



## A CHR-BASED SOLVER FOR WEAK MEMORY BEHAVIORS

CSTVA 2016 — **Allan Blanchard**<sup>1,2</sup>

Nikolai Kosmatov<sup>1</sup>

Frédéric Loulergue<sup>2</sup>



<sup>1</sup> CEA LIST - Software Reliability Laboratory

<sup>2</sup> Univ Orléans, INSA Centre Val de Loire, LIFO

# Memory model

## Reasoning about concurrency

Memory models define :

- how threads interact with memory
- in particular on shared data

They can be defined by :

- processor architectures
- languages

# Sequential Consistency

## Lamport 1979

The semantics of the parallel composition of two programs is given by the interleavings of the executions of both programs.

### A simple program

$$x := \Omega; y := \Omega$$

Thread 0		Thread 1
$x := 1;$		$y := 1;$
$r_0 := y;$		$r_1 := x;$

Possible results :

- $r_0 = 1 \wedge r_1 = 1$
- $r_0 = \Omega \wedge r_1 = 1$
- $r_0 = 1 \wedge r_1 = \Omega$

Impossible result :

- $r_0 = \Omega \wedge r_1 = \Omega$

# Weak behaviors

## Modern architectures

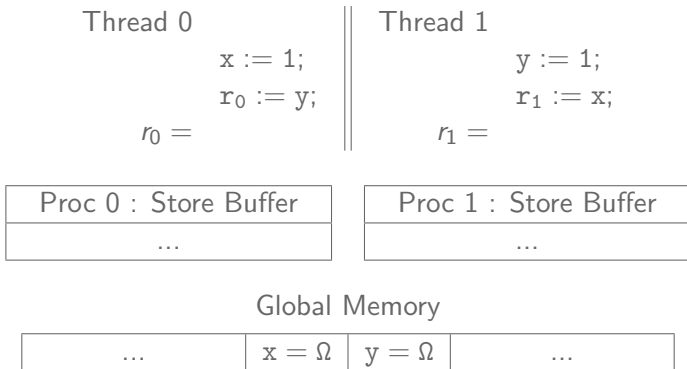
For example, x86-TSO or ARM do not respect SC :

- hard to ensure such synchronization in hardware,
- and far too costly.

## Examples of weak behaviors

- out-of-order execution,
- store buffering,
- speculative execution,
- ...

## Example with store buffering



## Example with store buffering

Thread 0

 $\rightarrow x := 1;$   
 $r_0 := y;$  $r_0 =$ 

Thread 1

 $y := 1;$   
 $r_1 := x;$  $r_1 =$ 

Proc 0 : Store Buffer

...  **$x = 1$**  ...

Proc 1 : Store Buffer

...

Global Memory

...  $x = \Omega$   $y = \Omega$  ...

## Example with store buffering

Thread 0

 $x := 1;$   
 $r_0 := y;$ 
 $r_0 =$ 

Thread 1

 $\rightarrow y := 1;$   
 $r_1 := x;$ 
 $r_1 =$ 

Proc 0 : Store Buffer

...	$x = 1$	...
-----	---------	-----

Proc 1 : Store Buffer

...	$y = 1$	...
-----	---------	-----

Global Memory

...	$x = \Omega$	$y = \Omega$	...
-----	--------------	--------------	-----

## Example with store buffering

Thread 0

```
x := 1;  
→ r0 := y;  
r0 = Ω
```

Thread 1

```
y := 1;  
r1 := x;  
r1 =
```

Proc 0 : Store Buffer

...	x = 1	...
-----	-------	-----

Proc 1 : Store Buffer

...	y = 1	...
-----	-------	-----

Global Memory

...	x = Ω	y = Ω	...
-----	-------	-------	-----



## Example with store buffering

Thread 0

$$\begin{aligned}
 &x := 1; \\
 &r_0 := y; \\
 &r_0 = \Omega
 \end{aligned}$$

Thread 1

$$\begin{aligned}
 &y := 1; \\
 &\longrightarrow r_1 := x; \\
 &r_1 = \Omega
 \end{aligned}$$

