

Logic in Automatic Verification

Javier Esparza

Software Reliability and Security Group
Institute for Formal Methods in Computer Science
University of Stuttgart

Many thanks to Abdelwaheb Ayari, David Basin,
Armin Biere, Paul Gastin and Denis Oddoux

Automatic verification

The dream:

feed a machine with a system and a specification

push a button

get 'yes' or 'no, because ...'

In this talk: three **small** examples of application of decision procedures for logics to this problem

SAT / QBF

Temporal logics


Monadic second order logics

Verifying adders with boolean logic

Modelling circuits with QBL

Gates as boolean formulas

Stable states as satisfying truth assignments




$$\mathbf{not}(a, b) \equiv \neg a \leftrightarrow b$$

$$\mathbf{and}(a, b, c) \equiv (a \wedge b) \leftrightarrow c$$

$$\mathbf{or}(a, b, c) \equiv (a \vee b) \leftrightarrow c$$



$$\mathbf{xor}(a, b, c) \equiv ((\neg a \wedge b) \vee (a \wedge \neg b)) \leftrightarrow c$$

Combine gates with \wedge , \exists (and renaming of variables)



$$R(x, y, q, r, s) = \exists w. R_1(x, y, w, q) \wedge R_2(y, w, r, s)$$

A full adder




full_adder(*a*, *b*, *s*, *cin*, *cout*) \equiv

$$\exists w_1, w_2, w_3. \mathbf{xor}(a, b, w_1) \wedge \mathbf{xor}(w_1, cin, s) \wedge \mathbf{and}(a, b, w_2) \wedge \\ \mathbf{and}(cin, w_1, w_3) \wedge \mathbf{or}(w_3, w_2, cout)$$

An n -bit ripple-carry adder

$$\begin{array}{r} c_2 \quad c_1 \quad cin \quad (= 0) \\ a_2 \quad a_1 \quad a_0 \\ + \quad b_2 \quad b_1 \quad b_0 \\ \hline cout \quad s_2 \quad s_1 \quad s_0 \end{array}$$

Wire together n 1-bit adders where i th carry-out is $i+1$ st carry-in, first carry is the carry-in and last is the carry-out.



We obtain the formula

$$\begin{aligned} \mathbf{adder}_n(a_0, \dots, a_{n-1}, b_0, \dots, b_{n-1}, s_0, \dots, s_{n-1}, \mathit{cin}, \mathit{cout}) \equiv \\ \exists c_0, \dots, c_n. (c_0 \leftrightarrow \mathit{cin}) \wedge (c_n \leftrightarrow \mathit{cout}) \wedge \\ \bigwedge_{i=1}^{n-1} \mathbf{full_adder}(a_i, b_i, s_i, c_i, c_{i+1}) \end{aligned}$$

Problem: **too slow!!**

Each c_i can only be computed after all of c_{i-1}, \dots, c_0 have been computed

Delay: $2n + 2$ time units for n -bit numbers

A carry-look-ahead n -adder

Compute all of c_{n-1}, \dots, c_0 (and c_{out}) **concurrently**

First step: given $a_{n-1} \dots a_0$ and $b_{n-1} \dots b_0$, identify the $i \in [0, n - 1]$ that are

- **Generating:** $c_{i+1} \equiv 1$ independently of c_i .
These are the positions such that $1 = g_i \equiv \mathbf{and}(a_i, b_i)$.
- **Propagating:** $c_{i+1} \equiv c_i$, i.e., c_i is ‘propagated’ to c_{i+1} .
These are the positions such that $1 = p_i \equiv \mathbf{xor}(a_i, b_i)$

