

Model Checking Algorithms and Reactive Systems

Satu Virtanen, satu@cs.hut.fi

12 November 2001

Branching time logics

$\mathbf{ML} ::= \mathcal{P} \mid \perp \mid (\mathbf{ML} \rightarrow \mathbf{ML}) \mid \langle \mathcal{R} \rangle \mathbf{ML}$

The validity of a multimodal formula φ is evaluated according to the following rule set.

- (i) $\mathcal{M} \models \mathbf{p}$ iff $w_0 \in \mathcal{I}(\mathbf{p})$
- (ii) $\mathcal{M} \not\models \perp$
- (iii) $\mathcal{M} \models (\varphi \rightarrow \psi)$ iff \mathcal{M}_φ implies $\mathcal{M} \models \psi$
- (iv) $\mathcal{M} \models \langle \mathcal{R} \rangle \varphi$ iff there exists a $w_1 \in U$ such that $(w_0, w_1) \in \mathcal{I}(\mathcal{R})$ and $(U, \mathcal{I}, w_1) \models \varphi$

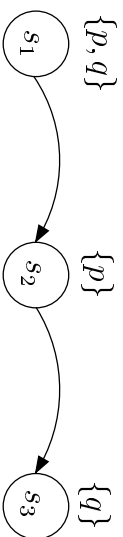
- *initial satisfiability* : whether $(U, \mathcal{I}, w_0) \models \varphi$
- *universal satisfiability*: whether $(U, \mathcal{I}) \models \varphi$, i.e. $\forall w_0 \in U : (U, \mathcal{I}, w_0) \models \varphi$
- *global* model checking algorithms iterate on the structure of φ and traverse the entire \mathcal{M} per each step
- *local* model checking algorithms extend the portion of \mathcal{M} to be examined per each round in an iterative manner, examining the entire φ per each step

The set of states satisfying φ is denoted with $\varphi^{\mathcal{F}}$ and calculated recursively:

- (i) for an *atomic proposition* \mathbf{p} , $\mathbf{p}^{\mathcal{F}} \triangleq \mathcal{I}(\mathbf{p})$
- (ii) $\perp^{\mathcal{F}} \triangleq \emptyset$
- (iii) $(\varphi \rightarrow \psi)^{\mathcal{F}} \triangleq (U \setminus \varphi^{\mathcal{F}}) \cup \psi^{\mathcal{F}}$
- (iv) $(\langle \mathcal{R} \rangle \psi)^{\mathcal{F}} \triangleq \{w \in U \mid \exists w' \in \psi^{\mathcal{F}}, (w, w') \in \mathcal{I}(\mathcal{R})\}$

Computational complexity: $\mathcal{O}(|\mathcal{M}| \times |\varphi|)$

Verifying a modal logic formula $\varphi = (q \rightarrow (R)p)$

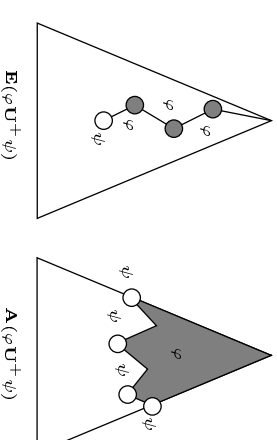


$w_0 \models (\phi \mathbf{U}^+ \psi)$ iff $\exists w_1 \in U$ such that $w_0 < w_1$ and $w_1 \models \psi$, also
 $\forall w_2 \in U$ such that $w_0 < w_2 < w_1$, applies that $w_2 \models \phi$.

$$\mathbf{E}(\psi_2 \mathbf{U}^+ \psi_1) \Leftrightarrow \mathbf{EX}(\psi_1 \vee (\psi_2 \wedge \mathbf{E}(\psi_2 \mathbf{U}^+ \psi_1)))$$

$$\mathbf{A}(\psi_2 \mathbf{U}^+ \psi_1) \Leftrightarrow \mathbf{AX}(\psi_1 \vee (\psi_2 \wedge \mathbf{A}(\psi_2 \mathbf{U}^+ \psi_1)))$$

$$\mathbf{CTL} ::= \mathcal{P} \mid \perp \mid (\mathbf{CTL} \rightarrow \mathbf{CTL}) \mid \mathbf{E}(\mathbf{CTL} \mathbf{U}^+ \mathbf{CTL}) \mid \mathbf{A}(\mathbf{CTL} \mathbf{U}^+ \mathbf{CTL})$$



```

procedure CTLcheck (model  $\mathcal{M} = (U, \mathcal{I}, w_0)$ , formula  $\varphi$ ) {
  if  $w_0 \in \text{eval}(\varphi)$ 
    print " $\varphi$  is satisfied at  $w_0$  of the frame  $\mathcal{F} = (U, \mathcal{I})$ "
  else
    print " $\varphi$  is not satisfied at  $w_0$  of the frame"
}

