc.loopback`.

The property to be checked may be specified using the `-n idx` or the `-p "formula"` options. If you need to generate a dimacs dump file of all generated problems, you must use the option `-o "filename"`.

Command options:

`-n index`

index is the numeric index of a valid LTL specification formula actually located in the properties database.

`-p "formula" [IN context]`

Checks the `formula` specified on the command-line.

`context` is the module instance name which the variables in `formula` must be evaluated in.

`-k max_length`

max_length is the maximum problem bound must be reached. Only natural number are valid values for this option. If no value is given the environment variable `bmc.length` is considered instead.

`-l loopback`

loopback value may be:

- a natural number in $(0, \text{max_length}-1)$. Positive sign ('+') can be also used as prefix of the number. Any invalid combination of length and loopback will be skipped during the generation/solving process.

- a negative number in $(-1, -bmc_length)$. In this case *loopback* is considered a value relative to *max_length*. Any invalid combination of length and loopback will be skipped during the generation/solving process.
- the symbol 'X', which means "no loopback"
- the symbol '*', which means "all possible loopback from zero to *length-1*"

-o *filename*

filename is the name of the dumped dimacs file. It may contain special symbols which will be macro-expanded to form the real file name. Possible symbols are:

- @F: model name with path part
- @f: model name without path part
- @k: current problem bound
- @l: current loopback value
- @n: index of the currently processed formula in the properties database
- @@: the '@' character

check_ltlspec_bmc_onepb - Checks the given LTL specification, or all LTL specifications if no formula is given. Checking parameters are the single problem bound and the loopback values Command

`check_ltlspec_bmc_onepb [-h | -n idx | -p "formula" [IN context]] [-k length] [-l loopback] [-o filename]`

As command `check_ltlspec_bmc` but it produces only one single problem with fixed bound and loopback values, with no iteration of the problem bound from zero to *max_length*.

Command options:

-n *index*

index is the numeric index of a valid LTL specification formula actually located in the properties database.

The validity of *index* value is checked out by the system.

-p "formula [IN context]"

Checks the formula specified on the command-line.

context is the module instance name which the variables in **formula** must be evaluated in.

-k *length*

length is the problem bound used when generating the single problem. Only natural number are valid values for this option. If no value is given the environment variable *bmc_length* is considered instead.

-l *loopback*

loopback value may be:

- a natural number in $(0, max_length-1)$. Positive sign ('+') can be also used as prefix of the number. Any invalid combination of length and loopback will be skipped during the generation/solving process.
- a negative number in $(-1, -bmc_length)$. In this case *loopback* is considered a value relative to *length*. Any invalid combination of length and loopback will be skipped during the generation/solving process.
- the symbol 'X', which means "no loopback"
- the symbol '*', which means "all possible loopback from zero to *length-1*"

-o *filename*

filename is the name of the dumped dimacs file. It may contain special symbols which will be macro-expanded to form the real file name. Possible symbols are:

- @F: model name with path part
- @f: model name without path part
- @k: current problem bound
- @l: current loopback value
- @n: index of the currently processed formula in the properties database
- @@: the '@' character

gen_ltlspec_bmc - *Dumps into one or more dimacs files the given LTL specification, or all LTL specifications if no formula is given. Generation and dumping parameters are the maximum bound and the loopback values* Command

`gen_ltlspec_bmc [-h | -n idx | -p "formula" [IN context]] [-k max_length] [-l loopback] [-o filename]`

This command generates one or more problems, and dumps each problem into a dimacs file. Each problem is related to a specific problem bound, which increases from zero (0) to the given maximum problem bound. In this short description "*length*" is the bound of the problem that system is going to dump out.

In this context the maximum problem bound is represented by the *max_length* parameter, or by its default value stored in the environment variable *bmc_length*.

Each dumped problem also depends on the loopback you can explicitly specify by the *-l* option, or by its default value stored in the environment variable *bmc_loopback*.

The property to be checked may be specified using the *-n idx* or the *-p "formula"* options. You may specify dimacs file name by using the option *-o "filename"*, otherwise the default value stored in the environment variable *bmc_dimacs_filename* will be considered.

Command options:

-n index

index is the numeric index of a valid LTL specification formula actually located in the properties database.

The validity of *index* value is checked out by the system.

-p "formula" [IN context]"

Checks the *formula* specified on the command-line.

context is the module instance name which the variables in *formula* must be evaluated in.

-k max_length

max_length is the maximum problem bound used when increasing problem bound starting from zero. Only natural number are valid values for this option. If no value is given the environment variable *bmc_length* value is considered instead.

-l loopback

loopback value may be:

- a natural number in (0, *max_length-1*). Positive sign ('+') can be also used as prefix of the number. Any invalid combination of bound and loopback will be skipped during the generation and dumping process.

- a negative number in (-1, *-bmc_length*). In this case *loopback* is considered a value relative to *max_length*. Any invalid combination of bound and loopback will be skipped during the generation process.

- the symbol 'X', which means "no loopback"

- the symbol '*', which means "all possible loopback from zero to *length-1*"

-o filename

filename is the name of dumped dimacs files. If this options is not specified, variable *bmc_dimacs_filename* will be considered. The file name string may contain special symbols which will be macro-expanded to form the real file

name. Possible symbols are:

- @F: model name with path part
- @f: model name without path part
- @k: current problem bound
- @l: current loopback value
- @n: index of the currently processed formula in the properties database
- @@: the '@' character

gen_ltlspec_bmc_onepb - *Dumps into one dimacs file the problem generated for the given LTL specification, or for all LTL specifications if no formula is explicitly given. Generation and dumping parameters are the problem bound and the loopback values* Command

```
gen_ltlspec_bmc_onepb [-h | -n idx | -p "formula" [IN context]] [-k length]
[-l loopback] [-o filename]
```

As the *gen_ltlspec_bmc* command, but it generates and dumps only one problem given its bound and loopback.

Command options:

-n *index*

index is the numeric index of a valid LTL specification formula actually located in the properties database.
The validity of *index* value is checked out by the system.

-p "formula [IN context]"

Checks the **formula** specified on the command-line.
context is the module instance name which the variables in **formula** must be evaluated in.