Observe: all g_i, p_i can be computed simultaneously

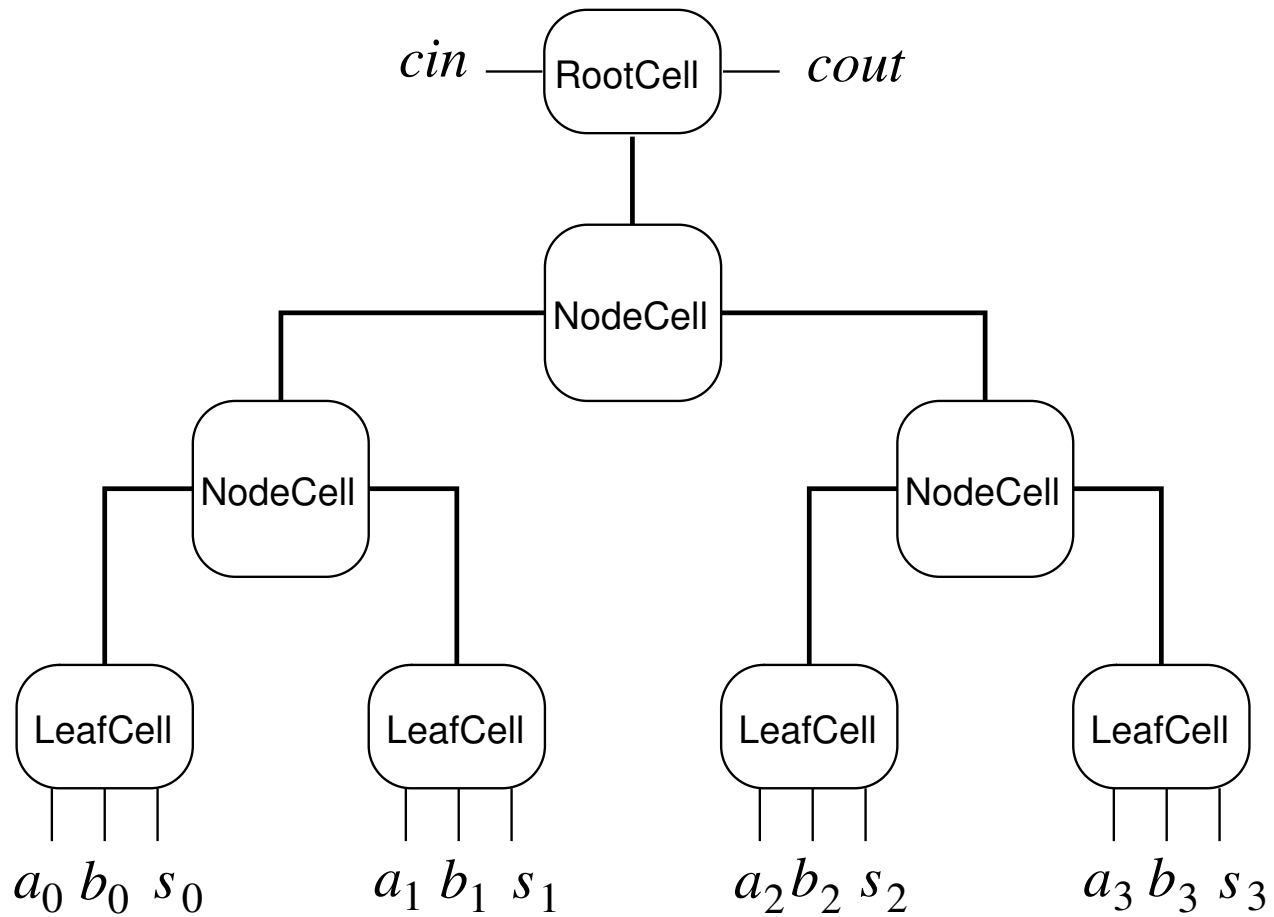
Second step: compute the c_i 's by

$$c_i \equiv g_i \vee (p_i \wedge g_{i-1}) \vee (p_i \wedge p_{i-1} \wedge g_{i-2}) \vee \dots \vee (p_i \wedge p_{i-1} \wedge \dots \wedge g_0)$$


Logarithmic delay in n using balanced \vee -trees and \wedge -trees

Delay depends on tree structure. For 64 bits: 23-56 units (instead of 130)


Description of the circuit (for 4 bits)




Description of the circuit II




LeafCell circuit



\otimes circuit



NodeCell circuit



RootCell circuit

Verification of the carry-look-ahead n -adder

Check if

adder $_n(a_0, \dots, a_{n-1}, b_0, \dots, b_{n-1}, s_0, \dots, s_{n-1}, cin, cout)$

