

A Pattern Recognition Approach for Speculative Firing Prediction in Distributed Saturation State-Space Generation

Ming-Ying Chung and Gianfranco Ciardo

Department of Computer Science and Engineering

University of California at Riverside

Riverside, CA 92521, USA

{chung,ciardo}@cs.ucr.edu

- **Introduction**

- *Parallel and distributed state-space generation*

- **Background**

- *Multi-valued Decision Diagrams (MDDs)* encoding of the state-space
- *Kronecker* encoding of the next state function
- *Saturation-style* fixed point iteration and its distributed version

- **Speculative firing prediction**

- Idea and the naïve approach
- The *pattern recognition* approach

- **Experimental results**

- **Conclusions**

Introduction

- *Formal verification* : for *quality assurance*
- *Model checking* : a, model base, automatic verification approach
E. Clarke and E. Emerson. *Synthesis of synchronization skeletons for branching time temporal logic*, Logic of Programs 1981
- *State-space generation* : the first step in model checking (**memory-intensive**)
- *Binary decision diagrams* (BDDs) : *symbolic* state-space construction
R. Bryant, *Graph-based algorithm for boolean function manipulation*, IEEE TC 1986
- *Parallel and distributed model checking* :
uses computation resources on a network or different computer architecture
- *Saturation NOW* : a *message-passing* scheme based on *saturation*
M.-Y. Chung and G. Ciardo, *Saturation NOW*, QEST 2004

- On *shared-memory multi-processor* : **special hardware**
S. Kimura and E. M. Clarke, *A parallel algorithm for constructing BDDs*, ICCD 1990
- On *distributed shared memory* : **special software**
Y. Parasuram, E. Stabler, and S.-K. Chin, *Parallel implementation of BDD algorithm using a DSM*, HICSS 1994
- On *network of workstations* :
 - **Job-based slicing** : **overlap (duplicate) image computation**
T. Stornetta and F. Brewer, *An efficient parallel BDD package*, DAC 1996
K. Milvang-Jensen and A. Hu, *BDDNOW*, FMCAD 1998
T. Heyman, D. Geist, O. Grumberg, T. Heyman, and A. Schuster, *Achieving scalability in parallel reachability analysis of very large circuits*, CAV 2000
O. Grumberg, T. Heyman, and A. Schuster, *A work-efficient distributed algorithm for reachability analysis*, CAV 2003
 - **Level-based slicing** : **scaling problem, no memory load balancing**
R. Ranjan, J. Snaghavi, R. Brayton, and A. Sangiovanni-Vincentelli, *BDDs on NOWs*, ICCD 1996

Background

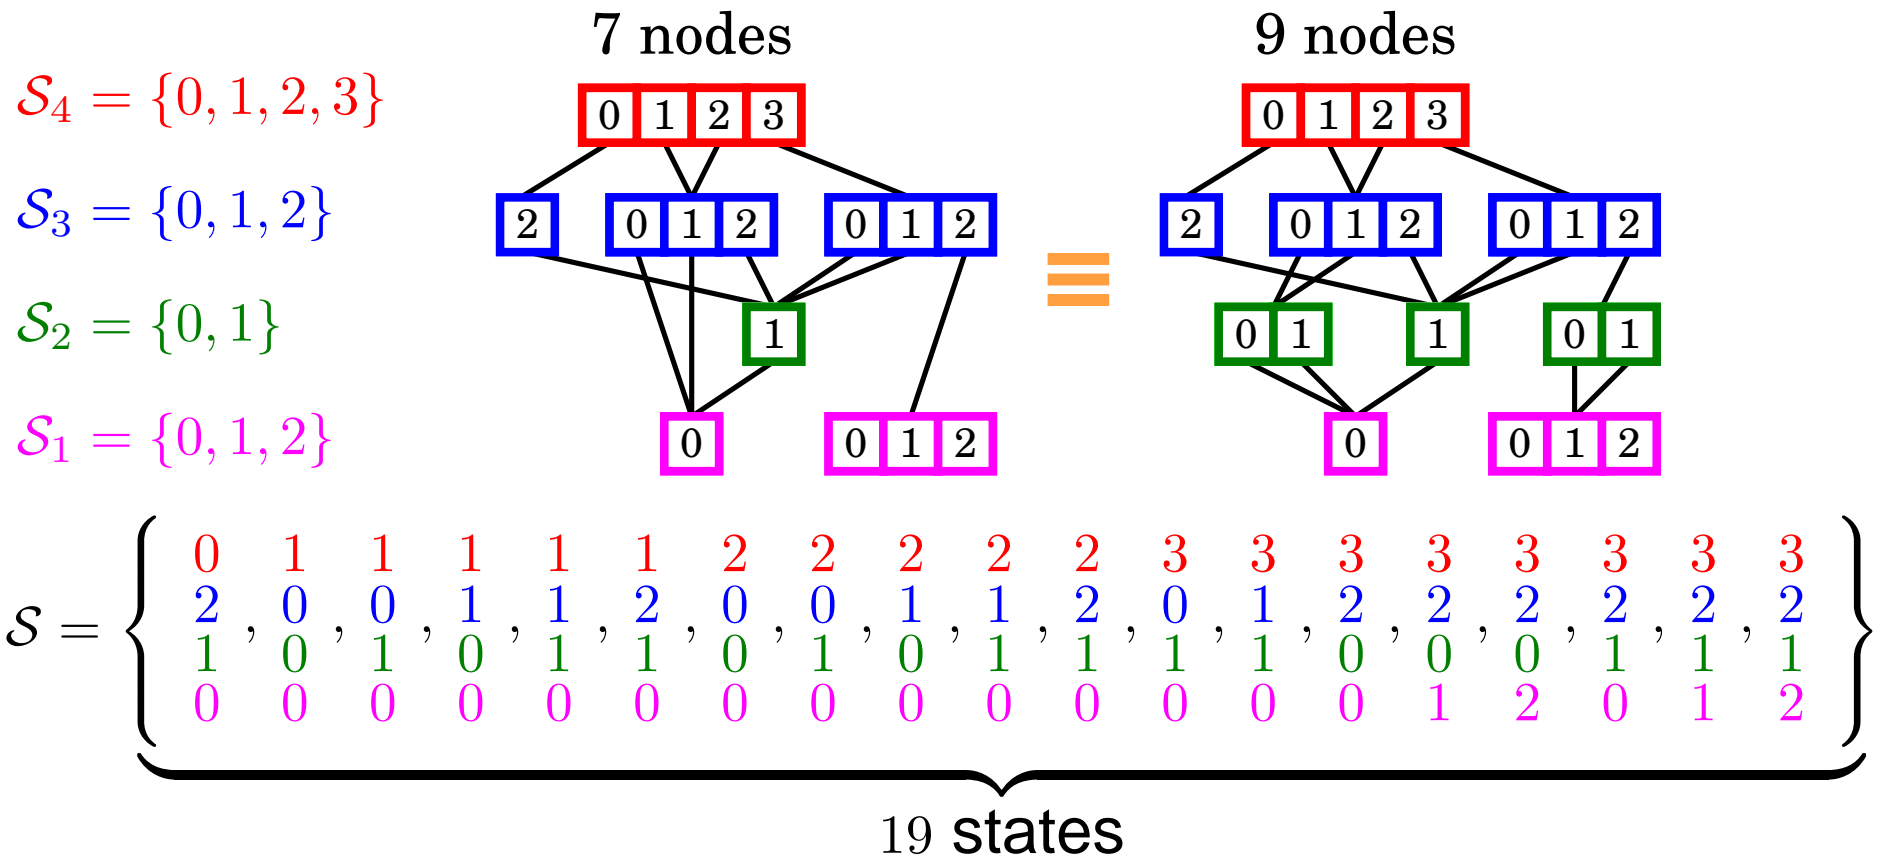
- A *structured discrete state model* is a triple $(\widehat{\mathcal{S}}, \mathcal{S}^{init}, \mathcal{N})$ where
 - $\widehat{\mathcal{S}}$ is the set of *potential states* of the model
 - $\mathcal{S}^{init} \subseteq \widehat{\mathcal{S}}$ is the set of *initial states*
 - $\mathcal{N} : \widehat{\mathcal{S}} \rightarrow 2^{\widehat{\mathcal{S}}}$ is the *next-state* function

- The *reachable state-space* $\mathcal{S} \in \widehat{\mathcal{S}}$ is the smallest set
 - containing \mathcal{S}^{init}
 - closed with respect to \mathcal{N}

$$\mathcal{S} = \mathcal{S}^{init} \cup \mathcal{N}(\mathcal{S}^{init}) \cup \mathcal{N}^2(\mathcal{S}^{init}) \cup \mathcal{N}^3(\mathcal{S}^{init}) \cup \dots = \mathcal{N}^*(\mathcal{S}^{init})$$

- We *structure the states by level* :

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



- We *structure the states by level* :