Proc 0 : Store Buffer

...	x = 1	...
-----	-------	-----

Proc 1 : Store Buffer

...	y = 1	...
-----	-------	-----

Global Memory

...	x = $\Omega$	y = $\Omega$	...
-----	--------------	--------------	-----

## Example with store buffering

Thread 0

 $x := 1;$   
 $r_0 := y;$   
 $r_0 = \Omega$ 

Thread 1

 $y := 1;$   
 $r_1 := x;$   
 $r_1 = \Omega$ 

Proc 0 : Store Buffer

... |  $x = 1$  | ...

Proc 1 : Store Buffer

... |  $y = 1$  | ...

Global Memory

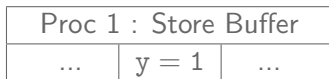
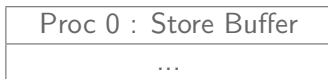
... |  $x = \Omega$  |  $y = \Omega$  | ...

## Example with store buffering

Thread 0

 $x := 1;$   
 $r_0 := y;$   
 $r_0 = \Omega$ 

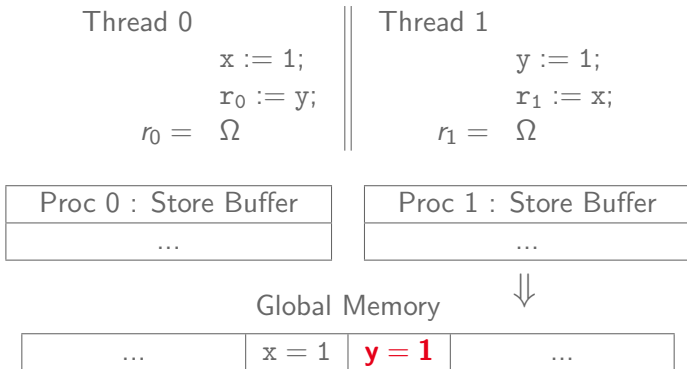
Thread 1

 $y := 1;$   
 $r_1 := x;$   
 $r_1 = \Omega$ 

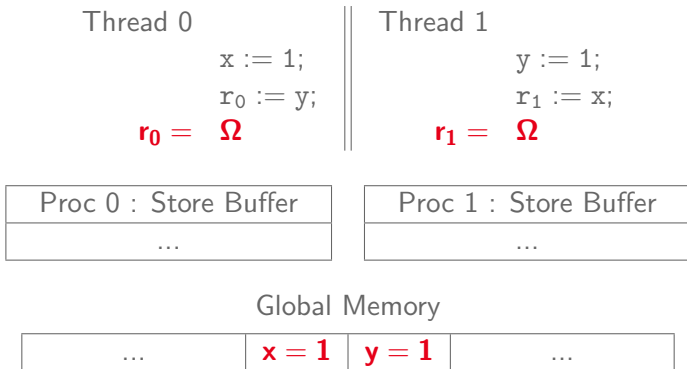
Global Memory



## Example with store buffering



## Example with store buffering



This behavior is allowed on ARM and TSO processors

# We need to understand weak behaviors

## Reasoning about programs

We have more and more multi-core software but

- it is hard to reason about them
- most analysis techniques are not aware of weak behaviors

## Existing dedicated tools

## CPPMem (Batty et al. 2010)

Program executions under C++11 model

## Herding cats (Alglave et al. 2014)

Generic framework for weak behaviors

- written in OCaml
- provides a language to specify memory models

## JMMSolve (Schrijvers 2004) based on CCMM (Saraswat 2004)

Program executions under Java Memory Model

- based on Concurrent Constraint-based Memory Machines
- written using Constraint Handling Rules (CHR)

# Prolog and CHR

## Prolog

Declarative language for logic programming

## Constraint Handling Rules

Declarative language for constraint programming

- maintains a store of constraints ( $\sim$  terms)
- handled by rules that will add or remove constraints