\Leftrightarrow

cla $_n(a_0, \dots, a_{n-1}, b_0, \dots, b_{n-1}, s_0, \dots, s_{n-1}, cin, cout)$

Use SAT solvers or QBF solvers

Results of the [SAT 2002 competition](#) on a variant of this problem:

- Task was to compare 2, 4, 8, ..., 256 bits adders (8 problems)
- From 26 variables and 70 3CNF clauses to 4584 variables and 13226 clauses
- Fastest solver (Zchaff) checked all 8 problems in 14 seconds
- More info at www.satlive.org/SATCompetition/2002/index.jsp

Rule-of-thumb: circuits with some hundreds of gates are routinely solved

Verifying mutual exclusion algorithms
with propositional temporal logics

The mutual exclusion problem

The setting:

Two computers connected to a database (e.g., of plane bookings)

Can communicate with each other through **shared variables** (i.e., variables that both can read and write)

Both computers run a program having a **critical section**, from which the program can update the database records

The problem: design the program run by the computers so that

At any time, at most one computer can be in the critical section

If a computer wishes to enter the critical section, it eventually will

These properties still hold if any of the computers breaks down in the non-critical section

Observe: not an input/output system!

A solution due to Peterson

```
var flag[0], flag[1] : {true, false} init false;  
var turn : {0,1};
```

```
while true do
```

```
  s0 non-critical section
```

```
  s1 flag[0] := true;
```

```
  s2 turn := 1;
```

```
  s3 while (flag[1] and turn=1) skip ;
```

```
  s4 critical section
```

```
  s5 flag[0] := false;
```

```
od
```

```
while true do
```

```
  r0 non-critical section
```

```
  r1 flag[1] := true;
```

```
  r2 turn := 0;
```

```
  r3 while (flag[0] and turn=0) skip ;
```

```
  r4 critical section
```

```
  r5 flag[1] := false;
```

```
od
```

Linear-time temporal logic (LTL)

Built on top of a set AP of **atomic propositions**

World: valuation of the atomic propositions over $\{\text{true}, \text{false}\}$

Formulas of LTL interpreted over **runs**: infinite sequences of worlds

Notation: $\text{run } \rho = \rho_0\rho_1\rho_2 \dots$
 $\text{suffix } \rho|_i = \rho_i\rho_{i+1}\rho_{i+2} \dots$

Type	Formula	$\rho \models \varphi$ iff ...	Intuition
atomic	p	p is true at ρ_0	p holds now
boolean	$\neg\varphi$	$\rho \not\models \varphi$	
	$\varphi \vee \psi$	$\rho \models \varphi$ or $\rho \models \psi$	
temporal	$\mathbf{X}\varphi$	$\rho _1 \models \varphi$	φ holds next
	$\mathbf{F}\varphi$	$\rho _i \models \varphi$ for some $i \in \mathbb{N}$	eventually φ
	$\mathbf{G}\varphi$	$\rho _i \models \varphi$ for all $i \in \mathbb{N}$	always φ
	$\varphi \mathbf{U} \psi$	there is $i \in \mathbb{N}$ such that $\rho _i \models \psi$ and $\rho _j \models \varphi$ for all $0 \leq j < i$	φ until ψ

Application to the mutex algorithm

Atomic propositions: $\text{flag}[0] = \mathbf{true}$, at_s_4 , ...

The program satisfies a property if **all** its runs (executions) satisfy it

The mutual exclusion property:

$$G(\neg \text{at_s}_4 \vee \neg \text{at_r}_4)$$

If computer 0 wants to enter the critical section, it eventually will:

$$G(\text{flag}[0] = \mathbf{true} \Rightarrow F \text{ at_s}_4)$$

But this property does not take breakdowns **out of the non-critical section** into account ...