```

```

function succ (state  $w$ ) : set of states
  return  $\{w' \mid (w, w') \in \mathcal{I}(\prec)\}$  // RETURNS STATES REACHABLE BY " $\prec$ "

```

Computational complexity: $\mathcal{O}(|\varphi| \cdot |U|^3)$

function eval (formula ϕ) : set of states

case ϕ of

p : return $\mathcal{I}(\phi)$

\perp : return \emptyset

$(\psi_1 \rightarrow \psi_2)$: return $((U \setminus \text{eval}(\psi_1)) \cup \text{eval}(\psi_2))$

$\mathbf{E}(\psi_2 \text{ U}^+ \psi_1)$: $E_1 := \text{eval}(\psi_1)$, $E_2 := \text{eval}(\psi_2)$, $E := \emptyset$

repeat until stabilization // UNTIL E DOES NOT GROW

$E := E \cup \{w \mid (\text{succ}(w) \cap (E_1 \cup (E_2 \cap E))) \neq \emptyset\}$

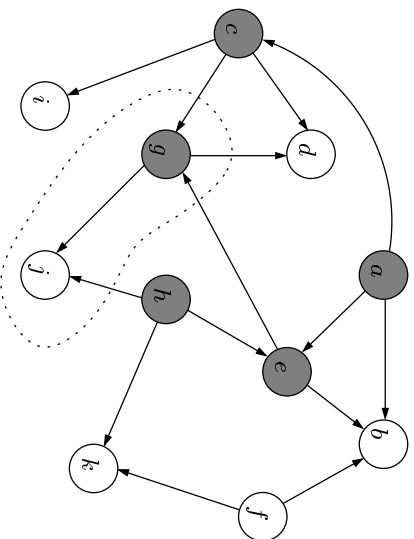
return E

$\mathbf{A}(\psi_2 \text{ U}^+ \psi_1)$: $E_1 := \text{eval}(\psi_1)$, $E_2 := \text{eval}(\psi_2)$, $E := \emptyset$

repeat until stabilization // UNTIL E DOES NOT GROW

$E := E \cup \{w \mid \emptyset \neq \text{succ}(w) \subseteq (E_1 \cup (E_2 \cap E))\}$

return E



Inverse reachability problem

Computational complexity: $\mathcal{O}(|E|)$.

function inverseReachability(set of points T) : set of points

<set> $S := \emptyset$ // THE RESULTING SET OF NODES IS INITIALLY EMPTY

<set> $T' = T$ // A WORKING SET FOR THE ALGORITHM

while $T' \neq \emptyset$ do

$T' := \text{predecessors}(T') \setminus S$

$S := S \cup T'$

return S ;

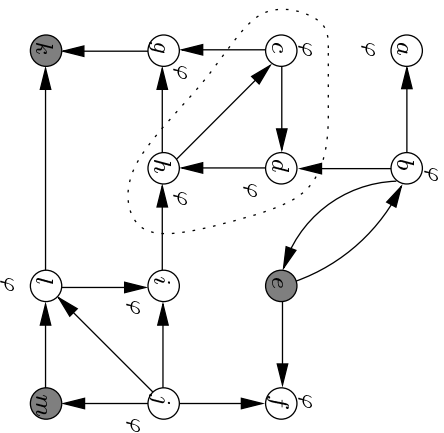
function predecessors(point w) : set of points

return $\{w' \mid (w', w) \in \mathcal{I}(\prec)\}$

A procedure for marking where $\text{EG}^+ \varphi$ holds

- (i) exclude all states in which φ does not hold
- (ii) mark the remaining set with $V_\varphi \subseteq U$
- (iii) mark all states $w \in U$ from which a state $w' \in V_\varphi$ without any successors at all that can be reached traversing only through states of V_φ
- (iv) find the maximal strongly connected components (SCCs) of V_φ
- (v) mark all states that are members of SCCs
- (vi) mark all states from which a non-trivial SCC of V_φ can be reached by a path in V_φ

$\psi^{\mathcal{F}} = \{b, c, d, e, h, i, j, l, m\}$ for $\psi = \mathbf{EG}^+\varphi$ for the \mathcal{M} below:



A set m of subformulas of φ is said to be *maximal* if for $\psi \in SF(\varphi)$: for each element, either the element or its negation is included in the set m , i.e. $\{\psi \mid \psi \in m\} \cup \{\neg\psi \mid \psi \notin m\}$. Such a set is denoted as $m \in SF(\varphi)$