-k *length*

length is the single problem bound used to generate and dump it. Only natural number are valid values for this option. If no value is given the environment variable *bmc_length* is considered instead.

-l *loopback*

loopback value may be:

- a natural number in (0, *length-1*). Positive sign ('+') can be also used as prefix of the number. Any invalid combination of length and loopback will be skipped during the generation and dumping process.
- a negative number in (-1, -*length*). Any invalid combination of length and loopback will be skipped during the generation process.
- the symbol 'X', which means "no loopback"
- the symbol '*', which means "all possible loopback from zero to *length-1*"

-o *filename*

filename is the name of the dumped dimacs file. If this options is not specified, variable *bmc_dimacs_filename* will be considered. The file name string may contain special symbols which will be macro-expanded to form the real file name. Possible symbols are:

- @F: model name with path part
- @f: model name without path part
- @k: current problem bound
- @l: current loopback value
- @n: index of the currently processed formula in the properties database
- @@: the '@' character

bmc_length Environment Variable

Sets the generated problem bound. Possible values are any natural number, but must be compatible with the current value held by the variable *bmc_loopback*. The default value is 10.

bmc_loopback Environment Variable

Sets the generated problem loop. Possible values are:

- Any natural number, but lesser than the current value of the variable *bmc_length*. In this case the loop point is absolute.
- Any negative number, but greater or equal than *-bmc_length*. In this case specified loop is the loop length.
- The symbol 'X', which means "no loopback".
- The symbol '*', which means "any possible loopbacks".

The default value is *.

bmc_dimacs_filename Environment Variable

This is the default file name used when generating DIMACS problem dumps. This variable may be taken into account by all commands which belong to the *gen_ltlspec_bmc* family. DIMACS file name can contain special symbols which will be expanded to represent the actual file name. Possible symbols are:

- **@F** The currently loaded model name with full path.
- **@f** The currently loaded model name without full path.
- **@n** The numerical index of the currently processed formula in the properties database.
- **@k** The currently generated problem length.
- **@l** The currently generated problem loopback value.

The default value is "**@f_k@k_l@l_n@n.dimacs**".

check_invar_bmc - *Generates and solve the given invariant, or all invariants if no formula is given* Command

`check_invar_bmc [-h | -n idx | -p "formula" [IN context]] [-o filename]`

Command options:

`-n index`

index is the numeric index of a valid INVAR specification formula actually located in the properties database.

The validity of *index* value is checked out by the system.

`-p "formula [IN context]"`

Checks the *formula* specified on the command-line.

context is the module instance name which the variables in *formula* must be evaluated in.

`-o filename`

filename is the name of the dumped dimacs file. It may contain special symbols which will be macro-expanded to form the real file name. Possible symbols are:

- **@F**: model name with path part
- **@f**: model name without path part
- **@n**: index of the currently processed formula in the properties database
- **@@**: the '@' character

gen_invar_bmc - Generates the given invariant, or all invariants if no formula is given Command

`gen_invar_bmc [-h | -n idx | -p "formula" [IN context]] [-o filename]`

Command options:

`-n index`

index is the numeric index of a valid INVAR specification formula actually located in the properties database.

The validity of *index* value is checked out by the system.

`-p "formula" [IN context]`

Checks the `formula` specified on the command-line.

`context` is the module instance name which the variables in `formula` must be evaluated in.

`-o filename`

filename is the name of the dumped dimacs file. If you do not use this option the dimacs file name is taken from the environment variable `bmc_invar_dimacs_filename`.

File name may contain special symbols which will be macro-expanded to form the real dimacs file name. Possible symbols are:

- @F: model name with path part
- @f: model name without path part
- @n: index of the currently processed formula in the properties database
- @@: the '@' character

bmc_invar_dimacs_filename Environment Variable

This is the default file name used when generating DIMACS invar dumps. This variable may be taken into account by the command `gen_invar_bmc`. DIMACS file name can contain special symbols which will be expanded to represent the actual file name. Possible symbols are:

- @F The currently loaded model name with full path.
- @f The currently loaded model name without full path.
- @n The numerical index of the currently processed formula in the properties database.

The default value is "`@f_invar_n@n.dimacs`".

sat_solver Environment Variable

The SAT solver's name actually to be used. Default SAT solver is SIM. Depending on the NuSMV configuration, also the Zchaff SAT solver can be available or not. Notice that Zchaff is for non-commercial purposes only.

bmc_simulate - Generates a trace of the model from 0 (zero) to *k* Command

`bmc_simulate [-h | -k]`

`bmc_simulate` does not require a specification to build the problem, because only the model is used to build it. The problem length is represented by the `-k` command parameter, or by its default value stored in the environment variable `bmc_length`.

Command options:

`-k length`

length is the length of the generated simulation.

4.4 Simulation Commands

In this section we describe the commands that allow to simulate a NuSMV specification.

pick_state - *Picks a state from the set of initial states* Command
pick_state [-h] [-v] [-r | -i [-a]] [-c "constraints"]

Chooses an element from the set of initial states, and makes it the **current state** (replacing the old one). The chosen state is stored as the first state of a new trace ready to be lengthened by **steps** states by the **simulate** command. The state can be chosen according to different policies which can be specified via command line options. By default the state is chosen in a deterministic way.

Command Options:

-v

Verbosely prints out chosen state (all state variables, otherwise it prints out only the label **t.1** of the state chosen, where **t** is the number of the new trace, that is the number of traces so far generated plus one).

-r

Randomly picks a state from the set of initial states.

-i

Enables the user to interactively pick up an initial state. The user is requested to choose a state from a list of possible items (every item in the list doesn't show state variables unchanged with respect to a previous item). If the number of possible states is too high, then the user has to specify some further constraints as "simple expression".

-a

Displays all state variables (changed and unchanged with respect to a previous item) in an interactive picking. This option works only if the **-i** options has been specified.

-c "constraints"

Uses **constraints** to restrict the set of initial states in which the state has to be picked. **constraints** must be enclosed between double quotes " ".

show_traces - *Shows the traces generated in a NuSMV session* Command

show_traces [[-h] [-v] [-m | -o output-file] -t | -a | trace_number]

Shows the traces currently stored in system memory, if any. By default it shows the last generated trace, if any.