Dealing with breakdowns

Introduce propositions $last_0$, $last_1$ saying which computer did the last step

No breakdowns for computer 0:

$$G F last_0$$

No breakdowns for computer 0 but possibly in the non-critical section:

$$G F last_0 \vee F G at_s_0$$

The final property to be checked:

$$\begin{aligned} & (G F last_0 \vee F G at_s_0) \wedge (G F last_1 \vee F G at_r_0) \\ & \implies \\ & G(flag[0] = \mathbf{true} \Rightarrow F at_r_4) \wedge G(flag[1] = \mathbf{true} \Rightarrow F at_s_4) \end{aligned}$$

Automatic verification

The **model-checking** problem: whether all runs of the algorithm satisfy a given LTL formula

Can be algorithmically solved in three steps (Vardi, Wolper 85):

Construct a **Büchi automaton** for the negation of the formula
(decision procedure for satisfiability)

Construct the product of this automaton and the state space of the system

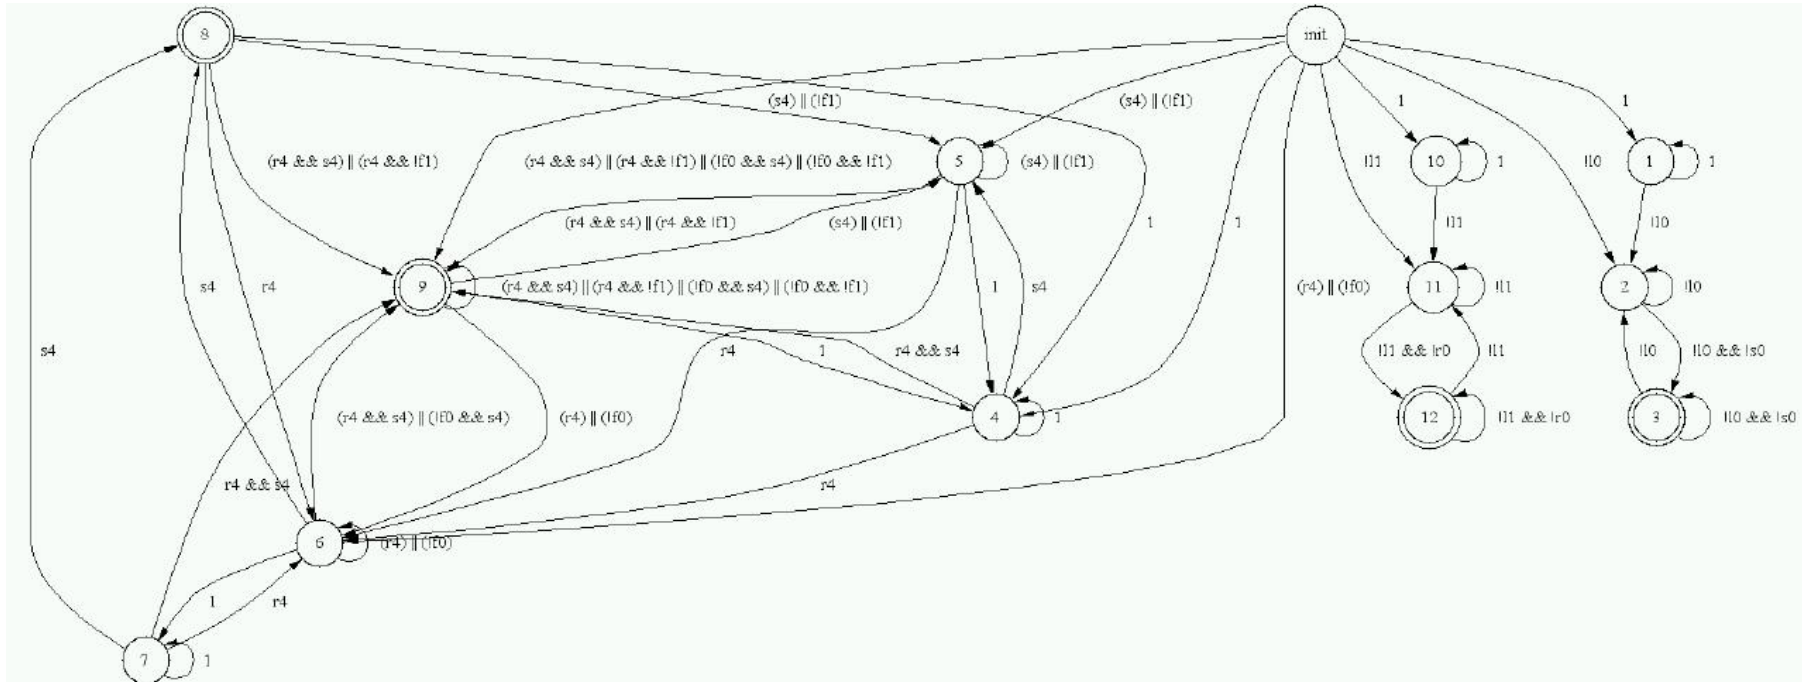
Check emptiness of the product

Linear complexity in the number of states of the program, **exponential complexity** in the size of the formula

Formula verified in less than one second with Holzmann's SPIN checker
(<http://spinroot.com/>)

Automaton for the formula

LTL2BA by Gastin and Oddoux (www.liafa.jussieu.fr/~oddoux/ltl2ba/)



Quite sophisticated: formula \rightarrow alternating Büchi \rightarrow generalized Büchi \rightarrow Büchi, with simplification heuristics at each step

The automaton for the negation of the formula has 36 states

Verifying parameterized adders
with monadic second order logics

WS1S : weak MSO logic of one successor

First order variables p, q, \dots interpreted over \mathbb{N}

Second-order variables X, Y, \dots interpreted over **finite subsets** of \mathbb{N}

$\phi ::= \mathbf{s}(q) \mid X(p) \mid \neg\phi \mid \phi \vee \phi \mid \exists p. \phi \mid \exists X. \phi$

