

Testing Distributed Systems through Symbolic Model Checking

G. Kalyon T. Massart C. Meuter L. Van Begin

Université Libre de Bruxelles
Département d'Informatique

FORTE'07

June 28, 2007

Need for Validation

Distributed control systems

- **concurrent** processes
- running on physically **distributed** hardware
- **hard** to design in nature
- **critical** systems (e.g. plant control system, ...)

Need for Validation

Distributed control systems

- **concurrent** processes
- running on physically **distributed** hardware
- **hard** to design in nature
- **critical** systems (e.g. plant control system, ...)

Validation Techniques

Model-Checking:

- **exhaustive** verification
- **but**: not (yet) scalable to real-sized systems
- works on a **model** of the system

Testing and Monitoring:

- **non-exhaustive** verification
- **scalable** for real-sized systems
- **widely used** in industry
- works on the **real** system

Testing

The System

- **distributed** and **asynchronous**
- instrumented to emit **relevant** events
(e.g. variable assignments, message transfer)
- execution **traces** are collected

Testing

The System

- **distributed** and **asynchronous**
- instrumented to emit **relevant** events
(e.g. variable assignments, message transfer)
- execution **traces** are collected

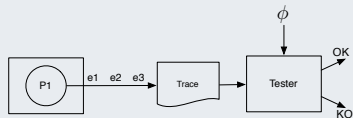
The Tester

- **analyses** the collected traces
- checks whether a certain **property** holds

Centralized v.s. Distributed Systems

Centralized System

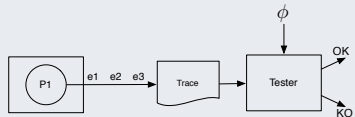
- only **one** process
- events are **totally** ordered



Centralized v.s. Distributed Systems

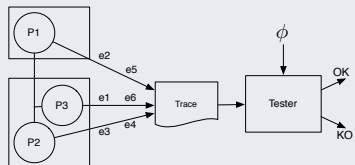
Centralized System

- only **one** process
- events are **totally** ordered



Distributed System

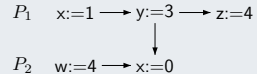
- **multiple** processes
- events are **not totally** ordered
- but a **partial order** can be obtained using **vector clocks** [Lam78, Mat89]



Partial Order Traces

Model

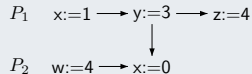
- **partially** ordered set of events
- events are labelled with **assignments**
- two **restrictions**:
 - 1 two events of the **same** process **must** be ordered
 - 2 two events assigning the **same** variable **must** be ordered



Partial Order Traces

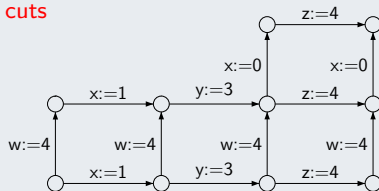
Model

- partially ordered set of events
- events are labelled with assignments
- two restrictions:
 - two events of the same process must be ordered
 - two events assigning the same variable must be ordered



Semantics

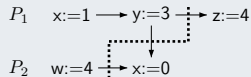
- a lattice of configurations called cuts



Partial Order Traces

Model

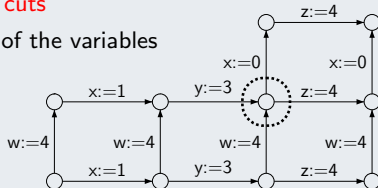
- partially ordered set of events
- events are labelled with assignments
- two restrictions:
 - two events of the same process must be ordered
 - two events assigning the same variable must be ordered



Semantics

- a lattice of configurations called cuts
- each cut has a unique valuation of the variables

$$v : \begin{cases} w & \mapsto & 4 \\ x & \mapsto & 1 \\ y & \mapsto & 3 \\ z & \mapsto & 0 \end{cases}$$



Global Predicate Detection

So far...

Efficient techniques have been developed for several classes of **predicate**:

- **Stable** predicates [CL85] (such that $\phi \implies AG\phi$)
- **Disjunctive** predicates (of the form $EF(p_1 \vee p_2 \vee \dots \vee p_n)$)
- **Conjunctive** predicates [GW94, GW96] (of the form $EF(p_1 \wedge p_2 \wedge \dots \wedge p_n)$)
- **Observer independent** predicates [CDF95] (such that $AF\phi \iff EF\phi$)
- **Linear, semi-linear** predicates [CG98]
- **Regular** predicates (RCTL logic) [GM01, SG03]

$$\phi ::= \top \mid p \mid \neg p \mid \phi \wedge \phi \mid EF\phi \mid EG\phi \mid AG\phi \mid EX[i]\phi$$

Global Predicate Detection

So far...