# Goals of our solver

## To identify allowed executions

- for a given parallel program
- according to a given memory model

## Additional goals

- possibility to add new memory models
- support of specific instructions

## Basic relations

 $(st, x, \Omega)$  $(st, y, \Omega)$  $(st, x, 1)$  $(st, y, 1)$  $(ld, y, R0 = 1)$  $(ld, x, R1 = 1)$ 

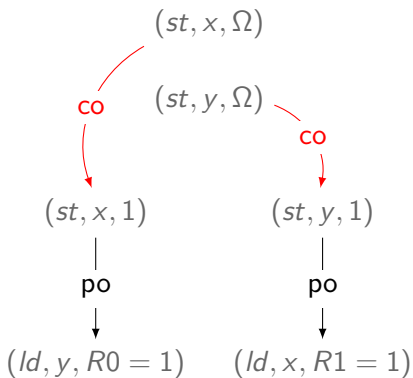
- Program-Order : PO
- Coherency-Order : CO
- Reads-From : RF

## Basic relations

 $(st, x, \Omega)$  $(st, y, \Omega)$  $(st, x, 1)$ |  
po  
↓ $(ld, y, R0 = 1)$  $(st, y, 1)$ |  
po  
↓ $(ld, x, R1 = 1)$ 

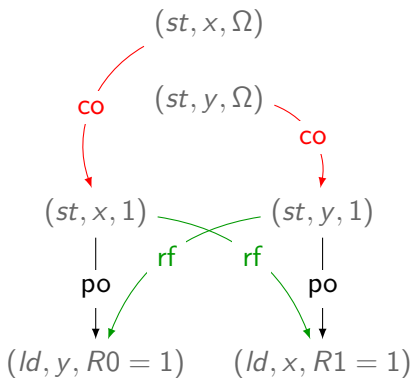
- Program-Order : PO
- Coherency-Order : CO
- Reads-From : RF

## Basic relations



- Program-Order : PO
- Coherency-Order : **CO**
- Reads-From : RF

## Basic relations



- Program-Order : PO
- Coherency-Order : CO
- Reads-From : RF

# Chosen approach

## Generate all candidate executions

An execution is represented by ordering relations :

- CO: for each location  $l$ , a total ordering of every store to  $l$
- RF: for each load, a store having written the value being read
- we combine all permutations of CO and RF using backtracking

## Filter out forbidden executions

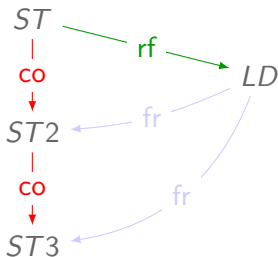
- apply model rules to deduce more ordering relations
- incoherent execution: an action must happen before itself (which means that some relations exhibits a cycle)

## Express and derive relations with CHR

Relation between 2 instructions	CHR constraint
$(st, x, \Omega) \xrightarrow{CO} (st, x, 1)$	$co((st, x, \text{undefined}), (st, x, 1))$
$(st, x, 2) \xrightarrow{RF} (ld, x, R)$	$rf((st, x, 2), (ld, x, 2))$

CHR rules to derive new relations :

- $rf(ST, LD), co(ST, ST2) \Rightarrow fr(LD, ST2)$ .
- $fr(LD, ST2), co(ST2, ST3) \Rightarrow fr(LD, ST3)$ .

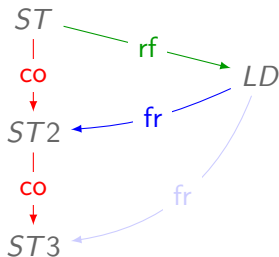


## Express and derive relations with CHR

Relation between 2 instructions	CHR constraint
$(st, x, \Omega) \xrightarrow{CO} (st, x, 1)$	$co((st, x, undefined), (st, x, 1))$
$(st, x, 2) \xrightarrow{RF} (ld, x, R)$	$rf((st, x, 2), (ld, x, 2))$

CHR rules to derive new relations :

- $rf(ST, LD), co(ST, ST2) \Rightarrow fr(LD, ST2)$ .
- $fr(LD, ST2), co(ST2, ST3) \Rightarrow fr(LD, ST3)$ .