Command Options:

-v

Verbosely prints traces content (all state variables, otherwise it prints out only those variables that have changed their value from previous state).

-t

Prints only the total number of currently stored traces.

-a

Prints all the currently stored traces.

-m

Pipes the output through the program specified by the **PAGER** shell variable if defined, else through the UNIX command "more".

-o output-file
Writes the output generated by the command to **output-file**

trace_number
The (ordinal) identifier number of the trace to be printed.

showed_states Environment Variable
Controls the maximum number of states showed during an interactive simulation session. Possible values are integers from 1 to 100. The default value is 25.

simulate - *Performs a simulation from the current selected state* Command

simulate [-h] [-p | -v] [-r | -i [-a]] [-c "constraints"] steps

Generates a sequence of at most **steps** states (representing a possible execution of the model), starting from the *current state*. The current state must be set via the *pick_state* or *goto_state* commands.

It is possible to run the simulation in three ways (according to different command line policies): deterministic (the default mode), random and interactive.

The resulting sequence is stored in a trace indexed with an integer number taking into account the total number of traces stored in the system. There is a different behavior in the way traces are built, according to how *current state* is set: *current state* is always put at the beginning of a new trace (so it will contain at most **steps + 1** states) except when it is the last state of an existent old trace. In this case the old trace is lengthened by at most **steps** states.

Command Options:

-p
Prints current generated trace (only those variables whose value changed from the previous state).

-v
Verbosely prints current generated trace (changed and unchanged state variables).

-r
Picks a state from a set of possible future states in a random way.

-i
Enables the user to interactively choose every state of the trace, step by step. If the number of possible states is too high, then the user has to specify some constraints as simple expression. These constraints are used only for a single simulation step and are *forgotten* in the following ones. They are to be intended in an opposite way with respect to those constraints eventually entered with the *pick_state* command, or during an interactive simulation session (when the number of future states to be displayed is too high), that are *local* only to a single step of the simulation and are *forgotten* in the next one.

-a
Displays all the state variables (changed and unchanged) during every step of an interactive session. This option works only if the **-i** option has been specified.

-c "constraints"
Performs a simulation in which computation is restricted to states satisfying those **constraints**. The desired sequence of states could not exist if such constraints were too strong or it may happen that at some point of the simulation

a future state satisfying those constraints doesn't exist: in that case a trace with a number of states less than `steps` trace is obtained. Note: `constraints` must be enclosed between double quotes " ".

`steps`

Maximum length of the path according to the constraints. The length of a trace could contain less than `steps` states: this is the case in which simulation stops in an intermediate step because it may not exist any future state satisfying those constraints.

4.5 Traces Inspection Commands

A trace is a sequence of states corresponding to a possible execution of the model. Traces are created by NuSMV when a formula is found to be false; they are also generated by the simulation feature (Section 4.4 [Simulation Commands], page 46). Each trace has a number, and the states are numbered within the trace. Trace *n* has states *n.1*, *n.2*, *n.3*, "...".

The trace inspection commands of NuSMV allow to navigate along the traces produced by NuSMV. During the navigation, there is a *current state*, and the *current trace* is the trace the *current state* belongs to. The commands are the following:

`goto_state` - Goes to a given state of a trace Command

`goto_state [-h] state`

Makes `state` the *current state*. This command is used to navigate alongs traces produced by NuSMV. During the navigation, there is a *current state*, and the *current trace* is the trace the *current state* belongs to.

`print_current_state` - Prints out the current state Command

`print_current_state [-h] [-v]`

Prints the name of the *current state* if defined.

Command options:

`-v`

Prints the value of all the state variables of the *current state*.

4.6 Interface to the DD Package

NuSMV uses the state of the art BDD package CUDD [Som98]. Control over the BDD package can very important to tune the performance of the system. In particular, the order of variables is critical to control the memory and the time required by operations over BDDs. Reordering methods can be activated to determine better variable orders, in order to reduce the size of the existing BDDs. Reordering methods can be activated either

Reordering of the variables can be triggered in two ways: by the user, or by the BDD package. In the first way, reordering is triggered by the interactive shell command `dynamic_var_ordering` with the `-f` option.

Reordering is triggered by the BDD package when the number of nodes reaches a given threshold. The threshold is initialized and automatically adjusted after each reordering by the package. This is called dynamical reordering, and can be enabled or disabled by the user. Dynamic reordering is enabled with the shell command `dynamic_var_ordering` with the option `-e`, and disabled with the `-d` option.

`enable_reorder` Environment Variable

Specifies whether dynamic reordering is enabled (when value is '0') or disabled (when value is '1').

reorder_method

Environment Variable

Specifies the ordering method to be used when dynamic variable reordering is fired. The possible values, corresponding to the reordering methods available with the CUDD package, are listed below. The default value is **sift**.

sift: Moves each variable throughout the order to find an optimal position for that variable (assuming all other variables are fixed). This generally achieves greater size reductions than the window method, but is slower.

random: Pairs of variables are randomly chosen, and swapped in the order. The swap is performed by a series of swaps of adjacent variables. The best order among those obtained by the series of swaps is retained. The number of pairs chosen for swapping equals the number of variables in the diagram.

random_pivot:

Same as **random**, but the two variables are chosen so that the first is above the variable with the largest number of nodes, and the second is below that variable. In case there are several variables tied for the maximum number of nodes, the one closest to the root is used.

sift_converge:

The **sift** method is iterated until no further improvement is obtained.

symmetry_sift:

This method is an implementation of symmetric sifting. It is similar to sifting, with one addition: Variables that become adjacent during sifting are tested for symmetry. If they are symmetric, they are linked in a group. Sifting then continues with a group being moved, instead of a single variable.

symmetry_sift_converge:

The **symmetry_sift** method is iterated until no further improvement is obtained.

window{2,3,4}:

Permutates the variables within windows of n adjacent variables, where n can be either 2, 3 or 4, so as to minimize the overall BDD size.

window{2,3,4}_converge:

The **window{2,3,4}** method is iterated until no further improvement is obtained.

group_sift:

This method is similar to **symmetry_sift**, but uses more general criteria to create groups.

group_sift_converge:

The **group_sift** method is iterated until no further improvement is obtained.