Definitions (Sample)

$$\phi_1 \wedge \phi_2 \equiv \neg(\neg\phi_1 \vee \neg\phi_2)$$

$$\forall p. \phi \equiv \neg\exists p. \neg\phi$$

$$X(0) \equiv \exists p. (\forall q. p \neq \mathbf{s}(q)) \wedge X(p)$$

$$x = y \equiv \forall X. X(x) \leftrightarrow X(y)$$

$$X(p + n) \equiv \exists p_1, \dots, p_n. p_1 = \mathbf{s}(p) \wedge \dots \wedge p_n = \mathbf{s}(p_{n-1}) \wedge X(p_n)$$

$$x \leq y \equiv \forall X. (X(y) \wedge \forall z, w. (X(z) \wedge \mathbf{s}(w) = z \rightarrow X(w)) \rightarrow X(x))$$

$$x < y \equiv x \leq y \wedge \neg(x = y)$$

WS1S as a logic of binary strings

Second-order variables interpreted as strings over $\{0, 1\}$

First-order variables interpreted as positions in the string

‘ $X(p)$ holds iff string X has a 1 at position p ’

Formula ϕ with free variables determines a language $\mathcal{L}(\phi)$

$$1101 \in \mathcal{L}(X(1) \wedge X(3)) \quad 1011 \notin \mathcal{L}(X(1) \wedge X(3))$$

n free variables in ϕ determine language over $\{0, 1\}^n$

$$\forall p. p < 4 \rightarrow (X(p) \leftrightarrow \neg Y(p))$$

$$\begin{array}{c} X \\ Y \end{array} \begin{array}{|c|c|c|c|} \hline 0 & 1 & 1 & 0 \\ \hline 1 & 0 & 0 & 1 \\ \hline \end{array} \in L(\phi) \quad \text{and} \quad \begin{array}{c} X \\ Y \end{array} \begin{array}{|c|c|c|} \hline 0 & 1 & 1 \\ \hline 0 & 0 & 0 \\ \hline \end{array} \notin L(\phi)$$

Examples

$$\exists p, q. p \neq q \wedge X(p) \wedge X(q)$$

– X is a string with a 1 in at least 2 positions, e.g., 010100

$$\exists p. (\forall q. p \neq \mathbf{s}(q)) \wedge X(p)$$

– X is a string whose initial letter is 1

$$\forall p. X(p) \leftrightarrow Y(\mathbf{s}(p))$$

– Y is X ‘right-shifted’ 1 position, e.g.,

0	1	1	0
0	0	1	1

Well-known results

Satisfiability of WS1S is decidable in non-elementary time
(each quantifier alternation adds one exponential)