Efficient techniques have been developed for several classes of **predicate**:

- **Stable** predicates [CL85] (such that $\phi \implies \text{AG}\phi$)
- **Disjunctive** predicates (of the form $\text{EF}(p_1 \vee p_2 \vee \dots \vee p_n)$)
- **Conjunctive** predicates [GW94, GW96] (of the form $\text{EF}(p_1 \wedge p_2 \wedge \dots \wedge p_n)$)
- **Observer independent** predicates [CDF95] (such that $\text{AF}\phi \iff \text{EF}\phi$)
- **Linear, semi-linear** predicates [CG98]
- **Regular** predicates (RCTL logic) [GM01, SG03]

$$\phi ::= \top \mid p \mid \neg p \mid \phi \wedge \phi \mid \text{EF}\phi \mid \text{EG}\phi \mid \text{AG}\phi \mid \text{EX}[i]\phi$$

Our goal

Provide an **efficient** technique for **full CTL** (Computational Tree Logic):

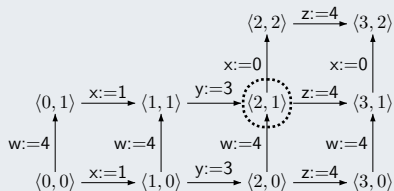
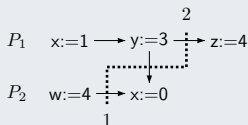
$$\phi ::= \top \mid p \mid \neg\phi \mid \phi \vee \phi \mid \phi \wedge \phi \mid \text{EX}\phi \mid \text{AX}\phi \mid \text{EG}\phi \mid \text{AG}\phi \mid \text{E}[\phi\text{U}\phi] \mid \text{A}[\phi\text{U}\phi]$$

Symbolic Representation

Multirectangles

Cuts can be viewed as **tuples** of naturals:

- in a tuple $\langle x_1, \dots, x_n \rangle$, x_i = the **number of events** process i has executed



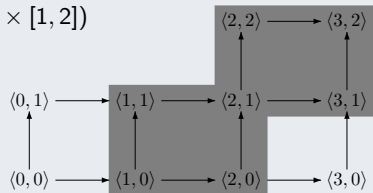
Symbolic Representation

Multirectangles

Cuts can be viewed as **tuples** of naturals:

- in a tuple $\langle x_1, \dots, x_n \rangle$, x_i = the **number of events** process i has executed
- sets of cut can be represented as a union of **multirectangles**

$$\blacksquare = ([1, 2] \times [0, 1]) \cup ([2, 3] \times [1, 2])$$



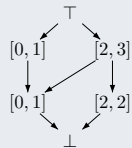
Symbolic Representation (cont'd)

Interval Sharing Tree (IST)

Symbolic data structure for representing sets of tuples:

- directed acyclic **graph**
- nodes are labelled with **intervals**
- each path from \top to \perp encodes a **multirectangle**

$$([0, 1] \times [0, 1]) \cup ([2, 3] \times [0, 1]) \cup ([2, 3] \times [2, 2])$$



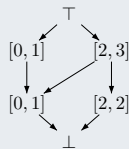
Symbolic Representation (cont'd)

Interval Sharing Tree (IST)

Symbolic data structure for representing sets of tuples:

- directed acyclic **graph**
- nodes are labelled with **intervals**
- each path from \top to \perp encodes a **multirectangle**

$$([0, 1] \times [0, 1]) \cup ([2, 3] \times [0, 1]) \cup ([2, 3] \times [2, 2])$$



Using IST for CTL

For each CTL formula ϕ , we build **symbolically** an IST \mathcal{I}_ϕ representing $\llbracket \phi \rrbracket$

Symbolic Computation

Tautology

If $\phi = \top$, we build \mathcal{I}_{\top} **iteratively**:

$$\begin{array}{l} x:=1 \longrightarrow y:=3 \longrightarrow z:=4 \\ \qquad \qquad \qquad \downarrow \\ w:=4 \longrightarrow x:=0 \end{array}$$

Symbolic Computation

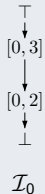
Tautology

If $\phi = \top$, we build \mathcal{I}_\top **iteratively**:

- start as if the trace had **no communications**

$x:=1 \longrightarrow y:=3 \longrightarrow z:=4$

$w:=4 \longrightarrow x:=0$



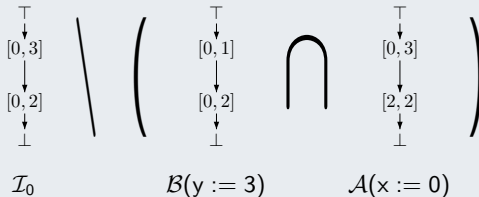
Symbolic Computation

Tautology

If $\phi = \top$, we build \mathcal{I}_{\top} **iteratively**:

