

Petri net analysis using decision diagrams

Gianfranco Ciardo

Department of Computer Science and Engineering

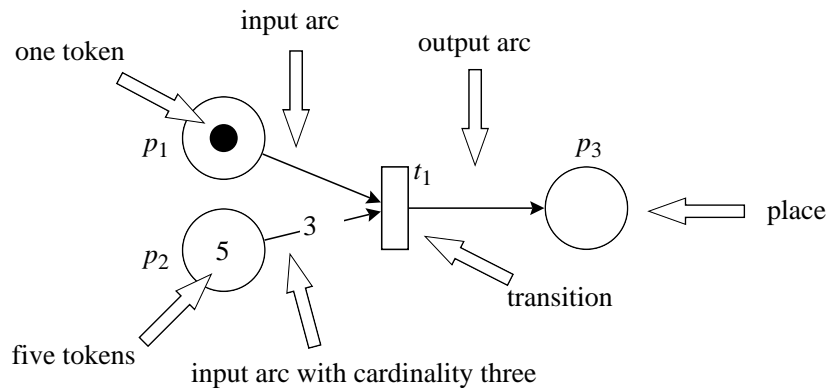
University of California, Riverside



State-space generation of Petri nets

Graphical representation of a Petri net

3



Petri nets

4

A Petri net is a tuple $(\mathcal{P}, \mathcal{E}, \mathbf{D}^-, \mathbf{D}^+, \mathbf{x}_{init})$ where:

- \mathcal{P} set of places, drawn as circles
- \mathcal{E} set of transitions drawn as rectangles
- $\mathbf{D}^- : \mathcal{P} \times \mathcal{E} \rightarrow \mathbb{N}$ input arc cardinalities
- $\mathbf{D}^+ : \mathcal{P} \times \mathcal{E} \rightarrow \mathbb{N}$ output arc cardinalities
- $\mathbf{x}_{init} \in \mathbb{N}^{|\mathcal{P}|}$ initial state, or marking


with $\mathcal{P} \cap \mathcal{E} = \emptyset$

Condition for transition α to be enabled in state $\mathbf{i} \in \mathbb{N}^{|\mathcal{P}|}$: $\alpha \in \mathcal{E}(\mathbf{i}) \Leftrightarrow \forall p \in \mathcal{P}, \mathbf{D}_{p,\alpha}^- \leq \mathbf{i}_p$


A transition α enabled in state \mathbf{i} can fire: $\mathbf{i} \xrightarrow{\alpha} \mathbf{j} \Leftrightarrow \forall p \in \mathcal{P}, \mathbf{j}_p = \mathbf{i}_p - \mathbf{D}_{p,\alpha}^- + \mathbf{D}_{p,\alpha}^+$

The next-state function \mathcal{N} satisfies $\mathbf{j} \in \mathcal{N}(\mathbf{i}) \Leftrightarrow \exists \alpha \in \mathcal{E}, \mathbf{j} \in \mathcal{N}_\alpha(\mathbf{i}) \Leftrightarrow \exists \alpha \in \mathcal{E}, \mathbf{i} \xrightarrow{\alpha} \mathbf{j}$

The state space, or reachability set, \mathcal{X}_{reach} is defined as usual



If the initial state is $\mathbf{x}_{init} = (N, 0, 0, 0, 0)$, \mathcal{X}_{reach} contains $\frac{(N+1)(N+2)(2N+3)}{6}$ states



\mathcal{X}_{reach} contains an infinite number of states regardless of the initial state $\mathbf{x}_{init} = (N)$

Why is state-space generation important?

- State-space generation is an essential step in logical model analysis
 - Can be used to discover the potential states (the possible range of a variable)
 - Discovers the actual states of the system (the possible combinations of variable values)
- State-space generation is enough to answer safety queries
 - Can we reach a "bad" state?
 - Is it true that $VAL < 0$ whenever $FLAG$ is raised?
- State-space generation is similar to other temporal logic queries
 - Is it possible to reach a state where $VAL < 0$ from a state where $VAL > 0$? (EF)
 - Is it true that, if $VAL > 0$, it must become 0 before it can become negative? (EU)

symbolic methods based on decision diagrams help immensely...
 ...but we still are (and always will be) memory and time bound

State-by-state (explicit) generation algorithm

```


ExploreExplicit( $\mathcal{X}_{init}$  : set of states,  $\mathcal{N}$  : next-state function) : set of states
local  $\mathcal{U}, \mathcal{X}_{reach}$  : set of states;
local  $\mathbf{i}, \mathbf{j}$  : state;
1  $\mathcal{X}_{reach} \leftarrow \emptyset;$ 
2  $\mathcal{U} \leftarrow \mathcal{X}_{init};$ 
3 while  $\mathcal{U} \neq \emptyset$  do
4   choose a state  $\mathbf{i}$  in  $\mathcal{U}$  and move it to  $\mathcal{X}_{reach};$ 
5   for each  $\mathbf{j} \in \mathcal{N}(\mathbf{i})$  do
6     if  $\mathbf{j} \notin \mathcal{X}_{reach} \cup \mathcal{U}$  then
7        $\mathcal{U} \leftarrow \mathcal{U} \cup \{\mathbf{j}\};$ 
8     end if;
9   end for;
10 end while;
11 return  $\mathcal{X}_{reach};$ 
    
```

\mathcal{X}_{reach} contains the known states already explored
 \mathcal{U} contains the known states not yet explored
search to determine whether \mathbf{j} is a new state
remember to explore \mathbf{j} later

the memory requirements are $O(|\mathcal{X}_{reach}|)$

most time is spent searching for a state (line 6)

Explicit generation of \mathcal{X}_{reach} adds **one state** at a time



with an explicit data structure

memory requirements increase monotonically during generation

they are proportional to $|\mathcal{X}_{reach}|$ at the end

Several informal concepts of “explicit” and “implicit” (a.k.a. “symbolic”) have been used

We adopt the following informal definitions


- **Explicit data structure:** each state requires a different memory location (bit, byte, word, array, etc.)
 $\Rightarrow O(|\mathcal{X}_{reach}|)$ **memory**
- **Explicit algorithm:** states are manipulated one by one
 $\Rightarrow O(|\mathcal{X}_{reach}|)$ or maybe $O(|\mathcal{X}_{reach}| \cdot \log |\mathcal{X}_{reach}|)$ **time**
Memory requirements increase linearly as new states are found
- **Implicit data structure:** each memory location **may** store information about multiple states
 $\Rightarrow O(|\mathcal{X}_{reach}|)$ **memory only in the worst case**
- **Implicit algorithm:** states are manipulated one set at a time
 $\Rightarrow O(|\mathcal{X}_{reach}|)$ or maybe $O(|\mathcal{X}_{reach}| \cdot \log |\mathcal{X}_{reach}|)$ **time only in the worst case**
Memory requirements grow and shrink as new states are found, peak not usually at the end

“Graph-based algorithms for boolean function manipulation”

Randy Bryant (Carnegie Mellon University)

IEEE Transactions on Computers, 1986

CiteSeer most cited document!



BDDs are a **canonical** representation of boolean functions $f : \mathbb{B}^L \rightarrow \mathbb{B}$

For the **root** node, $f(\mathbf{x}_4=0, \mathbf{x}_3=1, \mathbf{x}_2=1, \mathbf{x}_1=0) = f(\mathbf{x}_4=0, \mathbf{x}_3=1, \mathbf{x}_2=1, \mathbf{x}_1=1) = 1$

A BDD is an acyclic directed edge-labeled graph where:

- The only **terminal nodes** can be **0** and **1**, and are at **level 0** $\mathbf{0}.lvl = \mathbf{1}.lvl = 0$
- A **nonterminal node** p is at a **level** k , with $L \geq k \geq 1$ $p.lvl = k$
- A nonterminal node p has two outgoing edges labelled 0 and 1, pointing to children $p[0]$ and $p[1]$
- The level of the children is lower than that of p ; $p[0].lvl < p.lvl, p[1].lvl < p.lvl$
- A node p at level k encodes the function $v_p : \mathbb{B}^L \rightarrow \mathbb{B}$ defined recursively by

$$v_p(x_L, \dots, x_1) = \begin{cases} p & \text{if } k = 0 \\ v_{p[x_k]}(x_L, \dots, x_1) & \text{if } k > 0 \end{cases}$$

Instead of levels, we can also talk of **variables**:

- The terminal nodes are associated with the **range variable** x_0
- A nonterminal node is associated with a **domain variable** x_k , with $L \geq k \geq 1$

For canonical BDDs, we further require that

- There are no **duplicates**: if $p.lvl = q.lvl$ and $p[0] = q[0]$ and $p[1] = q[1]$, then $p = q$

Then, if the BDD is **quasi-reduced**, there is **no level skipping**:

- The only **root nodes** with no incoming arcs are at level L
- The children $p[0]$ and $p[1]$ of a node p are at level $p.lvl - 1$

Or, if the BDD is **fully-reduced**, there is **maximum level skipping**:


- There are no **redundant nodes** p satisfying $p[0] = p[1]$

Both versions are **canonical**, thus, if functions f and g are encoded using BDDs,

- Satisfiability, $f \neq 0$, or equivalence, $f = g$ $O(1)$
- Conjunction, $f \wedge g$, disjunction, $f \vee g$, relational product: $O(\|f\| \times \|g\|)$, if fully-reduced
 $\sum_{L \geq k \geq 1} O(\|f\|_k \times \|g\|_k)$, if quasi-reduced

$\|f\|$ = number of nodes in the BDD encoding f

$\|f\|_k$ = number of nodes at level k in the BDD encoding f



Fully-reduced BDDs: each node in the BDD encodes a different function

Quasi-reduced BDDs: each node at a given level of the BDD encodes a different function


Using BDDs to encode sets

We can encode a set $\mathcal{Y} \subseteq \mathbb{B}^L$ as a BDD p through its **characteristic function**:

$$\mathbf{i} = (i_L, \dots, i_1) \in \mathcal{Y} \Leftrightarrow v_p(i_L, \dots, i_1) = 1$$

The **size of the set** encoded by BDD p is not directly related to the **size of the BDD** itself

Indeed, any set requires as many nodes as its complement:



Union (Or) and Intersection (And) for fully-reduced BDDs

```

bdd Union(bdd p, bdd q) is fully-reduced version
local bdd r;
1 if p = 0 or q = 1 then return q;
2 if q = 0 or p = 1 then return p;
3 if p = q then return p;
4 if Cache contains entry <UnionCODE, {p, q} : r> then return r;
5 if p.lvl = q.lvl then
6   r ← UniqueTableInsert(p.lvl, Union(p[0], q[0]), Union(p[1], q[1]));
7 else if p.lvl > q.lvl then
8   r ← UniqueTableInsert(p.lvl, Union(p[0], q), Union(p[1], q));
9 else since p.lvl < q.lvl then
10  r ← UniqueTableInsert(q.lvl, Union(p, q[0]), Union(p, q[1]));
11 enter <UnionCODE, {p, q} : r> in Cache;
12 return r;
    
```

$Intersection(p, q)$ differs from $Union(p, q)$ only in the terminal cases:

- Union*: if $p = 0$ or $q = 1$ then return q ;
 if $q = 0$ or $p = 1$ then return p ;
- Intersection*: if $p = 1$ or $q = 0$ then return q ;
 if $q = 1$ or $p = 0$ then return p ;

complexity $O(\text{product of the numbers of nodes in } p \text{ and } q)$

Given an L -level BDD on (x_L, \dots, x_1) rooted at p_* encoding a set $\mathcal{Y} \subseteq \widehat{\mathcal{X}}$

Given a $2L$ -level BDD on $(x_L, x'_L, \dots, x_1, x'_1)$ rooted at r_* encoding a function $\mathcal{N} : \widehat{\mathcal{X}} \rightarrow 2^{\widehat{\mathcal{X}}}$

$RelationalProduct(p_*, r_*)$ returns the root of the BDD encoding the set

$$\{j : \exists i \in \mathcal{Y} \wedge j \in \mathcal{N}(i)\}$$

```

bdd RelationalProduct(bdd p, bdd2 r) is quasi-reduced version
local bdd q, q1, q2;
1 if p = 0 or r = 0 then return 0;
2 if p = 1 and r = 1 then return 1;
3 if Cache contains entry (RelationalProductCODE, p, r : q) then return q;
4 q0 ← Union(RelationalProduct(p[0], r[0][0]), RelationalProduct(p[1], r[1][0]));
5 q1 ← Union(RelationalProduct(p[0], r[0][1]), RelationalProduct(p[1], r[1][1]));
6 q ← UniqueTableInsert(p.lvl, q0, q1);
7 enter (RelationalProductCODE, p, r : q) in Cache;
8 return q;
    
```

The efficient manipulation of BDDs relies on the idea of dynamic programming

More specifically, on the use of an operation cache

Given enough memory for cache entries, we never recompute an operation on the same operands

The operation cache is implemented as a hash table with entries of the form $\langle \text{key} : \text{result} \rangle$

- given the **key** (*operator, operand, ..., operand*)
- we can retrieve the **result**, if it was previously computed

In practice, we can store

- either $\langle \vee, \{a, b\} : c \rangle$ in the **Operation cache** boolean OR is commutative
- or $\langle \{a, b\} : c \rangle$ in the **OR (or UNION) cache**
- either $\langle \Rightarrow, a, b : c \rangle$ in the **Operation cache** boolean IMPLIES is not commutative
- or $\langle a, b : c \rangle$ in the **IMPLIES cache**

- Given a boolean expression, or a function, $f : \mathbb{B}^L \rightarrow \mathbb{B}$, there is a unique BDD encoding it (for a fixed variable order x_L, \dots, x_1)
- Many functions have a **very compact BDD encoding**
- The constant functions 0 and 1 are represented by the nodes **0** and **1**, respectively
- Given the BDD encoding boolean expression f : test whether $f \equiv 0$ or $f \equiv 1$ in $O(1)$ time
- Given the BDDs encoding boolean expressions f and g : test whether $f \equiv g$ in $O(1)$ time
- The variable ordering affects the size of the BDD, consider $x_L = y_L \wedge \dots \wedge x_1 = y_1$
 - with the order $(x_L, y_L, \dots, x_1, y_1)$ $O(L)$ nodes
 - with the order $(x_L, \dots, x_1, y_L, \dots, y_1)$ $O(2^L)$ nodes
- The BDD encoding of some functions is exponentially large for any order
 - the expression for bit 32 of the 64-bit result of the multiplication of two 32-bit integers
- Finding the optimal ordering that minimizes the BDD size is an **NP-complete problem**

We can store

- any set of markings $\mathcal{Y} \subseteq \widehat{\mathcal{X}} = \mathbb{B}^{|\mathcal{P}|}$ of a safe PN with a $|\mathcal{P}|$ -level BDD
- any relation over $\widehat{\mathcal{X}}$, or function $\widehat{\mathcal{X}} \rightarrow 2^{\widehat{\mathcal{X}}}$, such as \mathcal{N} , with a $2|\mathcal{P}|$ -level BDD

We can encode \mathcal{N} using $4 \cdot |\mathcal{E}|$ boolean functions, each corresponding to a very simple BDD

- $APM_\alpha = \prod_{p: D-p, \alpha=1} (x_p = 1)$ (all predecessor places of α are marked)
- $NPM_\alpha = \prod_{p: D-p, \alpha=1} (x_p = 0)$ (no predecessor place of α is marked)
- $ASM_\alpha = \prod_{p: D+p, \alpha=1} (x_p = 1)$ (all successor places of α are marked)
- $NSM_\alpha = \prod_{p: D+p, \alpha=1} (x_p = 0)$ (no successor place of α is marked)

The **topological image computation** for a transition α on a set of states \mathcal{U} can be expressed as

$$\mathcal{N}_\alpha(\mathcal{U}) = (((\mathcal{U} \div APM_\alpha) \cdot NPM_\alpha) \div NSM_\alpha) \cdot ASM_\alpha$$

where “ \div ” indicates the **cofactor** operator and “ \cdot ” indicates boolean conjunction

Given

- a boolean function f over (x_L, \dots, x_1)
- a literal $x_k = i_k$, with $L \geq k \geq 1$ and $i_k \in \mathbb{B}$

the cofactor $f \div (x_k = i_k)$ is defined as

- $f(x_L, \dots, x_{k+1}, i_k, x_{k-1}, \dots, x_1)$

The extension to multiple literals, $f \div (x_{k_c} = i_{k_c}, \dots, x_{k_1} = i_{k_1})$, is recursively defined as

- $f(x_L, \dots, x_{k_c+1}, i_{k_c}, x_{k_c-1}, \dots, x_1) \div (x_{k_{c-1}} = i_{k_{c-1}}, \dots, x_{k_1} = i_{k_1})$

Thus, \mathcal{N} is stored in a **disjunctively partition** form as $\mathcal{N} = \bigcup_{\alpha \in \mathcal{E}} \mathcal{N}_\alpha$

An L -level BDD encodes a set of states \mathcal{Y} as a subset of the **potential state space** $\widehat{\mathcal{X}} = \mathbb{B}^L$

$$\mathbf{i} \equiv (\mathbf{i}_L, \dots, \mathbf{i}_1) \in \mathcal{Y} \Leftrightarrow \text{the corresponding path from the root leads to terminal 1}$$

A $2L$ -level BDD encodes the **next-state function** $\mathcal{N} : \widehat{\mathcal{X}} \rightarrow 2^{\widehat{\mathcal{X}}}$

$$\mathbf{j} \in \mathcal{N}(\mathbf{i}) \Leftrightarrow \text{the system can go from } \mathbf{i} \text{ to } \mathbf{j} \text{ in one step}$$

The state space \mathcal{X}_{reach} is the fixpoint of the iteration

$$\mathcal{X}_{init} \quad \mathcal{N}(\mathcal{X}_{init}) \quad \mathcal{N}(\mathcal{N}(\mathcal{X}_{init})) \quad \mathcal{N}(\mathcal{N}(\mathcal{N}(\mathcal{X}_{init}))) \quad \dots$$

The main operation is a repeated application of the relational product operator

$BfsGen(\mathcal{X}_{init}, \mathcal{N})$ is

```

1  $\mathcal{Y} \leftarrow \mathcal{X}_{init};$            known states
2  $\mathcal{U} \leftarrow \mathcal{X}_{init};$        unexplored states
3 while  $\mathcal{U} \neq \emptyset$  do
4    $\mathcal{W} \leftarrow \mathcal{N}(\mathcal{U});$    potentially new states
5    $\mathcal{U} \leftarrow \mathcal{W} \setminus \mathcal{Y};$  truly new states
6    $\mathcal{Y} \leftarrow \mathcal{Y} \cup \mathcal{U};$ 
7 return  $\mathcal{Y};$ 

```

sets and relations are encoded using BDDs

runtime is proportional to the BDD sizes

If $|\mathcal{X}_k| > 2$, use multiple boolean levels to encode \mathbf{i}_k

Ordered multiway decision diagrams (MDDs)

Assume a **domain** $\widehat{\mathcal{X}} = \mathcal{X}_L \times \dots \times \mathcal{X}_1$, where $\mathcal{X}_k = \{0, 1, \dots, n_k - 1\}$, for some $n_k \in \mathbb{N}$

Assume the **range** $\mathcal{X}_0 = \mathbb{B}$

An MDD is an acyclic directed edge-labeled graph where:

- The only **terminal** nodes can be **0** and **1**, and are at level 0 $0.lvl = 1.lvl = 0$
- A **nonterminal** node p is at a **level** k , with $L \geq k \geq 1$ $p.lvl = k$
- For each $i_k \in \mathcal{X}_k$, a nonterminal node p at level k has an outgoing edge pointing to **child** $p[i_k]$
- The level of a child is lower than that of p $p[i_k].lvl < p.lvl$
- A node p at level k encodes the **function** $v_p : \widehat{\mathcal{X}} \rightarrow \mathbb{B}$ defined recursively by

$$v_p(x_L, \dots, x_1) = \begin{cases} p & \text{if } k = 0 \\ v_{p[x_k]}(x_L, \dots, x_1) & \text{if } k > 0 \end{cases}$$

Instead of levels, we can also talk of **variables**:

- The terminal nodes are associated with the **range variable** x_0
- A nonterminal node is associated with a **domain variable** x_k , with $L \geq k \geq 1$

Canonical versions of MDDs

For **canonical** MDDs, we further require that




- There are no **duplicates**: if $p.lvl = q.lvl = k$ and $p[i_k] = q[i_k]$ for all $i_k \in \mathcal{X}_k$, then $p = q$

Then, if the MDD is **quasi-reduced**, there is **no level skipping**:

- The only **root** nodes with no incoming arcs are at level L
- If a node p is at level k , each child $p[i_k]$ is at level $k - 1$

Or, if the MDD is **fully-reduced**, there is **maximum level skipping**:

- There are no **redundant** nodes p at level k satisfying $p[i_k] = q$ for all $i_k \in \mathcal{X}_k$



What if we don't know the range of each \mathcal{X}_k ?