A set $m \subseteq SF(\varphi)$ is said to be *propositionally consistent* if both of the requirements below hold:

- (i) $\perp \notin m$
- (ii) if $(\psi_1 \rightarrow \psi_2) \in SF(\varphi)$, then the it must hold that $(\psi_1 \rightarrow \psi_2) \in m$ iff $\psi_1 \notin m$ or $\psi_2 \in m$

Validity for linear time logics

- the set of *maximal sequences* generated from $\mathcal{M} = (U, T, w_0)$ starting at the w_0 satisfies φ
- whether $\neg\varphi$ is valid for even one sequence

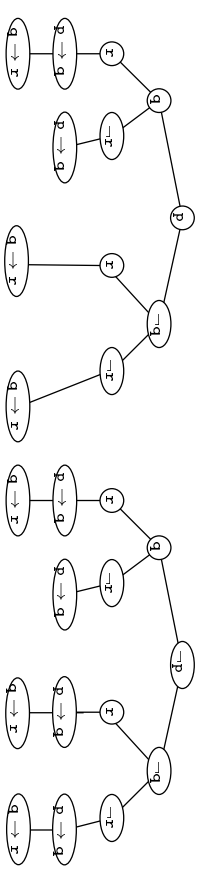
A simplified modal logic:

$$\mathbf{ML} ::= \mathcal{P} \mid \perp \mid (\mathbf{ML} \rightarrow \mathbf{ML}) \mid \langle R \rangle \mathbf{ML}$$

For given \mathcal{M} and φ : can a maximal sequence be generated from \mathcal{M} starting at w_0 such that $\neg\varphi$ is satisfied at w_0 ?

Finding propositionally consistent sets

$$\varphi = (q \rightarrow (R)(p \rightarrow q)), SF(\varphi) = \{q, \langle R \rangle(p \rightarrow q), (p \rightarrow q), p, \varphi\}$$



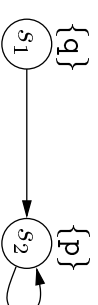
Maximal propositionally consistent sets

$$m \in \{ \{p, q, r, (p \rightarrow q), (q \rightarrow r)\}, \\ \{p, q, \neg r, (p \rightarrow q)\} \\ \{p, \neg q, r, (q \rightarrow r)\} \\ \{p, \neg q, \neg r, (q \rightarrow r)\} \\ \{\neg p, q, r, (p \rightarrow q), (q \rightarrow r)\}, \\ \{\neg p, q, \neg r, (p \rightarrow q)\} \\ \{\neg p, \neg q, r, (p \rightarrow q), (q \rightarrow r)\}, \\ \{\neg p, \neg q, \neg r, (p \rightarrow q), (q \rightarrow r)\} \}$$

Shorthand: $(q \rightarrow r) \triangleq \varphi$

Constructing a forest of admissible atoms

- *atom*: any pair (w, m) , where $w \in U$ and $m \subseteq SF(\varphi)$
- *admissible atom*: w and m agree on the interpretation of *propositions*
- *initial atom*: (w_0, m_0) where $\varphi \in m_0$
- *open leaf*: $\alpha = (w, m)$ with no children, is called *open* that has no formulas in m that begins with $\langle R \rangle$ in m
- *closed leaf*: other atoms in the forest
- *accepting path*: a path in the forest that is either infinite or ends in an open leaf (a counterexample to φ)



Admissible atoms for $\varphi = (q \rightarrow \langle R \rangle (p \rightarrow q))$ and above \mathcal{M} :

α	w	$m \in SF(\varphi)$
α_1	s_1	$\{p, \neg q, \langle R \rangle (p \rightarrow q), \varphi\}$
α_2	s_1	$\{p, \neg q, \varphi\}$
α_3	s_2	$\{\neg p, q, \langle R \rangle (p \rightarrow q), (p \rightarrow q), \varphi\}$
α_4	s_2	$\{\neg p, q, (p \rightarrow q)\}$

α_1 and α_2 are initial atoms, X_R contains (α_1, α_3) and (α_3, α_4) .