- start as if the trace had **no communications**
- treat communications **one at a time**

$x:=1 \longrightarrow y:=3 \longrightarrow z:=4$
 $ $
 $ \downarrow$
 $w:=4 \longrightarrow x:=0$

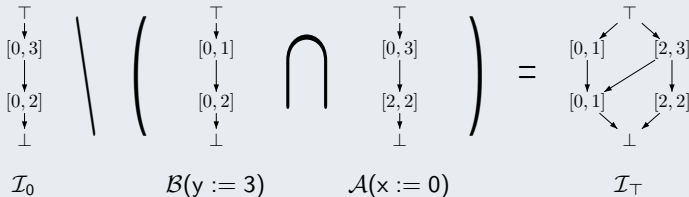
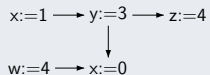


Symbolic Computation

Tautology

If $\phi = \top$, we build \mathcal{I}_{\top} **iteratively**:

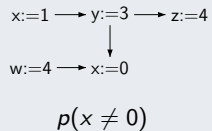
- start as if the trace had **no communications**
- treat communications **one at a time**
- stop when **all communications** have been treated



Symbolic Computation (cont'd)

Predicate

If $\phi = p \equiv (x \bullet c)$ with $\bullet \in \{=, \neq, <, \leq, >, \geq\}$, we build \mathcal{I}_p as follows:

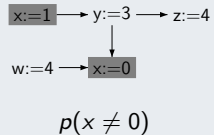


Symbolic Computation (cont'd)

Predicate

If $\phi = p \equiv (x \bullet c)$ with $\bullet \in \{=, \neq, <, \leq, >, \geq\}$, we build \mathcal{I}_p as follows:

- we know that all events assigning x are **totally ordered**

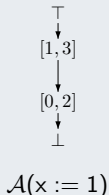
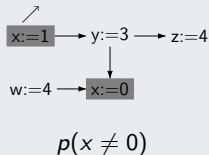


Symbolic Computation (cont'd)

Predicate

If $\phi = p \equiv (x \bullet c)$ with $\bullet \in \{=, \neq, <, \leq, >, \geq\}$, we build \mathcal{I}_p as follows:

- we know that all events assigning x are **totally ordered**
- some events sets p to **true**

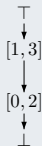
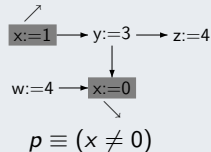


Symbolic Computation (cont'd)

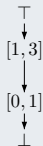
Predicate

If $\phi = p \equiv (x \bullet c)$ with $\bullet \in \{=, \neq, <, \leq, >, \geq\}$, we build \mathcal{I}_p as follows:

- we know that all events assigning x are **totally ordered**
- some events sets p to **true**
- some events sets p to **false**



$\mathcal{A}(x := 1)$



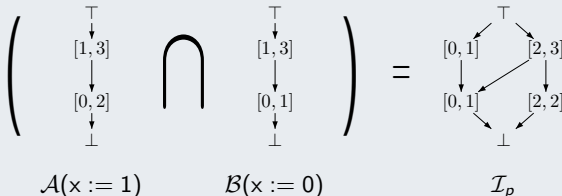
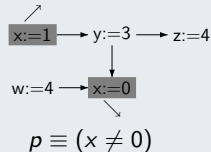
$\mathcal{B}(x := 0)$

Symbolic Computation (cont'd)

Predicate

If $\phi = p \equiv (x \bullet c)$ with $\bullet \in \{=, \neq, <, \leq, >, \geq\}$, we build \mathcal{I}_p as follows:

- we know that all events assigning x are **totally ordered**
- some events sets p to **true**
- some events sets p to **false**
- compute the union of all **slices** where:
 - p was set to **true**,
 - but not yet to **false**



Symbolic Computation (cont'd)

Boolean Operators

We can use **standard** set operations on IST:

$$\begin{aligned} \llbracket \phi_1 \wedge \phi_2 \rrbracket &= \llbracket \phi_1 \rrbracket \cap \llbracket \phi_2 \rrbracket \\ \llbracket \phi_1 \vee \phi_2 \rrbracket &= \llbracket \phi_1 \rrbracket \cup \llbracket \phi_2 \rrbracket \\ \llbracket \neg \phi \rrbracket &= \overline{\llbracket \phi \rrbracket} \end{aligned}$$

Symbolic Computation (cont'd)