We can simply assume $\hat{\mathcal{X}} = \mathbb{N}^L$

All but a finite number of edges point to 0

Only nodes encoding \emptyset are redundant

Fully-reduced MDDs: each node in the MDD encodes a different function

Quasi-reduced MDDs: each node at a given level of the MDD encodes a different function

Given an event $\alpha \in \mathcal{E}$, consider the subset of the state variables $\{x_L, \dots, x_1\}$ that:

- can be modified by α : $\mathcal{V}_M(\alpha) = \{x_k : \exists i, i' \in \hat{\mathcal{X}}, i' \in \mathcal{N}_\alpha(i) \wedge i[k] \neq i'[k]\}$
- can disable α : $\mathcal{V}_D(\alpha) = \{x_k : \exists i, j \in \hat{\mathcal{X}}, \forall h \neq k, i[h] = j[h] \wedge \mathcal{N}_\alpha(i) \neq \emptyset \wedge \mathcal{N}_\alpha(j) = \emptyset\}$

If $x_k \notin \mathcal{V}_M \cup \mathcal{V}_D$, we say that event α and variable x_k , or level k , are independent

Most events in a globally-asynchronous locally-synchronous model are highly localized:

- Let $Top(\alpha) = \max\{k : x_k \in \mathcal{V}_M(\alpha) \cup \mathcal{V}_D(\alpha)\}$ be the highest level dependent on α
- Let $Bot(\alpha) = \min\{k : x_k \in \mathcal{V}_M(\alpha) \cup \mathcal{V}_D(\alpha)\}$ be the lowest level dependent on α
- The span (of levels) $\{Top(\alpha), \dots, Bot(\alpha)\}$ for event α is often much smaller than $\{L, \dots, 1\}$

fully/quasi-reduced $2L$ -level MDD encoding does not exploit locality
 need **Kronecker, identity-reduced** $2L$ -level MDD, or **MxD** encoding

Kronecker description of the next-state function

$\mathcal{N} : \hat{\mathcal{X}} \rightarrow 2^{\hat{\mathcal{X}}}$ can be thought of as a boolean matrix $\mathbf{N} \in \mathbb{B}^{|\hat{\mathcal{X}}| \times |\hat{\mathcal{X}}|}$

The model is Kronecker-consistent if $\mathbf{N} = \sum_{\alpha \in \mathcal{E}} \left(\bigotimes_{L \geq k \geq 1} \mathbf{N}_{k,\alpha} \right)$ (using boolean operations)

In other words, $\mathcal{N} = \bigvee_{\alpha \in \mathcal{E}} \mathcal{N}_\alpha$ and each $\mathcal{N}_\alpha = \bigwedge_{L \geq k \geq 1} \mathcal{N}_{k,\alpha}$ where

$\mathcal{N}_{k,\alpha} : \mathcal{X}_k \rightarrow 2^{\mathcal{X}_k}$ is encoded by the boolean matrix $\mathbf{N}_{k,\alpha} \in \mathbb{B}^{|\mathcal{X}_k| \times |\mathcal{X}_k|}$

Locality: If the k^{th} local state does not affect and is not affected by event α , then $\mathbf{N}_{k,\alpha} = \mathbf{I}$


encode a huge \mathbf{N} with $L \cdot |\mathcal{E}|$ "small" matrices

Using structural information to encode \mathcal{N} ($L = 5$)

$\mathcal{X}_5 = ?$ $\mathcal{X}_4 = ?$ $\mathcal{X}_3 = ?$ $\mathcal{X}_2 = ?$ $\mathcal{X}_1 = ?$

EVENTS →					
LEVELS ↓	$\mathbf{N}_{5,a} : ?$	I	I	I	$\mathbf{N}_{5,e} : ?$
	$\mathbf{N}_{4,a} : ?$	$\mathbf{N}_{4,b} : ?$	$\mathbf{N}_{4,c} : ?$	I	I
	I	$\mathbf{N}_{3,b} : ?$	$\mathbf{N}_{3,c} : ?$	I	$\mathbf{N}_{3,e} : ?$
	$\mathbf{N}_{2,a} : ?$	I	I	$\mathbf{N}_{2,d} : ?$	I
	I	I	I	$\mathbf{N}_{1,d} : ?$	$\mathbf{N}_{1,e} : ?$


$Top(a) : 5$ $Top(b) : 4$ $Top(c) : 4$ $Top(d) : 2$ $Top(e) : 5$
 $Bot(a) : 2$ $Bot(b) : 3$ $Bot(c) : 3$ $Bot(d) : 1$ $Bot(e) : 1$



we determine a priori from the model whether $\mathbf{N}_{k,\alpha} = \mathbf{I}$

$\mathcal{X}_5: \{p^1, p^0\} \equiv \{0,1\}$ $\mathcal{X}_4: \{q^0, q^1\} \equiv \{0,1\}$ $\mathcal{X}_3: \{r^0, r^1\} \equiv \{0,1\}$ $\mathcal{X}_2: \{s^0, s^1\} \equiv \{0,1\}$ $\mathcal{X}_1: \{t^0, t^1\} \equiv \{0,1\}$


EVENTS →					
LEVELS ↓	$\mathbf{N}_{5,a}: \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$	I	I	I	$\mathbf{N}_{5,e}: \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$
	$\mathbf{N}_{4,a}: \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$	$\mathbf{N}_{4,b}: \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$	$\mathbf{N}_{4,c}: \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$	I	I
	I	$\mathbf{N}_{3,b}: \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$	$\mathbf{N}_{3,c}: \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}$	I	$\mathbf{N}_{3,e}: \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$
	$\mathbf{N}_{2,a}: \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$	I	I	$\mathbf{N}_{2,d}: \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$	I
	I	I	I	$\mathbf{N}_{1,d}: \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$	$\mathbf{N}_{1,e}: \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$



$Top(a):5$ $Top(b):4$ $Top(c):4$ $Top(d):2$ $Top(e):5$
 $Bot(a):2$ $Bot(b):3$ $Bot(c):3$ $Bot(d):1$ $Bot(e):1$

$\mathcal{X}_4 = ?$ $\mathcal{X}_3 = ?$ $\mathcal{X}_2 = ?$ $\mathcal{X}_1 = ?$

EVENTS →					
LEVELS ↓	$\mathbf{N}_{4,a}: ?$	I	I	I	$\mathbf{N}_{4,e}: ?$
	$\mathbf{N}_{3,a}: ?$	$\mathbf{N}_{3,b}: ?$	$\mathbf{N}_{3,c}: ?$	I	$\mathbf{N}_{3,e}: ?$
	$\mathbf{N}_{2,a}: ?$	I	I	$\mathbf{N}_{2,d}: ?$	I
	I	I	I	$\mathbf{N}_{1,d}: ?$	$\mathbf{N}_{1,e}: ?$




$Top(a):4$ $Top(b):3$ $Top(c):3$ $Top(d):2$ $Top(e):4$
 $Bot(a):2$ $Bot(b):3$ $Bot(c):3$ $Bot(d):1$ $Bot(e):1$


we determine automatically from the model whether $\mathbf{N}_{k,\alpha} = \mathbf{I}$

$\mathcal{X}_4: \{p^1, p^0\} \equiv \{0,1\}$ $\mathcal{X}_3: \{q^0 r^0, q^1 r^0, q^0 r^1\} \equiv \{0,1,2\}$ $\mathcal{X}_2: \{s^0, s^1\} \equiv \{0,1\}$ $\mathcal{X}_1: \{t^0, t^1\} \equiv \{0,1\}$

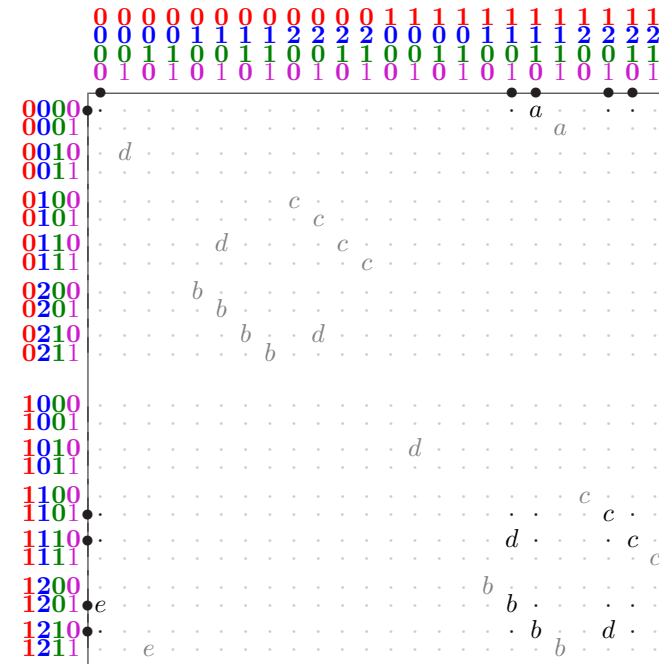
EVENTS →					
LEVELS ↓	$\mathbf{N}_{4,a}: \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$	I	I	I	$\mathbf{N}_{4,e}: \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$
	$\mathbf{N}_{3,a}: \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$	$\mathbf{N}_{3,b}: \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$	$\mathbf{N}_{3,c}: \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}$	I	$\mathbf{N}_{3,e}: \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}$
	$\mathbf{N}_{2,a}: \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$	I	I	$\mathbf{N}_{2,d}: \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$	I
	I	I	I	$\mathbf{N}_{1,d}: \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$	$\mathbf{N}_{1,e}: \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$




$Top(a):4$ $Top(b):3$ $Top(c):3$ $Top(d):2$ $Top(e):4$
 $Bot(a):2$ $Bot(b):3$ $Bot(c):3$ $Bot(d):1$ $Bot(e):1$



$\{p^1, p^0\} \equiv \{0,1\}$
 $\{q^0 r^0, q^1 r^0, q^0 r^1\} \equiv \{0,1,2\}$
 $\{s^0, s^1\} \equiv \{0,1\}$
 $\{t^0, t^1\} \equiv \{0,1\}$




To model software we must be able to model assignments, e.g., with self-modifying Petri nets



To enforce Kronecker consistency, places $a, b,$ and c must belong to the same submodel

However, each assignment may cause further grouping of variables



The local state space for $\{a, b, c, x, y, z, w, m, r\}$ may be too large

[Miner QEST 2004] [Ciardo and Yu CHARME 2005]

A self-modifying Petri net with inhibitor arcs is a tuple $(\mathcal{P}, \mathcal{E}, \mathbf{D}^-, \mathbf{D}^+, \mathbf{D}^\circ, \mathbf{i}^{init})$ where:

- \mathcal{P} and \mathcal{E} places and transitions
- $\mathbf{D}^-, \mathbf{D}^+ : \mathcal{P} \times \mathcal{E} \times \mathbb{N}^{|\mathcal{P}|} \rightarrow \mathbb{N}$ marking-dependent input, output arc cardinalities
- $\mathbf{D}^\circ : \mathcal{P} \times \mathcal{E} \times \mathbb{N}^{|\mathcal{P}|} \rightarrow \mathbb{N} \cup \{\infty\}$ marking-dependent inhibitor arc cardinalities
- $\mathbf{i}^{init} : \mathbb{N}^{|\mathcal{P}|}$ initial marking


Transition α is enabled in marking $\mathbf{i} \in \mathbb{N}^{|\mathcal{P}|}$ iff $\forall p \in \mathcal{P}, \mathbf{D}_{p,\alpha}^-(\mathbf{i}) \leq i_p \wedge \mathbf{D}_{p,\alpha}^\circ(\mathbf{i}) > i_p$

If α is enabled in \mathbf{i} , it can fire and lead to marking \mathbf{j} $\forall p \in \mathcal{P}, j_p = i_p - \mathbf{D}_{p,\alpha}^-(\mathbf{i}) + \mathbf{D}_{p,\alpha}^+(\mathbf{i})$

The effect of α is deterministic, so we can write $\mathbf{i} \xrightarrow{\alpha} \mathbf{j}$ or use the general notation $\mathbf{j} \in \mathcal{N}_\alpha(\mathbf{i})$

$\hat{\mathcal{X}} \equiv \mathbb{N}^{|\mathcal{P}|}$ $\mathcal{X}_{init} \equiv \{\mathbf{i}^{init}\}$ $\mathcal{N} \equiv \bigcup_{\alpha \in \mathcal{E}} \mathcal{N}_\alpha$


Problem with pre-generating the sets \mathcal{X}_k



each one of these four submodels is unbounded (in isolation)

A poor solution to the problem

Modifying the model to enforce known bounds is difficult (especially if we want the smallest \mathcal{X}_k)



more importantly, it's dangerous!

Another way express Kronecker-consistency:

$\mathcal{N} : \hat{\mathcal{X}} \rightarrow 2^{\hat{\mathcal{X}}}$ tells us which transitions between potential states are possible
 $\mathcal{N} = \bigvee_{\alpha \in \mathcal{E}} \mathcal{N}_\alpha$, $\mathcal{N}_\alpha : \hat{\mathcal{X}} \rightarrow 2^{\hat{\mathcal{X}}}$ asynchronous, disjunctive, decomposition of \mathcal{N}
 $\mathcal{N}_\alpha = (\bigwedge_{k \in \mathcal{D}_\alpha} \mathcal{N}_{k,\alpha}) \wedge (\bigwedge_{k \in \overline{\mathcal{D}}_\alpha} \mathcal{I}_k)$ synchronous, conjunctive, decomposition by level of \mathcal{N}_α


$\mathcal{D}_\alpha \subseteq \{L, \dots, 1\}$	set of levels on which event α depends/affects
$\overline{\mathcal{D}}_\alpha = \{L, \dots, 1\} \setminus \mathcal{D}_\alpha$	set of levels on which event α does not depend/affect
\mathcal{I}_k	identity transformation for the states of submodel k
$\mathcal{N}_{k,\alpha} : \mathcal{X}_k \rightarrow 2^{\mathcal{X}_k}$	next-state function restricted to level k only

The disjunctive-then-conjunctive decomposition of \mathcal{N} can be applied to arbitrary models:

$\mathcal{N}_\alpha = (\bigwedge_{c=1}^{m_\alpha} \mathcal{N}_{\mathcal{D}_{c,\alpha}}) \wedge (\bigwedge_{k \in \overline{\mathcal{D}}_\alpha} \mathcal{I}_k)$ general conjunctive decomposition of \mathcal{N}_α

$\bigcup_{c=1}^{m_\alpha} \mathcal{D}_{c,\alpha} = \mathcal{D}_\alpha \subseteq \{L, \dots, 1\}$	$\mathcal{N}_{\mathcal{D}_{c,\alpha}}$ depends on a set of levels
$\mathcal{N}_{\mathcal{D}_{c,\alpha}} : (\prod_{k \in \mathcal{D}_{c,\alpha}} \mathcal{X}_k) \rightarrow 2^{(\prod_{k \in \mathcal{D}_{c,\alpha}} \mathcal{X}_k)}$	

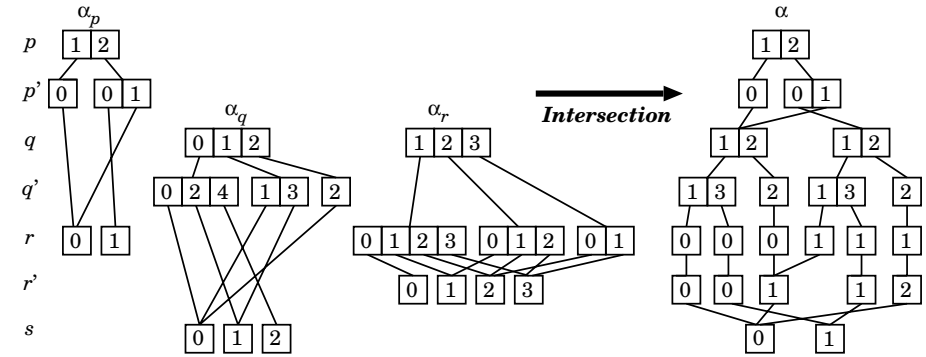
$Top(\alpha) = \max \mathcal{D}_\alpha$ $Bot(\alpha) = \min \mathcal{D}_\alpha$



Equivalent pseudocode (simultaneous statements)

if $p > r \wedge q + 2s \leq 4 \wedge q > 0 \wedge r + q \leq 4$ then
 $p \leftarrow p - (r + 1);$
 $q \leftarrow q + 2s;$
 $r \leftarrow r + q - 1;$

Assume $\mathcal{X}_p = \{0, 1, 2\}$, $\mathcal{X}_q = \{0, 1, 2, 3, 4\}$, $\mathcal{X}_r = \{0, 1, 2, 3\}$, $\mathcal{X}_s = \{0, 1, 2\}$



special reduction rule and interpretation of skipped levels

To confirm a new local state $\mathbf{i}_h \in \mathcal{X}_h$:

- 1 for $k = L$ down to h do
- 2 for each α such that $Top(\alpha) = k$ and $h \in \mathcal{D}_\alpha$ do
- 3 for each $\mathcal{D}_{c,\alpha}$ containing h do
- 4 explicitly build the set \mathcal{Y} of potential transitions from $\{\mathbf{i}_h\} \times (\prod_{k \in \mathcal{D}_{c,\alpha} \setminus \{h\}} \mathcal{X}_k)$;
- 5 $\mathcal{N}_{\mathcal{D}_{c,\alpha}} \leftarrow \mathcal{N}_{\mathcal{D}_{c,\alpha}} \cup \mathcal{Y}$; *build the conjunct*
- 6 $\mathcal{N}_\alpha \leftarrow (\bigwedge_{c=1}^{m_\alpha} \mathcal{N}_{\mathcal{D}_{c,\alpha}}) \wedge (\bigwedge_{l \in \overline{\mathcal{D}}_\alpha} \mathcal{I}_l)$; *build the disjunct*
- 7 $\mathcal{N}_k \leftarrow \mathcal{N}_k \cup \mathcal{N}_\alpha$; $\mathcal{N}_k = \bigcup_{Top(\alpha)=k} \mathcal{N}_\alpha$


the explicit enumeration of the elements of $\prod_{h \in \mathcal{D}_{c,\alpha}} \mathcal{X}_h$
is the reason for seeking the smallest possible $\mathcal{D}_{c,\alpha}$

In addition to the

- quasi-reduced form
- reduced form

we need an

- identity-reduced form to be used for "to", or "primed", levels



Identity-reduced level k'

Identity-reduced level k' and fully-reduced level k

canonical even if we use different reduction rules for each level

Accelerated fixpoint computation

An L -level MDD encodes a set of states $\mathcal{Y} \subseteq \widehat{\mathcal{X}} = \mathcal{X}_L \times \dots \times \mathcal{X}_1$

$\mathbf{i} \equiv (i_L, \dots, i_1) \in \mathcal{Y} \Leftrightarrow$ the path from the root corresponding to \mathbf{i} leads to terminal $\mathbf{1}$

A $2L$ -level MDD encodes the next-state function $\mathcal{N} : \widehat{\mathcal{X}} \rightarrow 2^{\widehat{\mathcal{X}}}$

$\mathbf{j} \in \mathcal{N}(\mathbf{i}) \Leftrightarrow$ the path from the root corresponding to the interleaving of \mathbf{i} and \mathbf{j} leads to terminal $\mathbf{1}$

The state space \mathcal{X}_{reach} is the fixpoint of the iteration

$$\mathcal{X}_{init} \cup \mathcal{N}(\mathcal{X}_{init}) \cup \mathcal{N}(\mathcal{N}(\mathcal{X}_{init})) \cup \mathcal{N}(\mathcal{N}(\mathcal{N}(\mathcal{X}_{init}))) \cup \dots$$

Standard method

Alternative *All* method

ExploreMdd($\mathcal{X}_{init}, \mathcal{N}$) is

```

1  $\mathcal{Y} \leftarrow \mathcal{X}_{init};$  known states
2  $\mathcal{U} \leftarrow \mathcal{X}_{init};$  unexplored states
3 repeat
4    $\mathcal{W} \leftarrow \mathcal{N}(\mathcal{U});$  potentially new states
5    $\mathcal{U} \leftarrow \mathcal{W} \setminus \mathcal{Y};$  truly new states
6    $\mathcal{Y} \leftarrow \mathcal{Y} \cup \mathcal{U};$ 
7 until  $\mathcal{U} = \emptyset;$ 
8 return  $\mathcal{Y};$ 

```

AllExploreMdd($\mathcal{X}_{init}, \mathcal{N}$) is

```

1  $\mathcal{Y} \leftarrow \mathcal{X}_{init};$ 
2 repeat
3    $\mathcal{O} \leftarrow \mathcal{Y};$  old states
4    $\mathcal{Y} \leftarrow \mathcal{O} \cup \mathcal{N}(\mathcal{O});$  new states
5 until  $\mathcal{O} = \mathcal{Y};$ 
6 return  $\mathcal{Y};$ 

```

Chaining the next-state function \mathcal{N} [Roig95]

If \mathcal{N} is stored in a disjointly partitioned form as $\mathcal{N} = \bigcup_{\alpha \in \mathcal{E}} \mathcal{N}_\alpha$, using $|\mathcal{E}|$ MDDs, the effect of

```

 $\mathcal{W} \leftarrow \mathcal{N}(\mathcal{U});$  potentially new states
 $\mathcal{U} \leftarrow \mathcal{W} \setminus \mathcal{Y};$  truly new states