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

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

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

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

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

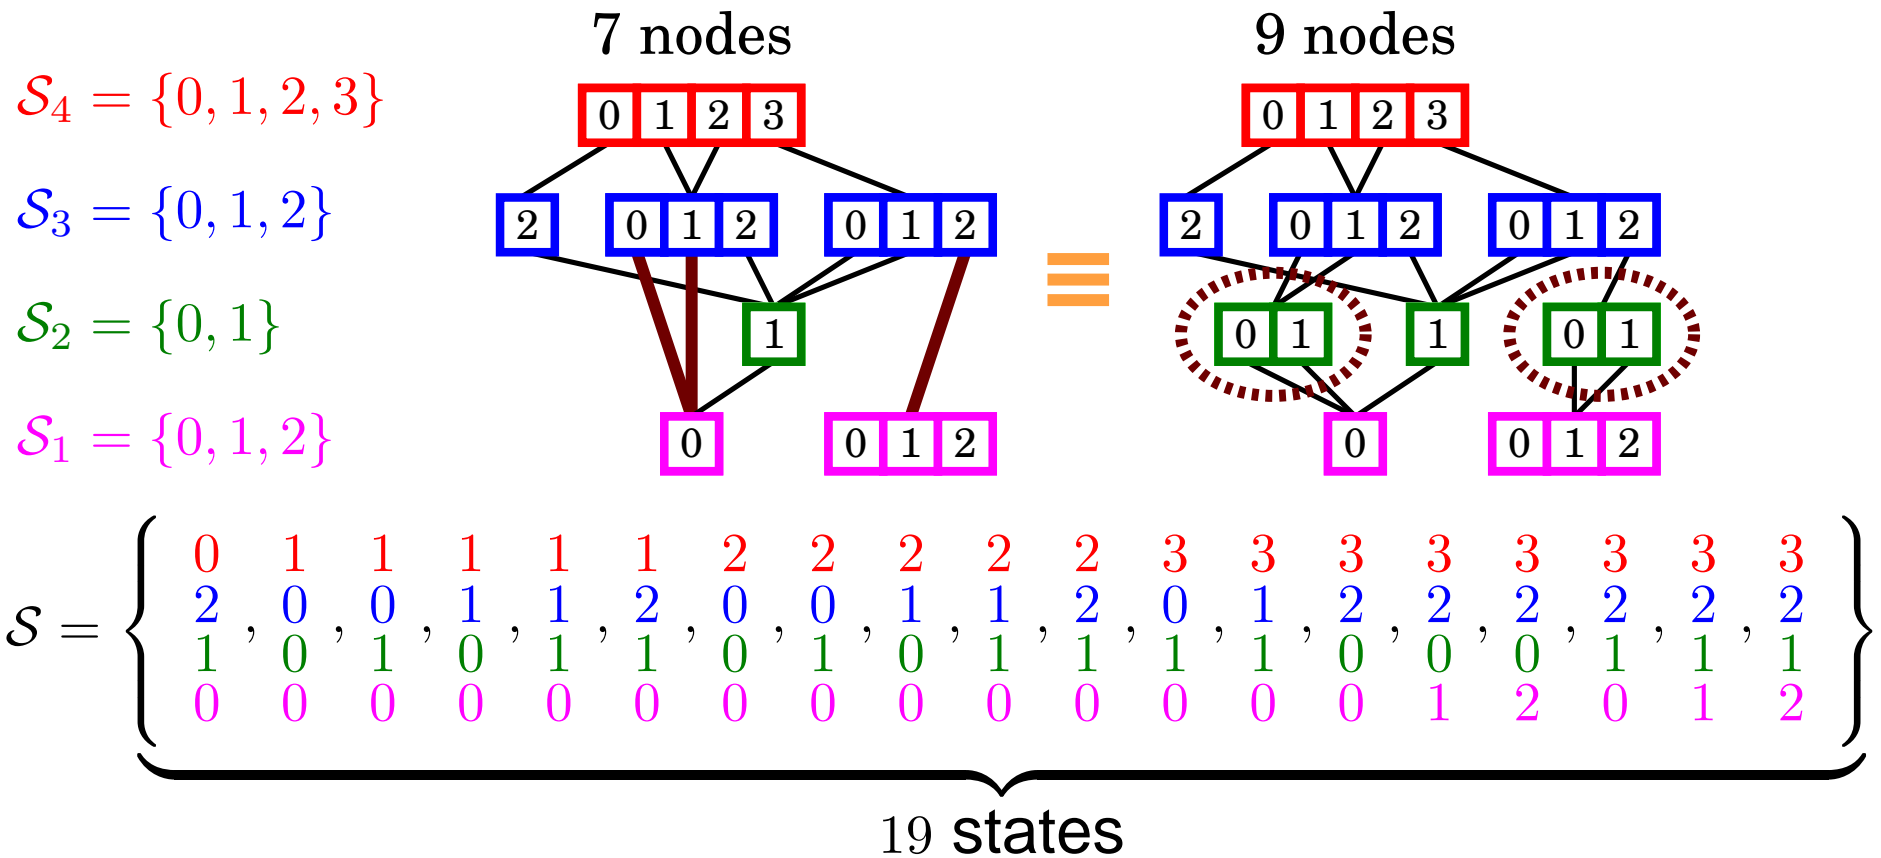


$$\mathcal{S} = \left\{ \begin{array}{cccccccc|c|cccccccc} 0 & 1 & 1 & 1 & 1 & 1 & 2 & 2 & 2 & 2 & 2 & 3 & 3 & 3 & 3 & 3 & 3 & 3 & 3 \\ 2 & 0 & 0 & 1 & 1 & 2 & 0 & 0 & 1 & 1 & 2 & 0 & 1 & 2 & 2 & 2 & 2 & 2 & 2 \\ 1 & 0 & 1 & 0 & 1 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 2 & 0 & 1 & 2 \end{array} \right\}$$

19 states

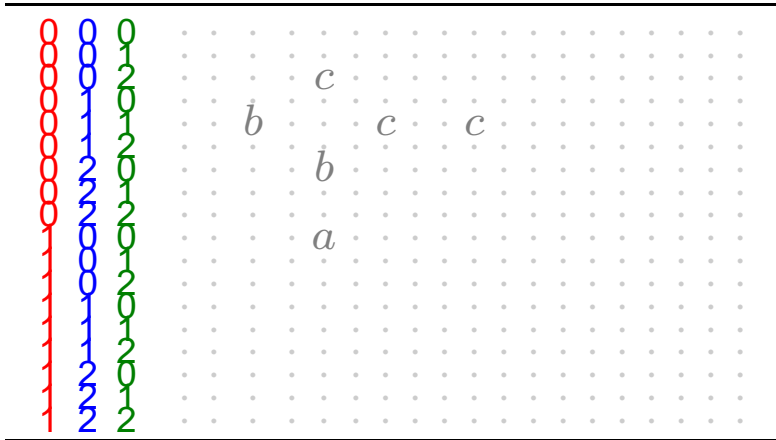
- We *structure the states by level* :

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



$$\mathbf{N} = \sum_{e \in \mathcal{E}} \mathbf{N}_e = \sum_{e \in \mathcal{E}} \bigotimes_{K \geq k \geq 1} \mathbf{N}_{k,e}$$

00 0 0 0 0 0 0 0 1 1 1 1 1 1 1 1 1
 00 0 1 1 1 2 2 2 0 0 0 0 1 1 1 2 2 2
 0 1 2 0 1 2 0 1 2 0 1 2 0 1 2 0 1 2



	a	b	c	d
x	$1 \rightarrow 0$	\mathbf{I}	\mathbf{I}	$0 \rightarrow 1$
y	$0 \rightarrow 1$	$1 \rightarrow 0$ $2 \rightarrow 1$	$0 \rightarrow 1$ $1 \rightarrow 2$	$1 \rightarrow 0$
z	$0 \rightarrow 1$	$0 \rightarrow 1$ $1 \rightarrow 2$	$2 \rightarrow 1$ $1 \rightarrow 0$	$1 \rightarrow 0$

$Top(\alpha)$: the highest MDD level affected by event α


$$Top(a) = Top(d) = x \quad Top(b) = Top(c) = y$$

$$\mathbb{N} = \sum_{e \in \mathcal{E}} \mathbb{N}_e = \sum_{e \in \mathcal{E}} \bigotimes_{K \geq k \geq 1} \mathbb{N}_{k,e}$$

00 0 0000001111111111

00 0 1112220000111222

01 2 0120120120012012



	a	b	c	d
x	$1 \rightarrow 0$	I	I	$0 \rightarrow 1$
y	$0 \rightarrow 1$	<div style="display: flex; flex-direction: column; align-items: center;"> <div style="border: 1px solid blue; padding: 2px;">1</div> \rightarrow <div style="border: 1px solid blue; padding: 2px;">0</div> </div> <div style="display: flex; flex-direction: column; align-items: center; margin-top: 5px;"> <div style="color: blue;">2</div> \rightarrow <div style="color: blue;">1</div> </div>	<div style="display: flex; flex-direction: column; align-items: center;"> <div style="color: blue;">$0 \rightarrow 1$</div> <div style="color: blue;">$1 \rightarrow 2$</div> </div>	$1 \rightarrow 0$
z	$0 \rightarrow 1$	<div style="display: flex; flex-direction: column; align-items: center;"> <div style="color: green;">$0 \rightarrow 1$</div> <div style="border: 1px solid green; padding: 2px;">1</div> \rightarrow <div style="border: 1px solid green; padding: 2px;">2</div> </div>	<div style="display: flex; flex-direction: column; align-items: center;"> <div style="color: green;">$2 \rightarrow 1$</div> <div style="color: green;">$1 \rightarrow 0$</div> </div>	$1 \rightarrow 0$