$X_R((w, m), (w', m'))$ holds between admissible atoms iff

- (i) $(w, w') \in \mathcal{I}(R)$
- (ii) if $\langle R \rangle \psi \in SF(\varphi)$ and $\psi \in m'$, then $\langle R \rangle \psi \in m$
- (iii) if $\langle R \rangle \psi \in m$, then $\psi \in m'$
- (iv) some $\langle R \rangle \psi \in m$

LTL

Formulas in positive conjunctive normal form:

$$\mathbf{F}^* p_1 \wedge \dots \wedge \mathbf{F}^* p_n \wedge \mathbf{G}^*(p_1 \rightarrow \mathbf{X}\mathbf{G}^* \neg p_1) \wedge \dots \wedge \mathbf{G}^*(p_n \rightarrow \mathbf{X}\mathbf{G}^* \neg p_n)$$

Definition of $\mathcal{M} = (U, \mathcal{I}, w_0)$:

- the set of states $U = V \cup \{s, t\}$, where $s, t \notin V$
- $w_0 = s$
- $\mathcal{I}(R) = E \cup \{(s, v_i) \mid v_i \in V\} \cup \{(v_i, t) \mid v_i \in V\} \cup \{(t, t)\}$
- $\mathcal{I}(\mathbf{p})$ for all atomic propositions $\mathbf{p} \in \mathcal{P}$ is such that
 - $v_i \in \mathcal{I}(\mathbf{p}_i)$ iff $1 \leq i \leq |V|$
 - $v_i \notin \mathcal{I}(\mathbf{p}_j)$ iff $1 \leq i, j \leq |V|, i \neq j$
 - $\forall 1 \leq i \leq |V| : s \notin \mathcal{I}(\mathbf{p}_i)$
 - $\forall 1 \leq i \leq |V| : t \notin \mathcal{I}(\mathbf{p}_i)$

Atoms $\alpha = (w, c)$ where $w \in U$ and c as defined below:

- $\forall \mathbf{p} \in \mathcal{P} : \mathbf{p} \in c_i$ if and only if $w_i \in \mathcal{I}(\mathbf{p})$
 - $\forall \psi \in CL(\varphi) : \psi \in c_i$ if and only if $\neg \psi \notin c_i$
 - $\forall (\psi_1 \vee \psi_2) \in CL(\varphi) : (\psi_1 \vee \psi_2) \in c_i$ if and only if either $\psi_1 \in c_i$ or $\psi_2 \in c_i$
 - $\forall \neg \mathbf{X}\psi \in CL(\varphi) : \neg \mathbf{X}\psi \in c_i$ if and only if $\mathbf{X}\neg \psi \in c_i$
 - $\forall (\psi_2 \mathbf{U}^* \psi_1) \in CL(\varphi), (\psi_2 \mathbf{U}^* \psi_1) \in c_i$ if and only if either $\psi_1 \in c_i$ or both $\psi_2 \in c_i$ and $\mathbf{X}(\psi_2 \mathbf{U}^* \psi_1) \in c_i$
- $(\alpha, \alpha') \in E$ for $\alpha = (w, c)$ and $\alpha' = (w', c')$ iff
- $(\alpha, \alpha') \in E$ if and only if $(w, w') \in \mathcal{I}(R)$ and
 - $\forall \mathbf{X}\psi \in CL(\varphi) : \mathbf{X}\psi \in c$ if and only if $\psi \in c'$

The *closure* of a formula φ , $CL(\varphi)$, is defined as the minimal set of formulas containing φ for which

- (i) $\neg \psi \in CL(\varphi)$ if and only if $\psi \in CL(\varphi)$
- (ii) if $\psi_1 \vee \psi_2 \in CL(\varphi)$, then both $\psi_1, \psi_2 \in CL(\varphi)$
- (iii) if $\mathbf{X}\psi \in CL(\varphi)$, then $\psi \in CL(\varphi)$
- (iv) if $\neg \mathbf{X}\psi \in CL(\varphi)$, then $\mathbf{X}\neg \psi \in CL(\varphi)$
- (v) if $(\psi_2 \mathbf{U}^* \psi_1) \in CL(\varphi)$, then all $\psi_1, \psi_2, \mathbf{X}(\psi_2 \mathbf{U}^* \psi_1) \in CL(\varphi)$

$\mathcal{M}, w_0 \models \mathbf{E}\varphi$ iff there exists an eventuality sequence starting at an initial atom $\alpha = (w_0, c_0)$.

An eventuality sequence starts from an atom $\alpha = (w, c)$ iff there there is a path in the atom graph G such that a self-fulfilling SCC can be reached by it from α .