```

is exactly achieved with the statements

```

 $\mathcal{W} \leftarrow \emptyset;$ 
for each  $\alpha \in \mathcal{E}$  do
   $\mathcal{W} \leftarrow \mathcal{W} \cup \mathcal{N}_\alpha(\mathcal{U});$ 
 $\mathcal{U} \leftarrow \mathcal{W} \setminus \mathcal{Y};$ 

```

However, if we do not require strict breadth-first order, we can use chaining and do

```

for each  $\alpha \in \mathcal{E}$  do
   $\mathcal{U} \leftarrow \mathcal{U} \cup \mathcal{N}_\alpha(\mathcal{U});$ 
 $\mathcal{U} \leftarrow \mathcal{U} \setminus \mathcal{Y};$ 

```

Symbolic *SsGen*: breadth-first vs. chaining, new vs. all states

BfSsGen($\mathcal{X}_{init}, \{\mathcal{N}_\alpha : \alpha \in \mathcal{E}\}$)

```

1  $\mathcal{Y} \leftarrow \mathcal{X}_{init};$  known states
2  $\mathcal{U} \leftarrow \mathcal{X}_{init};$  unexplored known states
3 repeat
4    $\mathcal{W} \leftarrow \emptyset;$ 
5   for each  $\alpha \in \mathcal{E}$  do
6      $\mathcal{W} \leftarrow \mathcal{W} \cup \mathcal{N}_\alpha(\mathcal{U});$ 
7    $\mathcal{U} \leftarrow \mathcal{W} \setminus \mathcal{Y};$  truly new states
8    $\mathcal{Y} \leftarrow \mathcal{Y} \cup \mathcal{U};$ 
9 until  $\mathcal{U} = \emptyset;$ 
10 return  $\mathcal{Y};$ 

```

ChSsGen($\mathcal{X}_{init}, \{\mathcal{N}_\alpha : \alpha \in \mathcal{E}\}$)

```

1  $\mathcal{Y} \leftarrow \mathcal{X}_{init};$  known states
2  $\mathcal{U} \leftarrow \mathcal{X}_{init};$  unexplored known states
3 repeat
4   for each  $\alpha \in \mathcal{E}$  do
5      $\mathcal{U} \leftarrow \mathcal{U} \cup \mathcal{N}_\alpha(\mathcal{U});$ 
6    $\mathcal{U} \leftarrow \mathcal{U} \setminus \mathcal{Y};$  truly new states
7    $\mathcal{Y} \leftarrow \mathcal{Y} \cup \mathcal{U};$ 
8 until  $\mathcal{U} = \emptyset;$ 
9 return  $\mathcal{Y};$ 

```

AllBfSsGen($\mathcal{X}_{init}, \{\mathcal{N}_\alpha : \alpha \in \mathcal{E}\}$)

```

1  $\mathcal{Y} \leftarrow \mathcal{X}_{init};$  known states
2 repeat
3    $\mathcal{O} \leftarrow \mathcal{Y};$  save old state space
4    $\mathcal{W} \leftarrow \emptyset;$ 
5   for each  $\alpha \in \mathcal{E}$  do
6      $\mathcal{W} \leftarrow \mathcal{W} \cup \mathcal{N}_\alpha(\mathcal{O});$ 
7    $\mathcal{Y} \leftarrow \mathcal{O} \cup \mathcal{W};$ 
8 until  $\mathcal{O} = \mathcal{Y};$ 
9 return  $\mathcal{Y};$ 

```

AllChSsGen($\mathcal{X}_{init}, \{\mathcal{N}_\alpha : \alpha \in \mathcal{E}\}$)

```

1  $\mathcal{Y} \leftarrow \mathcal{X}_{init};$  known states
2 repeat
3    $\mathcal{O} \leftarrow \mathcal{Y};$  save old state space
4   for each  $\alpha \in \mathcal{E}$  do
5      $\mathcal{Y} \leftarrow \mathcal{Y} \cup \mathcal{N}_\alpha(\mathcal{Y});$ 
6 until  $\mathcal{O} = \mathcal{Y};$ 
7 return  $\mathcal{Y};$ 

```

N	X _{reach}	Time (sec)				Memory (MB)				final
		Bf	AllBf	Ch	AllCh	Bf	AllBf	Ch	AllCh	

Dining Philosophers: $L = N/2$, $|\mathcal{X}_k| = 34$ for all k

50	2.2×10^{31}	37.6	36.8	1.3	1.3	146.8	131.6	2.2	2.2	<0.1
100	5.0×10^{62}	644.1	630.4	5.4	5.3	>999.9	>999.9	8.9	8.9	<0.1
1000	9.2×10^{626}	—	—	895.4	915.5	—	—	895.2	895.0	0.3

Slotted Ring Network: $L = N$, $|\mathcal{X}_k| = 15$ for all k

5	5.3×10^4	0.2	0.3	0.1	0.1	0.8	1.1	0.3	0.2	<0.1
10	8.3×10^9	21.5	24.1	2.1	1.2	39.0	45.0	5.7	3.3	<0.1
15	1.5×10^{15}	745.4	771.5	18.5	8.9	344.3	375.4	35.1	20.2	<0.1

Round Robin Mutual Exclusion: $L = N + 1$, $|\mathcal{X}_k| = 10$ for all k except $|\mathcal{X}_1| = N + 1$

10	2.3×10^4	0.2	0.3	0.1	0.1	0.6	1.2	0.1	0.1	<0.1
20	4.7×10^7	2.7	4.4	0.3	0.3	5.9	12.8	0.5	0.5	<0.1
50	1.3×10^{17}	263.2	427.6	2.9	2.8	126.7	257.7	4.3	3.8	0.1

FMS: $L = 19$, $|\mathcal{X}_k| = N + 1$ for all k except $|\mathcal{X}_{17}| = 4$, $|\mathcal{X}_{12}| = 3$, $|\mathcal{X}_7| = 2$


5	2.9×10^6	0.7	0.7	0.1	0.1	2.6	2.2	0.4	0.2	<0.1
10	2.5×10^9	7.0	5.8	0.5	0.3	18.2	14.7	2.3	1.3	<0.1
25	8.5×10^{13}	677.2	437.9	12.9	5.1	319.7	245.3	42.7	21.2	0.1

MDD node p at level k is **saturated** if it encodes a fixed point w.r.t. any event α s.t. the highest MDD level it depends on, $Top(\alpha)$, is at most $k \Rightarrow$ all MDD nodes reachable from p are also saturated

- build the L -level MDD encoding of \mathcal{X}_{init} if $|\mathcal{X}_{init}| = 1$, there is one node per level
- saturate each node at level 1: fire in them all events α s.t. $Top(\alpha) = 1$
- saturate each node at level 2: fire in them all events α s.t. $Top(\alpha) = 2$ (if this creates nodes at level 1, saturate them immediately upon creation)
- saturate each node at level 3: fire in them all events α s.t. $Top(\alpha) = 3$ (if this creates nodes at levels 2 or 1, saturate them immediately upon creation)
- ...
- saturate the root node at level L : fire in it all events α s.t. $Top(\alpha) = L$ (if this creates nodes at levels $L-1, L-2, \dots, 1$, saturate them immediately upon creation)


states are **not** discovered in breadth-first order

enormous time and memory savings for asynchronous systems




a	bc	d	e
	I	I	
		I	
	I		I
I	I		

- <4|2> 0 $\mathcal{X}_4 = \{p^1\} \equiv \{0\}$
- <3|2> 0 $\mathcal{X}_3 = \{q^0 r^0\} \equiv \{0\}$
- <2|2> 0 $\mathcal{X}_2 = \{s^0\} \equiv \{0\}$
- <1|2> 0 $\mathcal{X}_1 = \{t^0\} \equiv \{0\}$




a	bc	d	e
0 : 1	I	I	0 : 2
0 : 1	0 : -	I	0 : -
0 : 1	I	0 : -	I
I	I	0 : 1	0 : -

- <4|2> 0 $\mathcal{X}_4 = \{p^1, p^0, p^2\} \equiv \{0, 1, 2\}$
- <3|2> 0 $\mathcal{X}_3 = \{q^0 r^0, q^1 r^0\} \equiv \{0, 1\}$
- <2|2> 0 $\mathcal{X}_2 = \{s^0, s^1\} \equiv \{0, 1\}$
- <1|2> 0 $\mathcal{X}_1 = \{t^0, t^1\} \equiv \{0, 1\}$




<i>a</i>	<i>bc</i>	<i>d</i>	<i>e</i>
0 : 1	I	I	0 : 2
0 : 1	0 : -	I	0 : -
0 : 1	I	0 : -	I
I	I	0 : 1	0 : -

$\langle 4|2\rangle$ 0
 $\mathcal{X}_4 = \{p^1, p^0, p^2\} \equiv \{0, 1, 2\}$
 $\langle 3|2\rangle$ 0
 $\mathcal{X}_3 = \{q^0 r^0, q^1 r^0\} \equiv \{0, 1\}$
 $\langle 2|2\rangle$ 0
 $\mathcal{X}_2 = \{s^0, s^1\} \equiv \{0, 1\}$
 $\langle 1|2\rangle$ 0
 $\mathcal{X}_1 = \{t^0, t^1\} \equiv \{0, 1\}$




<i>a</i>	<i>bc</i>	<i>d</i>	<i>e</i>
0 : 1	I	I	0 : 2
0 : 1	0 : -	I	0 : -
0 : 1	I	0 : -	I
I	I	0 : 1	0 : -

$\langle 4|2\rangle$ 0 1
 $\mathcal{X}_4 = \{p^1, p^0, p^2\} \equiv \{0, 1, 2\}$
 $\langle 3|2\rangle$ 0 $\langle 3|3\rangle$ 1
 $\mathcal{X}_3 = \{q^0 r^0, q^1 r^0\} \equiv \{0, 1\}$
 $\langle 2|2\rangle$ 0 $\langle 2|3\rangle$ 1
 $\mathcal{X}_2 = \{s^0, s^1\} \equiv \{0, 1\}$
 $\langle 1|2\rangle$ 0
 $\mathcal{X}_1 = \{t^0, t^1\} \equiv \{0, 1\}$




<i>a</i>	<i>bc</i>	<i>d</i>	<i>e</i>
0 : 1	I	I	0 : 2
0 : 1	0 : -	I	0 : -
0 : 1 1 : 2	I	0 : - 1 : 0	I
I	I	0 : 1	0 : -

$\langle 4|2\rangle$ 0 1
 $\mathcal{X}_4 = \{p^1, p^0, p^2\} \equiv \{0, 1, 2\}$
 $\langle 3|2\rangle$ 0 $\langle 3|3\rangle$ 1
 $\mathcal{X}_3 = \{q^0 r^0, q^1 r^0\} \equiv \{0, 1\}$
 $\langle 2|2\rangle$ 0 $\langle 2|3\rangle$ 1
 $\mathcal{X}_2 = \{s^0, s^1, s^2\} \equiv \{0, 1, 2\}$
 $\langle 1|2\rangle$ 0
 $\mathcal{X}_1 = \{t^0, t^1\} \equiv \{0, 1\}$




<i>a</i>	<i>bc</i>	<i>d</i>	<i>e</i>
0 : 1	I	I	0 : 2
0 : 1	0 : -	I	0 : -
0 : 1 1 : 2	I	0 : - 1 : 0	I
I	I	0 : 1	0 : -

$\langle 4|2\rangle$ 0 1
 $\mathcal{X}_4 = \{p^1, p^0, p^2\} \equiv \{0, 1, 2\}$
 $\langle 3|2\rangle$ 0 $\langle 3|3\rangle$ 1
 $\mathcal{X}_3 = \{q^0 r^0, q^1 r^0\} \equiv \{0, 1\}$
 $\langle 2|2\rangle$ 0 $\langle 2|3\rangle$ 0 1
 $\mathcal{X}_2 = \{s^0, s^1, s^2\} \equiv \{0, 1, 2\}$
 $\langle 1|2\rangle$ 0 $\langle 1|3\rangle$ 1
 $\mathcal{X}_1 = \{t^0, t^1\} \equiv \{0, 1\}$




<i>a</i>	<i>bc</i>	<i>d</i>	<i>e</i>
0 : 1	I	I	0 : 2
0 : 1	0 : -	I	0 : -
0 : 1 1 : 2	I	0 : - 1 : 0	I
I	I	0 : 1 1 : 2	0 : - 1 : 0

$\langle 4|2 \rangle \begin{bmatrix} 0 & 1 \end{bmatrix} \quad \mathcal{X}_4 = \{p^1, p^0, p^2\} \equiv \{0, 1, 2\}$
 $\langle 3|2 \rangle \begin{bmatrix} 0 & & & 1 \end{bmatrix} \quad \mathcal{X}_3 = \{q^0 r^0, q^1 r^0\} \equiv \{0, 1\}$
 $\langle 2|2 \rangle \begin{bmatrix} 0 & & & 0 & 1 \end{bmatrix} \quad \mathcal{X}_2 = \{s^0, s^1, s^2\} \equiv \{0, 1, 2\}$
 $\langle 1|2 \rangle \begin{bmatrix} 0 & & & & 1 \end{bmatrix} \quad \mathcal{X}_1 = \{t^0, t^1, t^2\} \equiv \{0, 1, 2\}$




<i>a</i>	<i>bc</i>	<i>d</i>	<i>e</i>
0 : 1	I	I	0 : 2
0 : 1	0 : -	I	0 : -
0 : 1 1 : 2	I	0 : - 1 : 0	I
I	I	0 : 1 1 : 2	0 : - 1 : 0

$\langle 4|2 \rangle \begin{bmatrix} 0 & 1 \end{bmatrix} \quad \mathcal{X}_4 = \{p^1, p^0, p^2\} \equiv \{0, 1, 2\}$
 $\langle 3|2 \rangle \begin{bmatrix} 0 & & & 1 \end{bmatrix} \quad \mathcal{X}_3 = \{q^0 r^0, q^1 r^0\} \equiv \{0, 1\}$
 $\langle 2|2 \rangle \begin{bmatrix} 0 & & & 0 & 1 \end{bmatrix} \quad \mathcal{X}_2 = \{s^0, s^1, s^2\} \equiv \{0, 1, 2\}$
 $\langle 1|2 \rangle \begin{bmatrix} 0 & & & & 1 \end{bmatrix} \quad \mathcal{X}_1 = \{t^0, t^1, t^2\} \equiv \{0, 1, 2\}$




<i>a</i>	<i>bc</i>	<i>d</i>	<i>e</i>
0 : 1	I	I	0 : 2
0 : 1 1 : 2	0 : - 1 : 3	I	0 : - 1 : -
0 : 1 1 : 2	I	0 : - 1 : 0	I
I	I	0 : 1 1 : 2	0 : - 1 : 0

$\langle 4|2 \rangle \begin{bmatrix} 0 & 1 \end{bmatrix} \quad \mathcal{X}_4 = \{p^1, p^0, p^2\} \equiv \{0, 1, 2\}$
 $\langle 3|2 \rangle \begin{bmatrix} 0 & & & 1 \end{bmatrix} \quad \mathcal{X}_3 = \{q^0 r^0, q^1 r^0, q^2 r^0, q^0 r^1\} \equiv \{0, 1, 2, 3\}$
 $\langle 2|2 \rangle \begin{bmatrix} 0 & & & 0 & 1 \end{bmatrix} \quad \mathcal{X}_2 = \{s^0, s^1, s^2\} \equiv \{0, 1, 2\}$
 $\langle 1|2 \rangle \begin{bmatrix} 0 & & & & 1 \end{bmatrix} \quad \mathcal{X}_1 = \{t^0, t^1, t^2\} \equiv \{0, 1, 2\}$




<i>a</i>	<i>bc</i>	<i>d</i>	<i>e</i>
0 : 1	I	I	0 : 2
0 : 1 1 : 2	0 : - 1 : 3	I	0 : - 1 : -
0 : 1 1 : 2	I	0 : - 1 : 0	I
I	I	0 : 1 1 : 2	0 : - 1 : 0

$\langle 4|2 \rangle \begin{bmatrix} 0 & 1 \end{bmatrix} \quad \mathcal{X}_4 = \{p^1, p^0, p^2\} \equiv \{0, 1, 2\}$
 $\langle 3|2 \rangle \begin{bmatrix} 0 & & & 1 & 3 \end{bmatrix} \quad \mathcal{X}_3 = \{q^0 r^0, q^1 r^0, q^2 r^0, q^0 r^1\} \equiv \{0, 1, 2, 3\}$
 $\langle 2|2 \rangle \begin{bmatrix} 0 & & & 0 & 1 \end{bmatrix} \quad \mathcal{X}_2 = \{s^0, s^1, s^2\} \equiv \{0, 1, 2\}$
 $\langle 1|2 \rangle \begin{bmatrix} 0 & & & & 1 \end{bmatrix} \quad \mathcal{X}_1 = \{t^0, t^1, t^2\} \equiv \{0, 1, 2\}$




<i>a</i>	<i>bc</i>	<i>d</i>	<i>e</i>
0 : 1	I	I	0 : 2
0 : 1 1 : 2 3 : 4	0 : - 1 : 3 3 : 1	I	0 : - 1 : - 3 : 0
0 : 1 1 : 2	I	0 : - 1 : 0	I
I	I	0 : 1 1 : 2	0 : - 1 : 0

$\langle 4|2 \rangle$ $\boxed{0 \ 1}$ $\mathcal{X}_4 = \{p^1, p^0, p^2\} \equiv \{0, 1, 2\}$
 $\langle 3|2 \rangle$ $\boxed{0}$ $\langle 3|3 \rangle$ $\boxed{1 \ 3}$ $\mathcal{X}_3 = \{q^0r^0, q^1r^0, q^2r^0, q^0r^1, q^1r^1\} \equiv \{0, 1, 2, 3, 4\}$
 $\langle 2|2 \rangle$ $\boxed{0}$ $\langle 2|3 \rangle$ $\boxed{0 \ 1}$ $\mathcal{X}_2 = \{s^0, s^1, s^2\} \equiv \{0, 1, 2\}$
 $\langle 1|2 \rangle$ $\boxed{0}$ $\langle 1|3 \rangle$ $\boxed{1}$ $\mathcal{X}_1 = \{t^0, t^1, t^2\} \equiv \{0, 1, 2\}$




<i>a</i>	<i>bc</i>	<i>d</i>	<i>e</i>
0 : 1 1 : -	I	I	0 : 2 1 : 0
0 : 1 1 : 2 3 : 4	0 : - 1 : 3 3 : 1	I	0 : - 1 : - 3 : 0
0 : 1 1 : 2	I	0 : - 1 : 0	I
I	I	0 : 1 1 : 2	0 : - 1 : 0

$\langle 4|2 \rangle$ $\boxed{0 \ 1}$ $\mathcal{X}_4 = \{p^1, p^0, p^2\} \equiv \{0, 1, 2\}$
 $\langle 3|2 \rangle$ $\boxed{0}$ $\langle 3|3 \rangle$ $\boxed{1 \ 3}$ $\mathcal{X}_3 = \{q^0r^0, q^1r^0, q^2r^0, q^0r^1, q^1r^1\} \equiv \{0, 1, 2, 3, 4\}$
 $\langle 2|2 \rangle$ $\boxed{0}$ $\langle 2|3 \rangle$ $\boxed{0 \ 1}$ $\mathcal{X}_2 = \{s^0, s^1, s^2\} \equiv \{0, 1, 2\}$
 $\langle 1|2 \rangle$ $\boxed{0}$ $\langle 1|3 \rangle$ $\boxed{1}$ $\mathcal{X}_1 = \{t^0, t^1, t^2\} \equiv \{0, 1, 2\}$



<i>a</i>	<i>bc</i>	<i>d</i>	<i>e</i>
0 : 1 1 : -	I	I	0 : 2 1 : 0
0 : 1 1 : 2 3 : 4	0 : - 1 : 3 3 : 1	I	0 : - 1 : - 3 : 0
0 : 1 1 : 2	I	0 : - 1 : 0	I
I	I	0 : 1 1 : 2	0 : - 1 : 0

$\langle 4|2 \rangle$ $\boxed{0 \ 1}$ $\mathcal{X}_4 = \{p^1, p^0, p^2\} \equiv \{0, 1, 2\}$
 $\langle 3|2 \rangle$ $\boxed{0}$ $\langle 3|3 \rangle$ $\boxed{1 \ 3}$ $\mathcal{X}_3 = \{q^0r^0, q^1r^0, q^2r^0, q^0r^1, q^1r^1\} \equiv \{0, 1, 2, 3, 4\}$
 $\langle 2|2 \rangle$ $\boxed{0}$ $\langle 2|3 \rangle$ $\boxed{0 \ 1}$ $\mathcal{X}_2 = \{s^0, s^1, s^2\} \equiv \{0, 1, 2\}$
 $\langle 1|2 \rangle$ $\boxed{0}$ $\langle 1|3 \rangle$ $\boxed{1}$ $\mathcal{X}_1 = \{t^0, t^1, t^2\} \equiv \{0, 1, 2\}$