Boolean Operators

We can use **standard** set operations on IST:

$$\begin{aligned} \llbracket \phi_1 \wedge \phi_2 \rrbracket &= \llbracket \phi_1 \rrbracket \cap \llbracket \phi_2 \rrbracket \\ \llbracket \phi_1 \vee \phi_2 \rrbracket &= \llbracket \phi_1 \rrbracket \cup \llbracket \phi_2 \rrbracket \\ \llbracket \neg \phi \rrbracket &= \overline{\llbracket \phi \rrbracket} \end{aligned}$$

Existential Modalities

- **symbolic** computation for $\text{EX}\phi$ using $\llbracket \text{EX}\phi \rrbracket = \bigcup_{i \in [1, k]} \llbracket \text{EX}[i]\phi \rrbracket$:

$$\mathcal{I}_{\text{EX}\phi} = \bigcup_{i \in [1, k]} (\mathcal{I}_{\phi}^{[x_i \leftarrow (x_i - 1)]} \cap \mathcal{I}_{\text{T}})$$

- for $\text{EG}\phi$ and $\text{E}[\phi_1 \text{U}\phi_2]$, use classical **fixed point**:

$$\begin{aligned} \llbracket \text{EG}\phi \rrbracket &= \mathbf{gfp} \lambda \psi \cdot \llbracket \phi \rrbracket \cap \llbracket \text{EX}\psi \rrbracket \\ \llbracket \text{E}[\phi_1 \text{U}\phi_2] \rrbracket &= \mathbf{lfp} \lambda \psi \cdot \llbracket \phi_2 \rrbracket \cup (\llbracket \phi_1 \rrbracket \cap \llbracket \text{EX}\psi \rrbracket) \end{aligned}$$

- for $\text{EF}\phi \stackrel{\text{def}}{=} \text{E}[\text{TU}\phi]$, we can compute $\mathcal{I}_{\text{EF}\phi} = \downarrow \mathcal{I}_{\phi} \cap \mathcal{I}_{\text{T}}$

Symbolic Computation (cont'd)

Universal Modalities

- **symbolic** computation for $AX\phi$ using $\llbracket AX\phi \rrbracket = \llbracket \neg EX\neg\phi \rrbracket$:

$$\mathcal{I}_{AX\phi} = \overline{\mathcal{I}_{EX\neg\phi}}$$

- for $AG\phi$ and $A[\phi_1 U \phi_2]$, use classical **fixed point**:

$$\begin{aligned} \llbracket AG\phi \rrbracket &= \mathbf{gfp} \lambda\psi \cdot \llbracket \phi \rrbracket \cap \llbracket AX\psi \rrbracket \\ \llbracket A[\phi_1 U \phi_2] \rrbracket &= \mathbf{lfp} \lambda\psi \cdot \llbracket \phi_2 \rrbracket \cup (\llbracket \phi_1 \rrbracket \cap \llbracket AX\psi \rrbracket) \end{aligned}$$

- for $AG\phi$, since $\llbracket AG\phi \rrbracket = \llbracket \neg EF\neg\phi \rrbracket$, we can compute $\mathcal{I}_{AG\phi} = \overline{\downarrow \mathcal{I}_{\neg\phi} \cap \mathcal{I}_T}$

Experimental Results

Model	#proc	#events	IST (in sec.)	NuSMV (in sec.)
Peterson	2	2000	0.46	349.57
	2	5000	7.53	↑↑
	2	15000	189.65	↑↑
Peterson Generalized	2	2000	0.20	294.46
	2	5000	6.44	↑↑
	2	20000	390.90	↑↑
	5	1000	2.04	13.74
	5	1500	6.82	↑↑
	5	5000	176.62	↑↑
	10	1500	7.53	150.23
	10	2000	27.01	↑↑
	10	5000	147.89	↑↑

↑↑ indicates (> 10 min.)

Experimental Results

Model	#proc	#events	IST (in sec.)	NuSMV (in sec.)
Alternating Bit Protocol	2	1000	13.60	297.28
	2	2000	27.56	↑↑
	2	5000	257.29	↑↑
Dining Philosophers	3	100	0.15	6.36
	3	200	1.11	↑↑
	3	2000	366.22	↑↑
	5	100	0.25	↑↑
	5	200	27.05	↑↑
	5	500	125.56	↑↑
	10	100	1.67	↑↑
	10	200	26.94	↑↑
	10	500	↑↑	↑↑

↑↑ indicates (> 10 min.)

Conclusion

What has been done?

- we proposed a **symbolic** technique for testing full CTL properties on distributed systems
- has been **implemented** and works **well** in practice

Future Work

- integration of the methods in our **tool** TraX
- interface with our **industrial** design environment $\mathcal{d}SL$ [DMM03,DGMM05]
- investigate possible further **improvements** of our technique
- investigate **similar** models (MSC...)
- comparison with other symbolic **data structures** (IDD,...)

Questions?