$Top(\alpha)$: the highest MDD level affected by event α

$$Top(a) = Top(d) = x \qquad Top(b) = Top(c) = y$$

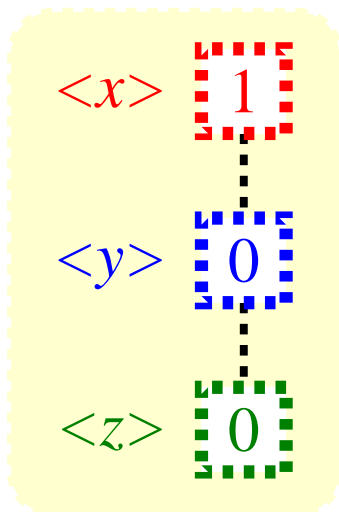
An MDD node at level k is **saturated** if the state set it encodes is a fixed point with respect to any event α such that $Top(\alpha) \leq k$ (α is independent of any MDD level higher than k)

- build the MDD encoding of initial state set
- From $k = 1$ to K ,
 - saturate each node at level k by firing all events α s.t. $Top(\alpha) = k$
(if this creates nodes below level k , saturate them **immediately upon creation**)
- When the root node is saturated, all reachable states have been discovered

Saturation has Enormous time and memory efficiency compared to traditional breath-first iteration for asynchronous models



	a	b	c	d
x	$1 \rightarrow 0$	I	I	$0 \rightarrow 1$
y	$0 \rightarrow 1$	$1 \rightarrow 0$ $2 \rightarrow 1$	$0 \rightarrow 1$ $1 \rightarrow 2$	$1 \rightarrow 0$
z	$0 \rightarrow 1$	$0 \rightarrow 1$ $1 \rightarrow 2$	$2 \rightarrow 1$ $1 \rightarrow 0$	$1 \rightarrow 0$




$$\mathcal{S}_p = \{p^0, p^1\} \equiv \{0, 1\}$$

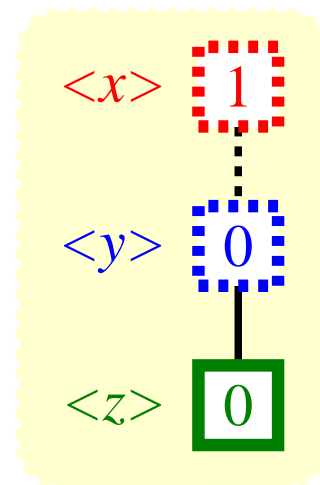

$$\mathcal{S}_q = \{q^0, q^1, q^2\} \equiv \{0, 1, 2\}$$

$$\mathcal{S}_r = \{r^0, r^1, r^2\} \equiv \{0, 1, 2\}$$


Saturate node **0** at level $\langle z \rangle$





	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
<i>x</i>	$1 \rightarrow 0$	I	I	$0 \rightarrow 1$
<i>y</i>	$0 \rightarrow 1$	$1 \rightarrow 0$ $2 \rightarrow 1$	$0 \rightarrow 1$ $1 \rightarrow 2$	$1 \rightarrow 0$
<i>z</i>	$0 \rightarrow 1$	$0 \rightarrow 1$ $1 \rightarrow 2$	$2 \rightarrow 1$ $1 \rightarrow 0$	$1 \rightarrow 0$




Saturate node $\boxed{0}$ at level $\langle y \rangle$





	a	\boxed{b}	\boxed{c}	d
x	$1 \rightarrow 0$	I	I	$0 \rightarrow 1$
y	$0 \rightarrow 1$	$1 \rightarrow 0$ $2 \rightarrow 1$	$\underline{0} \rightarrow \boxed{1}$ $1 \rightarrow 2$	$1 \rightarrow 0$
z	$0 \rightarrow 1$	$\underline{0} \rightarrow \boxed{1}$ $1 \rightarrow 2$	$2 \rightarrow 1$ $1 \rightarrow 0$	$1 \rightarrow 0$




Saturate node **1** at level $\langle x \rangle$



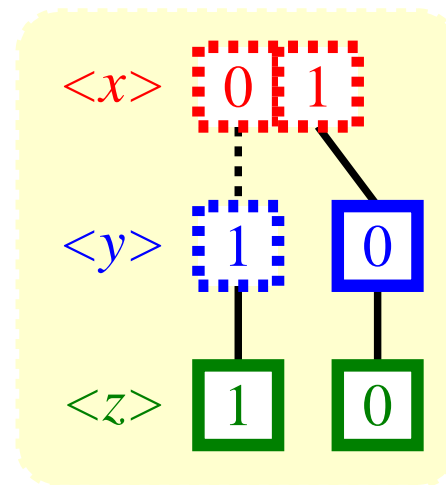

	a	<i>b</i>	<i>c</i>	<i>d</i>
<i>x</i>	<u>1</u> → 0	I	I	<u>0</u> → 1
<i>y</i>	<u>0</u> → 1	1 → 0 2 → 1	0 → 1 1 → 2	1 → 0
<i>z</i>	<u>0</u> → 1	0 → 1 1 → 2	2 → 1 1 → 0	1 → 0




Saturate node **1** at level $\langle z \rangle$





	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
<i>x</i>	1 → 0	I	I	0 → 1
<i>y</i>	0 → 1	1 → 0 2 → 1	0 → 1 1 → 2	1 → 0
<i>z</i>	0 → 1	0 → 1 1 → 2	2 → 1 1 → 0	1 → 0




Saturate node **1** at level $\langle y \rangle$



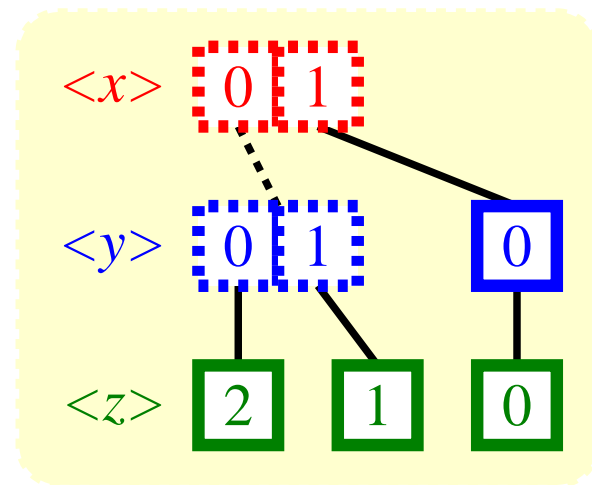

	<i>a</i>	1	<i>c</i>	<i>d</i>
<i>x</i>	1 → 0	I	I	0 → 1
<i>y</i>	0 → 1	1 → 0 2 → 1	0 → 1 1 → 2	1 → 0
<i>z</i>	0 → 1	0 → 1 1 → 2	2 → 1 1 → 0	1 → 0




Saturate node **2** at level $\langle z \rangle$




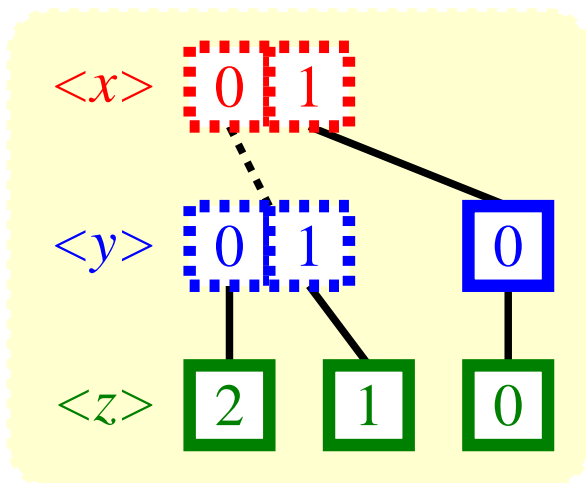
	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
<i>x</i>	1 → 0	I	I	0 → 1
<i>y</i>	0 → 1	1 → 0 2 → 1	0 → 1 1 → 2	1 → 0
<i>z</i>	0 → 1	0 → 1 1 → 2	2 → 1 1 → 0	1 → 0




Saturate node **0 1** at level $\langle y \rangle$





	<i>a</i>	<i>b</i>	c	<i>d</i>
<i>x</i>	1 → 0	I	I	0 → 1
<i>y</i>	0 → 1	1 → 0 2 → 1	0 → 1 <u>1</u> → 2	1 → 0
<i>z</i>	0 → 1	0 → 1 1 → 2	2 → 1 <u>1</u> → 0	1 → 0