annealing:

This method is an implementation of simulated annealing for variable ordering. This method is potentially very slow.

genetic:

This method is an implementation of a genetic algorithm for variable ordering. This method is potentially very slow.

exact:

This method implements a dynamic programming approach to exact reordering. It only stores a BDD at a time. Therefore, it is relatively efficient in terms of memory. Compared to other reordering strategies, it is very slow, and is not recommended for more than 16 boolean variables.

linear:

This method is a combination of sifting and linear transformations.

linear_conv:

The **linear** method is iterated until no further improvement is obtained.

dynamic_var_ordering - Deals with the dynamic variable ordering. Command

`dynamic_var_ordering [-d] [-e <method>] [-f <method>] [-h]`

Controls the application and the modalities of (dynamic) variable ordering. Dynamic ordering is a technique to reorder the BDD variables to reduce the size of the existing BDDs. When no options are specified, the current status of dynamic ordering is displayed. At most one of the options `-e`, `-f`, and `-d` should be specified.

Dynamic ordering may be time consuming, but can often reduce the size of the BDDs dramatically. A good point to invoke dynamic ordering explicitly (using the `-f` option) is after the commands `build_model`, once the transition relation has been built. It is possible to save the ordering found using `write_order` in order to reuse it (using `build_model -i order-file`) in the future.

Command options:

`-d`

Disable dynamic ordering from triggering automatically.

`-e <method>`

Enable dynamic ordering to trigger automatically whenever a certain threshold on the overall BDD size is reached. `<method>` must be one of the following:

sift: Moves each variable throughout the order to find an optimal position for that variable (assuming all other variables are fixed). This generally achieves greater size reductions than the window method, but is slower.

random: Pairs of variables are randomly chosen, and swapped in the order. The swap is performed by a series of swaps of adjacent variables. The best order among those obtained by the series of swaps is retained. The number of pairs chosen for swapping equals the number of variables in the diagram.

random_pivot: Same as **random**, but the two variables are chosen so that the first is above the variable with the largest number of nodes, and the second is below that variable. In case there are several variables tied for the maximum number of nodes, the one closest to the root is used.

sift_converge: The **sift** method is iterated until no further improvement is obtained.

symmetry_sift: This method is an implementation of symmetric sifting. It is similar to sifting, with one addition: Variables that become adjacent during sifting are tested for symmetry. If they are symmetric, they are linked in a group. Sifting then continues with a group being moved, instead of a single variable.

symmetry_sift_converge: The **symmetry_sift** method is iterated until no further improvement is obtained.

window2,3,4: Permutes the variables within windows of "n" adjacent variables, where "n" can be either 2, 3 or 4, so as to minimize the overall BDD size.

window2,3,4_converge: The **window2,3,4** method is iterated until no further improvement is obtained.

group_sift: This method is similar to **symmetry_sift**, but uses more general criteria to create groups.

group_sift_converge: The **group_sift** method is iterated until no further improvement is obtained.

annealing: This method is an implementation of simulated annealing for variable ordering. This method is potentially very slow.

genetic: This method is an implementation of a genetic algorithm for variable ordering. This method is potentially very slow.

exact: This method implements a dynamic programming approach to exact reordering. It only stores a BDD at a time. Therefore, it is relatively efficient in terms of memory. Compared to other reordering strategies, it is very slow, and is not recommended for more than 16 boolean variables.

linear: This method is a combination of sifting and linear transformations.

linear.converge: The **linear** method is iterated until no further improvement is obtained.

-f <method>

Force dynamic ordering to be invoked immediately. The values for **<method>** are the same as in option **-e**.

print_bdd_stats - *Prints out the BDD statistics and parameters* Command

print_bdd_stats [-h]

Prints the statistics for the BDD package. The amount of information depends on the BDD package configuration established at compilation time. The configuration parameters are printed out too. More information about statistics and parameters can be found in the documentation of the CUDD Decision Diagram package.

set_bdd_parameters - *Creates a table with the value of all currently active NuSMV flags and change accordingly the configurable parameters of the BDD package.* Command

set_bdd_parameters [-h] [-s]

Applies the variables table of the NuSMV environment to the BDD package, so the user can set specific BDD parameters to the given value. This command works in conjunction with the **print_bdd_stats** and **set** commands.

print_bdd_stats first prints a report of the parameters and statistics of the current **bdd_manager**. By using the command **set**, the user may modify the value of any of the parameters of the underlying BDD package. The way to do it is by setting a value in the variable **BDD.parameter name** where **parameter name** is the name of the parameter exactly as printed by the **print_bdd_stats** command.

Command options:

-s

Prints the BDD parameter and statistics after the modification.

4.7 Administration Commands

This section describes the administrative commands offered by the interactive shell of NuSMV.

! shell_command Command

Executes a shell command. The *shell_command* is executed by calling **bin/sh -c shell_command**. If the command does not exist or you have not the right to execute it, then an error message is printed.

alias - *Provides an alias for a command* Command

alias [-h] [<name> [<string>]]

The "alias" command, if given no arguments, will print the definition of all current aliases. Given a single argument, it will print the definition of that alias (if any). Given two arguments, the keyword **<name>** becomes an alias for the command string **<string>**, replacing any other alias with the same name.

Command options:

<name>

Alias

<string>

Command string

It is possible to create aliases that take arguments by using the history substitution mechanism. To protect the history substitution character ‘%’ from immediate expansion, it must be preceded by a ‘\’ when entering the alias.

For example:

```
NuSMV> alias read "read_model -i \\\%:1.smv ; set input_order_file \\\%:1.ord"
```

```
NuSMV> read short
```

will create an alias ‘read’, execute “read_model -i short.smv; set input_order_file short.ord”.

And again:

```
NuSMV> alias echo2 "echo Hi ; echo \\\%* !"
```

```
NuSMV> echo2 happy birthday
```

will print:

```
Hi
```

```
happy birthday !
```

CAVEAT: Currently there is no check to see if there is a circular dependency in the alias definition. e.g.

```
NuSMV> alias foo "echo print_bdd_stats; foo"
```

creates an alias which refers to itself. Executing the command `foo` will result an infinite loop during which the command `print_bdd_stats` will be executed.

echo - *Merely echoes the arguments* Command

```
echo [-h] <args>
```

Echoes its arguments to standard output.

help - *Provides on-line information on commands* Command

```
help [-a] [-h] [<command>]
```

If invoked with no arguments “help” prints the list of all commands known to the command interpreter. If a command name is given, detailed information for that command will be provided.