<i>a</i>	<i>bc</i>	<i>d</i>	<i>e</i>
0 : 1 1 : -	I	I	0 : - 1 : 0
0 : 1 1 : - 2 : -	0 : - 1 : 2 2 : 1	I	0 : - 1 : - 2 : 0
0 : 1 1 : -	I	0 : - 1 : 0	I
I	I	0 : 1 1 : -	0 : - 1 : 0

$\langle 4|2 \rangle$ $\boxed{0 \ 1}$ $\mathcal{X}_4 = \{p^1, p^0\} \equiv \{0, 1\}$
 $\langle 3|2 \rangle$ $\boxed{0}$ $\langle 3|3 \rangle$ $\boxed{1 \ 2}$ $\mathcal{X}_3 = \{q^0r^0, q^1r^0, q^0r^1\} \equiv \{0, 1, 2\}$
 $\langle 2|2 \rangle$ $\boxed{0}$ $\langle 2|3 \rangle$ $\boxed{0 \ 1}$ $\mathcal{X}_2 = \{s^0, s^1\} \equiv \{0, 1\}$
 $\langle 1|2 \rangle$ $\boxed{0}$ $\langle 1|3 \rangle$ $\boxed{1}$ $\mathcal{X}_1 = \{t^0, t^1\} \equiv \{0, 1\}$

N	Reachable states	Final memory (KB)			Peak memory (KB)			Time (sec)		
		OTF	PRE	NuSMV	OTF	PRE	NuSMV	OTF	PRE	NuSMV
Dining Philosophers: $L = N/2$, $ \mathcal{X}_k = 34$ for all k										
20	3.46×10^{12}	4	3	4,178	5	4	4,192	0.01	0.01	0.4
50	2.23×10^{31}	11	10	8,847	14	12	8,863	0.03	0.02	13.1
100	4.97×10^{62}	24	20	8,891	28	25	15,256	0.06	0.05	990.8
200	2.47×10^{125}	48	40	21,618	57	50	59,423	0.15	0.11	18,129.3
5,000	6.53×10^{3134}	1,210	1,015	—	1,445	1,269	—	65.55	51.29	—
Slotted Ring Network: $L = N$, $ \mathcal{X}_k = 15$ for all k										
5	5.39×10^4	1	1	502	5	5	507	0.01	0.01	0.1
10	8.29×10^9	5	5	4,332	28	27	8,863	0.06	0.04	6.1
15	1.46×10^{15}	10	9	771	80	77	11,054	0.18	0.13	2,853.1
100	2.60×10^{105}	434	398	—	15,753	14,486	—	41.72	25.78	—
Round Robin Mutual Exclusion: $L = N + 1$, $ \mathcal{X}_k = 10$ for all k except $ \mathcal{X}_1 = N + 1$										
10	2.30×10^4	5	5	917	6	7	932	0.01	0.01	0.2
20	4.72×10^7	18	17	5,980	20	21	5,985	0.04	0.03	1.4
30	7.25×10^{10}	37	36	2,222	41	41	8,716	0.09	0.07	5.6
100	2.85×10^{32}	357	355	13,789	372	372	21,814	2.11	1.55	2,836.5
150	4.82×10^{47}	784	781	—	807	807	—	7.04	5.07	—
FMS: $L = 19$, $ \mathcal{X}_k = N + 1$ for all k except $ \mathcal{X}_{17} = 4$, $ \mathcal{X}_{12} = 3$, $ \mathcal{X}_7 = 2$										
5	1.92×10^4	5	6	2,113	6	9	2,126	0.01	0.01	1.0
10	2.50×10^9	16	19	1,152	26	31	8,928	0.02	0.02	41.6
25	8.54×10^{13}	86	135	17,045	163	239	152,253	0.16	0.11	17,321.9
150	4.84×10^{23}	6,291	15,459	—	16,140	29,998	—	18.50	10.92	—

Traditional approaches apply the global next-state function \mathcal{N} once to each node at each iteration and make extensive use of the unique table and operation caches

- We exhaustively fire each event α in each node p at level $k = Top(\alpha)$, from $k = 1$ up to L
- We must consider redundant nodes as well, thus we prefer quasi-reduced MDDs
- Once node p at level k is saturated, we never fire an event α with $k = Top(\alpha)$ on p again
- The recursive *Fire* calls stop at level $Bot(\alpha)$, although the *Union* calls can go deeper
- Only saturated nodes are placed in the unique table and in the union and firing caches
- Many (most?) nodes we insert in the MDD will still be present in the final MDD
- Firing α in p benefits from having saturated the nodes below p (finds more states)

usually enormous memory and time savings
but Saturation is **not** optimal for all models

Saturation pseudocode (for Kronecker-consistent models)

63

Generate(in s : array[1..L] of lcl): idx

```

1  $p \leftarrow \mathbf{1}$ ;
2 for  $k = 1$  to  $L$  do
3    $r \leftarrow NewNode(k)$ ;
4    $r[s[k]] \leftarrow p$ ;
5   Saturate( $k, r$ );
6   UniqueTableInsert( $k, r$ );
7    $p \leftarrow r$ ;
8   return  $r$ ;
```

Saturate(in k : lvl, p : idx)

```

1 repeat
2    $pCng \leftarrow false$ ;
3   foreach  $\alpha \in \mathcal{E}_k$  do
4      $\mathcal{L} \leftarrow Locals(k, \alpha, p)$ ;
5     while  $\mathcal{L} \neq \emptyset$  do
6        $i \leftarrow Pick(\mathcal{L})$ ;
7        $f \leftarrow RecFire(k-1, \alpha, p[i])$ ;
8       if  $f \neq \mathbf{0}$  then
9         foreach  $j \in \mathcal{N}_{k,\alpha}(i)$  do
10           $u \leftarrow Union(k-1, f, p[j])$ ;
11          if  $u \neq p[j]$  then
12             $p[j] \leftarrow u$ ;
13             $pCng \leftarrow true$ ;
14            if  $\mathcal{N}_{k,\alpha}(j) \neq \emptyset$  then
15               $\mathcal{L} \leftarrow \mathcal{L} \cup \{j\}$ ;
16 until  $pCng = false$ ;
```

RecFire(in h : lvl, α : evt, q : idx): idx

```

1 if  $h < Bot(\alpha)$  then return  $q$ ;
2 if Find(FC[ $h$ ], ( $q, \alpha$ ),  $s$ ) then return  $s$ ;
3  $s \leftarrow NewNode(h)$ ;
4  $sCng \leftarrow false$ ;
5  $\mathcal{L} \leftarrow Locals(h, \alpha, q)$ ;
6 while  $\mathcal{L} \neq \emptyset$  do
7    $i \leftarrow Pick(\mathcal{L})$ ;
8    $f \leftarrow RecFire(h-1, \alpha, q[i])$ ;
9   if  $f \neq \mathbf{0}$  then
10    foreach  $j \in \mathcal{N}_{h,\alpha}(i)$  do
11      $u \leftarrow Union(h-1, f, s[j])$ ;
12     if  $u \neq s[j]$  then
13        $s[j] \leftarrow u$ ;
14        $sCng \leftarrow true$ ;
15 if  $sCng$  then
16   Saturate( $h, s$ );
17 UniqueTableInsert( $h, s$ );
18 Insert(FC[ $h$ ], ( $q, \alpha$ ),  $s$ );
19 return  $s$ ;
```

FC is the firing cache

Saturation pseudocode (for arbitrary models)

64

mdd Saturate(level k , mdd p) is

quasi-reduced version

```

local mdd  $r, r_0, \dots, r_{n_k-1}$ ;
1 if  $k = 0$  then return  $p$ ;
2 if Cache contains entry  $\langle SaturateCODE, p : r \rangle$  then return  $r$ ;
3 foreach  $i_k \in \mathcal{X}_k$  do
4    $r_{i_k} \leftarrow Saturate(k-1, p[i_k])$ ; first, be sure that the children are saturated
5 repeat
6   choose  $\alpha \in \mathcal{E}$ ,  $i_k, j_k \in \mathcal{X}_k$  s.t.  $Top(\alpha) = k$  and  $r_{i_k} \neq \mathbf{0}$  and  $\mathcal{N}_\alpha[i_k][j_k] \neq \mathbf{0}$ ;
7    $r_{j_k} \leftarrow Union(r_{j_k}, RelProdSat(k-1, r_{i_k}, \mathcal{N}_\alpha[i_k][j_k]))$ ;
8 until  $r_0, \dots, r_{n_k-1}$  do not change;
9  $r \leftarrow UniqueTableInsert(k, r_0, \dots, r_{n_k-1})$ ;
10 enter  $\langle SaturateCODE, p : r \rangle$  in Cache;
11 return  $r$ ;
```


mdd RelProdSat(level k , mdd q , mdd2 f) is

```

local mdd  $r, r_0, \dots, r_{n_k-1}$ ;
1 if  $k = 0$  then return  $q \wedge f$ ;
2 if Cache contains entry  $\langle RelProdSatCODE, q, f : r \rangle$  then return  $r$ ;
3 foreach  $i_k, j_k \in \mathcal{X}_k$  s.t.  $q[i_k] \neq \mathbf{0}$  and  $f[i_k][j_k] \neq \mathbf{0}$  do
4    $r_{j_k} \leftarrow Union(r_{j_k}, RelProdSat(k-1, q[i_k], f[i_k][j_k]))$ ;
5  $r \leftarrow Saturate(k, UniqueTableInsert(k, r_0, \dots, r_{n_k-1}))$ ;
6 enter  $\langle RelProdSatCODE, q, f : r \rangle$  in Cache;
7 return  $r$ .
```



CTL model checking of Petri nets

Given $(\hat{\mathcal{X}}, \mathcal{X}_{init}, \mathcal{N})$, we can unwind the graph into a computation tree rooted at each $i \in \mathcal{X}_{init}$



A Kripke structure is specified by $(\hat{\mathcal{X}}, \mathcal{X}_{init}, \mathcal{N}, \mathcal{A}, \mathcal{L})$

- \mathcal{A} is a set of atomic properties
- $\mathcal{L} : \hat{\mathcal{X}} \rightarrow 2^{\mathcal{A}}$ is a labeling function listing the atomic properties that hold in each state



CTL: computation tree logic

Given a Kripke structure $(\hat{\mathcal{X}}, \mathcal{X}_{init}, \mathcal{N}, \mathcal{A}, \mathcal{L})$

CTL has state formulas and path formulas

- State formulas:
 - if $a \in \mathcal{A}$, a is a state formula (a is an atomic proposition, true or false in each state)
 - if p and p' are state formulas, $\neg p$, $p \vee p'$, $p \wedge p'$ are state formulas
 - if q is a path formula, $E q$, $A q$ are state formulas
- Path formulas:
 - if p and p' are state formulas, $X p$, $F p$, $G p$, $p U p'$, $p R p'$ are path formulas
 - Note: unlike CTL*, a state formula is **not** also a path formula

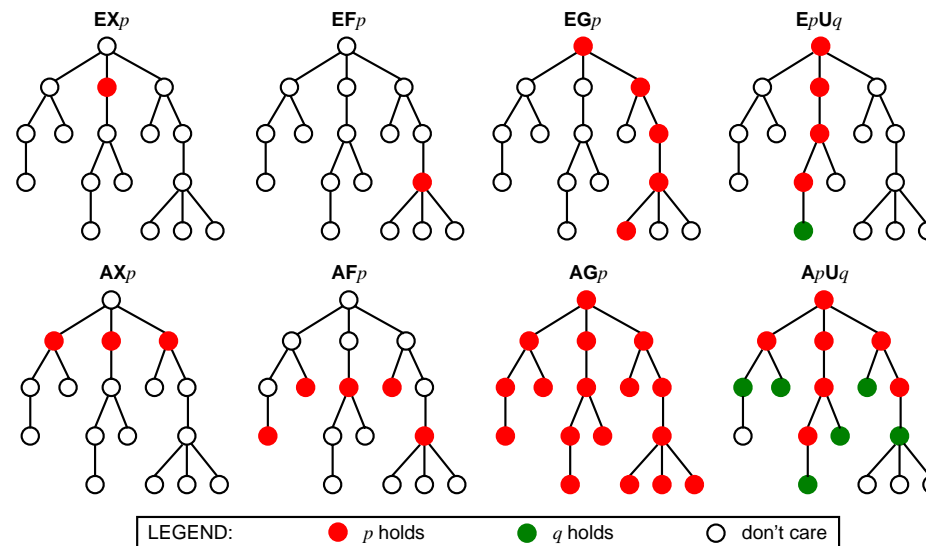
In CTL, operators occur in pairs:

- a path quantifier, E or A, must always immediately precede a temporal operator, X, F, G, U, R

CTL expressions can be nested: $p \vee E \neg p U (\neg p \wedge A X p)$

A CTL formula p identifies a set of model states (those satisfying p)

CTL semantics



EX, EU, and EG form a complete set of CTL operators, since:

- $AXp = \neg EX \neg p$
- $AFP = \neg EGF \neg p$
- $AGp = \neg EFG \neg p$
- $ApRq = \neg A \neg p U \neg q$
- $EFp = E true U p$
- $ApUq = \neg (E \neg q U \neg p \wedge \neg q) \wedge \neg EG \neg q$
- $EPRq = \neg A \neg p U q$
- $AFp = \neg EG \neg p$
- $ApUq = \neg (E \neg q U \neg p \wedge \neg q) \wedge \neg EG \neg q$
- $APRq = \neg E \neg p U q$

The system always reaches a stable state and remains in stable states after an initial startup period

- $initial \Rightarrow AF (AG \text{ stable})$ or $initial \Rightarrow AF (A \text{ stable} \cup \text{shutdown})$

At any point in the execution, it is possible to return to a reset state

- $AG \text{ EF reset}$

If a process asks access to the critical region, it eventually obtains it

- $AG \text{ request_critical} \Rightarrow AF \text{ access_critical}$

An algorithm to label all states that satisfy Exp

We assume that all states satisfying p have been correctly labeled already

```

LabelEX(p) is
1  $\mathcal{Y} \leftarrow \{\mathbf{i} \in \mathcal{X}_{reach} : p \in labels(\mathbf{i})\};$  initialize  $\mathcal{Y}$  with the states satisfying  $p$ 
2 while  $\mathcal{Y} \neq \emptyset$  do
3   pick and remove a state  $\mathbf{j}$  from  $\mathcal{Y}$ ;
4   for each  $\mathbf{i} \in \mathcal{N}^{-1}(\mathbf{j})$  do state  $\mathbf{i}$  can transition to state  $\mathbf{j}$ 
5      $labels(\mathbf{i}) \leftarrow labels(\mathbf{i}) \cup \{Exp\};$ 

```

If we have $\mathcal{X}_{reach} \subseteq \hat{\mathcal{X}}$, we can assume $\mathcal{N} : \mathcal{X}_{reach} \rightarrow 2^{\mathcal{X}_{reach}}$ instead of $\mathcal{N} : \hat{\mathcal{X}} \rightarrow 2^{\hat{\mathcal{X}}}$

An algorithm to label all states that satisfy $EpUq$

We assume that all states satisfying p and all states satisfying q have been correctly labeled already

```

LabelEU(p, q) is
1  $\mathcal{Y} \leftarrow \{\mathbf{i} \in \mathcal{X}_{reach} : q \in labels(\mathbf{i})\};$  initialize  $\mathcal{Y}$  with the states satisfying  $q$ 
2 for each  $\mathbf{i} \in \mathcal{Y}$  do
3    $labels(\mathbf{i}) \leftarrow labels(\mathbf{i}) \cup \{EpUq\};$ 
4 while  $\mathcal{Y} \neq \emptyset$  do
5   pick and remove a state  $\mathbf{j}$  from  $\mathcal{Y}$ ;
6   for each  $\mathbf{i} \in \mathcal{N}^{-1}(\mathbf{j})$  do state  $\mathbf{i}$  can transition to state  $\mathbf{j}$ 
7     if  $EpUq \notin labels(\mathbf{i})$  and  $p \in labels(\mathbf{i})$  then
8        $labels(\mathbf{i}) \leftarrow labels(\mathbf{i}) \cup \{EpUq\};$ 
9        $\mathcal{Y} \leftarrow \mathcal{Y} \cup \{\mathbf{i}\};$ 

```

An algorithm to label all states that satisfy EGp

We assume that all states satisfying p have been correctly labeled already

The algorithm relies on finding the (nontrivial) strongly connected components (SCCs) of a graph

```

LabelEG(p) is
1  $\mathcal{Y} \leftarrow \{\mathbf{i} \in \mathcal{X}_{reach} : p \in labels(\mathbf{i})\};$  initialize  $\mathcal{Y}$  with the states satisfying  $p$ 
2 build the set  $\mathcal{C}$  of SCCs in the subgraph of  $\mathcal{N}$  induced by  $\mathcal{Y}$ ;
3  $\mathcal{W} \leftarrow \{\mathbf{i} : \mathbf{i} \text{ is a state in a SCC of } \mathcal{C}\};$ 
4 for each  $\mathbf{i} \in \mathcal{W}$  do
5    $labels(\mathbf{i}) \leftarrow labels(\mathbf{i}) \cup \{EGp\};$ 
6 while  $\mathcal{W} \neq \emptyset$  do
7   pick and remove a state  $\mathbf{j}$  from  $\mathcal{W}$ ;
8   for each  $\mathbf{i} \in \mathcal{N}^{-1}(\mathbf{j})$  do state  $\mathbf{i}$  can transition to state  $\mathbf{j}$ 
9     if  $EGp \notin labels(\mathbf{i})$  and  $p \in labels(\mathbf{i})$  then
10       $labels(\mathbf{i}) \leftarrow labels(\mathbf{i}) \cup \{EGp\};$ 
11       $\mathcal{W} \leftarrow \mathcal{W} \cup \{\mathbf{i}\};$ 

```

All sets of states and relations over sets of states are encoded using DDs

An algorithm to build the DD encoding the set of states that satisfy EXp

Assume that the DD encoding the set \mathcal{P} of states satisfying p has been built already

BuildEXsymbolic(\mathcal{P}) is

```

1 return RelationalProduct( $\mathcal{P}, \mathcal{N}^{-1}$ );           perform one backward step in the transition relation
    
```

Where

- \mathcal{N}^{-1} is the inverse or backwards transition relation:

$$\mathbf{i} \in \mathcal{N}^{-1}(\mathbf{j}) \Leftrightarrow \mathbf{j} \in \mathcal{N}(\mathbf{i})$$

- given a relation $\mathcal{R} : \mathcal{A} \rightarrow 2^{\mathcal{B}}$ and a set $\mathcal{Y} \subseteq \mathcal{A}$:

$$\text{RelationalProduct}(\mathcal{Y}, \mathcal{R}) = \mathcal{R}(\mathcal{Y}) = \bigcup_{\mathbf{i} \in \mathcal{Y}} \mathcal{R}(\mathbf{i}) \subseteq \mathcal{B}$$

Two algorithms to build the DD encoding the set of states that satisfy $EpUq$

Assume that the DDs encoding the sets \mathcal{P} and \mathcal{Q} of states satisfying p and q have been built already

BuildEUsymbolic(\mathcal{P}, \mathcal{Q}) is

```

1  $\mathcal{Y} \leftarrow \emptyset$ ;
2  $\mathcal{U} \leftarrow \mathcal{Q}$ ;                               initialize the unexplored set  $\mathcal{U}$  with the states satisfying  $q$ 
3 repeat
4    $\mathcal{Y} \leftarrow \text{Union}(\mathcal{Y}, \mathcal{U})$ ;           currently known states satisfying  $EpUq$ 
5    $\mathcal{W} \leftarrow \text{RelationalProduct}(\mathcal{U}, \mathcal{N}^{-1})$ ;   perform one backward step in the transition relation
6    $\mathcal{Z} \leftarrow \text{Intersection}(\mathcal{W}, \mathcal{P})$ ;         discard the states that do not satisfy  $p$ 
7    $\mathcal{U} \leftarrow \text{Difference}(\mathcal{Z}, \mathcal{Y})$ ;         discard the states that are not new
8 until  $\mathcal{U} = \emptyset$ ;
9 return  $\mathcal{Y}$ ;
    
```

BuildEUsymbolicAll(\mathcal{P}, \mathcal{Q}) is

```

1  $\mathcal{Y} \leftarrow \mathcal{Q}$ ;                               initialize the currently known result with the states satisfying  $q$ 
2 repeat
3    $\mathcal{O} \leftarrow \mathcal{Y}$ ;                               save the old set of states
4    $\mathcal{W} \leftarrow \text{RelationalProduct}(\mathcal{Y}, \mathcal{N}^{-1})$ ;   perform one backward step in the transition relation
5    $\mathcal{Z} \leftarrow \text{Intersection}(\mathcal{W}, \mathcal{P})$ ;         discard the states that do not satisfy  $p$ 
6    $\mathcal{Y} \leftarrow \text{Union}(\mathcal{Z}, \mathcal{Y})$ ;           add to the currently known result
7 until  $\mathcal{O} = \mathcal{Y}$ ;
8 return  $\mathcal{Y}$ ;
    
```

An algorithm to build the DD encoding the set of states that satisfy EGp

Assume that the DDs encoding the set \mathcal{P} of states satisfying p has been built already

BuildEGsymbolic(\mathcal{P}) is

```

1  $\mathcal{Y} \leftarrow \mathcal{P}$ ;                               initialize  $\mathcal{Y}$  with the states satisfying  $p$ 
2 repeat
3    $\mathcal{O} \leftarrow \mathcal{Y}$ ;                               save the old set of states
4    $\mathcal{W} \leftarrow \text{RelationalProduct}(\mathcal{Y}, \mathcal{N}^{-1})$ ;   perform one backward step in the transition relation
5    $\mathcal{Y} \leftarrow \text{Intersection}(\mathcal{Y}, \mathcal{W})$ ;
6 until  $\mathcal{O} = \mathcal{Y}$ ;
7 return  $\mathcal{Y}$ ;
    
```

This algorithm starts with a larger set of states and reduces it

This algorithm is not based on finding the strongly connected components of \mathcal{N}

Traditional symbolic CTL model checking (EF, EU, EG) uses a breadth-first fixed-point iteration

Just like for state-space generation, breadth-first can require huge peak memory, hence runtime

Using the model structure results in better algorithms for symbolic CTL model checking:

- exploit locality
- employ a Saturation-based algorithm for EF and EU
- greatly reduced memory and time requirements for asynchronous systems
- implemented in our tool [SMART](#)
 - can we apply Saturation to EG?
 - can we extend this to fair CTL?