The language $\mathcal{L}(\phi)$ is regular

A finite automaton accepting $\mathcal{L}(\phi)$ can be computed directly from ϕ

Modelling the family of **ALL** ripple-carry adders

Recall the formula for a ripple carry n -adder

$$\begin{aligned} \mathbf{adder}_n(a_0, \dots, a_{n-1}, b_0, \dots, b_{n-1}, s_0, \dots, s_{n-1}, cin, cout) \equiv \\ \exists c_0, \dots, c_n. (c_0 \leftrightarrow cin) \wedge (c_n \leftrightarrow cout) \wedge \\ \bigwedge_{i=0}^{n-1} \mathbf{full_adder}(a_i, b_i, s_i, c_i, c_{i+1}) \end{aligned}$$

We construct the WS1S formula

$$\begin{aligned} \mathbf{adder}(n, A, B, S, cin, cout) \equiv \\ \exists C. (C(\mathbf{0}) \leftrightarrow cin) \wedge (C(n) \leftrightarrow cout) \wedge \\ \forall p. p < n \rightarrow \mathbf{full_adder}(A(p), B(p), S(p), C(p), C(p + \mathbf{1})) \wedge \\ \forall p. p \geq n \rightarrow (\neg A(p) \wedge \neg B(p) \wedge \neg S(p)) \end{aligned}$$

A model of **adder**(*A*, *B*, *S*, *cin*, *cout*)

	0	1	2	3	4	5			
A	1	1	1	0	0	0	...	7	n = 4
B	1	0	0	1	0	0	...	9	cin = 0
S	0	0	0	0	1	0	...	16	cout = 0

The set of models of **adder** is 'the union' of all the sets of models of **adder**_{*n*}

WS2S : weak MSO logic of two successors

Seen as a logic over **binary trees**

Second-order variables interpreted as **trees over $\{0, 1\}$**

First-order variables interpreted as **positions (nodes) in the tree**

Example:

$$X(\epsilon) \wedge (\forall p. X(\mathbf{s}_0(p)) \leftrightarrow X(\mathbf{s}_1(p))) \wedge \forall p. \neg Y(\mathbf{s}_0(p)) \vee \neg Y(\mathbf{s}_1(p))$$

‘ X contains the root node ϵ , and


a node p is in X iff its brother is also in X , and

for any node p , Y contains at most one of p 's successors’

Models

A model of a formula with n free variables is a ‘superposition’ of trees over \mathcal{B} , i.e., a tree whose nodes are labelled with elements of $\{0, 1\}^n$