Modified Tarjan's algorithm

```

procedure LTLcheck(model  $\mathcal{M}$ , formula  $\varphi$ )
  <integer> recursionDepth := 0 // NUMBER OF RECURSIVE CALLS MADE
  <stack> SCC :=  $\emptyset$ 
  <hashable> storage // FOR STORING THE TRAVERSED "PATH"
  <set> initial := { $\alpha$  |  $\alpha$  is an initial atom of  $\mathcal{M}$  and  $\varphi$ }
  for all  $\alpha \in$  initial do dfs( $\alpha$ )
  print " $\varphi$  is not satisfiable in  $\mathcal{M}$ "

  • exponential with respect to the number of  $\mathbf{U}^+$  formulas
  • linear with respect to the size of  $\mathcal{M}$ ,  $\mathcal{O}(|U|^2)$ 
  • generally: with past-temporal operators is complete in
    PSPACE with respect to  $|\varphi|$  and NLOGSPACE in size of the
    model  $\mathcal{M}$ 

```

```

procedure processRoot (atom  $\alpha$ )
  <set> required :=  $\emptyset$  // SET OF REQUIRED EVENTUALITIES
  <set> fulfilled :=  $\emptyset$  // SET OF FULFILLED EVENTUALITIES
  repeat
     $\beta :=$  pop(SCC)
    storage[ $\beta$ ] :=  $\infty$  // PUT "BELOW" EVERYTHING ELSE
    required := required  $\cup$  { $\psi_1$  | ( $\psi_2 \mathbf{U}^+ \psi_1$ )  $\in \beta$ }
    fulfilled := fulfilled  $\cup$  { $\psi$  |  $\psi \in \beta$ }
  until ( $\alpha = \beta$ ) // ALL ELEMENTS IN THE SCC OF  $\alpha$  ARE PROCESSED
  if (required  $\subseteq$  fulfilled) // THE SCC OF  $\alpha$  IS SELF-FULFILLING
    print " $\varphi$  is satisfiable in  $\mathcal{M}$ "; exit

```

The function children returns the set of atoms accessible from the given atom in the constructed graph.

```

procedure dfs(atom  $\alpha$ ) // DEPTH FIRST SEARCH
  if (storage does not contain  $\alpha$ ) // A "FRESH" ATOM
    <integer> currentDepth = recursionDepth
    recursionDepth++ // INCREMENT DEPTH BY ONE
    // INITIAL VALUE AT CURRENT RECURSION DEPTH
    storage[ $\alpha$ ] = currentDepth
    push(SCC,  $\alpha$ ) // TO THE TOP OF THE STACK
    <set> successorAtoms := children( $\alpha$ )
    for all ( $\beta \in$  successorAtoms) do
      dfs( $\beta$ )
      // CHECK WHETHER  $\beta$  IS "ABOVE"  $\alpha$ 
      storage[ $\alpha$ ] = min(storage[ $\alpha$ ], storage[ $\beta$ ])
    // IF NOTHING WAS ABOVE  $\Rightarrow \alpha$  IS A ROOT OF AN SCC
    if (storage[ $\alpha$ ] = currentDepth)
      processRoot( $\alpha$ )

```

$$\mu\mathbf{TL} ::= \mathcal{P} \mid \mathcal{Q} \mid \perp \mid (\mu\mathbf{TL} \rightarrow \mu\mathbf{TL}) \mid (\mathcal{R})\mu\mathbf{TL} \mid \nu\mathcal{Q} \mu\mathbf{TL}$$

- the quantifier ν is a *restricted existential quantifier* on sets of points
- the quantifier μ is defined by ν as $\mu q \varphi \triangleq \neg \nu q \neg(\varphi \{q := \neg q\})$
- Knaster-Tarski:

$$(U, \mathcal{I}, w) \models \nu q \leftrightarrow w \in \bigcup \{Q \mid Q \subseteq \varphi^{\mathcal{F}} \{q := Q\}\}$$

$$(U, \mathcal{I}, w) \models \mu q \leftrightarrow w \in \bigcap \{Q \mid \varphi^{\mathcal{F}} \{q := Q\} \subseteq Q\}$$

If φ^i is union-continuous, the fixed points $\nu q \varphi$ and $\mu q \varphi$ can be obtained with the following limits:

$$\nu q \varphi = \lim_{i \rightarrow \infty} \varphi^i(\top)$$

$$\mu q \varphi = \lim_{i \rightarrow \infty} \varphi^i(\perp).$$

For a finite universe U , every monotonic function is also union-continuous:

$$\nu q \varphi = \lim_{i \leq |U|} \varphi^i(\top)$$

$$\mu q \varphi = \lim_{i \leq |U|} \varphi^i(\perp)$$

- the cycles stabilize after at most $|U|$ rounds: $\mathcal{O}(|\varphi| \cdot |U|^{qd(\varphi)})$
- $qd(\varphi)$ is the *nesting depth* of fixed-point operators in φ
- the computation of an inner fixed-point formula needs to be restarted for each iteration on the enclosing formula

Example: $\psi(q_1, q_2) = \mu q_1(\mathbf{p}_1 \wedge \mu q_2(\mathbf{X}q_1 \vee \mathbf{X}q_2 \vee \mathbf{p}_2))$.