Substantial time and memory improvements for EX and EG (thanks to locality)

Enormous time and memory improvements for EF and EU (thanks to locality and saturation)

S (depends on parameter N)	NuSMV				SMART				NuSMV				SMART			
	after SS		alone		EUsat		EUsat		after SS		alone		EGtrud			
	sec	kB	sec	kB	iter	sec	kB	iter	sec	kB	sec	kB	sec	kB		
Phils $E[(phil_1 \neq eat) \cup (phil_0 = eat)]$ $EG(phil_0 \neq eat)$ starvation																
2.23×10^{31}	1.2	46	39.7	46	100	0.17	1	4	0.06	1	0.9	46	132.3	50	0.02	1
4.96×10^{62}	7.9	316	1121.8	316	200	0.67	3	4	0.14	3	9.0	316	2525.3	358	0.05	3
3.03×10^{313}	—	—	—	—	1000	19.09	78	4	0.77	60	—	—	—	—	0.28	58
FMS $E[(M_1 > 0) \cup (P_1s = P_2s = P_3s = N)]$ $EG\neg(P_1s = P_2s = P_3s = N)$																
3.44×10^3	0.2	17	318.1	43	31	0.04	<.5	6	0.01	<.5	0.2	17	128.9	18	<.005	<.5
4.86×10^4	1.0	127	—	—	46	0.16	<.5	8	0.02	<.5	1.0	127	—	—	0.01	<.5
8.54×10^{13}	—	—	—	—	376	—	—	52	1010.85	293	—	—	—	—	50.38	251
Round robin $E[(p_1 \neq load) \cup (p_0 = send)]$ $EG(true)$ find all cycles																
2.30×10^5	0.2	11	85.0	11	39	0.01	<.5	11	0.01	<.5	0.3	11	78.5	13	<.005	<.5
1.10×10^6	0.6	40	4922.7	40	59	0.03	<.5	16	0.01	<.5	1.2	40	4739.5	44	0.01	<.5
2.85×10^{32}	—	—	—	—	399	13.32	32	101	4.67	19	—	—	—	—	1.29	20
Leader $E[(pref_1 = 0) \cup (status_0 = leader)]$ $EG(status_0 \neq leader)$																
1.15×10^4	2.3	11	8104.7	371	62	0.36	1	38	0.27	1	232.8	12	1189.1	235	0.11	2
1.50×10^5	52.0	33	—	—	81	3.74	7	52	3.09	7	18023.6	104	—	—	0.44	9
2.39×10^7	—	—	—	—	121	690.85	116	85	416.85	101	—	—	—	—	7.15	128
Slotted ring $E[(slot_1 \neq bf) \cup (slot_0 = ag)]$ $EG(slot_0 \neq hg)$																
8.29×10^9	0.2	10	0.4	3	63	0.01	<.5	9	0.01	<.5	0.6	10	0.1	1	0.01	<.5
1.46×10^{15}	1.8	15	2.0	10	93	0.37	1	9	0.02	<.5	4.7	15	0.2	2	0.01	<.5
3.03×10^{105}	—	—	—	—	603	—	—	9	1.60	62	—	—	—	—	0.62	62

Decision diagrams for integer-valued functions

Ordered multiterminal multiway decision diagrams (MTMDDs) 79

Assume a domain $\hat{\mathcal{X}} = \mathcal{X}_L \times \dots \times \mathcal{X}_1$, where $\mathcal{X}_k = \{0, 1, \dots, n_k - 1\}$, for some $n_k \in \mathbb{N}$

Assume a range $\mathcal{X}_0 = \{0, 1, \dots, n_0 - 1\}$, for some $n_0 \in \mathbb{N}$ (any set \mathcal{X}_0 will actually do)

An MTMDD is an acyclic directed edge-labeled graph where:

- The only terminal nodes are values from \mathcal{X}_0 and are at level 0 $\forall i_0 \in \mathcal{X}_0, i_0.lvl = 0$
- A nonterminal node p is at a level k , with $L \geq k \geq 1$ $p.lvl = k$
- A nonterminal node p at level k has n_k outgoing edges pointing to children $p[i_k]$, for $i_k \in \mathcal{X}_k$
- The level of the children is lower than that of p ; $p[0].lvl < p.lvl, p[1].lvl < p.lvl$
- A node p at level k encodes the function $v_p : \hat{\mathcal{X}} \rightarrow \mathcal{X}_0$ defined recursively by

$$v_p(x_L, \dots, x_1) = \begin{cases} p & \text{if } k = 0 \\ v_{p[x_k]}(x_L, \dots, x_1) & \text{if } k > 0 \end{cases}$$

Instead of levels, we can also talk of variables:

- The terminal nodes are associated with the range variable x_0
- A nonterminal node is associated with a domain variable x_k , with $L \geq k \geq 1$

Canonical versions of MTMDDs 80

For canonical MTMDDs, we further require that


- There are no duplicates: if $p.lvl = q.lvl = k$ and $p[i_k] = q[i_k]$ for all $i_k \in \mathcal{X}_k$, then $p = q$

Then, if the MTMDD is quasi-reduced, there is no level skipping:

- The only root nodes with no incoming arcs are at level L
- Each child $p[i_k]$ of a node p is at level $p.lvl - 1$

Or, if the MTMDD is fully-reduced, there is maximum level skipping:

- There are no redundant nodes p satisfying $p[i_k] = q$ for all $i_k \in \mathcal{X}_k$



$\mathcal{X}_4 = \{0, 1, 2, 3\}$

$\mathcal{X}_3 = \{0, 1, 2\}$

$\mathcal{X}_2 = \{0, 1\}$

$\mathcal{X}_1 = \{0, 1, 2\}$

These MTMDDs encode a function $\hat{\mathcal{X}} \rightarrow \mathbb{N}$ (or \mathbb{Z} , or \mathbb{R} , or any arbitrary set)

Given a model $(\hat{\mathcal{X}}, \mathcal{X}_{init}, \mathcal{N})$, we can define the distance function $\delta : \hat{\mathcal{X}} \rightarrow \mathbb{N} \cup \{\infty\}$

$\delta(\mathbf{i}) = \min\{d : \mathbf{i} \in \mathcal{N}^d(\mathcal{X}_{init})\}$ thus $\delta(\mathbf{i}) = \infty \Leftrightarrow \mathbf{i} \notin \mathcal{X}_{reach}$

Build $\mathcal{X}^{[d]} = \{\mathbf{i} : \delta(\mathbf{i}) = d\}$,
for $d = 0, 1, \dots, d_{max}$


Build $\mathcal{Y}^{[d]} = \{\mathbf{i} : \delta(\mathbf{i}) \leq d\}$,
for $d = 0, 1, \dots, d_{max}$

```
DistanceMddForestEQ( $\mathcal{X}_{init}, \mathcal{N}$ ) is
1  $d \leftarrow 0$ ;
2  $\mathcal{X}_{reach} \leftarrow \mathcal{X}_{init}$ ;
3  $\mathcal{X}^{[0]} \leftarrow \mathcal{X}_{init}$ ;
4 repeat
5    $\mathcal{X}^{[d+1]} \leftarrow \mathcal{N}(\mathcal{X}^{[d]}) \setminus \mathcal{X}_{reach}$ ;
6    $d \leftarrow d + 1$ ;
7    $\mathcal{X}_{reach} \leftarrow \mathcal{X}_{reach} \cup \mathcal{X}^{[d]}$ ;
8 until  $\mathcal{X}^{[d]} = \emptyset$ ;
```

```
DistanceMddForestLE( $\mathcal{X}_{init}, \mathcal{N}$ ) is
1  $d \leftarrow 0$ ;
2  $\mathcal{Y}^{[0]} \leftarrow \mathcal{X}_{init}$ ;
3 repeat
4    $\mathcal{Y}^{[d+1]} \leftarrow \mathcal{N}(\mathcal{Y}^{[d]}) \cup \mathcal{Y}^{[d]}$ ;
5    $d \leftarrow d + 1$ ;
6 until  $\mathcal{Y}^{[d]} = \mathcal{Y}^{[d-1]}$ ;
```

This is breadth-first symbolic state space generation

$\mathcal{X}_{reach} = \{\mathbf{i} \in \hat{\mathcal{X}} : \delta(\mathbf{i}) < \infty\} = \bigcup_{d=0}^{d_{max}} \mathcal{X}^{[d]} = \mathcal{Y}^{[d_{max}]}$ is a by-product of this process!



With an MDD forest: node merging can be poor at the top

With an MTMDD: node merging can be poor at the bottom

Both approaches are explicit in the number of distinct distance values

Assume a domain $\hat{\mathcal{X}} = \mathcal{X}_L \times \dots \times \mathcal{X}_1$, where $\mathcal{X}_k = \{0, 1, \dots, n_k - 1\}$, for some $n_k \in \mathbb{N}$

Assume the range \mathbb{Z} and the combinator "+" (addition over the integers)

An EVMDD is an acyclic directed edge-labeled graph where:

- The only terminal node is Ω and is at level 0 $\Omega.lvl = 0$
- A nonterminal node p is at a level k , with $L \geq k \geq 1$ $p.lvl = k$
- A nonterminal node p at level k has n_k outgoing edges
- For $i_k \in \mathcal{X}_k$, edge $p[i_k]$ points to child $p[i_k].child$, and has value $p[i_k].val \in \mathbb{Z}$
- The level of the children is lower than that of p $p[i_k].child.lvl < p.lvl$
- An edge $\langle \sigma, p \rangle$, with $p.lvl = k$ encodes the function $v_{\langle \sigma, p \rangle} : \hat{\mathcal{X}} \rightarrow \mathbb{Z}$ defined recursively by

$$v_{\langle \sigma, p \rangle}(x_L, \dots, x_1) = \begin{cases} \sigma & \text{if } k = 0, \text{ i.e., } p = \Omega \\ \sigma + v_{p[x_k]}(x_L, \dots, x_1) & \text{if } k > 0, \text{ i.e., } p \neq \Omega \end{cases}$$

For canonical EVMDDs, we first normalize each node p at level $k \geq 1$ in one of two ways:

- $p[0].val = 0$, or EVMDDs
- $p[i_k].val \geq 0$ for all $i_k \in \mathcal{X}_k$, and $p[j_k].val = 0$ for at least one $j_k \in \mathcal{X}_k$ EV⁺MDDs

Then, the usual reduction requirements apply:

- There are no **duplicates**: if $p.lvl = q.lvl = k$ and $p[i_k] = q[i_k]$ for all $i_k \in \mathcal{X}_k$, then $p = q$

And, if the MDD is **quasi-reduced**, there is no level skipping:

- The only **root** nodes with no incoming arcs are at level L , and have **root edge values** in \mathbb{Z}
- Each child $p[i_k].child$ of a node p is at level $p.lvl - 1$

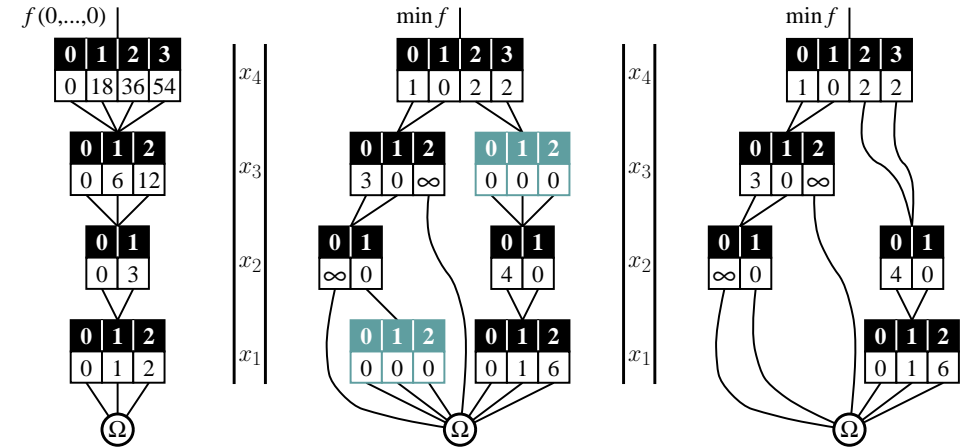
Or, if the MDD is **fully-reduced**, there is maximum level skipping:

- There are no **redundant** nodes p satisfying $p[i_k].child = q$ and $p[i_k].val = 0$ for all $i_k \in \mathcal{X}_k$


For EVMDDs, the value of the incoming root edge is $f(0, \dots, 0)$

For EV⁺MDDs, the value of the incoming root edge is $\min f$

The EV⁺MDDs normalization allows to store **partial functions** $\hat{\mathcal{X}} \rightarrow \mathbb{Z} \cup \{\infty\}$



Example: the *Minimum* operator for quasi-reduced EV⁺MDDs 87



The *Minimum* operator for quasi-reduced EV⁺MDDs 88

edge $Minimum(\text{level } k, \text{edge } \langle \alpha, p \rangle, \text{edge } \langle \beta, q \rangle)$ edge is a pair $\langle \text{int}, \text{node} \rangle$

local node p', q', r ;
 local int μ, α', β' ;
 local local i_k ;


- 1 if $\alpha = \infty$ then return $\langle \beta, q \rangle$;
- 2 if $\beta = \infty$ then return $\langle \alpha, p \rangle$;
- 3 $\mu \leftarrow \min\{\alpha, \beta\}$;
- 4 if $k = 0$ then return $\langle \mu, \Omega \rangle$;
- 5 if *Cache* contains entry $\langle MinimumCODE, k, p, q, \alpha - \beta : \gamma, r \rangle$ then return $\langle \gamma + \mu, r \rangle$;
- 6 $r \leftarrow NewNode(k)$;
- 7 foreach $i_k \in \mathcal{X}_k$ do the only node at level 0 is Ω
- 8 $p' \leftarrow p.child[i_k]$;
- 9 $\alpha' \leftarrow \alpha - \mu + p.val[i_k]$;
- 10 $q' \leftarrow q.child[i_k]$;
- 11 $\beta' \leftarrow \beta - \mu + q.val[i_k]$;
- 12 $r[i_k] \leftarrow Minimum(k-1, \langle \alpha', p' \rangle, \langle \beta', q' \rangle)$;
- 13 *UniqueTableInsert*(k, r); continue downstream
- 14 enter $\langle MinimumCODE, k, p, q, \alpha - \beta : \mu, r \rangle$ in *Cache*;
- 15 return $\langle \mu, r \rangle$;

The distance function δ is the fixed-point of the iteration $\delta^{[m+1]} = \Phi(\delta^{[m]})$ where

$$\delta^{[m+1]}(\mathbf{i}) = \min \left(\delta^{[m]}(\mathbf{i}), \min \left\{ 1 + \delta^{[m]}(\mathbf{j}) \mid \exists \alpha \in \mathcal{E} : \mathbf{i} \in \mathcal{T}_\alpha(\mathbf{j}) \right\} \right)$$

initialized with

$$\delta^{[0]}(\mathbf{i}) = \begin{cases} 0 & \text{if } \mathbf{i} \in \mathcal{X}_{init} \\ \infty & \text{otherwise} \end{cases}$$




- At each iteration, we monotonically reduce the value of $\delta^{[m]}(\mathbf{i})$ for at least one state \mathbf{i}
- The traditional breadth-first iteration does this in an all-or-nothing fashion:

$$\delta^{[m]}(\mathbf{i}) = \infty \text{ for } m < d, \text{ the actual distance } d \text{ of } \mathbf{i}, \text{ then } \delta^{[m]}(\mathbf{i}) = d \text{ for } m \geq d$$

- The Saturation approach may instead reduce $\delta^{[m]}(\mathbf{i})$ multiple times, until it reaches d

effect of event $e : \mathbf{i}_3 \xrightarrow{e} \mathbf{j}_3$ (synchronously)
 $\mathbf{i}_2 \xrightarrow{e} \mathbf{j}_2$



$$\text{where } (\beta, u) = \text{Min}((6, z), (3, w))$$

Results: time and memory to generate and store δ

N	S	Time (in seconds)					Final nodes			Peak nodes				
		E_s	E_b	M_b	T_s	T_b	$E_s E_b$	$M_b T_s T_b$	E_s	E_b	M_b	T_s	T_b	
Dining philosophers: $d_{max} = 2N, L = N/2, \mathcal{X}_k = 34$ for all k														
10	1.9×10^6	0.01	0.06	0.05	0.12	0.46	21	255	170	21	605	644	238	4022
30	6.4×10^{18}	0.02	0.86	0.70	7.39	56.80	71	2545	1710	71	7225	7364	2788	140262
1000	9.2×10^{626}	0.48	—	—	—	—	2496	—	—	2496	—	—	—	—
Kanban system: $d_{max} = 14N, L = 4, \mathcal{X}_k = (N+3)(N+2)(N+1)/6$ for all k														
5	2.5×10^6	0.02	0.14	0.12	0.24	1.55	9	444	133	57	1132	1156	776	13241
12	5.5×10^9	0.34	4.34	3.45	11.08	129.46	16	2368	518	218	5633	5805	5585	165938
50	1.0×10^{16}	179.48	—	—	—	—	58	—	—	2802	—	—	—	—
Flex. manuf. syst.: $d_{max} = 14N, L = 19, \mathcal{X}_k = N+1$ for all k except $ \mathcal{X}_{17} = 4, \mathcal{X}_{12} = 3, \mathcal{X}_2 = 2$														
5	2.9×10^6	0.01	0.42	0.34	0.88	11.78	149	5640	2989	211	15205	15693	4903	179577
10	2.5×10^9	0.04	2.96	2.40	5.79	608.92	354	28225	11894	536	76676	78649	17885	1681625
140	2.0×10^{23}	20.03	—	—	—	—	32012	—	—	52864	—	—	—	—
Round-robin mutex protocol: $d_{max} = 8N - 6, L = N + 1, \mathcal{X}_k = 10$ for all k except $ \mathcal{X}_1 = N + 1$														
10	2.3×10^4	0.01	0.06	0.05	0.22	0.50	92	1038	1123	107	1898	1948	1210	9245
30	7.2×10^{10}	0.05	0.95	0.89	16.04	224.83	582	12798	19495	637	24122	24566	20072	376609
200	7.2×10^{62}	1.63	—	—	—	—	20897	—	—	21292	—	—	—	—


E_s : EV⁺MDD & Saturation E_b : EV⁺MDD & breadth-first M_b : multiple MDDs & breadth-first
 T_s : MTMDD & Saturation T_b : MTMDD & breadth-first

Generating an EF trace using EV⁺MDDs

INPUT: the MDD x encoding a set of states \mathcal{X} , the EV⁺MDD $\langle \rho, r \rangle$ encoding δ

OUTPUT: a (minimum) μ -length trace $\mathbf{j}^{[0]}, \dots, \mathbf{j}^{[\mu]}$ from a state in \mathcal{X}_{init} to a state in \mathcal{X}

1. Build the EV⁺MDD $\langle 0, x \rangle$ encoding $\delta_x(\mathbf{i}) = 0$ if $\mathbf{i} \in \mathcal{X}$ and $\delta_x(\mathbf{i}) = \infty$ if $\mathbf{i} \in \hat{\mathcal{X}} \setminus \mathcal{X}$
2. Compute the EV⁺MDD $\langle \mu, m \rangle$ encoding $\text{Max}(\langle \rho, r, \cdot \rangle, \langle 0, x \rangle)$
 μ is the length of one of the shortest-paths we are seeking
3. If $\mu = \infty$, exit: \mathcal{X} does not contain reachable states
4. Otherwise, extract from $\langle \mu, m \rangle$ a state $\mathbf{j}^{[\mu]} = (j_L^{[\mu]}, \dots, j_1^{[\mu]})$ on a 0-labelled path from m to Ω
 $\mathbf{j}^{[\mu]}$ is a reachable state in \mathcal{X} at the desired minimum distance μ from \mathcal{X}_{init}
5. Initialize ν to μ and iterate until $\nu = 0$:
 - (a) For each state $\mathbf{i} \in \hat{\mathcal{X}}$ such that $\mathbf{j}^{[\nu]} \in \mathcal{N}(\mathbf{i})$ (use the backward function \mathcal{N}^{-1})
 - compute $\delta(\mathbf{i})$ using $\langle \rho, r \rangle$ and stop on the first \mathbf{i} such that $\delta(\mathbf{i}) = \nu - 1$
there exists at least one such state \mathbf{i}^*
 - (b) Decrement ν
 - (c) Let $\mathbf{j}^{[\nu]}$ be \mathbf{i}^*



$$\mathcal{Y} = \left\{ \begin{array}{cccccccccccccccccccc} 0 & 1 & 1 & 1 & 1 & 1 & 2 & 2 & 2 & 2 & 2 & 3 & 3 & 3 & 3 & 3 & 3 & 3 & 3 \\ 2 & 0 & 0 & 1 & 1 & 1 & 2 & 0 & 0 & 1 & 1 & 2 & 0 & 1 & 2 & 2 & 2 & 2 & 2 \\ 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 2 & 0 & 1 & 2 \end{array} \right\}$$