Command options:

-a

Provides a list of all internal commands, whose names begin with the underscore character (‘_’) by convention.

history - *list previous commands and their event numbers* Command

```
history [-h] [<num>]
```

Lists previous commands and their event numbers. This is a UNIX-like history mechanism inside the NuSMV shell.

Command options:

<num>

Lists the last `<num>` events. Lists the last 30 events if `<num>` is not specified.

History Substitution:

The history substitution mechanism is a simpler version of the `cs`h history substitution mechanism. It enables you to reuse words from previously typed commands.

The default history substitution character is the ‘%’ (‘!’ is default for shell escapes, and ‘#’ marks the beginning of a comment). This can be changed using the "set" command. In this description ‘%’ is used as the history_char. The ‘%’ can appear anywhere in a line. A line containing a history substitution is echoed to the screen after the substitution takes place. ‘%’ can be preceded by a ‘\’ in order to escape the substitution, for example, to enter a ‘%’ into an alias or to set the prompt.

Each valid line typed at the prompt is saved. If the "history" variable is set (see help page for "set"), each line is also echoed to the history file. You can use the "history" command to list the previously typed commands.

Substitutions:

At any point in a line these history substitutions are available.

<code>!:0</code>	Initial word of last command.
<code>!:n</code>	n-th argument of last command.
<code>!\$</code>	Last argument of last command.
<code>!*</code>	All but initial word of last command.
<code>!!</code>	Last command.
<code>!stuf</code>	Last command beginning with "stuf".
<code>!n</code>	Repeat the n-th command.
<code>!-n</code>	Repeat the n-th previous command.
<code>~old~new</code>	Replace "old" with "new" in previous command. Trailing spaces are significant during substitution. Initial spaces are not significant.

print_usage - *Prints processor and BDD statistics.* Command

`print_usage [-h]`

Prints a formatted dump of processor-specific usage statistics, and BDD usage statistics. For Berkeley Unix, this includes all of the information in the `getrusage()` structure.

Command options:

`-h`
Prints the command usage.

quit - *exits NuSMV* Command

`quit [-h] [-s]`

Stops the program. Does not save the current network before exiting.

Command options:

`-s`
Frees all the used memory before quitting. This is slower, and it is used for finding memory leaks.

reset - *Resets the whole system.* Command

`reset [-h]`

Resets the whole system, in order to read in another model and to perform verification on it.

Command options:

`-h`
Prints the command usage.

set - *Sets an environment variable* Command

set [-h] [<name>] [<value>]

A variable environment is maintained by the command interpreter. The "set" command sets a variable to a particular value, and the "unset" command removes the definition of a variable. If "set" is given no arguments, it prints the current value of all variables.

Command options:

-h

Prints the command usage.

<name>

Variable name

<value>

Value to be assigned to the variable.

Interpolation of variables is allowed when using the set command. The variables are referred to with the prefix of '\$'. So for example, what follows can be done to check the value of a set variable:

```
NuSMV> set foo bar
NuSMV> echo $foo
bar
```

The last line "bar" will be the output produced by NuSMV.

Variables can be extended by using the character ':' to concatenate values. For example:

```
NuSMV> set foo bar
NuSMV> set foo $foo:foobar
NuSMV> echo $foo
bar:foobar
```

The variable `foo` is extended with the value `foobar`.

Whitespace characters may be present within quotes. However, variable interpolation lays the restriction that the characters ':' and '/' may not be used within quotes. This is to allow for recursive interpolation. So for example, the following is allowed

```
NuSMV> set "foo bar" this
NuSMV> echo $"foo bar"
this
```

The last line will be the output produced by NuSMV.

But in the following, the value of the variable `foo/bar` will not be interpreted correctly:

```
NuSMV> set "foo/bar" this
NuSMV> echo $"foo/bar"
foo/bar
```

If a variable is not set by the "set" command, then the variable is returned unchanged.

Different commands use environment information for different purposes. The command interpreter makes use of the following parameters:

autoexec

Defines a command string to be automatically executed after every command processed by the command interpreter. This is useful for things like timing commands, or tracing the progress of optimization.

open_path

"open_path" (in analogy to the shell-variable PATH) is a list of colon-separated strings giving directories to be searched whenever a file is opened for read. Typically the current directory (.) is the first item in this list. The standard system library (typically \$NuSMV_LIBRARY_PATH) is always implicitly appended to the current path. This provides a convenient short-hand mechanism for reaching standard library files.

nusmv_stderr

Standard error (normally stderr) can be re-directed to a file by setting the variable nusmv_stderr.

nusmv_stdout

Standard output (normally stdout) can be re-directed to a file by setting the variable nusmv_stdout.

source - *Executes a sequence of commands from a file* Command

`source [-h] [-p] [-s] [-x] <file> [<args>]`

Reads and executes commands from a file.

Command options:

`-p`

Prints a prompt before reading each command.

`-s`

Silently ignores an attempt to execute commands from a nonexistent file.

`-x`

Echoes each command before it is executed.

`<file>`

File name

Arguments on the command line after the filename are remembered but not evaluated. Commands in the script file can then refer to these arguments using the history substitution mechanism.

EXAMPLE:

Contents of test.scr:

```
read_model -i %:2
flatten_hierarchy
build_variables
build_model
compute_fairness
```

Typing "source test.scr short.smv" on the command line will execute the sequence

```
read_model -i short.smv
flatten_hierarchy
build_variables
build_model
compute_fairness
```

(In this case %:0 gets "source", %:1 gets "test.scr", and %:2 gets "short.smv".) If you type "alias st source test.scr" and then type "st short.smv bozo", you will execute

```
read_model -i bozo
flatten_hierarchy
build_variables
build_model
compute_fairness
```

because "bozo" was the second argument on the last command line typed. In other words,

command substitution in a script file depends on how the script file was invoked. Switches passed to a command are also counted as positional parameters. Therefore, if you type "st -x short.smv bozo", you will execute

```
read_model -i short.smv
flatten_hierarchy
build_variables
build_model
compute_fairness
```

To pass the "-x" switch (or any other switch) to "source" when the script uses positional parameters, you may define an alias. For instance, "alias srcx source -x".

time - *Provides a simple CPU elapsed time value* Command
time [-h]

Prints the processor time used since the last invocation of the "time" command, and the total processor time used since NuSMV was started.

unalias - *Removes the definition of an alias.* Command
unalias [-h] <alias-names>

Removes the definition of an alias specified via the alias command.