For each iteration of q_1 , the inner formula $\mu q_2(\mathbf{X}q_1 \vee \mathbf{X}q_2 \vee \mathbf{p}_2)$ is re-evaluated.

function eval (formula φ) : set of states

case φ of

p: **return** $\mathcal{I}(\mathbf{p})$ // ATOMIC PROPOSITION

q : **return** $v(q)$ // VALUATION OF A PROPOSITION VARIABLE

\perp : **return** \emptyset

$(\psi_1 \rightarrow \psi_2)$: **return** $((U \setminus \text{eval}(\psi_1)) \cup \text{eval}(\psi_2))$

$\langle R \rangle \psi$: **return** $R^{-1}(\text{eval}(\psi))$

$\nu q \psi$: $H_1 := U$

repeat until stabilization // UNTIL H_1 DOES NOT SHRINK

$H_1 := \text{eval}(\psi\{q := H_1\})$

return H_1

$\mu q \psi$: $H_2 := \emptyset$

repeat until stabilization // UNTIL H_2 DOES NOT GROW

$H_2 := \text{eval}(\psi\{q := H_2\})$

return H_2

Local model checking for μ TL

- a tableau method, which explores a portion of the model using depth-first search
- the nodes of the tableau are sequences of the form $\Delta, w \models \psi$, where $w \in U$ and ψ is a sub-formula of φ
- Δ is a list of definitions that contains a sequence of declarations $(q_1 = \psi_1, \dots, q_n = \psi_n)$ where each proposition variable q_i appears only once and each ψ_i may only contain variables q_j such that $j < i$

Steps for decomposition of the formula φ into a tableau:

- (i) a node $\Delta, w \models (\psi_1 \wedge \psi_2)$ has two children: $\Delta, w \models \psi_1$ and $\Delta, w \models \psi_2$
- (ii) a node $\Delta, w \models (\psi_1 \vee \psi_2)$ has one child: either $\Delta, w \models \psi_1$ or $\Delta, w \models \psi_2$
- (iii) a node $\Delta, w \models \langle R \rangle \psi$ has one child: $\Delta, w' \models \psi$
- (iv) a node $\Delta, w \models [R] \psi$ has n children: $\Delta, w_1 \models \psi, \dots, \Delta, w_n \models \psi$
- (v) a node $\Delta, w \models \mu q \psi$ has one child: $\Delta', w \models \psi$
- (vi) a node $\Delta, w \models \nu q \psi$ has one child: $\Delta', w \models \psi$
- (vii) a node $\Delta, w \models q$ has one child: $\Delta, w \models \psi$

When no rule cannot be applied for any leaf, the tableau is said to be *maximal*.

A leaf $\Delta, w \models \psi$ of a maximal tableau is said to be *successful* if the following holds:

- (1) $(\psi = \mathbf{p} \in \mathcal{P} \wedge w \in \mathcal{I}(\mathbf{p})) \vee (\psi = \neg \mathbf{p} \wedge w \notin \mathcal{I}(\mathbf{p}))$
- (2) $\psi = q \in \mathcal{Q}, q \notin \Delta, w \in v(q)$, or $\psi = \neg q, q \notin \Delta, w \notin v(q)$, or
- (3) rule (iv) produces no children, that is: $\psi = [R] \psi' \wedge R(w) = \emptyset$
- (4) $\psi = q \in \mathcal{Q}$ and q was included in Δ by rule (vi)

Theorem: $w \in \varphi^{\mathcal{F}}$ iff there exists a successful tableau with root $\emptyset, w \models \varphi$.

Restrictions for the construction:

$\vee, \wedge, \langle R \rangle, [R], \mu$ and ν are used as basic operations, negations are assumed to appear only in literals.

Each μ or ν quantification in the formula φ is expected to bind a different propositional variable.

- rule (iii) can only be applied when $w' \in R(w)$
- when applying rule (iv), it must hold that $R(w) = \{w_1, \dots, w_n\}$
- for rules (v) and (vi), $\Delta' = \Delta \cup \{q = \psi\}$
- rule (vii) is only applicable when $(q = \psi) \in \Delta$ and no ancestor node is of the form $\Delta', w \models \psi$ with the same w and ψ

Message passing

- processes send each other messages, where the messages and the order of sending them are specified by some protocol
- each process may *send* or *receive* a message.
- messages may be directed to specific *channels* with specified recipients or *broadcasted* to any process within the messaging space

In *synchronized* communication, the *send* and *receive* primitives are *blocking* operations, i.e. the sending process must wait for the recipient to react before proceeding further.