Saturate node **0 1 2** at level $\langle y \rangle$





	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
<i>x</i>	1 → 0	I	I	0 → 1
<i>y</i>	0 → 1	1 → 0 2 → 1	0 → 1 1 → 2	1 → 0
<i>z</i>	0 → 1	0 → 1 1 → 2	2 → 1 1 → 0	1 → 0



Saturate node **0 1** at level $\langle x \rangle$



	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
<i>x</i>	1 → 0	I	I	0 → 1
<i>y</i>	0 → 1	1 → 0 2 → 1	0 → 1 1 → 2	1 → 0
<i>z</i>	0 → 1	0 → 1 1 → 2	2 → 1 1 → 0	1 → 0



- **Problem :**

We target relatively large models where state-space generation with a single workstation must rely on virtual memory or runs out of memory

- **Solution :**

- Apply the saturation algorithm on a network of workstations to increase the available memory
- *Horizontal MDDs slicing for distributed state-space construction :*
 - no duplicate MDDs
 - no global synchronization
- *Nested approach of dynamic memory load balancing scheme :*
 - pairwise communication to achieve nearly ideal memory scalability


Pattern Recognition Approach of Speculative Firing Prediction

- **Problem :**

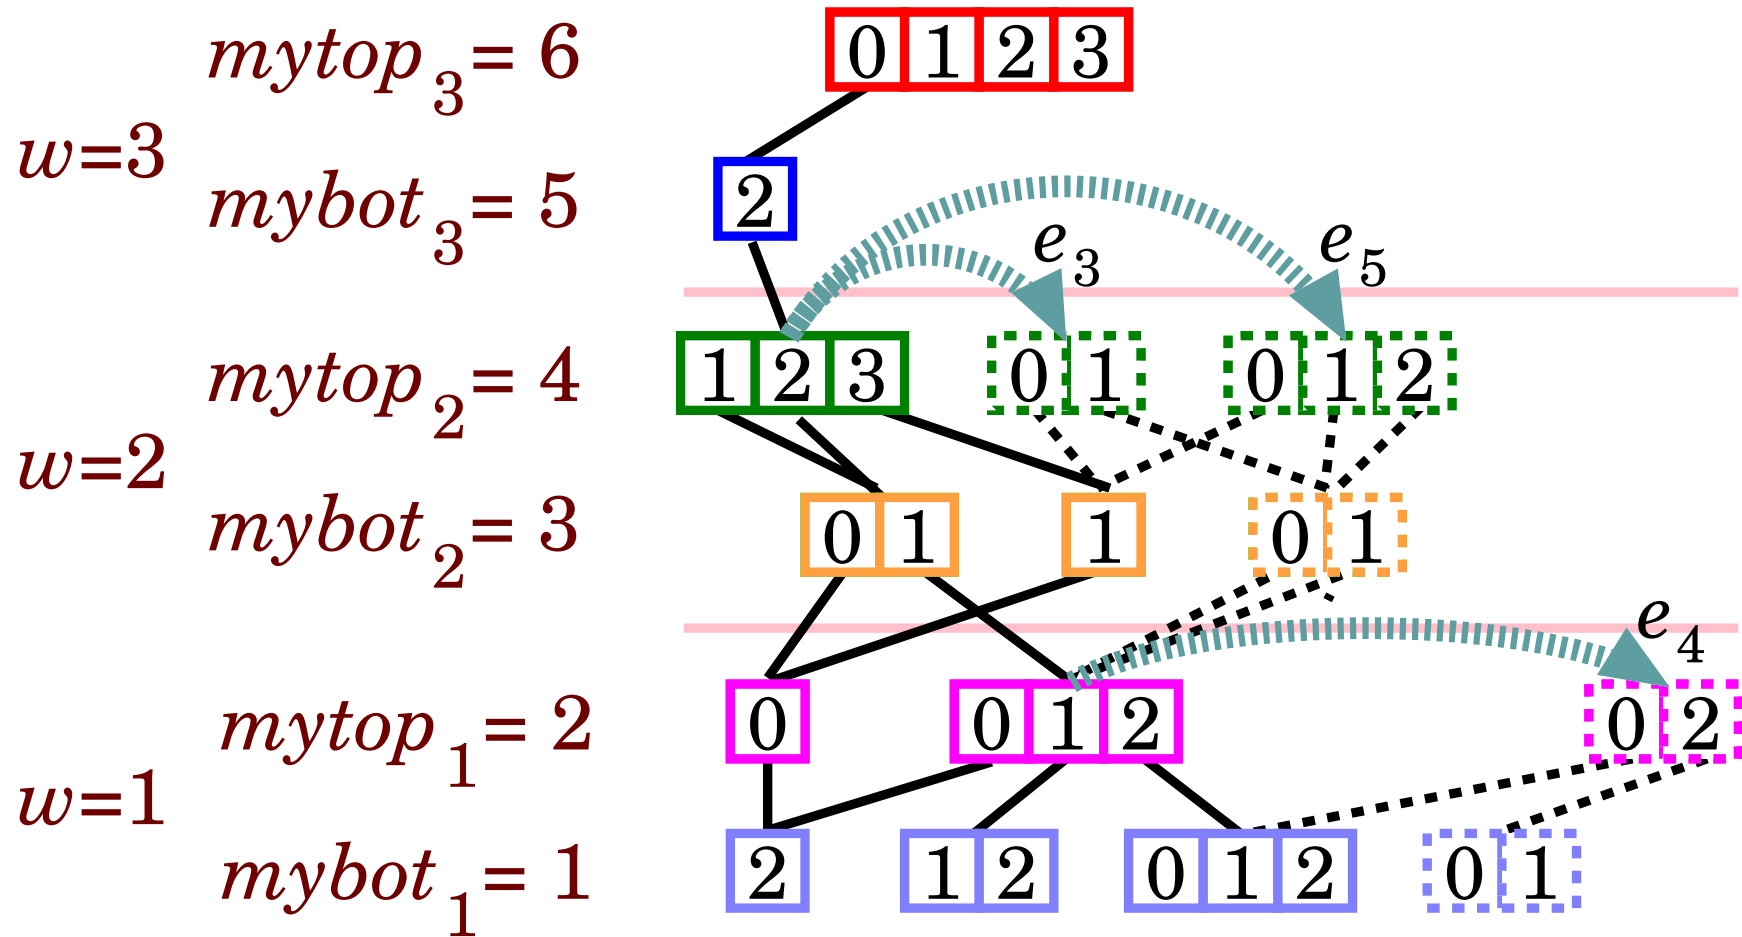
- *Only one workstation is active* at any time :
no theoretical speedup is achieved
- The horizontal slicing scheme *sequentializes the distributed computation* :
parallelization becomes a challenge now