Command options:

<alias-names>
Aliases to be removed

unset - *Unsets an environment variable* Command
unset [-h] <variables>

A variable environment is maintained by the command interpreter. The "set" command sets a variable to a particular value, and the "unset" command removes the definition of a variable.

Command options:

-h
Prints the command usage.

<variables>
Variables to be unset

usage - *Provides a dump of process statistics* Command
usage [-h]

Prints a formatted dump of processor-specific usage statistics. For Berkeley Unix, this includes all of the information in the getrusage() structure.

which - *Looks for a file called "file_name"* Command
which [-h] <file_name>

Looks for a file in a set of directories which includes the current directory as well as those in the NuSMV path. If it finds the specified file, it reports the found file's path. The searching path is specified through the "set open_path" command in ".nusmvr".

Command options:

<file_name>
File to be searched

4.8 Other Environment Variables

The behavior of the system depends on the value of some environment variables. For instance, an environment variable specifies the partitioning method to be used in building the transition relation. The value of environment variables can be inspected and modified with the ‘set’ command. Environment variables can be either logical or utility.

autoexec Environment Variable
Defines a command string to be automatically executed after every command processed by the command interpreter. This may be useful for timing commands, or tracing the progress of optimization.

filec Environment Variable
Enables file completion a la "csh". If the system has been compiled with the "readline" library, the user is able to perform file completion by typing the `<TAB>` key (in a way similar to the file completion inside the "bash" shell). If the system has not been compiled with the "readline" library, a built-in method to perform file completion a la "csh" can be used. This method is enabled with the ‘set filec’ command. The "csh" file completion method can be also enabled if the "readline" library has been used. In this case the features offered by readline will be disabled.

shell_char Environment Variable
`shell_char` specifies a character to be used as shell escape. The default value of this environment variable is ‘!’.

history_char Environment Variable
`history_char` specifies a character to be used in history substitutions. The default value of this environment variable is ‘%’.

open_path Environment Variable
`open_path` (in analogy to the shell-variable `PATH`) is a list of colon-separated strings giving directories to be searched whenever a file is opened for read. Typically the current directory (‘.’) is first in this list. The standard system library (`NuSMV_LIBRARY_PATH`) is always implicitly appended to the current path. This provides a convenient short-hand mechanism for reaching standard library files.

nusmv_stderr Environment Variable
Standard error (normally `stderr`) can be re-directed to a file by setting the variable `nusmv_stderr`.

nusmv_stdout Environment Variable
Standard output (normally `stdout`) can be re-directed to a file by setting the internal variable `nusmv_stdout`.

nusmv_stdin Environment Variable
Standard input (normally `stdin`) can be re-directed to a file by setting the internal variable `nusmv_stdin`.

5 Running NuSMV batch

When the `-int` option is not specified, NuSMV runs as a batch program, in the style of SMV, performing (some of) the steps described in previous section in a fixed sequence.

```
system_prompt> NuSMV [command line options] input-file RET
```

The program described in *input-file* is processed, and the corresponding finite state machine is built. Then, if *input-file* contains formulas to verify, their truth in the specified structure is evaluated. For each formula which is not true a counterexample is printed.

The batch mode can be controlled with the following command line options:

```
NuSMV [-s] [-ctt] [-lp] [-n idx] [-v vl] [-cpp] [-is] [-ic] [-ils]
      [-ii] [-r] [-f] [-int] [-h | -help] [-i iv_file] [-o ov_file]
      [-AG] [-load script_file] [-reorder] [-dynamic] [-m method]
      [[-mono] | [-thresh cp_t] | [-cp cp_t] | [-iwls95 cp_t]] [-coi]
      [-noaffinity] [-iwls95preorder] [-bmc] [-bmc_length k]
      [-ofm fm_file] [-obm bm_file] [input-file]
```

where the meaning of the options is described below. If *input-file* is not provided in batch mode, then the model is read from standard input.

- `-s` Avoids to load the NuSMV commands contained in ‘~/`.nusmvrc`’ or in ‘`.nusmvrc`’ or in ‘`$$${NuSMV_LIBRARY_PATH}/master.nusmvrc`’.
- `-ctt` Checks whether the transition relation is total.
- `-lp` Lists all properties in SMV model
- `-n idx` Specifies which property of SMV model should be checked
- `-v verbose-level`
 Enables printing of additional information on the internal operations of NuSMV. Setting *verbose-level* to 1 gives the basic information. Using this option makes you feel better, since otherwise the program prints nothing until it finishes, and there is no evidence that it is doing anything at all. Setting the *verbose-level* higher than 1 enables printing of much extra information.
- `-cpp` Runs preprocessor on SMV files
- `-is` Does not check SPEC
- `-ic` Does not check COMPUTE
- `-ils` Does not check LTLSPEC
- `-ii` Does not check INVARSPEC
- `-r` Prints the number of reachable states before exiting. If the `-f` option is not used, the set of reachable states is computed.
- `-f` Computes the set of reachable states before evaluating CTL expressions.
- `-int` Starts interactive shell.
- `-help`
- `-h` Prints the command line help.
- `-i iv_file` Reads the variable ordering from file *iv_file*.
- `-o ov_file` Reads the variable ordering from file *ov_file*.
- `-AG` Verifies only AG formulas using an ad hoc algorithm (see documentation for the `ag_only_search` environment variable).