To compute the index of a state, use edge values:


- Sum the values found on the corresponding path:

$$\psi(2, 1, 1, 0) = 0 + 6 + 2 + 1 + 0 = 9$$

- A state is unreachable if the path is not complete:

$$\psi(0, 2, 0, 0) = 0 + 0 + 0 + \infty = \infty$$

(a missing edge has the default value of ∞)



Continuous-time Markov chains

A stochastic process $\{X(t) : t \geq 0\}$ is a collection of r.v.'s indexed by a time parameter t

We say that $X(t)$ is the state of the process at time t

The possible values $X(t)$ can ever assume for any t is (a subset of) the state space \mathcal{X}_{reach}

$\{X(t) : t \geq 0\}$ over a discrete \mathcal{X}_{reach} is a continuous-time Markov chain (CTMC) if

$$\Pr \{X(t^{[n+1]}) = \mathbf{i}^{[n+1]} \mid X(t^{[n]}) = \mathbf{i}^{[n]} \wedge X(t^{[n-1]}) = \mathbf{i}^{[n-1]} \wedge \dots \wedge X(t^{[0]}) = \mathbf{i}^{[0]}\}$$

$$= \Pr \{X(t^{[n+1]}) = \mathbf{i}^{[n+1]} \mid X(t^{[n]}) = \mathbf{i}^{[n]}\}$$

for any $0 \leq t^{[0]} \leq \dots \leq t^{[n-1]} \leq t^{[n]} \leq t^{[n+1]}$ and $\mathbf{i}^{[0]}, \dots, \mathbf{i}^{[n-1]}, \mathbf{i}^{[n]}, \mathbf{i}^{[n+1]} \in \mathcal{X}_{reach}$

Markov property:

“given the present state, the future is independent of the past”

“the most recent knowledge about the state is all we need”

Numerical analysis of stochastic Petri nets

Markov chain description and analysis

A continuous-time Markov chain (CTMC) $\{X(t) : t \geq 0\}$ with state space \mathcal{X}_{reach} is described by

- its infinitesimal generator $\mathbf{Q} = \mathbf{R} - \text{diag}(\mathbf{R} \cdot \mathbf{1}) = \mathbf{R} - \text{diag}(\mathbf{h})^{-1} \in \mathbb{R}^{|\mathcal{X}_{reach}| \times |\mathcal{X}_{reach}|}$
- its initial probability vector $\boldsymbol{\pi}(0) \in \mathbb{R}^{|\mathcal{X}_{reach}|}$

where

- \mathbf{R} is the transition rate matrix: $\mathbf{R}[\mathbf{i}, \mathbf{j}]$ is the rate of going to state \mathbf{j} when in state \mathbf{i}
- \mathbf{h} is the expected holding time vector: $\mathbf{h}[\mathbf{i}] = 1 / \sum_{\mathbf{j} \in \mathcal{X}_{reach}} \mathbf{R}[\mathbf{i}, \mathbf{j}]$
- $\boldsymbol{\pi}(0)[\mathbf{i}] = \Pr \{\text{chain is in state } \mathbf{i} \text{ at time } 0, \text{ i.e., initially}\}$

Transient probability vector $\boldsymbol{\pi}(t) \in \mathbb{R}^{|\mathcal{X}_{reach}|}$: $\boldsymbol{\pi}(t)[\mathbf{i}] = \Pr \{X(t) = \mathbf{i}\}$

- $\boldsymbol{\pi}(t)$ is the solution of $\frac{d\boldsymbol{\pi}(t)}{dt} = \boldsymbol{\pi}(t) \cdot \mathbf{Q}$ with initial condition $\boldsymbol{\pi}(0)$

Steady-state probability vector $\boldsymbol{\pi} \in \mathbb{R}^{|\mathcal{X}_{reach}|}$: $\boldsymbol{\pi}[\mathbf{i}] = \lim_{t \rightarrow \infty} \Pr \{X(t) = \mathbf{i}\}$

- $\boldsymbol{\pi}$ is the solution of $\boldsymbol{\pi} \cdot \mathbf{Q} = \mathbf{0}$ subject to $\sum_{\mathbf{i} \in \mathcal{X}_{reach}} \boldsymbol{\pi}[\mathbf{i}] = 1$ \mathbf{Q} must be ergodic

Explore(in: $\mathcal{X}_{init}, \mathcal{N}$; out: $\mathcal{X}_{reach}, \mathbf{R}, \psi$) is

```

1  $n \leftarrow 0$ ; state indices start at 0
2  $\mathcal{X}_{reach} \leftarrow \emptyset$ ;  $\mathcal{X}_{reach}$  contains the states explored so far
3  $\mathcal{U} \leftarrow \mathcal{X}_{init}$ ;  $\mathcal{U}$  contains the unexplored states known so far
4 for each  $\mathbf{i} \in \mathcal{X}_{init}$  do
5    $\psi(\mathbf{i}) \leftarrow n++$ ; assign to  $\mathbf{i}$  the next available index and increment  $n$ 
6 end for
7 while  $\mathcal{U} \neq \emptyset$  do
8   choose a state  $\mathbf{i}$  in  $\mathcal{U}$  and move it from  $\mathcal{U}$  to  $\mathcal{X}_{reach}$ ;
9   for each event  $\alpha \in \mathcal{E}$  and each state  $\mathbf{j} \in \mathcal{N}_\alpha(\mathbf{i})$  do
10    if  $\mathbf{j} \notin \mathcal{X}_{reach} \cup \mathcal{U}$  then search to determine whether  $\mathbf{j}$  is a new state
11      $\psi(\mathbf{j}) \leftarrow n++$ ; assign to  $\mathbf{j}$  the next available index and increment  $n$ 
12      $\mathcal{U} \leftarrow \mathcal{U} \cup \{\mathbf{j}\}$ ; remember to explore  $\mathbf{j}$  later
13   end if;
14    $\mathbf{R}[\psi(\mathbf{i}), \psi(\mathbf{j})] \leftarrow \mathbf{R}[\psi(\mathbf{i}), \psi(\mathbf{j})] + \lambda_\alpha(\mathbf{i})\Delta_\alpha(\mathbf{i}, \mathbf{j})$ ;  $\psi$  is used to index  $\mathbf{R}$ 
15 end for;
16 end while;
```

$\psi : \hat{\mathcal{X}} \rightarrow \{0, \dots, |\mathcal{X}_{reach}| - 1\} \cup \{\text{null}\}$ is a state indexing function (e.g., discovery order)

$\lambda_\alpha(\mathbf{i})$ is the rate at which event α fires in state \mathbf{i}

$\Delta_\alpha(\mathbf{i}, \mathbf{j})$ is the probability that, if event α fires in state \mathbf{i} , the next state is \mathbf{j}

A decomposition of a discrete-state model describing a CTMC is Kronecker-consistent if:

- the potential transition rate matrix $\hat{\mathbf{R}}$ is additively partitioned

$$\hat{\mathbf{R}} = \sum_{\alpha \in \mathcal{E}} \hat{\mathbf{R}}_\alpha$$

- $\hat{\mathcal{S}} = \mathcal{X}_L \times \dots \times \mathcal{X}_1$, a global state \mathbf{i} consists of L local states

$$\mathbf{i} = (\mathbf{i}_L, \dots, \mathbf{i}_1)$$

- and, most importantly, we can multiplicatively partition each $\hat{\mathbf{R}}_\alpha$, that is, we can write

$$\lambda_\alpha(\mathbf{i}) = \lambda_{L,\alpha}(\mathbf{i}_L) \cdots \lambda_{1,\alpha}(\mathbf{i}_1) \quad \text{and} \quad \Delta_\alpha(\mathbf{i}, \mathbf{j}) = \Delta_{L,\alpha}(\mathbf{i}_L, \mathbf{j}_L) \cdots \Delta_{1,\alpha}(\mathbf{i}_1, \mathbf{j}_1)$$

$$\hat{\mathbf{R}}_\alpha = \mathbf{R}_{L,\alpha} \otimes \dots \otimes \mathbf{R}_{1,\alpha}$$

We encode the potential transition rate matrix $\hat{\mathbf{R}}$ with $|\mathcal{E}| \times L$ small matrices $\mathbf{R}_{k,\alpha} \in \mathbb{R}^{n_k \times n_k}$

for stochastic Petri nets with transition rates depending on at most one place, any partition of the places into L subsets is consistent (even with inhibitor, reset, or probabilistic arcs)

Generalized stochastic Petri nets

99

A GSPN is a tuple $(\mathcal{P}, \mathcal{E}_T, \mathcal{E}_V, \mathbf{D}^-, \mathbf{D}^+, \mathbf{D}^\circ, \mathbf{G}, \succ, \mathbf{x}_{init}, \lambda, w)$ where:

- $\mathcal{E} = \mathcal{E}_T \cup \mathcal{E}_V$ is a set of transitions, or events
- $(\mathcal{P}, \mathcal{E}, \mathbf{D}^-, \mathbf{D}^+, \mathbf{D}^\circ, \mathbf{G}, \succ, \mathbf{x}_{init})$ is a self-modifying PN with inhibitor arcs, guards, priorities
- $\lambda : \mathcal{E}_T \times \mathbb{N}^{|\mathcal{P}|} \rightarrow [0, +\infty)$ are state-dependent firing rates
- $w : \mathcal{E}_V \times \mathbb{N}^{|\mathcal{P}|} \rightarrow [0, +\infty)$ are state-dependent firing weights

Events in \mathcal{E}_T are timed, events in \mathcal{E}_V are immediate

We require that $\lambda_\alpha(\mathbf{i}) = 0 \Leftrightarrow \alpha$ is not (state-) enabled in \mathbf{i}

We require that $w_\alpha(\mathbf{i}) = 0 \Leftrightarrow \alpha$ is not enabled in \mathbf{i}

The firing distribution of event α in state \mathbf{i} is

- Undefined if α is not enabled in \mathbf{i}
- $\text{Expo}(\lambda_\alpha(\mathbf{i}))$ if $\alpha \in \mathcal{E}_T(\mathbf{i})$
- $\text{Const}(0)$ if $\alpha \in \mathcal{E}_V(\mathbf{i})$

Enabling rule for GSPNs

100

Partition the reachability set \mathcal{X}_{reach} into

- vanishing states (drawn with dotted lines): $\mathcal{V} = \{\mathbf{i} : \exists \alpha \in \mathcal{E}_V, \alpha \text{ is enabled in } \mathbf{i}\}$
- tangible states (drawn with solid lines): $\mathcal{T} = \mathcal{X}_{reach} \setminus \mathcal{V} = \{\mathbf{i} : \forall \alpha \in \mathcal{E}_V, \alpha \text{ is disabled in } \mathbf{i}\}$

If a state \mathbf{i} enables an immediate event, no timed event can fire in \mathbf{i}

All timed events are (stochastically-) disabled in the vanishing states

The expected holding time in tangible state \mathbf{i} is:

$$\mathbf{h}[\mathbf{i}] = \frac{1}{\sum_{\alpha \in \mathcal{E}_T(\mathbf{i})} \lambda_\alpha(\mathbf{i})}$$

The firing probability of enabled event $\alpha \in \mathcal{E}_T$ in tangible state \mathbf{i} is:

$$\hat{w}_\alpha(\mathbf{i}) = \lambda_\alpha(\mathbf{i}) \cdot \mathbf{h}[\mathbf{i}]$$

The firing probability of enabled event $\alpha \in \mathcal{E}_V$ in vanishing state \mathbf{i} is:

$$\hat{w}_\alpha(\mathbf{i}) = \frac{w_\alpha(\mathbf{i})}{\sum_{\beta \in \mathcal{E}_V(\mathbf{i})} w_\beta(\mathbf{i})}$$

Consider the stochastic process $\{(\mathbf{i}^{[n]}, \alpha^{[n]}, t^{[n]}) : n \in \mathbb{N}\}$ where

- $\mathbf{i}^{[n]}$ is the n -th state entered by the GSPN
- $\alpha^{[n]}$ is the n -th event to fire
- $t^{[n]}$ is the time of the n -th firing

$\mathbf{i}^{[0]} = \mathbf{x}_{init}$
 $\alpha^{[0]}$ is undefined
 $t^{[0]} = 0$


$\{\mathbf{i}^{[n]} : n \in \mathbb{N}\}$ is a discrete-time Markov chain (DTMC) over $\mathcal{X}_{reach} = \mathcal{T} \cup \mathcal{V}$

Consider sojourns into tangible states:

- $\mathbf{i}(\theta) = \mathbf{i} \Leftrightarrow \mathbf{i} = \mathbf{i}^{[n]}$ and $\theta^{[n]} \leq \theta < \theta^{[n+1]}$

$\{\mathbf{i}(\theta) : \theta \geq 0\}$ is a continuous-time Markov chain (CTMC) over $\mathcal{X}_{reach} = \mathcal{T}$

$\mathcal{X}_4 : \{p^1, p^0\} \equiv \{0, 1\}$ $\mathcal{X}_3 : \{q^0 r^0, q^1 r^0, q^0 r^1\} \equiv \{0, 1, 2\}$ $\mathcal{X}_2 : \{s^0, s^1\} \equiv \{0, 1\}$ $\mathcal{X}_1 : \{t^0, t^1\} \equiv \{0, 1\}$




note the shaded identity patterns!!!

EV*MDDs by example

One way to think about EV*MDDs is "EV*MDD = $-\log(\text{EV*MDD})$ ":

- $0 \Leftrightarrow 1$
- edge values $\in [0, +\infty] \Leftrightarrow$ edge values $\in [0, 1]$
- root incoming edge $\in (-\infty, +\infty] \Leftrightarrow$ root incoming edge $\in [0, +\infty)$
- values add along the path \Leftrightarrow values multiply along the path

\mathbf{i}_3	0	0	0	0	1	1	1	1
\mathbf{i}_2	0	0	1	1	0	0	1	1
\mathbf{i}_1	0	1	0	1	0	1	0	1
f	3.5	7	2.1	0	0	2.8	7	1.4




Canonicity: all edge values are in $[0, 1]$ and at least one is 1

In canonical form, the root incoming edge has value $\max_{\mathbf{i} \in \hat{\mathcal{X}}} f(\mathbf{i})$

Encoding R with an EV*MDD: initial non-canonical EV*MDDs


We can store R with a $2K$ -level EV*MDD: consider the example of Kronecker encoding

$\mathbf{R} = \sum_{t \in \{a, l, d, e\}} \mathbf{R}_t = \sum_{t \in \{a, l, d, e\}} \bigotimes_{4 \geq k \geq 1} \mathbf{R}_{k,t}$




note the shaded identity patterns!!!

(assume that β_3 is the largest rate in \mathbf{R})



Use a recursive algorithm to compute

$$\mathbf{R} = \sum_{t \in \{a,l,d,e\}} \mathbf{R}_t$$



hidden identity patterns remain!!!

Matrix diagrams (MxDs)

Assume a domain $\hat{\mathcal{X}} = \mathcal{X}_L \times \dots \times \mathcal{X}_1$, where $\mathcal{X}_k = \{0, 1, \dots, n_k - 1\}$, for some $n_k \in \mathbb{N}$

Assume the range $\mathbb{R}^{\geq 0} = [0, +\infty)$ and the combinator “.” (multiplication over the reals)

An (edge-valued) MxD is an acyclic directed edge-labeled graph where:

- The only terminal node is Ω and is at level 0 $\Omega.lvl = 0$
- A nonterminal node p is at a level k , with $L \geq k \geq 1$ $p.lvl = k$
- A nonterminal node p at level k has $n_k \times n_k$ outgoing edges
- For $i_k, i'_k \in \mathcal{X}_k$, edge $p[i_k, i'_k]$ points to child $p[i_k, i'_k].child$, and has value $p[i_k, i'_k].val \geq 0$
- The level of the children is lower than that of p $p[i_k, i'_k].child.lvl < p.lvl$
- An edge $\langle \sigma, p \rangle$, with $p.lvl = k$ encodes the function $v_{\langle \sigma, p \rangle} : \hat{\mathcal{X}} \rightarrow \mathbb{Z}$ defined recursively by

$$v_{\langle \sigma, p \rangle}(x_L, x'_L, \dots, x_1, x'_1) = \begin{cases} \sigma & \text{if } k = 0, \text{ i.e., } p = \Omega \\ \sigma \cdot v_{p[x_k, x'_k]}(x_L, x'_L, \dots, x_1, x'_1) & \text{if } k > 0, \text{ i.e., } p \neq \Omega \end{cases}$$

This definition of $v_{\langle \sigma, p \rangle}$ applies when no edge skips a level, otherwise we have more choices...

Canonical versions of MxDs

For canonical MxDs, we first normalize each node p in one of two ways:

- $\max\{p[i_k, i'_k].val : i_k, i'_k \in \mathcal{X}_k\} = 1$, or
- $\min\{p[i_k, i'_k].val : i_k, i'_k \in \mathcal{X}_k, p[i_k, i'_k].val \neq 0\} = 1$

Then, the usual reduction requirements apply, there are no duplicates:

- If $p.lvl = q.lvl = k$ and $p[i_k] = q[i_k]$ for all $i_k \in \mathcal{X}_k$, then $p = q$

And, if the MxD is quasi-reduced, there is no level skipping:

- The only root nodes with no incoming arcs are at level L , and have root edge values in \mathbb{Z}
- Each child $p[i_k, i'_k].child$ of a node p is at level $p.lvl - 1$

Or, if the MxD is fully-reduced, there is no redundant node p satisfying:

- $p[i_k, i'_k].child = q$ and $p[i_k, i'_k].val = 1$ for all $i_k, i'_k \in \mathcal{X}_k$


Or, if the MxD is identity-reduced, there are no identity nodes p satisfying:

- $p[i_k, i_k].child = q$ and $p[i_k, i_k].val = 1$ for all $i_k \in \mathcal{X}_k$
- $p[i_k, i'_k].val = 0$ for all $i_k \neq i'_k$

Rows and columns of the matrix are indexed by $x_2 \cdot 2 + x_1$ and $x'_2 \cdot 2 + x'_1$


- 0 $\equiv (x_2 = 0, x_1 = 0)$
- 1 $\equiv (x_2 = 0, x_1 = 1)$
- 2 $\equiv (x_2 = 1, x_1 = 0)$
- 3 $\equiv (x_2 = 1, x_1 = 1)$
- 4 $\equiv (x_2 = 2, x_1 = 0)$
- 5 $\equiv (x_2 = 2, x_1 = 1)$

0	1	2	3	4	5
0	0	0	0	0	0
1	0	0	0	0	1
2	0	0	0	0	0
3	1	6	2	12	0
4	0	0	7	0	0
5	0	0	0	7	0



$R[001,210] = 1*5*2*9 = 90$

$R[102,101] = 1*4*11 = 44$



0	0	0	0	0	1	1	1	1	1	2	2	2	2	2
0	0	0	1	1	0	0	0	1	1	0	0	1	1	1
0	1	2	0	1	2	0	1	2	0	1	2	0	1	2
000					8									70
001					5	1							90	10
002					11									20
010						8			40					
011							5	1		25	30			
012							11			55				
100					32									
101					20	4								
102					44									
110							32			56			14	
111								20	4	72	8		18	2
112								44			16			4
200					42									28
201					54	6								36
202							12							8
210	24										16			
211		15	3									10	2	
212			33										22	

State indexing options: potential $\hat{\psi}$ vs. actual ψ

For Markov analysis, we can generate \mathcal{X}_{reach} first, using \mathcal{X}_{init} and $\mathcal{N} : \hat{\mathcal{X}} \rightarrow 2^{\hat{\mathcal{X}}}$

Once we know \mathcal{X}_{reach} :

- We can restrict \mathcal{N} to $\mathcal{N} : \mathcal{X}_{reach} \rightarrow 2^{\mathcal{X}_{reach}}$ (if needed for further logical analysis)
- We can store $\hat{\mathbf{R}} : \hat{\mathcal{X}} \times \hat{\mathcal{X}} \rightarrow \mathbb{R}$ or $\mathbf{R} : \mathcal{X}_{reach} \times \mathcal{X}_{reach} \rightarrow \mathbb{R}$
- We can choose algorithms that use $\hat{\pi} : \hat{\mathcal{X}} \rightarrow \mathbb{R}$ or $\pi : \mathcal{X}_{reach} \rightarrow \mathbb{R}$