- **Solution :**

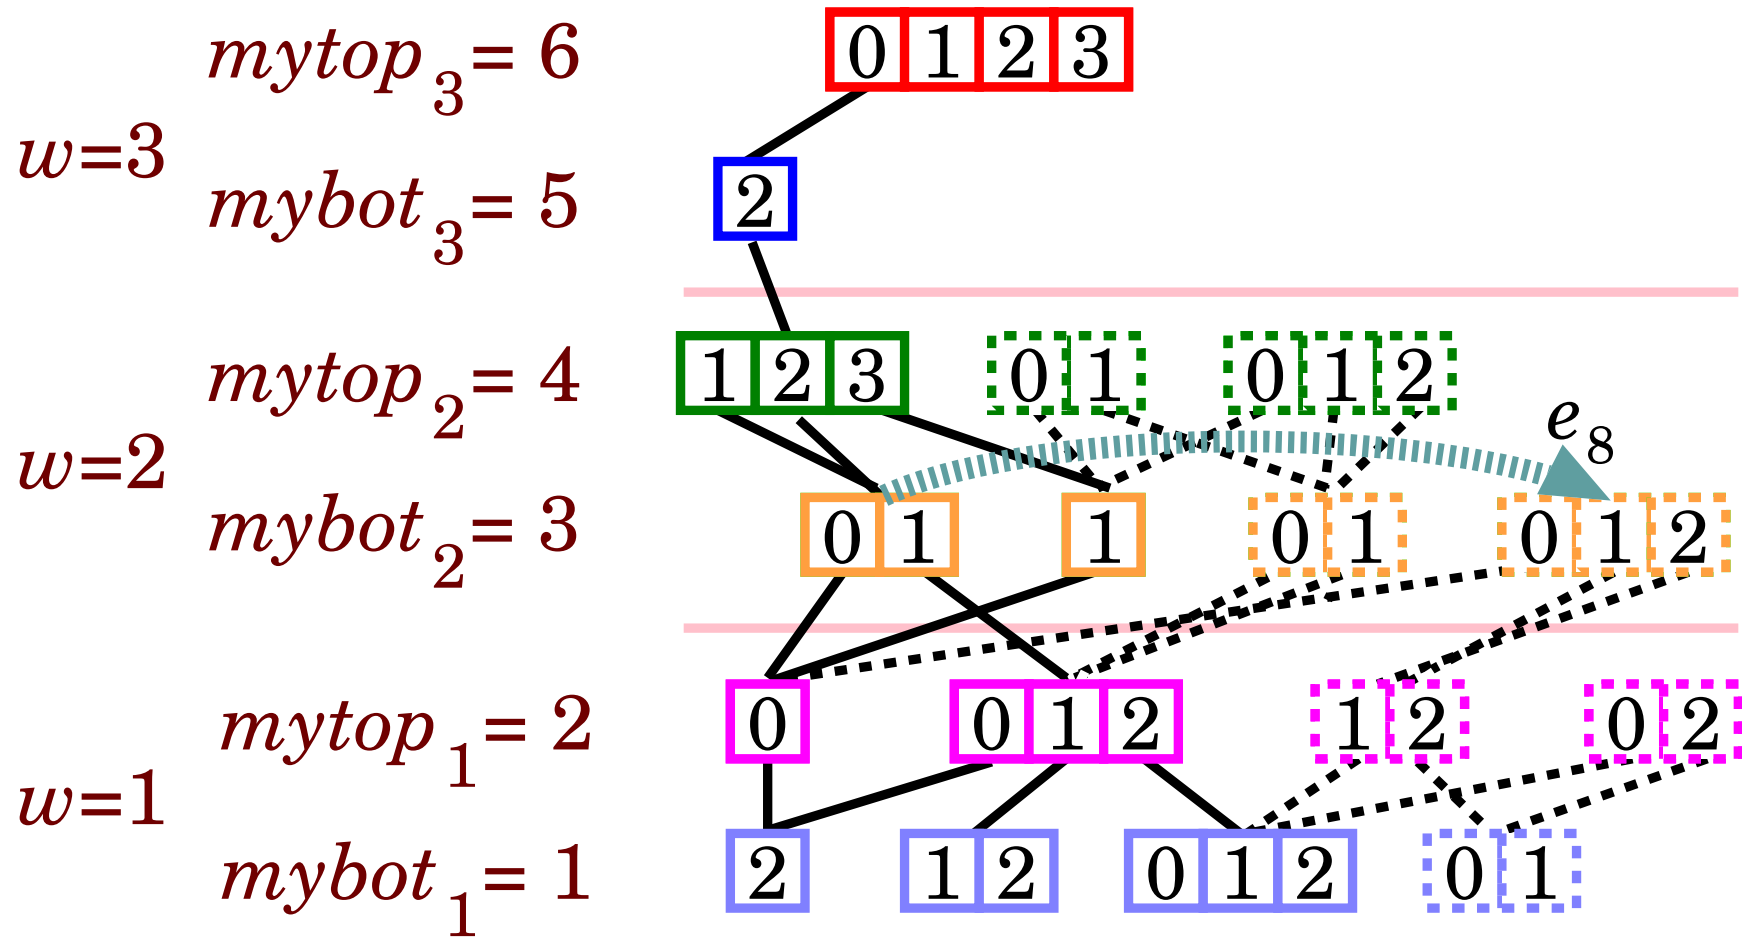
- *Idle workstations perform some event firings in advance* and cache the results of those firings
- Some later firing requests might retrieve those results from caches and be resolved quickly




A runtime snapshot




w_2 and w_1 perform speculative event firings individually




w_2 and w_1 perform some speculative event firing together




w_3 retrieves event firing result from operation cache




w_3 makes event firing requests remotely



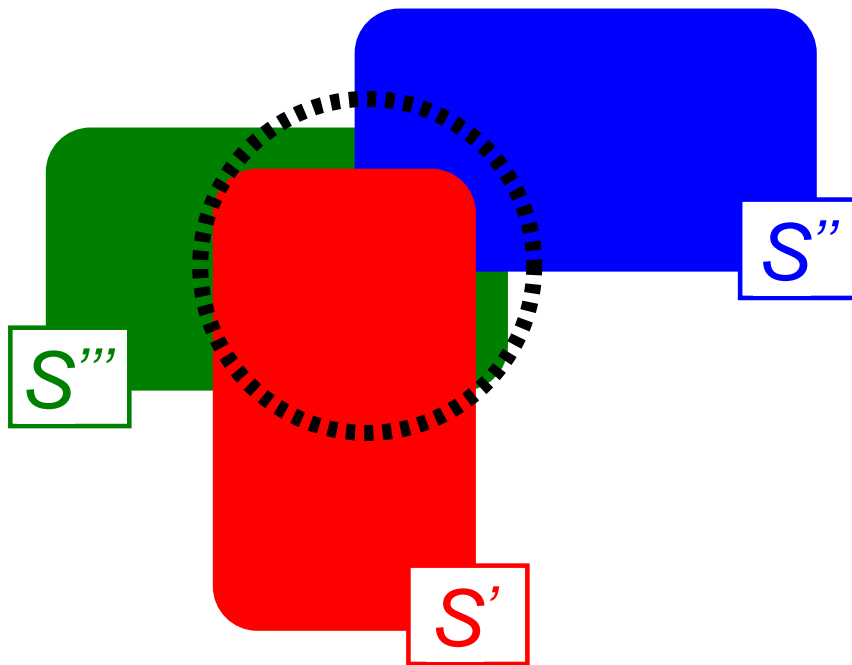
- Space for storing the disconnected MDDs
- Space for storing the additional caches entries



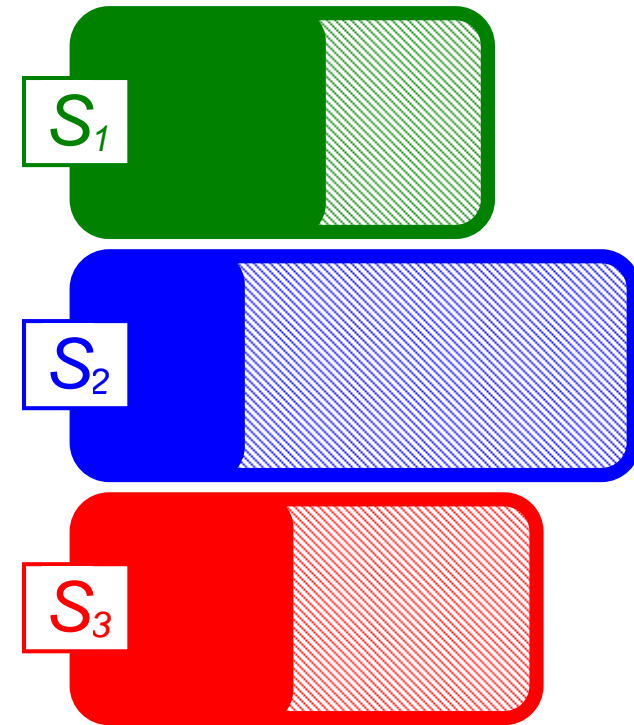
Overlapped BDDs
(Duplicate)



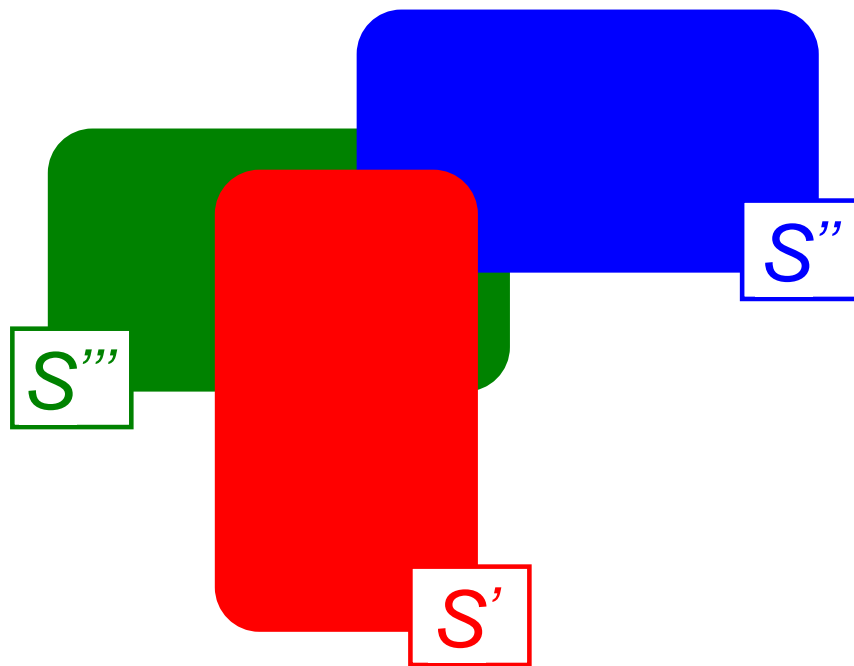
Disconnected MDDs
(Useless)