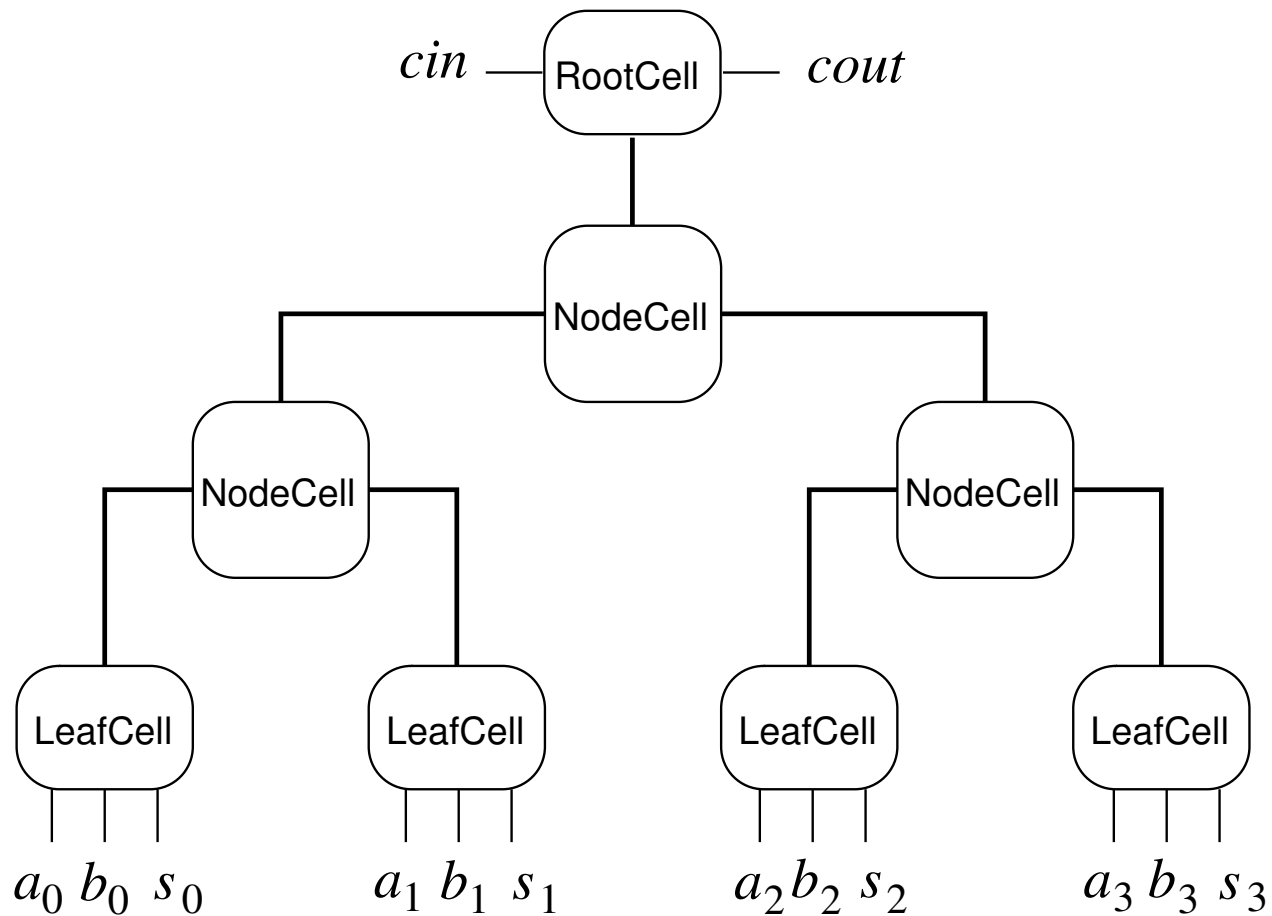
The tree



is a model of

$$X(\epsilon) \wedge (\forall p. X(\mathbf{s}_0(p)) \leftrightarrow X(\mathbf{s}_1(p))) \wedge \forall p. \neg Y(\mathbf{s}_0(p)) \vee \neg Y(\mathbf{s}_1(p))$$

Modelling the family of ALL carry-look-ahead adders



The family can be modelled by the formula

$$\begin{aligned} \mathbf{cla}(A, B, S, cin, cout) \equiv & \exists T, E_1, E_2, F_1, F_2 \\ & \mathbf{RootCell}(F_1, F_2, E_1, E_2, cin, cout) \wedge \\ & (\forall p. (\mathbf{leaf}(p, T) \rightarrow \mathbf{LeafCell}(A, B, S, F_1, F_2, E_1, E_2, p)) \wedge \\ & (\mathbf{node}(p, T) \rightarrow \mathbf{NodeCell}(F_1, F_2, E_1, E_2, p))) \wedge \\ & \mathbf{shape_cond}(A, B, S, T) \end{aligned}$$

Verification of a parameterized cla-adder

Check validity of

$$\forall A, B, S, cin, cout. \mathbf{adder}(A, B, S, cin, cout) \Leftrightarrow \mathbf{cla}(A, B, S, cin, cout)$$

(Requires to embed WS1S into WS2S)

Checked in 1 second by MONA (Mona at www.brics.dk/mona)

Restrictions:

- only array or tree structures
- only one parameter (two parameters \rightarrow quantification on binary relations)

Conclusions

Conclusions

No conclusions, just examples!