- load** *cmd-file*
Interprets NuSMV commands from file *cmd-file*.
- reorder** Enables variable reordering after having checked all the specification if any.
- dynamic** Enables dynamic reordering of variables
- m** *method* Uses *method* when variable ordering is enabled. Possible values for method are those allowed for the **reorder_method** environment variable (see Section 4.6 [Interface to the DD package], page 48).
- mono** Enables monolithic transition relation
- thresh** *cp_t*
conjunctive partitioning with threshold of each partition set to *cp_t* (DEFAULT, with *cp_t*=1000)
- cp** *cp_t* DEPRECATED: use -thresh instead.
- iwls95** *cp_t*
Enables Iwls95 conjunctive partitioning and sets the threshold of each partition to *cp_t*
- coi** Enables cone of influence reduction
- noaffinity**
Disables affinity clustering
- iwls95preoder**
Enables iwls95 preordering
- bmc** Enables BMC instead of BDD model checking
- bmc** *k* Sets *bmc_length* variable, used by BMC
- ofm** *fn_file*
prints flattened model to file *fn_file*
- obm** *bn_file*
Prints boolean model to file *bn_file*

Bibliography

- [BCCZ99] A. Biere, A. Cimatti, E. Clarke, and Y. Zhu. Symbolic Model Checking without BDDs. *Tools and Algorithms for Construction and Analysis of Systems*, In TACAS'99, March 1999.
- [CCG+02] A. Cimatti, E. Clarke, E. Giunchiglia, F. Giunchiglia, M. Pistore, M. Roveri, R. Sebastiani, and A. Tacchella. NuSMV 2: An OpenSource Tool for Symbolic Model Checking. *Proceedings of Computer Aided Verification (CAV 02)*, 2002.
- [CCGR00] A. Cimatti, E. Clarke, F. Giunchiglia, and M. Roveri. NuSMV: a new symbolic model checker.. *International Journal on Software Tools for Technology Transfer (STTT)*, 2(4), March 2000.
- [BCLM+94] J.R. Burch, E.M. Clarke, D.E. Long, K.L. McMillan, and D.L. Dill. Symbolic Model Checking for Sequential Circuit Verification. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 13(4):401–424, April 1994.
- [CGH97] E. Clarke, O. Grumberg and K. Hamaguchi. Another Look at LTL Model Checking. *Formal Methods in System Design*, 10(1):57–71, February 1997.
- [CMB90] O. Coudert, C. Berthet, and J. C. Madre. Verification of synchronous sequential machines based on symbolic execution. In J. Sifakis, editor, *Proceedings of the International Workshop on Automatic Verification Methods for Finite State Systems*, volume 407 of *LNCS*, pages 365–373, Berlin, June 1990. Springer.
- [Dil88] D. Dill. *Trace Theory for Automatic Hierarchical Verification of Speed-Independent Circuits*. ACM Distinguished Dissertations. MIT Press, 1988.
- [EMSS91] E. Allen Emerson, A. K. Mok, A. Prasad Sistla, and Jai Srinivasan. Quantitative temporal reasoning. In Edmund M. Clarke and Robert P. Kurshan, editors, *Proceedings of Computer-Aided Verification (CAV '90)*, volume 531 of *LNCS*, pages 136–145, Berlin, Germany, June 1991. Springer.
- [McMil92] K.L. McMillan. *The SMV system – DRAFT*. 1992. Available at <http://www.cs.cmu.edu/~modelcheck/smv/smvmanual.r2.2.ps>.
- [McMil93] K.L. McMillan. *Symbolic Model Checking*. Kluwer Academic Publ., 1993.
- [Mart85] A.J. Martin. The design of a self-timed circuit for distributed mutual exclusion. In H. Fuchs and W.H. Freeman, editors, *Proceedings of the 1985 Chapel Hill Conference on VLSI*, pages 245–260, New York, 1985.
- [MOON00] Moon, Hachtel, Somenzi, Border-Block Tringular Form and Conjunction Schedule in Image Computation, to appear at FMCAD00.
- [RAP+95] R. K. Ranjan, A. Aziz, B. Plessier, C. Pixley, and R. K. Brayton. Efficient BDD algorithms for FSM synthesis and verification. In *IEEE/ACM Proceedings International Workshop on Logic Synthesis*, Lake Tahoe (NV), May 1995.
- [Som98] F. Somenzi. CUDD: CU Decision Diagram package — release 2.2.0 Department of Electrical and Computer Engineering — University of Colorado at Boulder, May 1998.
- [Vis96] "VIS: A system for Verification and Synthesis", The VIS Group, In the Proceedings of the 8th International Conference on Computer Aided Verification, p428-432, Springer Lecture Notes in Computer Science, #1102, Edited by R. Alur and T. Henzinger, New Brunswick, NJ, July 1996

Appendix A Compatibility with CMU SMV

The NuSMV language is mostly source compatible with the original version of SMV distributed at Carnegie Mellon University from which we started. In this appendix we describe the most common problems that can be encountered when trying to use old CMU SMV programs with NuSMV.

The main problem is variable names in old programs that conflicts with new reserved words. The list of the new reserved words of NuSMV w.r.t. CMU SMV is the following:

F, G, X, U, V, W, H, O, Y, Z, S, T, B

These names are reserved for the LTL temporal operators.

LTLSPEC It is used to introduce LTL specifications.

INVARSPEC

It is used to introduce invariant specifications.

IVAR It is used to introduce input variables.

JUSTICE It is used to introduce "justice" fairness constraints.

COMPASSION

It is used to introduce "compassion" fairness constraints.

The **IMPLEMENTS**, **INPUT**, **OUTPUT** statements are not supported by NuSMV. They are parsed from the input file, but are internally ignored.

NuSMV differs from CMU SMV also in the controls that are performed on the input formulas. Several formulas that are valid for CMU SMV, but that have no clear semantics, are not accepted by NuSMV. In particular:

- It is no longer possible to write formulas containing nested 'next'.

TRANS

```
next(alpha & next(beta | next(gamma))) -> delta
```

- It is no longer possible to write formulas containing 'next' in the right hand side of "normal" and "init" assignments (they are allowed in the right hand side of "next" assignments), and with the statements 'INVAR' and 'INIT'.

INVAR

```
next(alpha) & beta
```

INIT

```
next(beta) -> alpha
```

ASSIGN

```
delta := alpha & next(gamma);      -- normal assignments
```

```
init(gamma) := alpha & next(delta); -- init assignments
```

- It is no longer possible to write 'SPEC', 'FAIRNESS' statements containing 'next'.