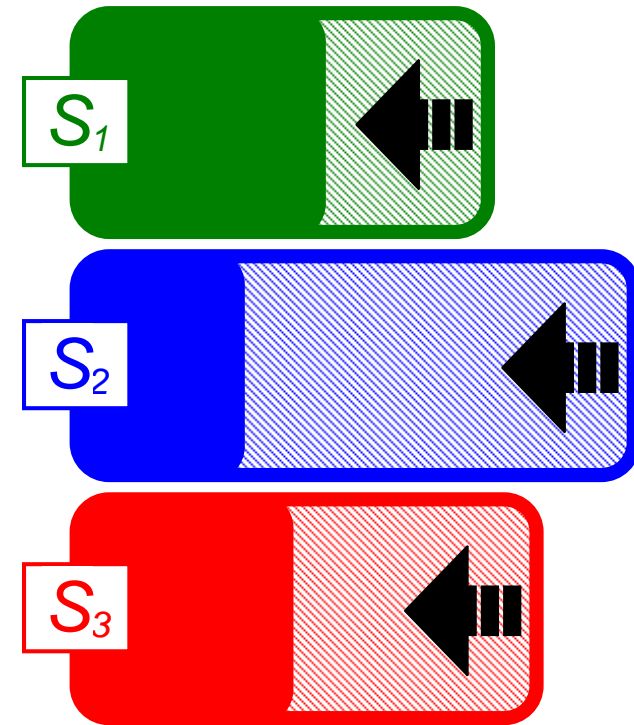
**Overlapped BDDs
(Duplicate)**



**Disconnected MDDs
(Useless)**



Overlapped BDDs
(Duplicate)



Disconnected MDDs
(Useless)

- **Problem :**

- Do not know a priori whether an event will be fired on some MDD node
- A NAÏVE approach : *idle workstations exhaust all possible firings*

$$\mathcal{E}_{all}(p) = \{e : Top(p.lvl) > k \geq Bot(e)\} \text{ for each node } p$$

- It might create many useless nodes and firing cache entries
- It will not stop when the prediction does not help much


- **Observation :**

- $\mathcal{E}_{patt}(p) = \{e : \text{event } e \text{ need to be fired on the node } p \text{ to saturate } p\}$
- There are always some MDD nodes which have similar or even the same set of events fired on.

- **Solution :**

- A more informed firing prediction based on *firing patterns*

$\alpha\beta\gamma\delta$ $\alpha\beta\gamma$ $\alpha\beta\gamma\delta$ $\alpha\beta\gamma\delta$ $\alpha\gamma$ $\gamma\delta$ $\alpha\beta\gamma\delta$ $\alpha\beta\gamma\delta$ $\alpha\beta\gamma\delta$




Final

Snapshot

HIST Final

Requests


(stack)




FirePredict(*Requests* : stack, *Class* : array)

```

while Requests ≠ ∅ do
  (e, p) ← Pop(Requests);
  Enqueue(e,  $\mathcal{E}_{hist}(p)$ );
  q ← Classp.lvl[e];
  if  $\mathcal{E}_{hist}(p) \supset \mathcal{E}_{hist}(q)$  then
    Classp.lvl[e] ← p;
    fire e on q and cache the result;
  else if  $\mathcal{E}_{hist}(p) \subset \mathcal{E}_{hist}(q)$  then
    foreach e' ∈  $\mathcal{E}_{hist}(q) \setminus \mathcal{E}_{hist}(p)$  do
      fire e' on p and cache the result;
  
```




Pop one element of *Requests* in each iteration




FirePredict(*Requests* : stack, *Class* : array)

```

while Requests ≠ ∅ do
  (e, p) ← Pop(Requests);
  Enqueue(e,  $\mathcal{E}_{hist}(p)$ );
  q ← Classp.lvl[e];
  if  $\mathcal{E}_{hist}(p) \supset \mathcal{E}_{hist}(q)$  then
    Classp.lvl[e] ← p;
    fire e on q and cache the result;
  else if  $\mathcal{E}_{hist}(p) \subset \mathcal{E}_{hist}(q)$  then
    foreach e' ∈  $\mathcal{E}_{hist}(q) \setminus \mathcal{E}_{hist}(p)$  do
      fire e' on p and cache the result;
    
```




For each *k*, *Class*_{*k*}[*e*] indicate the representative node w.r.t. *e*




```

FirePredict(Requests : stack, Class : array)
while Requests ≠ ∅ do
    (e, p) ← Pop(Requests);
    Enqueue(e,  $\mathcal{E}_{hist}(p)$ );
    q ←  $Class_{p.lvl}[e]$ ;
    if  $\mathcal{E}_{hist}(p) \supset \mathcal{E}_{hist}(q)$  then
         $Class_{p.lvl}[e] \leftarrow p$ ;
        fire e on q and cache the result;
    else if  $\mathcal{E}_{hist}(p) \subset \mathcal{E}_{hist}(q)$  then
        foreach  $e' \in \mathcal{E}_{hist}(q) \setminus \mathcal{E}_{hist}(p)$  do
            fire  $e'$  on p and cache the result;
    
```




Pop (δ, p) and enqueue δ into $\mathcal{E}_{hist}(p)$




FirePredict(Requests : stack, Class : array)

```

while Requests  $\neq$   $\emptyset$  do
    (e, p)  $\leftarrow$  Pop(Requests);
    Enqueue(e,  $\mathcal{E}_{hist}(p)$ );
    q  $\leftarrow$  Classp.lvl[e];
    if  $\mathcal{E}_{hist}(p) \supset \mathcal{E}_{hist}(q)$  then
        Classp.lvl[e]  $\leftarrow$  p;
        fire e on q and cache the result;
    else if  $\mathcal{E}_{hist}(p) \subset \mathcal{E}_{hist}(q)$  then
        foreach  $e' \in \mathcal{E}_{hist}(q) \setminus \mathcal{E}_{hist}(p)$  do
            fire  $e'$  on p and cache the result;
    
```




Set p as the new representative node of δ




FirePredict(Requests : stack, Class : array)

```

while Requests  $\neq$   $\emptyset$  do
  (e, p)  $\leftarrow$  Pop(Requests);
  Enqueue(e,  $\mathcal{E}_{hist}(p)$ );
  q  $\leftarrow$  Classp.lvl[e];
  if  $\mathcal{E}_{hist}(p) \supset \mathcal{E}_{hist}(q)$  then
    Classp.lvl[e]  $\leftarrow$  p;
    fire e on q and cache the result;
  else if  $\mathcal{E}_{hist}(p) \subset \mathcal{E}_{hist}(q)$  then
    foreach  $e' \in \mathcal{E}_{hist}(q) \setminus \mathcal{E}_{hist}(p)$  do
      fire  $e'$  on p and cache the result;
  
```



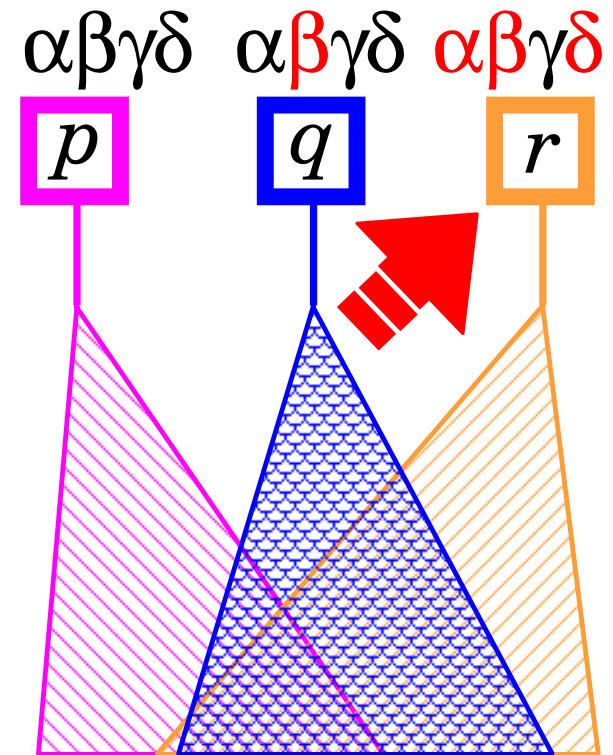
Speculatively fire β on q




FirePredict(Requests : stack, Class : array)

```

while Requests ≠ ∅ do
  (e, p) ← Pop(Requests);
  Enqueue(e,  $\mathcal{E}_{hist}(p)$ );
  q ←  $Class_{p.lvl}[e]$ ;
  if  $\mathcal{E}_{hist}(p) \supset \mathcal{E}_{hist}(q)$  then
     $Class_{p.lvl}[e] \leftarrow p$ ;
    fire e on q and cache the result;
  else if  $\mathcal{E}_{hist}(p) \subset \mathcal{E}_{hist}(q)$  then
    foreach  $e' \in \mathcal{E}_{hist}(q) \setminus \mathcal{E}_{hist}(p)$  do
      fire  $e'$  on p and cache the result;
  