Synchronous communication: a process may not proceed until all the communication partners defined for the particular task have expressed willingness to participate in that action.

Asynchronous communication: the processes deciding themselves whether they wish to wait, usually by using some sort of buffer for messages that are not immediately reacted to

When non-determinism is required for the inter-process communication, *guards* that consists of a Boolean expression and communication statements may be placed on the messaging operations.

Communication by shared variables

When two or more processes interact by reading from and writing to a shared memory space, they may exchange information by using *shared variables* for separate issues on which to communicate. The semantic of possible values for each of the variables must be defined, similarly to defining a grammar or a protocol for the message passing paradigm.

It must be ensured that the accesses to the shared variables are *legal*: two processes should not be able to modify the value of a variable at the same time, and no process should read a value after some other process has reserved it for modification.

When the system modules are said to be *synchronized*, for each of the discrete time steps, each of the modules performs a task. The tasks are generally independent on each other, but it can also be stated that only when certain conditions apply, certain transitions may be taken by the processes.

One process may have to stay in an *idle loop*, performing essentially a *no-op* (stands for *no-operation*) until the conditions for taking the next actual step are reached.

In *asynchronous* computation, for each time step, some but not necessarily all modules advance. If performing a *no-op* is possible for every process in every state of the computation, then the synchronous and asynchronous computation modes coincide. The idle step is also called a “stutter”.

Transition system (Σ, S, Δ, S_0)

- an *alphabet* Σ
- a non-empty and finite set of states S
- a subset $S_0 \subseteq S$ as *initial states*
- a *transition relation* $\Delta \subseteq S \times \Sigma \times S$ describes which states in S are reachable in one step from one another for the symbols of the alphabet Σ

Parallel transition system $\mathcal{T} = (T_1, \dots, T_n)$

n -tuple of transition systems in which $\forall i, j$ such that $i < j$, applies that $S_i \cap S_j = \emptyset$.

- the *global alphabet* is the union of all alphabets, $\Sigma = \bigcup_{i=1}^n \Sigma_i$
- the *global state space* is defined as $S = S_1 \times \dots \times S_n$
- the set of *global initial states* is defined as $S_0 = S_{10} \times \dots \times S_{n0}$
- for the *global transition relation*, a transition $\delta = ((s_1, \dots, s_n), a, (s'_1, \dots, s'_n)) \in \Delta$ iff $\forall T_i$
 - (i) if $a \in \Sigma_i$, then $(s_1, a, s'_1) \in \Delta_i$, and
 - (ii) if $a \notin \Sigma_i$, then $s_i = s'_i$

[E]lementary Petri net $N = (P, T, F, m_0)$

- a finite set P of *places*
- a set of *transitions* T such that $P \cap T = \emptyset$
- a *flow relation* F that describes the relations of the places and the transitions: $F \subseteq (P \times T) \cup (T \times P)$
- an *initial marking* $m_0 \subseteq P$ imposed on the net
- any subset of P is a *marking* of the Petri net

A marking function \mathcal{L} is a function from the set of places P to the set of P 's subsets 2^P

- *preset* : $\bullet t \doteq \{p \mid (p, t) \in F\}$
- *postset* : $t \bullet \doteq \{p \mid (t, p) \in F\}$.

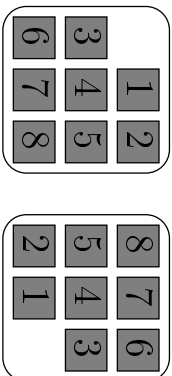
A transition is *enabled* at a marking $m \subseteq P$ if $\bullet t \subseteq m \wedge t \bullet \cap m \subseteq \bullet t$

A marking m' is said to be the *result of firing a transition* t from some marking m , if the transition t is enabled at m and the marking m' consists of the states $m' = (m \setminus \bullet t) \cup t \bullet$.

Shared variable program (V, D, T, s_0)

- a set of *program variables* $V = \{v_1, \dots, v_n\}$
- a *state space* $D = D_1 \times \dots \times D_n$ in which each $D_i = \{d_{i1}, \dots, d_{in_i}\}$ is a *finite domain* of the corresponding variable v_i
- a *transition relation* T is defined as $T \subseteq D \times D$
- one *initial state* $s_0 = (d_{11}, \dots, d_{n1})$