FAIRNESS

```
next(running)
```

SPEC

```
next(x) & y
```

- The check for circular dependencies among variables has been done more restrictive.

We say that variable x depends on variable y if $x := f(y)$. We say that there is a circular dependency in the definition of x if:

- x depends on itself (e.g. $x := f(x,y)$);
- x depends on y and y depends on x (e.g. $x := f(y)$ and $y := f(x)$ or $x := f(z)$, $z := f(y)$ and $y := f(x)$).

In the case of circular dependencies among variables there is no fixed order in which we can compute the involved variables. Avoiding circular dependencies among variables guarantee that there exists an order in which the variables can be computed. In NuSMV circular dependencies are not allowed.

In CMU SMV the test for circular dependencies is able to detect circular dependencies only in "normal" assignments, and not in "next" assignments. The circular dependencies check of NuSMV has been extended to detect circularities also in "next" assignments. For instance the following fragment of code is accepted by CMU SMV but discarded by NuSMV.

```
MODULE main
VAR
  y : boolean;
  x : boolean;
ASSIGN
  next(x) := x & next(y);
  next(y) := y & next(x);
```

Another difference between NuSMV and CMU SMV is in the variable order file. The variable ordering file accepted by NuSMV can be partial and can contain variables not declared in the model. Variables listed in the ordering file but not declared in the model are simply discarded. The variables declared in the model but not listed in the variable file provided in input are created at the end of the given ordering following the default ordering. All the ordering files generated by CMU SMV are accepted in input from NuSMV but the ordering files generated by NuSMV may be not accepted by CMU SMV. Notice that there is no guarantee that a good ordering for CMU SMV is also a good ordering for NuSMV. In the ordering files for NuSMV, identifier `_process_selector_` can be used to control the position of the variable that encodes process selection. In CMU SMV it is not possible to control the position of this variable in the ordering; it is hard-coded at the top of the ordering.

Command Index

!

! 51

A

add_property 39
alias 51

B

bmc_setup 40
bmc_simulate 45
build_model 34

C

check_invar 37
check_invar_bmc 44
check_ltlspec 38
check_ltlspec_bmc 40
check_ltlspec_bmc_onepb 41
check_spec 37
check_trans 36
compute 38
compute_reachable 36

D

dynamic_var_ordering 50

E

echo 52
encode_variables 34

F

flatten_hierarchy 33

G

gen_invar_bmc 45
gen_ltlspec_bmc 42
gen_ltlspec_bmc_onepb 43
go 35
go_bmc 40
goto_state 48

H

help 52
history 52

P

pick_state 46
print_bdd_stats 51
print_current_state 48
print_iwls95options 35
print_reachable_states 36
print_usage 53
process_model 36

Q

quit 53

R

read_model 33
reset 53

S

set 54
set_bdd_parameters 51
show_traces 46
show_vars 33
simulate 47
source 55

T

time 56

U

unalias 56
unset 56
usage 56

W

which 56
write_order 34

Variable Index

A

affinity 35
 ag_only_search 37
 autoexec 57

B

bmc_dimacs_filename 44
 bmc_invar_dimacs_filename 45
 bmc_length 44
 bmc_loopback 44

C

check_trans 37
 conj_part_threshold 35

E

enable_reorder 48

F

filec 57
 forward_search 37

H

history_char 57

I

image_cluster_size, image_W{1,2,3,4} 35
 image_verbosity 35
 input_file 33

input_order_file 34
 iwls95preorder 35

N

NuSMV_LIBRARY_PATH 57, 58
 nusmv_stderr 57
 nusmv_stdin 57
 nusmv_stdout 57

O

open_path 57
 output_order_file 34

P

partition_method 35

R

reorder_method 49

S

sat_solver 45
 shell_char 57
 showed_states 47

V

verbose_level 33

Index

-

-AG	58
-bmc	59
-bmc <i>k</i>	59
-coi	59
-cp <i>cp_t</i>	59
-cpp	58
-ctt	58
-dynamic	59
-f	58
-h	58
-help	58
-i <i>iv_file</i>	58
-ic	58
-ii	58
-ils	58
-int	58
-is	58
-iwls95 <i>cp_t</i>	59
-iwls95preorder	59
-load <i>cmd_file</i>	59
-lp	58
-m <i>method</i>	59
-mono	59
-n <i>idx</i>	58
-noaffinity	59
-o <i>ov_file</i>	58
-obm <i>bm_file</i>	59
-ofm <i>fm_file</i>	59
-r	58
-reorder	59
-s	58
-thresh <i>cp_t</i>	59
-v <i>verbose-level</i>	58

•

.nusmvrc	58
----------	----

-

_process_selector	62
-------------------	----

~

~/nusmvrc	58
-----------	----

A

administration commands	51
-------------------------	----

B

batch, running NuSMV	58
Bounded Model Checking	14

C

case expressions	22
Commands for Bounded Model Checking	39
comments in NuSMV language	21
compassion constraints	29
CTL model checking	10
CTL Specifications	29

D

DD package interface	48
DEFINE declaration	25

E

Examples	2
expressions	21

F

fair paths	29
fairness constraint	4
fairness constraints	29
fairness constraints declaration	29
FAIRNESS declarations	29

I

identifiers	27
INIT declaration	25
input variables	24
interactive shell	32
interactive, running NuSMV	32
interface to DD Package	48
introduction	1
INVAR declaration	25
ISA declarations	26
IVAR	24

J

justice constraints	29
---------------------	----

L

LTL model checking	13
LTL Specifications	30

M

main module	28
master.nusmvrc	58
model compiling	33
model parsing	33
model reading	33
MODULE declarations	26

N

next expressions 22

O

options 58

P

process keyword 28

processes 28

R

Real Time CTL Specifications and Computations
..... 31

running 29

S

self 28

set expressions 22

Shell configuration Variables 57

simple expressions 21

Simulation 6

Simulation Commands 46

state variables syntax 23

syntax 21

T

'temp.ord' 34

Traces Inspection Commands 48

TRANS declarations 25

Tutorial 2

type declaration 23

type specifiers 23

V

VAR declaration 23