```



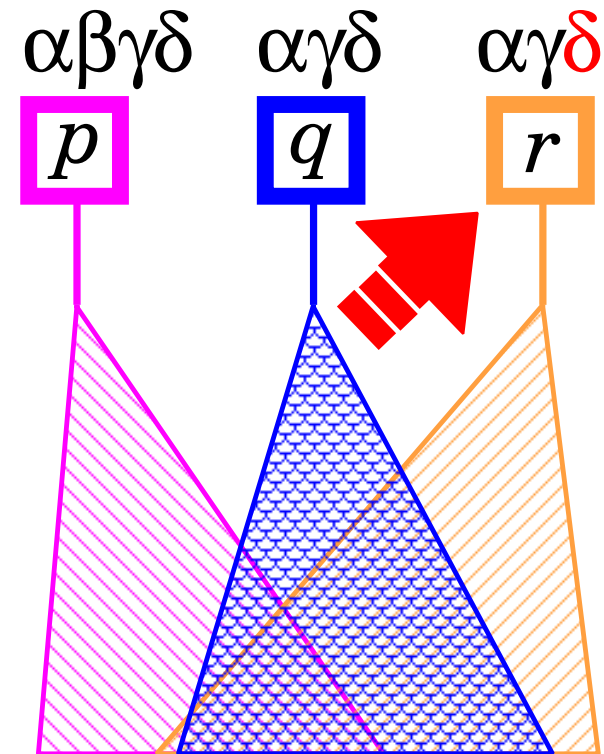
Aggressively update and predict the firings of α , β , and δ on r




FirePredict(Requests : stack, Class : array)

```

while Requests ≠ ∅ do
  (e, p) ← Pop(Requests);
  Enqueue(e,  $\mathcal{E}_{hist}(p)$ );
  q ←  $Class_{p.lvl}[e]$ ;
  if  $\mathcal{E}_{hist}(p) \supset \mathcal{E}_{hist}(q)$  then
     $Class_{p.lvl}[e] \leftarrow p$ ;
    fire e on q and cache the result;
  else if  $\mathcal{E}_{hist}(p) \subset \mathcal{E}_{hist}(q)$  then
    foreach  $e' \in \mathcal{E}_{hist}(q) \setminus \mathcal{E}_{hist}(p)$  do
      fire  $e'$  on p and cache the result;
  
```




Pop (α, r), enqueue α , and speculatively fire δ on r



```

FirePredict(Requests : stack, Class : array)
while Requests ≠ ∅ do
  (e, p) ← Pop(Requests);
  Enqueue(e, E_hist(p));
  q ← Class_{p.lvl}[e];
  if E_hist(p) ⊃ E_hist(q) then
    Class_{p.lvl}[e] ← p;
    fire e on q and cache the result;
  else if E_hist(p) ⊂ E_hist(q) then
    foreach e' ∈ E_hist(q) \ E_hist(p) do
      fire e' on p and cache the result;
  
```



Pop (β, r), enqueue β , and speculatively fire α and δ on r

Experimental Results

- Protocol for local area networks : N is the number of nodes within the network ($K = N$, $|\mathcal{S}_k| = 10$ for all k , $|\mathcal{E}| = 3N$)

E. Pastor, O. Roig, J. Cortadella, and R. Badia, *Petri net analysis using boolean manipulation*, ICATPN 1994

W	Time (sec)			Total Memory (MB)			Max Memory (MB)		
	SMART _N ^{OW}	NAÏVE	HIST	SMART _N ^{OW}	NAÏVE	HIST	SMART _N ^{OW}	NAÏVE	HIST
$N = 200$ $ \mathcal{S} = 8.38 \cdot 10^{211}$ SMART completes in 108 sec using 284MB									
2	119	-24%	-13%	286	+3%	+45%	197	+1%	+53%
4	139	-27%	-15%	286	+11%	+51%	127	+61%	+58%
8	182	-32%	-24%	286	+129%	+62%	69	+239%	+62%
$N = 300$ $ \mathcal{S} = 8.38 \cdot 10^{211}$ SMART does not complete in 5 hrs using 512MB									
2	^s 552	^s +5%	^s -5%	962	+25%	+11%	562	+8%	+7%
4	^d 490	> 5hrs	^d -16%	962	-	+34%	352	-	+12%
8	564	> 5hrs	-39%	962	-	+50%	252	-	+23%

- Both NAÏVE and HIST approaches outperform SMART_N^{OW}

- Protocol for local area networks : N is the number of nodes within the network ($K = N$, $|\mathcal{S}_k| = 10$ for all k , $|\mathcal{E}| = 3N$)

E. Pastor, O. Roig, J. Cortadella, and R. Badia, *Petri net analysis using boolean manipulation*, ICATPN 1994

W	Time (sec)			Total Memory (MB)			Max Memory (MB)		
	SMART _N ^{OW}	NAÏVE	HIST	SMART _N ^{OW}	NAÏVE	HIST	SMART _N ^{OW}	NAÏVE	HIST
$N = 200$ $ \mathcal{S} = 8.38 \cdot 10^{211}$ SMART completes in 108 sec using 284MB									
2	119	-24%	-13%	286	+3%	+45%	197	+1%	+53%
4	139	-27%	-15%	286	+11%	+51%	127	+61%	+58%
8	182	-32%	-24%	286	+129%	+62%	69	+239%	+62%
$N = 300$ $ \mathcal{S} = 8.38 \cdot 10^{211}$ SMART does not complete in 5 hrs using 512MB									
2	^s 552	^s +5%	^s -5%	962	+25%	+11%	562	+8%	+7%
4	^d 490	> 5hrs	^d -16%	962	-	+34%	352	-	+12%
8	564	> 5hrs	-39%	962	-	+50%	252	-	+23%

- Both NAÏVE and HIST approaches outperform SMART in some cases

- Protocol for local area networks : N is the number of nodes within the network
($K = N$, $|\mathcal{S}_k| = 10$ for all k , $|\mathcal{E}| = 3N$)

E. Pastor, O. Roig, J. Cortadella, and R. Badia, *Petri net analysis using boolean manipulation*, ICATPN 1994

W	Time (sec)			Total Memory (MB)			Max Memory (MB)		
	SMART _N ^{OW}	NAÏVE	HIST	SMART _N ^{OW}	NAÏVE	HIST	SMART _N ^{OW}	NAÏVE	HIST
$N = 200$ $ \mathcal{S} = 8.38 \cdot 10^{211}$ SMART completes in 108 sec using 284MB									
2	119	-24%	-13%	286	+3%	+45%	197	+1%	+53%
4	139	-27%	-15%	286	+11%	+51%	127	+61%	+58%
8	182	-32%	-24%	286	+129%	+62%	69	+239%	+62%
$N = 300$ $ \mathcal{S} = 8.38 \cdot 10^{211}$ SMART does not complete in 5 hrs using 512MB									
2	^s 552	^s +5%	^s -5%	962	+25%	+11%	562	+8%	+7%
4	^d 490	> 5hrs	^d -16%	962	-	+34%	352	-	+12%
8	564	> 5hrs	-39%	962	-	+50%	252	-	+23%

- The memory overhead is not trivial
- Overall idle time increases \Rightarrow overall memory consumption increases

- Protocol for local area networks : N is the number of nodes within the network ($K = N$, $|\mathcal{S}_k| = 10$ for all k , $|\mathcal{E}| = 3N$)

E. Pastor, O. Roig, J. Cortadella, and R. Badia, *Petri net analysis using boolean manipulation*, ICATPN 1994

W	Time (sec)			Total Memory (MB)			Max Memory (MB)		
	$\text{SMART}_N^{\text{OW}}$	NAÏVE	HIST	$\text{SMART}_N^{\text{OW}}$	NAÏVE	HIST	$\text{SMART}_N^{\text{OW}}$	NAÏVE	HIST
$N = 200$ $ \mathcal{S} = 8.38 \cdot 10^{211}$ SMART completes in 108 sec using 284MB									
2	119	-24%	-13%	286	+3%	+45%	197	+1%	+53%
4	139	-27%	-15%	286	+11%	+51%	127	+61%	+58%
8	182	-32%	-24%	286	+129%	+62%	69	+239%	+62%
$N = 300$ $ \mathcal{S} = 8.38 \cdot 10^{211}$ SMART does not complete in 5 hrs using 512MB									
2	^s 552	^s +5%	^s -5%	962	+25%	+11%	562	+8%	+7%
4	^d 490	> 5hrs	^d -16%	962	-	+34%	352	-	+12%
8	564	> 5hrs	-39%	962	-	+50%	252	-	+23%

- The excessive memory requirement could force to trigger the dynamic load balancing or memory swapping earlier
- The HIST approach releases the heavy memory overhead in the NAÏVE one

- System with three machines to process three different types of parts : N is the number of each type of parts

($K = 19$, $|\mathcal{S}_k| = N + 1$ for all k except $|\mathcal{S}_{17}| = 4$, $|\mathcal{S}_{12}| = 3$, and $|\mathcal{S}_7| = 2$, $|\mathcal{E}| = 20$)

A. Miner and G. Ciardo, *Efficient reachability set generation and storage using DDs*, ICATPN 1999

W	Time (sec)			Total Memory (MB)			Max Memory (MB)		
	SMART _N ^{OW}	NAÏVE	HIST	SMART _N ^{OW}	NAÏVE	HIST	SMART _N ^{OW}	NAÏVE	HIST
$N = 300 \quad \mathcal{S} = 3.64 \cdot 10^{27}$ SMART completes in 55 sec using 241MB									
2	79	-8%	-8%	243	+12%	+24%	121	+26%	+52%
4	91	^d +67%	-9%	243	+102%	+30%	119	+205%	+50%
8	260	> 5hrs	-30%	243	-	+42%	103	-	+47%
$N = 450 \quad \mathcal{S} = 6.90 \cdot 10^{29}$ SMART does not complete in 5 hrs using 512MB									
2	^s 257	^s +12%	^s -14%	826	+16%	+5%	512	+15%	+7%
4	^d 311	> 5hrs	^d -18%	826	-	+33%	372	-	+6%
8	959	> 5hrs	-25%	826	-	+61%	343	-	+6%