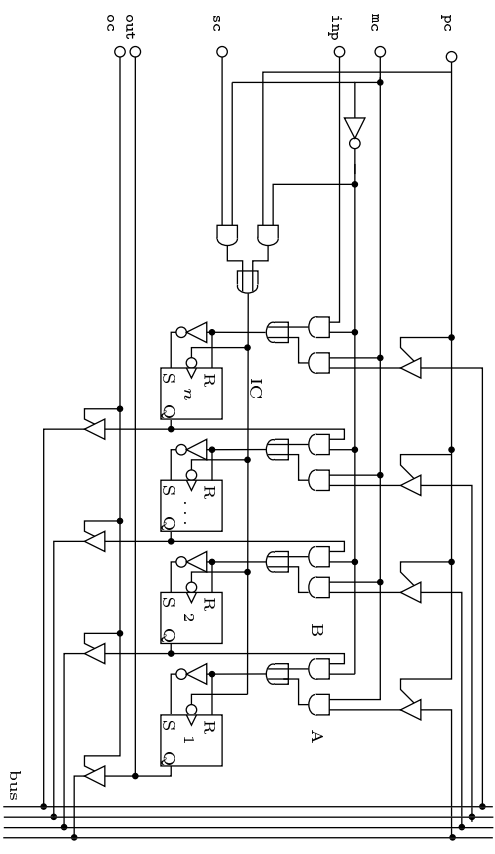
A transition (s, s') is in T iff the proposition φ_T is a logical implication of the valuation T , i.e. $T \models \varphi_T$, where $s = (d_1, \dots, d_n)$, $s' = (d'_1, \dots, d'_n)$, $\mathcal{I}(v_i) = d_i$, and $\mathcal{I}(v'_i) = d'_i$



The puzzle as a shared variables program: a variable for each tile $i = 1, \dots, (h \cdot v - 1)$ (a two-component vector $t_i = (v_{ih}, v_{iv})$) and one variable for the direction of the next move, transitions are swappings of a tile with the empty slot.

S	R	Q'	description
0	0	Q	maintain
1	0	1	set
0	1	0	reset
1	1	-	undefined

The undefined state can be modeled by making a non-deterministic choice between 0 and 1 for the value of the next state Q' internally. The output here changes only when the clock line changes from high to low.



The **correctness requirements** with n representing the width of the data bus are

$$AG^*(mc \wedge pc \rightarrow \bigvee_{i=1}^n (bus[i] \leftrightarrow A((oc \rightarrow AX(bus[i] U^{+ic}))))$$

$$AG^*(\neg mc \wedge sc \rightarrow \bigvee_{i=2}^n (q[i] \leftrightarrow A(q[i-1] U^{+ic}))),$$

ic is the result of the bit-wise operation $ic = ((\neg mc \wedge pc) \vee (mc \wedge sc))$.

```

MODULE main
VAR q, bus: array 1 .. n of boolean;
inp, mc, pc, sc, oc: boolean; -- INPUT LINES
DEFINE out := q[1]; ic := ((!mc & pc) | (mc & sc));
-- FOR ALL OF THE n BITS INDEXED WITH 0, ..., n-1:
A[0] := mc & pc & bus[0]; B[0] := !mc & q[1];
R[0] := !(A[0] | B[0]); S[0] := !R[0];
ASSIGN next(q[0]) := case ic: case
    !S[0] & !R[0]: q[0]; -- MAINTAIN
    S[0] & !R[0]: 1; -- SET
    !S[0] & R[0]: 0; -- RESET
    S[0] & R[0]: {0, 1}; esac; -- UNDEFINED
next(bus[?]) := case oc: q[0]; !oc: {0, 1}; esac;
FAIRNESS ic FAIRNESS oc

```

input	buffer	output	input	buffer	output
nil	$\langle \rangle$	nil	nil	$\langle \rangle$	nil
x	$\langle \rangle$	nil	nil	$\langle \rangle$	x
nil	$\langle x_1, \dots, x_n \rangle$	nil	nil	$\langle x_1, \dots, x_{n-1} \rangle$	x_n
x	$\langle x_1, \dots, x_n \rangle$	nil	nil	$\langle x, x_1, \dots, x_{n-1} \rangle$	x_n
nil	$\langle \rangle$	y	nil	$\langle \rangle$	
x	$\langle \rangle$	y	nil	$\langle x \rangle$	
x	$\langle x_1, \dots, x_n \rangle$	y	nil	$\langle x, x_1, \dots, x_n \rangle$	
x	$\langle x_1, \dots, x_n \rangle$	y	x	$\langle x_1, \dots, x_n \rangle$	

A **bounded buffer** used by a set of communicating processes within the operating system of a Siemens cellular phone

- should be deadlock-free: **AG*EF*** init
- five processes and the operating system kernel process
- size of the state space per process is app. 10–20 states
- there are app. 50 types of messages
- processes scheduled by the operating system based on priorities
- storage and delivery of the messages is also done by the OS