Strictly **explicit** methods: using **actual R** and π works best

Strictly **implicit** methods: decision diagrams usually don't work well to store $\hat{\pi}$ or π

Implicit methods have tradeoffs:

- Storing π instead of $\hat{\pi}$ is often unavoidable if we employ a full vector and $|\hat{\mathcal{X}}| \gg |\mathcal{X}_{reach}|$
- Symbolic storage of $\hat{\mathbf{R}}$ is usually cheaper than that of \mathbf{R} in terms of memory requirements
- However, using $\hat{\mathbf{R}}$ in conjunction with π complicates indexing...
- ...forcing us to store $\psi : \hat{\mathcal{X}} \rightarrow \{0, 1, \dots, |\mathcal{X}_{reach}| - 1\} \cup \{\text{null}\}$, hence \mathcal{X}_{reach}

MxD-based vector-matrix multiplication algorithm

```

real[n] VectorMatrixMult(real[n] x, mxdd_node A, evmdd_node  $\psi$ ) is
    local natural s;
    local real[n] y;
    local sparse_real c;
    1 s  $\leftarrow$  0;
    2 for each  $\mathbf{j} = (j_L, \dots, j_1) \in \mathcal{X}_{reach}$  in lexicographic order do
        3  $\mathbf{c} \leftarrow \text{GetCol}(L, A, \psi, j_L, \dots, j_1)$ ;
        4  $\mathbf{y}[s] \leftarrow \text{ElementWiseMult}(\mathbf{x}, \mathbf{c})$ ;
        5 s  $\leftarrow$  s + 1;
    6 return y;
    
```

```

sparse_real GetCol(level k, mxdd_node M, evmdd_node  $\phi$ , natural  $j_k, \dots, j_1$ ) is
    local sparse_real c, d;
    1 if k = 0 then return [1];
    2 if Cache contains entry  $\langle \text{GetColCODE}, M, \phi, j_k, \dots, j_1 : \mathbf{c} \rangle$  then return c;
    3  $\mathbf{c} \leftarrow \mathbf{0}$ ;
    4 for each  $i_k \in \mathcal{X}_k$  such that  $M[i_k, j_k].val \neq 0$  and  $\phi[i_k].val \neq \infty$  do
        5  $\mathbf{d} \leftarrow \text{GetCol}(k - 1, M[i_k, j_k].child, \phi[i_k].child, j_{k-1}, \dots, j_1)$ ;
        6 for each i such that  $\mathbf{d}[i] \neq 0$  do
            7  $\mathbf{c}[i + \phi[i_k].val] \leftarrow \mathbf{c}[i + \phi[i_k].val] + M[i_k, j_k].val \cdot \mathbf{d}[i]$ ;
    8 enter  $\langle \text{GetColCODE}, M, \phi, j_k, \dots, j_1 : \mathbf{c} \rangle$  in Cache;
    9 return c;
    
```

Memory consumption in bytes for:

\mathcal{X}_{reach} (MDD), \mathbf{R} (Sparse), $\hat{\mathbf{R}}$ (Kronecker), $\hat{\mathbf{R}}$ and \mathbf{R} (Pot/Act MxD), $\hat{\mathbf{R}}$ and \mathbf{R} (Pot/Act MTMDD)

Model	N	$ \hat{\mathcal{X}} $	$ \mathcal{X}_{reach} $	MDD	Sparse	Kron	Pot MxD	Act MxD	Pot MTMDD	Act MTMDD
qn4	2	324	324	333	14,256	772	586	722	22,784	22,784
	6	38,416	38,416	499	2,524,480	3,092	2,494	2,870	36,864	36,864
	10	527,076	527,076	905	38,524,464	7,076	5,778	6,522	62,720	62,720
qn8	2	6,561	324	681	14,256	1,204	738	1,688	43,776	49,152
	6	5,764,801	38,416	1,119	2,524,480	2,404	1,674	5,872	55,040	70,912
	10	214,358,881	527,076	1,953	38,524,464	3,604	2,610	12,040	66,304	98,560
mserv2	3	1,485	495	705	23,352	4,124	3,246	3,952	34,560	40,704
	6	6,345	2,115	3,176	111,408	17,468	13,998	16,432	111,104	135,168
	10	18,495	6,165	8,846	342,720	52,228	42,278	49,032	306,560	378,460
mserv4	3	14,256	495	1,174	23,352	5,568	4,098	4,916	68,864	79,616
	6	106,596	2,115	8,453	111,408	22,920	17,502	20,054	254,360	298,856
	10	488,268	6,165	33,739	342,720	67,560	52,342	58,934	873,896	998,552
mserv6	3	32,076	495	1,333	23,352	5,724	4,066	5,316	86,784	101,376
	6	239,841	2,115	8,614	111,408	23,076	17,470	20,238	298,596	347,956
	10	1,098,603	6,165	33,900	342,720	67,716	52,310	59,118	982,396	1,112,684

Model	N	$ \hat{\mathcal{X}} $	$ \mathcal{X}_{reach} $	MDD	Sparse	Kron	Pot MxD	Act MxD	Pot MTMDD	Act MTMDD
molloy4	5	4,536	91	660	4,204	1,316	1,148	2,534	23,552	28,160
	8	32,805	285	1,215	14,676	2,528	2,300	5,216	27,648	38,656
	10	87,846	506	1,766	27,104	3,556	3,288	7,504	31,232	47,360
molloy5	5	7,776	91	846	4,204	1,100	792	4,298	28,416	37,120
	8	59,049	285	1,545	14,676	1,592	1,188	9,356	31,232	50,944
	10	161,051	506	2,223	27,104	1,920	1,452	13,778	33,280	61,952
kan3	1	160	160	264	8,032	500	412	544	18,432	18,432
	3	58,400	58,400	937	5,590,400	7,572	6,786	8,134	66,816	67,072
	5	2,546,432	2,546,432	5,646	303,705,920	45,660	41,816	48,780	303,776	303,776
kan4	1	256	160	332	8,032	420	354	602	23,552	24,576
	3	160,000	58,400	628	5,590,400	2,500	2,216	3,284	44,032	50,176
	5	9,834,496	2,546,432	1,532	303,705,920	7,940	7,118	9,950	92,928	110,592
kan16	1	65,536	160	1,275	8,032	2,148	866	3,000	95,232	107,520
	3	—	58,400	1,902	5,590,400	3,236	1,746	10,566	115,456	151,808
	5	—	2,546,432	3,149	303,705,920	4,324	2,626	24,106	135,168	216,320
fms5	1	2,100	84	535	3,228	1,456	604	1,808	36,096	40,960
	3	9,432,500	20,600	3,294	1,554,080	8,304	5,224	24,320	151,296	247,040
	5	2,016,379,008	852,012	30,490	82,727,748	34,484	24,664	138,244	654,892	1,255,108
fms21	1	4,194,304	84	2,050	3,228	3,132	1,132	7,396	126,976	148,224
	3	—	20,600	6,777	1,554,080	5,028	2,328	68,762	176,896	437,760
	5	—	852,012	22,038	82,727,748	6,924	3,524	255,988	235,008	1,393,932

Results for the numerical solution

Matrix diagrams achieve the highest efficiency in the vector-matrix multiplications. . .
 . . . and can provide access by columns as required by Gauss–Seidel

Time requirements for the Kanban model

N	$ \mathcal{X}_{reach} $	number of nonzeros in \mathbf{R}	MxDs		Kronecker				Explicit	
			Gauss–Seidel lters	sec/iter	Gauss–Seidel lters	sec/iter	Jacobi lters	sec/iter	Gauss–Seidel lters	sec/iter
2	4,600	28,120	40	0.11	55	0.17	134	0.09	55	0.02
3	58,400	446,400	67	1.46	97	2.56	240	1.34	97	0.34
4	454,475	3,979,850	99	12.33	149	23.69	370	11.99	149	3.04
5	2,546,432	24,460,016	139	73.09	214	147.70	527	74.09	214	18.51
6	11,261,376	115,708,992	185	336.21	289	723.30	713	359.15	—	—
7	41,644,800	450,455,040	238	1,289.91	374	2,922.80	—	—	—	—

Invariant analysis of Petri nets

P-semiflows: Definitions and meaning. Farkas' algorithm. Zero-suppressed integer-range MDDs. A fully symbolic algorithm for p-semiflow computation.

Let $\mathbf{D} = \mathbf{D}^+ - \mathbf{D}^-$ be the flow matrix

A p-semiflow is a non-zero solution $\mathbf{w} \in \mathbb{N}^n$ to the set of linear flow equations $\mathbf{w} \cdot \mathbf{D} = \mathbf{0}$

P-semiflow \mathbf{w} defines the invariant constraint $\sum_{p \in \mathcal{P}} \mathbf{w}_p \cdot \mu_p = C$ on any reachable marking μ

The initial marking μ^{init} determines the constant $C = \sum_{p \in \mathcal{P}} \mathbf{w}_p \cdot \mu_p^{init}$

P-semiflows provide necessary, not sufficient, conditions on reachability

A linear combination of p-semiflows is a p-semiflow \Rightarrow either none or infinite number of p-semiflows

Support of p-semiflow \mathbf{w} : the set of places with positive weight $Supp(\mathbf{w}) = \{p \in \mathcal{P} : \mathbf{w}_p > 0\}$

A p-semiflow \mathbf{w} is minimal if it is scaled back (GCD of its entries is 1) and has minimal support

We seek a (minimal) generator set of minimal p-semiflows $\mathcal{W} = \{\mathbf{w}^{(1)}, \dots, \mathbf{w}^{(r)}\}$

Any p-semiflow \mathbf{w} can be derived from \mathcal{W} through non-negative integer linear combinations

The generator set \mathcal{W} is unique, but its size can be exponential in the number of places n

We manipulate a matrix $[\mathbf{T} | \mathbf{P}]$ stored as a set \mathcal{A} of integer row vectors of length $m + n$

Initially, $[\mathbf{T} | \mathbf{P}] = [\mathbf{D} | \mathbf{I}] \in \mathbb{Z}^{n \times m} \times \mathbb{N}^{n \times n}$ \mathbf{D} is the flow matrix, \mathbf{I} is the $n \times n$ identity matrix

We iteratively manipulate the set of rows, forcing zero entries in the first j columns, for $j = 1, \dots, m$

We substitute rows with positive or negative entry j with linear combinations having a zero entry j

At the end, any row $[\mathbf{t} | \mathbf{p}]$ in \mathcal{A} is such that $\mathbf{t} \in \mathbb{Z}^m$ is all zero and $\mathbf{p} \in \mathbb{N}^n$ is a p-semiflow

```

set of array[m+n] of int ExpPSemiflows(int n, int m, set of array[m+n] of int A) is
1 for j = 1 to m do
2   A_N ← ∅;                                     set of rows with negative entry j
3   A_P ← ∅;                                     set of rows with positive entry j
4   foreach a ∈ A do                             partition the rows of A according to entry j
5     if a[j] < 0 then A_N ← A_N ∪ {a};
6     if a[j] > 0 then A_P ← A_P ∪ {a};
7   A ← A \ (A_N ∪ A_P);                         remove from A rows with nonzero entry j
8   foreach a_N ∈ A_N and a_P ∈ A_P do          |A_N| · |A_P| linear combinations
9     v ← MinimumCommonMultiple(-a_N[j], a_P[j]);
10  A ← A ∪ {(-v/a_N[j]) · a_N + (v/a_P[j]) · a_P};  entry j of new row is 0
11 return A;
    
```

Compute a generator set (explicit version)

The output \mathcal{A} of Farkas' algorithm may contain **unscaled** and **non-minimal-support** p-semiflows

To scale back the set \mathcal{A} of p-semiflows:

- for each row $\mathbf{a} \in \mathcal{A}$, divide \mathbf{a} by the GCD of all entries in \mathbf{a} :

$$\mathcal{A} \leftarrow (\mathcal{A} \setminus \{\mathbf{a}\}) \cup \{\mathbf{a} / \gcd(\mathbf{a}[m+1], \dots, \mathbf{a}[m+n])\}$$

- requires examining each support in \mathcal{A}

$$\text{time complexity } O(|\mathcal{A}| \cdot n)$$

To eliminate from \mathcal{A} the non-minimal-support p-semiflows:

- for each pair of distinct rows \mathbf{a} and \mathbf{b} in \mathcal{A} , delete \mathbf{b} if its support is a superset of that of \mathbf{a} :

$$\text{if } Supp(\mathbf{a}) \subset Supp(\mathbf{b}) \text{ then } \mathcal{A} \leftarrow \mathcal{A} \setminus \{\mathbf{b}\}$$


- requires $|\mathcal{A}| \cdot (|\mathcal{A}| - 1) / 2$ support comparisons

$$\text{time complexity } O(|\mathcal{A}|^2 \cdot n)$$

Alternatively, we can avoid adding redundant rows to \mathcal{A} during Algorithm *ExpPSemiflows*:

- before iteration j , \mathcal{A} contains the minimal support p-semiflows ignoring transitions j, \dots, m
- during iteration j , a row is added to \mathcal{A} only if it does not contain the support of a row in \mathcal{A}
- worst-case complexity remains $O(|\mathcal{A}|^2 \cdot n)$, but this alternative is quite beneficial in practice

Example of p-semiflows: a traffic light controller



Initial set of rows \mathcal{A} describing matrix $[\mathbf{T} | \mathbf{P}] = [\mathbf{D} | \mathbf{I}]$

	T_{1a}	T_{1b}	T_{1c}	T_{2a}	T_{2b}	T_{2c}	G_1	Y_1	R_1	G_2	Y_2	R_2	S
G_1	1	-1					1						
Y_1		1	-1					1					
R_1			1						1				
G_2				1	-1					1			
Y_2					1	-1					1		
R_2					-1	1						1	
S							-1	1	-1				1

Final set of rows \mathcal{A}

$\mathbf{w}^{(1)}$							1	1	1				
$\mathbf{w}^{(2)}$										1	1	1	
$\mathbf{w}^{(3)}$							1	1		1	1		1

An MDD over variables $x_L \succ x_{L-1} \succ \dots \succ x_1$ is a directed acyclic edge-labeled multi-graph:

- A **nonterminal** node p is associated with a variable $p.var = x_k$, with $L \geq k \geq 1$, and has an infinite set of outgoing edges, each indexed by a different $i \in \mathbb{Z}$
- The only **terminal** nodes are $\mathbf{0}$ and $\mathbf{1}$ $\mathbf{0}.var = \mathbf{1}.var = x_0$, with $x_k \succ x_0$ for $L \geq k \geq 1$
- The edge with index $i \in \mathbb{Z}$ from node p points to a node q $p[i] = q$, with $p.var \succ q.var$

Using a **zero-suppressed** semantic, node p with $p.var = x_k$ encodes the set of k -tuples:

$$\mathcal{X}(p) = \begin{cases} \emptyset & \text{if } p = \mathbf{0} \\ \{\epsilon\} & \text{if } p = \mathbf{1} \\ \bigcup_{i:p[i] \neq \mathbf{0} \wedge p[i].var = x_h} \{i\} \cdot \{0^{k-h-1}\} \cdot \mathcal{X}(p[i]) & \text{otherwise} \end{cases}$$

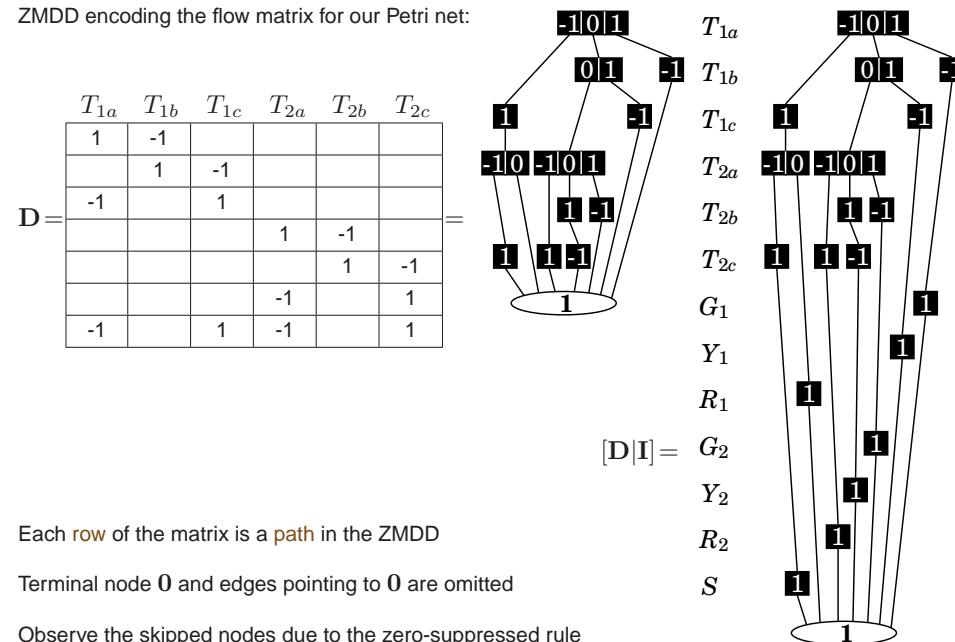
To enforce a **finite representation**

- we require that $\{i \in \mathbb{Z} : p[i] \neq \mathbf{0}\}$ be finite

To enforce **canonicity**

- we forbid nodes where all edges point to $\mathbf{0}$
- we forbid duplicate nodes if $p.var = q.var$ and $p[i] = q[i]$ for all $i \in \mathbb{Z}$, then $p = q$
- we force variable-skipping when possible no node p has $p[i] = \mathbf{0}$ for all $i \in \mathbb{Z} \setminus \{0\}$

ZMDD encoding the flow matrix for our Petri net:



Each row of the matrix is a **path** in the ZMDD

Terminal node $\mathbf{0}$ and edges pointing to $\mathbf{0}$ are omitted

Observe the skipped nodes due to the zero-suppressed rule

A symbolic approach to compute the generator set

The explicit algorithm to compute the p-semiflows manages **sets of row vectors of size $m + n$**

Traditional MDDs are ideal to **encode and manipulate large sets of same-length tuples**

Our ZMDDs **excel when most vectors contain mostly zeros**, often the case in p-semiflow computation

Our approach is thus:

- Build a ZMDD a over $v_1 \succ \dots \succ v_m \succ w_1 \succ \dots \succ w_n$ encoding the set \mathcal{A} of rows in $[D \mid I]$
- For $j = 1, \dots, m$, symbolically eliminate all rows with $v_j \neq 0$, using linear combinations
 - \Rightarrow After iteration j , all rows have $v_1 = \dots = v_j = 0$, thus the root of the ZMDD is below v_j
- Apply a symbolic algorithm to eliminate non-minimal support p-semiflows
 - \Rightarrow We propose **four variants**, differing on when and how rows are eliminated
- At the end, apply a symbolic algorithm to scale back the remaining invariants
 - \Rightarrow We use a symbolic brute force search for the scale-back factors

Pseudo-code for symbolic p-semiflow computation

Given MDD a encoding a set of rows, column $a.var = v_j$ is annulled through the following steps:

- Collect the rows with negative v_j and positive v_j into two new MDDs a_N and a_P
- Leave the rows with $v_j = 0$ in a (being a ZMDD, a becomes $a[0]$, i.e., its height decreases)
- Perform the pairwise linear combination between each $a_P[i_P]$ and $a_N[i_N]$ to annul v_j
- Finally, combine the resulting MDD with a using an ordinary *Union* operation

mdd *SymPSemiflows*(int m , mdd a) is

```

1 for  $j = 1$  to  $m$  do
2   if  $a.var \neq v_j$  skip; nothing to do if  $a$  encodes only rows with  $v_j = 0$ 
3    $a_N \leftarrow Intersection(a, Potential(v_j < 0));$  set of rows with negative  $v_j$ 
4    $a_P \leftarrow Intersection(a, Potential(v_j > 0));$  set of rows with positive  $v_j$ 
5    $a \leftarrow Intersection(a, Potential(v_j = 0));$  redefine  $a$  for the next iteration
6   foreach  $i_N$  s.t.  $a_N[i_N] \neq \mathbf{0}$  and  $i_P$  s.t.  $a_P[i_P] \neq \mathbf{0}$  do
7      $\rho \leftarrow MinimumCommonMultiple(-i_N, i_P);$ 
8      $\rho_N \leftarrow \rho / (-i_N);$ 
9      $\rho_P \leftarrow \rho / i_P;$ 
10     $a \leftarrow Union(a, SymLinComb(\rho_N, a_N[i_N], \rho_P, a_P[i_P]));$ 
11 return  $a;$ 

```

Given two ZMDDs with $x.var = y.var$, create a ZMDD r with $r.var = x.var = y.var$

For each pair of edges $x[i_x]$ and $y[i_y]$, add a new edge $r[j]$ to r , where $j = \rho_x i_x + \rho_y i_y$

This edge points to the ZMDD node encoding $SymLinComb(\rho_x, x[i_x], \rho_y, y[i_y])$