- It is a different type of model in comparison with the previous model

- System with three machines to process three different types of parts : N is the number of each type of parts

($K = 19$, $|\mathcal{S}_k| = N + 1$ for all k except $|\mathcal{S}_{17}| = 4$, $|\mathcal{S}_{12}| = 3$, and $|\mathcal{S}_7| = 2$, $|\mathcal{E}| = 20$)

A. Miner and G. Ciardo, *Efficient reachability set generation and storage using DDs*, ICATPN 1999

W	Time (sec)			Total Memory (MB)			Max Memory (MB)		
	SMART _N ^{OW}	NAÏVE	HIST	SMART _N ^{OW}	NAÏVE	HIST	SMART _N ^{OW}	NAÏVE	HIST
$N = 300$ $ \mathcal{S} = 3.64 \cdot 10^{27}$ SMART completes in 55 sec using 241MB									
2	79	-8%	-8%	243	+12%	+24%	121	+26%	+52%
4	91	^d +67%	-9%	243	+102%	+30%	119	+205%	+50%
8	260	> 5hrs	-30%	243	-	+42%	103	-	+47%
$N = 450$ $ \mathcal{S} = 6.90 \cdot 10^{29}$ SMART does not complete in 5 hrs using 512MB									
2	^s 257	^s +12%	^s -14%	826	+16%	+5%	512	+15%	+7%
4	^d 311	> 5hrs	^d -18%	826	-	+33%	372	-	+6%
8	959	> 5hrs	-25%	826	-	+61%	343	-	+6%

- The NAÏVE approach only outperforms SMART_N^{OW} in some small case

- System with three machines to process three different types of parts : N is the number of each type of parts

($K = 19$, $|\mathcal{S}_k| = N + 1$ for all k except $|\mathcal{S}_{17}| = 4$, $|\mathcal{S}_{12}| = 3$, and $|\mathcal{S}_7| = 2$, $|\mathcal{E}| = 20$)

A. Miner and G. Ciardo, *Efficient reachability set generation and storage using DDs*, ICATPN 1999

W	Time (sec)			Total Memory (MB)			Max Memory (MB)		
	SMART _N ^{OW}	NAÏVE	HIST	SMART _N ^{OW}	NAÏVE	HIST	SMART _N ^{OW}	NAÏVE	HIST
$N = 300 \quad \mathcal{S} = 3.64 \cdot 10^{27}$ SMART completes in 55 sec using 241MB									
2	79	-8%	-8%	243	+12%	+24%	121	+26%	+52%
4	91	^d +67%	-9%	243	+102%	+30%	119	+205%	+50%
8	260	> 5hrs	-30%	243	-	+42%	103	-	+47%
$N = 450 \quad \mathcal{S} = 6.90 \cdot 10^{29}$ SMART does not complete in 5 hrs using 512MB									
2	^s 257	^s +12%	^s -14%	826	+16%	+5%	512	+15%	+7%
4	^d 311	> 5hrs	^d -18%	826	-	+33%	372	-	+6%
8	959	> 5hrs	-25%	826	-	+61%	343	-	+6%

- The HIST approach outperforms SMART_N^{OW} and controls memory usage well
- HIST approach works on various types of models

- Round robin version of mutual exclusion algorithm : N is the number of processors involved

($K = N + 1$, $|\mathcal{S}_k| = 10$ for all k except $|\mathcal{S}_1| = N + 1$, $|\mathcal{E}| = 5N$).

S. Graf, B. Steffen, and G. Lüttgen, *Compositional minimisation of finite state systems using interface specification*, Formal Aspect of Computer 1996

W	Time (sec)			Total Memory (MB)			Max Memory (MB)		
	SMART _N ^{OW}	NAÏVE	HIST	SMART _N ^{OW}	NAÏVE	HIST	SMART _N ^{OW}	NAÏVE	HIST
$N = 800 \quad \mathcal{S} = 1.20 \cdot 10^{196}$ SMART completes in 27 sec using 290MB									
2	29	+37%	+6%	293	+110%	+85%	215	+52%	+63%
4	36	+33%	+8%	293	+348%	+109%	130	+186%	+65%
8	51	+33%	+5%	293	+807%	+148%	73	+433%	+73%
$N = 1100 \quad \mathcal{S} = 3.36 \cdot 10^{334}$ SMART does not complete in 5 hrs using 512MB									
2	^d 65	^s +62%	^s +18%	794	+46%	+6%	379	+79%	+30%
4	47	^s +131%	^d +10%	794	+119%	+38%	265	+104%	+40%
8	56	^d +164%	+7%	794	+299%	+50%	173	+126%	+38%

- The NAÏVE approach performs excessive fire prediction but most speculative firing results are not useful

- Round robin version of mutual exclusion algorithm : N is the number of processors involved

($K = N + 1$, $|\mathcal{S}_k| = 10$ for all k except $|\mathcal{S}_1| = N + 1$, $|\mathcal{E}| = 5N$).

S. Graf, B. Steffen, and G. Lüttgen, *Compositional minimisation of finite state systems using interface specification*, Formal Aspect of Computer 1996

W	Time (sec)			Total Memory (MB)			Max Memory (MB)		
	SMART _N ^{OW}	NAÏVE	HIST	SMART _N ^{OW}	NAÏVE	HIST	SMART _N ^{OW}	NAÏVE	HIST
$N = 800 \quad \mathcal{S} = 1.20 \cdot 10^{196}$ SMART completes in 27 sec using 290MB									
2	29	+37%	+6%	293	+110%	+85%	215	+52%	+63%
4	36	+33%	+8%	293	+348%	+109%	130	+186%	+65%
8	51	+33%	+5%	293	+807%	+148%	73	+433%	+73%
$N = 1100 \quad \mathcal{S} = 3.36 \cdot 10^{334}$ SMART does not complete in 5 hrs using 512MB									
2	^d 65	^s +62%	^s +18%	794	+46%	+6%	379	+79%	+30%
4	47	^s +131%	^d +10%	794	+119%	+38%	265	+104%	+40%
8	56	^d +164%	+7%	794	+299%	+50%	173	+126%	+38%

- The HIST approach controls the space usage well by not to perform prediction if no firing pattern exists

- Avionics system : monitors T targets with S speeds on a $X \times Y \times Z$ runway ($K = 5(T+1)$, $|\mathcal{S}_{5+5i}| = 3$, $|\mathcal{S}_{4+5i}| = 14$, $|\mathcal{S}_{3+5i}| = 1 + X(10 + 6(S-1))$, $|\mathcal{S}_{2+5i}| = 1 + Y(10 + 6(S-1))$, $|\mathcal{S}_{1+5i}| = 1 + Z(10 + 6(S-1))$, for $i = 0, \dots, T$, except $|\mathcal{S}_{4+5T}| = 7$, $|\mathcal{E}| = 49 + T(56 + (Y-2)(31 + (X-2)(13 + 4Z))) + 3(X-2)(1 + YZ) + 2X + 5Y + 3Z$).

R. Siminiceanu and G. Ciardo, *Formal verification of the NASA Runway Safety Monitor*, AVoCS 2004

W	Time (sec)			Total Memory (MB)			Max Memory (MB)		
	SMART _N ^{OW}	NAÏVE	HIST	SMART _N ^{OW}	NAÏVE	HIST	SMART _N ^{OW}	NAÏVE	HIST
$Z = 2 \quad \mathcal{S} = 1.51 \cdot 10^{15}$ SMART completes in 236 sec using 314MB									
2	731	> 10hrs	-2%	332	-	+39%	191	-	+48%
4	938	> 10hrs	-8%	332	-	+88%	190	-	+30%
8	1480	> 10hrs	-22%	332	-	+128%	173	-	+13%
$Z = 3 \quad \mathcal{S} = 5.07 \cdot 10^{15}$ SMART does not complete in 10 hrs using 512MB									
2	^s 11280	> 10hrs	^s -1%	962	-	+10%	595	-	+16%
4	^d 9762	> 10hrs	^d -15%	962	-	+31%	371	-	+8%
8	^d 14101	> 10hrs	^d -17%	962	-	+58%	359	-	+6%

- The HIST approach does work in practice

Conclusions

- **Conclusions :**

- Pattern recognition approach is capable to guide the speculative firing prediction to improve the runtime of SMART_N^{OW}.
- This approach does work on some real world model.
- It seems applicable to the breadth-first iteration.

- **Future research directions :**

- Explore more sophisticated, but still low-overhead, heuristics to improve the usefulness of the predicted events.
- Augment this approach to include information about the current memory consumption.

Thank You!