If r already contains an edge $r[j]$, perform a union instead of creating a new edge

```

mdd SymLinComb(int  $\rho_x$ , mdd  $x$ , int  $\rho_y$ , mdd  $y$ ) is
1 if  $x = \mathbf{1}$  and  $y = \mathbf{1}$  then return  $\mathbf{1}$ ;
2 if  $x = \mathbf{0}$  or  $y = \mathbf{0}$  then return  $\mathbf{0}$ ;
3 if  $InCache(C_{SymLinComb}, \rho_x, x, \rho_y, y, r)$  then return  $r$ ;
4 if  $x.var \succ y.var$  then  $y.var$  is skipped
5    $r \leftarrow NewNode(x.var)$ ;
6   foreach  $i_x$  s.t.  $x[i_x] \neq \mathbf{0}$  do  $r[\rho_x i_x] \leftarrow SymLinComb(\rho_x, x[i_x], \rho_y, y)$ ;
7 else if  $y.var \succ x.var$  then  $x.var$  is skipped
8    $r \leftarrow NewNode(y.var)$ ;
9   foreach  $i_y$  s.t.  $y[i_y] \neq \mathbf{0}$  do  $r[\rho_y i_y] \leftarrow SymLinComb(\rho_x, x, \rho_y, y[i_y])$ ;
10 else  $y.var = x.var$ 
11    $r \leftarrow NewNode(x.var)$ ;
12   foreach  $i_x$  s.t.  $x[i_x] \neq \mathbf{0}$  and  $i_y$  s.t.  $y[i_y] \neq \mathbf{0}$  do
13      $j \leftarrow \rho_x i_x + \rho_y i_y$ ;
14      $r[j] \leftarrow Union(r[j], SymLinComb(\rho_x, x[i_x], \rho_y, y[i_y]))$ 
15  $r \leftarrow UniqueTableInsert(r)$ ;
16  $CacheAdd(C_{SymLinComb}, \rho_x, x, \rho_y, y, r)$ ;
17 return  $r$ ;

```

The support of a new row is never a subset of the support of a row already in \mathcal{A}

new rows never eliminate old rows

Non-minimal-support p-semiflows can be eliminated periodically or at the end

Periodic “internal” elimination, or at the end

- Use a single ZMDD a
- $MinSuppInt$ eliminates p-semiflows that are a linear combination of multiple p-semiflows in a
- The result of $SymLinComb$ is immediately unioned with a
- $MinSuppInt$ can be applied any time to the intermediate result of the p-semiflow computation

Periodic “external” elimination

- ZMDD a encodes a set of minimal-support p-semiflows
- A second ZMDD b stores a single or multiple linear combinations, computed by $SymLinComb$
- $ElimNMSupp$ removes from b p-semiflows with non-minimal-support w.r.t. b
- $MinSuppExt$ removes from b p-semiflows with non-minimal-support w.r.t. a or a and b
- Only then, ZMDDs a and b can be safely unioned

Helper functions to remove non-minimal support p-semiflows 127

Prune given a ZMDD a , return the *Union* of all nodes p such that

- p is either a or a descendant of a , and
- either $p.var = w_1$ or $w_1 \succ p.var$ and p is pointed to by an edge from a node q s.t. $q.var \succ w_1$

MkBool given a ZMDD a , return the ZBDD encoding the set of boolean vectors

$$\{\mathbf{b} \in \mathbb{B}^n : \exists \mathbf{x} \in \mathcal{X}(Prune(a)), \forall i, 1 \leq i \leq n, \mathbf{b}[i] = 1 \Leftrightarrow \mathbf{x}[i] > 0\}$$

Filter given a ZMDD a and a ZBDD b , return the ZMDD encoding

$$\{\mathbf{x} \in \mathcal{X}(Prune(a)) : \exists \mathbf{b} \in \mathcal{X}(b), \forall i, 1 \leq i \leq n, \mathbf{b}[i] = 1 \Leftrightarrow \mathbf{x}[i] > 0\}$$

An unusual element-wise symbolic operator 128

EWOr given two ZBDDs p and q , return a ZBDD r of the *element-wise-or* of all pairs of tuples

$$\mathcal{X}(r) = \{\mathbf{i} \vee \mathbf{j} : \mathbf{i} \in \mathcal{X}(p), \mathbf{j} \in \mathcal{X}(q)\}$$

- quite unlike the much more familiar (non-element-wise) union of sets encoded by two BDDs
- nevertheless, efficient and elegant symbolic implementation

bdd $EWOr$ (bdd p , bdd q) is

```

1 if  $p = \mathbf{0}$  or  $q = \mathbf{0}$  then return  $\mathbf{0}$ ;
2 if  $p = q$  then return  $p$ ;
3 if  $InCache(C_{EWOr}, p, q, r)$  then return  $r$ ;
4  $r \leftarrow NewNode(p.var)$ ;
5 if  $p.var \succ q.var$  then  $r[0] \leftarrow EWOr(p[0], q)$ ;  $r[1] \leftarrow EWOr(p[1], q)$ ;  $q.var$  is skipped
6 if  $q.var \succ p.var$  then  $r[0] \leftarrow EWOr(p, q[0])$ ;  $r[1] \leftarrow EWOr(p, q[1])$ ;  $p.var$  is skipped
7 else  $p.var = q.var$ 
8    $r[0] \leftarrow EWOr(p[0], q[0])$ ;
9    $r_{01} \leftarrow EWOr(p[0], q[1])$ ;  $r_{10} \leftarrow EWOr(p[1], q[0])$ ;  $r_{11} \leftarrow EWOr(p[1], q[1])$ ;
10   $r[1] \leftarrow Union(r_{01}, Union(r_{10}, r_{11}))$ ;
11  $r \leftarrow UniqueTableInsert(r)$ ;
12  $CacheAdd(C_{EWOr}, p, q, r)$ ;
13 return  $r$ ;

```



```

mdd SymPSemiflows(int m, mdd a) is common portion to all variants
1 for j = 1 to m do
2   if a.var ≠ vj skip;
3   aN ← Intersection(a, Potential(vj < 0));
4   aP ← Intersection(a, Potential(vj > 0));
5   a ← Intersection(a, Potential(vj = 0));
6   newRows ← 0; only for V2
7   foreach iN s.t aN[iN] ≠ 0 and iP s.t. aP[iP] ≠ 0 do
8     ρ ← MinimumCommonMultiple(-iN, iP);
9     ρN ← ρ/(-iN); ρP ← ρ/iP;

V1: Minimize after each linear combination (using external comparisons)
101 linComb = SymLinComb(ρN, aN[iN], ρP, aP[iP]);
111 a ← Union(a, MinSuppExt(linComb, a));
121 return a;

V2: Minimize after annulling column (using external comparisons)
102 linComb ← SymLinComb(ρN, aN[iN], ρP, aP[iP]);
112 newRows ← Union(newRows, linComb);
122 a ← Union(a, MinSuppExt(newRows, a));
132 return a;

V3: Minimize after annulling column (using internal comparisons)
103 a ← Union(a, SymLinComb(ρN, aN[iN], ρP, aP[iP]));
113 a ← MinSuppInt(a);
123 return a;

V4: Minimize only at the end (using internal comparisons)
104 a ← Union(a, SymLinComb(ρN, aN[iN], ρP, aP[iP]));
114 return MinSuppInt(a);
    
```

SymScalePsemiflows repeatedly scales a by the primes in $\{2, \dots, \lfloor \sqrt{\text{largest value in } a} \rfloor\}$
 The MDDs returned by *ScaleByNumber* are unioned into the new scaled-back MDD

```

mdd SymScalePsemiflows(mdd a) is
1 γ ← maxi ∈ N{p[i] ≠ 0 : p is a node in the MDD a};
2 foreach μ ∈ {2, ..., ⌊√γ⌋ : μ is prime} do
3   repeat
4     ⟨s, u⟩ ← ScaleByNumber(a, μ); s encodes scaled paths, u encodes unscaled paths
5     a ← Union(s, u); update a by combining scaled and unscaled paths
6   until s = 0;
7 return a;
    
```

ScaleByNumber takes an MDD a and an integer μ and returns two MDDs:


- s encodes all scaled-back p-semiflows in a which could be scaled by μ
- u encodes those that could not be scaled

```

⟨mdd, mdd⟩ ScaleByNumber(mdd a, int μ) is
1 if a = 0 or a = 1 then return ⟨a, 0⟩;
2 if InCache(CScaleByNumber, a, μ, ⟨s, u⟩) then return ⟨s, u⟩;
3 s ← NewNode(a.var); s will encode paths that were scaled
4 u ← NewNode(a.var); u will encode paths that could not be scaled
5 foreach i s.t a[i] ≠ 0 do
6   if μ divides i then ⟨s[i/μ], u[i]⟩ ← ScaleByNumber(a[i], μ); else u[i] ← a[i];
7 s ← UniqueTableInsert(s);
8 u ← UniqueTableInsert(u);
9 CacheAdd(CScaleByNumber, a, μ, ⟨s, u⟩);
10 return ⟨s, u⟩;
    
```

The four variants are implemented in SMART and compared with the explicit algorithm in GreatSPN

Experiments run on a Pentium4 3.0GHz PC with 1.0GB of RAM running CentoOS Linux 2.6.9



- classic**: a classic net with m transitions and m stages of u places each (u^m minimal p-semiflows)
- classicX**: a modified version of the previous model, using multiple cardinality arcs (still u^m p-semiflows)
- trains**: a circular railway system with u trains and s rail trunks
- slot**: a local area network protocol with u nodes in the network
- robin**: a round robin solution to the mutual exclusion among u processes
- aloha**: the ALOHA networking protocol on u nodes
- mmarch**: a multi-threaded architecture with $u \times u$ processing nodes
- phil**: the classic dining philosophers problem, with u philosophers
- power**: power distribution system with u generators (one p-semiflow)

model	trans	p-semiflows	nodes	edges	mem	PS	MS	SB	time	
trains10	100	37	642	678	V1	6	5.02%	94.95%	0.03%	4.70
					V2	6	4.99%	94.98%	0.03%	4.81
					V3	7	4.76%	95.17%	0.07%	4.05
					V4	om				
					GS	0.06	-	-	-	0.03
trains15	255	99	3,533	3,620	V1	19	4.86%	95.13%	0.01%	210.60
					V2	19	4.80%	95.17%	0.03%	210.05
					V3	21	4.98%	95.01%	0.01%	216.30
					V4	om				
					GS	0.242	-	-	-	0.33
slot20	160	42	1,597	1,798	V1	58	0.20%	99.79%	0.01%	57.44
					V2	58	0.20%	99.79%	0.01%	57.46
					V3	om				
					V4	5	98.46%	0.01%	1.53%	0.29
					GS	3	-	-	-	0.08
slot2000	16,000	4,002	31,997	35,998	V1	om				
					V2	om				
					V3	om				
					V4	242	99.64%	0.01%	0.36%	126.12
					GS	876	-	-	-	189.02

Models requiring minimal-support elimination (cont.)

133

model	trans	p-semiflows	nodes	edges	mem	PS	MS	SB	time	
robin4	24	30	78	96	V1	0.198	15.11%	83.45%	1.44%	0.012
					V2	0.198	15.06%	83.51%	1.42%	0.012
					V3	0.187	10.76%	88.14%	1.10%	0.015
					V4	26	99.9%	0%	0%	11.06
					GS	3	-	-	-	0.01
robin90	540	1.24×10^{27}	1,798	2,160	V1	57	1.02%	99.97%	0.01%	64.89
					V2	57	0.36%	99.62%	0.01%	64.81
					V3	om				
					V4	om				
					GS	ot				
aloha15	60	32,771	78	96	V1	0.325	30.17%	68.99%	0.83%	0.02
					V2	0.326	30.31%	68.88%	0.81%	0.02
					V3	0.362	17.25%	82.26%	0.49%	0.03
					V4	om				
					GS	4	-	-	-	33.20
aloha100	400	1.27×10^{30}	503	606	V1	12	43.42%	56.43%	0.14%	1.00
					V2	12	43.91%	55.95%	0.14%	1.02
					V3	14	6.21%	93.76%	0.03%	4.96
					V4	om				
					GS	ot				

Models not requiring minimal-support elimination

134

model	trans	p-semiflows	nodes	edges	mem	PS	MS	SB	time	
classic10	10	1×10^{10}	100	190	V1	0.055	16.19%	75.17%	8.64%	0.0027
					V2	0.055	15.99%	75.69%	8.31%	0.0028
					V3	0.058	8.51%	87.15%	4.33%	0.0054
					V4	0.031	80.66%	0.81%	18.53%	0.0014
					GS	ot				
classic250	250	3.05×10^{599}	62,500	124,750	V1	613	0.58%	99.00%	0.37%	54.20
					V2	613	0.58%	99.05%	0.37%	53.59
					V3	om				
					V4	8	84.97%	0.01%	15.03%	0.92
					GS	ot				
mmarch10	1,400	404	1,200	1,603	V1	40	0.82%	99.12%	0.06%	4.76
					V2	40	0.83%	99.11%	0.06%	4.79
					V3	om				
					V4	2	95.85%	0.05%	4.93%	0.0056
					GS	3	-	-	-	0.01
mmarch20	5,600	1,604	4,800	6,403	V1	198	3.84%	96.15%	0.01%	122.15
					V2	198	3.91%	96.08%	0.01%	121.99
					V3	om				
					V4	6	96.00%	0%	3.99%	0.32
					GS	110	-	-	-	4.29

Models not requiring minimal-support elimination (cont.)

135

model	trans	p-semiflows	nodes	edges	mem	PS	MS	SB	time	
phil30	120	90	239	328	V1	11	2.35%	97.59%	0.06%	1.04
					V2	11	2.36%	97.59%	0.06%	1.04
					V3	11	0.43%	99.53%	0.04%	1.61
					V4	0.346	94.45%	0.17%	5.38%	0.011
					GS	0.1	-	-	-	0.01
phil100	400	300	799	1,098	V1	50	0.19%	99.79%	0.01%	30.90
					V2	50	0.19%	99.80%	0.01%	30.85
					V3	50	0.13%	99.86%	0.01%	47.64
					V4	2	96.82%	0.04%	3.13%	0.077
					GS	1	-	-	-	0.04
power50	2,600	1	51	51	V1	2	51.25%	48.68%	0.06%	0.21
					V2	2	51.17%	48.76%	0.06%	0.21
					V3	2	24.03%	75.94%	0.03%	0.45
					V4	2	99.84%	0.02%	0.13%	0.11
					GS	0.6	-	-	-	0.08
power100	10,200	1	101	101	V1	14	51.39%	48.58%	0.01%	2.05
					V2	14	51.37%	48.62%	0.01%	2.04
					V3	14	23.51%	76.49%	0.01%	4.50
					V4	14	99.98%	0.00%	0.01%	1.05
					GS	4	-	-	-	0.64
classicX8	8	1.67×10^6	5,193	9,402	V1	14	0.31%	99.87%	0.01%	273.42
					V2	2	68.60%	17.26%	13.94%	0.27
					V3	2	70.53%	16.33%	13.17%	0.29
					V4	4	86.47%	0.01%	13.53%	0.28
					GS	ot				
classicX12	12	8.92×10^{12}	61,584	114,648	V1	om				
					V2	18	84.43%	9.54%	6.02%	29.73
					V3	19	82.78%	11.05%	6.16%	29.07
					V4	18	93.71%	0%	6.29%	28.31
					GS	ot				

Conclusions from the numerical results

136

For models whose runtime is greater than 1 sec, less than 1% of runtime is spent scaling back Except for *classicX*, which requires many *ScaleByNumber* calls due to the large row entries Even on *classicX*, scaling back requires less than 14% of runtime

Best: V1,V2,V3 (*trains*), V4 (*slot,classic,mmarch,phil,power*) V1,V2 (*robin,aloha*) V2,V3,V4 (*classicX*)

At least one of either V2 or V4 was the most efficient (or very close) for each model

V2 works best for models that add many non-minimal support p-semiflows at each step

V4 works best for models where few non-minimal support p-semiflows are generated

We could start with V4 and switch to V2 if many non-minimal support p-semiflows are being generated Alternatively, we could run V2 and V4 on two independent workstations

GreatSPN tends to be more efficient for models with relatively few p-semiflows

For most models, at least one of our variants outperforms GreatSPN for large enough instances

Our new symbolic method offers vast time and space improvements

— the most dramatic example is *classic250*: 3.05×10^{599} p-semiflows in 1 sec using 8MBytes

— many Petri nets with unit arc cardinalities require a single call to *SymLinComb* per column

Two types of models have larger time and memory requirements with our symbolic method

— models with a dense flow matrix, they do not benefit as much from the ZMDD properties

— models with arc cardinalities greater than one, as revealed by comparing *classic* and *classicX* (still, the symbolic method can generate the 8.92×10^{12} p-semiflows of *classicX12* in 30 sec)

Relevant references from my work

- [1] P. Buchholz, G. Ciardo, S. Donatelli, and P. Kemper. Complexity of memory-efficient Kronecker operations with applications to the solution of Markov models. *INFORMS J. Comp.*, 12(3):203–222, 2000.
- [2] M.-Y. Chung, G. Ciardo, and A. J. Yu. A fine-grained fullness-guided chaining heuristic for symbolic reachability analysis. In S. Graf and W. Zhang, editors, *Proc. 4th International Symposium on Automated Technology for Verification and Analysis (ATVA)*, LNCS 4218, pages 51–66, Beijing, China, Oct. 2006. Springer-Verlag.
- [3] G. Ciardo. Data representation and efficient solution: a decision diagram approach. In M. Bernardo and J. Hillston, editors, *Formal Methods for Performance Evaluation*, LNCS 4486, pages 371–394, Bertinoro, Italy, May 2007. Springer-Verlag.
- [4] G. Ciardo, M. Forno, P. L. Grieco, and A. S. Miner. Comparing implicit representations of large CTMCs. In W. J. Stewart and A. N. Langville, editors, *Numerical Solution of Markov Chains*, pages 323–327, Urbana, IL, USA, Sept. 2003.
- [5] G. Ciardo, G. Lüttgen, and A. S. Miner. Exploiting interleaving semantics in symbolic state-space generation. *Formal Methods in System Design*, 31:63–100, 2007.
- [6] G. Ciardo, G. Lüttgen, and R. Siminiceanu. Saturation: An efficient iteration strategy for symbolic state space generation. In T. Margaria and W. Yi, editors, *Proc. Tools and Algorithms for the*

Construction and Analysis of Systems (TACAS), LNCS 2031, pages 328–342, Genova, Italy, Apr. 2001. Springer-Verlag.

- [7] G. Ciardo, R. Marmorstein, and R. Siminiceanu. Saturation unbound. In H. Gavel and J. Hatcliff, editors, *Proc. Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, LNCS 2619, pages 379–393, Warsaw, Poland, Apr. 2003. Springer-Verlag.
- [8] G. Ciardo, R. Marmorstein, and R. Siminiceanu. The saturation algorithm for symbolic state space exploration. *Software Tools for Technology Transfer*, 8(1):4–25, Feb. 2006.
- [9] G. Ciardo, G. Mecham, E. Paviot-Adet, and M. Wan. P-semiflow computation with decision diagrams. In G. Franceschinis and K. Wolf, editors, *Proc. 30th International Conference on Application and Theory of Petri nets and Other Models of Concurrency (ICATPN)*, LNCS 5606, pages 143–162, Paris, France, June 2007. Springer-Verlag.
- [10] G. Ciardo and A. S. Miner. Implicit data structures for logic and stochastic systems analysis. *ACM SIGMETRICS Perf. Eval. Rev.*, 32(4):4–9, 2005.
- [11] G. Ciardo and R. Siminiceanu. Using edge-valued decision diagrams for symbolic generation of shortest paths. In M. D. Aagaard and J. W. O’Leary, editors, *Proc. Fourth International Conference on Formal Methods in Computer-Aided Design (FMCAD)*, LNCS 2517, pages 256–273, Portland, OR, USA, Nov. 2002. Springer-Verlag.
- [12] G. Ciardo and R. Siminiceanu. Structural symbolic CTL model checking of asynchronous systems. In W. Hunt, Jr. and F. Somenzi, editors, *Computer Aided Verification (CAV’03)*, LNCS 2725,

pages 40–53, Boulder, CO, USA, July 2003. Springer-Verlag.

- [13] G. Ciardo and A. J. Yu. Saturation-based symbolic reachability analysis using conjunctive and disjunctive partitioning. In D. Borriore and W. Paul, editors, *Proc. CHARME*, LNCS 3725, pages 146–161, Saarbrücken, Germany, Oct. 2005. Springer-Verlag.
- [14] M. Wan and G. Ciardo. Symbolic reachability analysis of integer timed Petri nets. In M. Nielsen et al., editors, *Proc. 35th Int. Conf. Current Trends in Theory and Practice of Computer Science (SOFSEM)*, LNCS 5404, pages 595–608, Špindlerův Mlýn, Czech Republic, Feb. 2009. Springer-Verlag.
- [15] M. Wan and G. Ciardo. Symbolic state-space generation of asynchronous systems using extensible decision diagrams. In M. Nielsen et al., editors, *Proc. 35th Int. Conf. Current Trends in Theory and Practice of Computer Science (SOFSEM)*, LNCS 5404, pages 582–594, Špindlerův Mlýn, Czech Republic, Feb. 2009. Springer-Verlag.