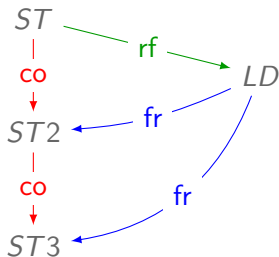


## Express and derive relations with CHR

Relation between 2 instructions	CHR constraint
$(st, x, \Omega) \xrightarrow{CO} (st, x, 1)$	$co((st, x, undefined), (st, x, 1))$
$(st, x, 2) \xrightarrow{RF} (ld, x, R)$	$rf((st, x, 2), (ld, x, 2))$

CHR rules to derive new relations :

- $rf(ST, LD), co(ST, ST2) \Rightarrow fr(LD, ST2).$
- $fr(LD, ST2), co(ST2, ST3) \Rightarrow fr(LD, ST3).$



## Detection of incoherent execution

Reminder : incoherent execution = cycle in the perserved relations

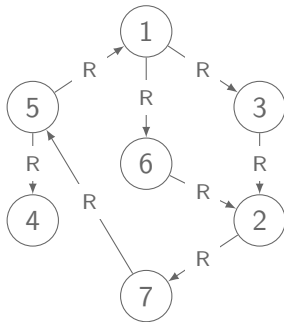
```
:- chr_constraint r/2, tc/2, cycle/1.
```

```
tc(B,E) \ tc(B,E) <=>  
  true.
```

```
tc(B,E), r(E,B) <=>  
  cycle(B).
```

```
tc(B,E), r(E,N) ==>  
  inf(B,N) | tc(B,N).
```

```
r(I,J) ==>  
  inf(I,J) | tc(I,J).
```



## Detection of incoherent execution

Reminder : incoherent execution = cycle in the perserved relations

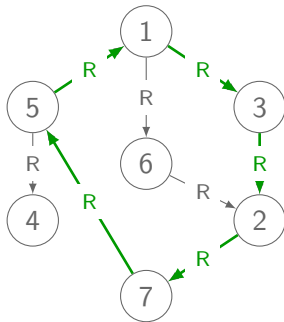
```
:- chr_constraint r/2, tc/2, cycle/1.
```

```
tc(B,E) \ tc(B,E) <=>  
  true.
```

```
tc(B,E), r(E,B) <=>  
  cycle(B).
```

```
tc(B,E), r(E,N) ==>  
  inf(B,N) | tc(B,N).
```

```
r(I,J) ==>  
  inf(I,J) | tc(I,J).
```



## Detection of incoherent execution

Reminder : incoherent execution = cycle in the perserved relations

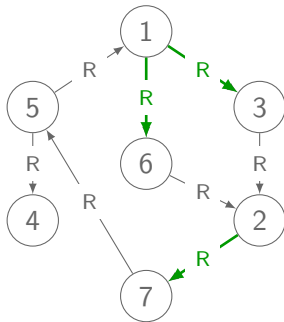
```
:- chr_constraint r/2, tc/2, cycle/1.
```

```
tc(B,E) \ tc(B,E) <=>  
  true.
```

```
tc(B,E), r(E,B) <=>  
  cycle(B).
```

```
tc(B,E), r(E,N) ==>  
  inf(B,N) | tc(B,N).
```

```
r(I,J) ==>  
  inf(I,J) | tc(I,J).
```



## Detection of incoherent execution

Reminder : incoherent execution = cycle in the perserved relations

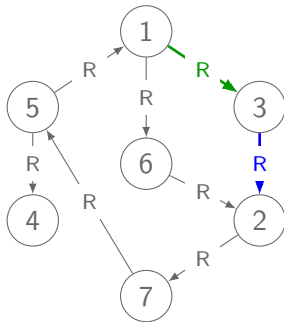
```
:- chr_constraint r/2, tc/2, cycle/1.
```

```
tc(B,E) \ tc(B,E) <=>  
  true.
```

```
tc(B,E), r(E,B) <=>  
  cycle(B).
```

```
tc(B,E), r(E,N) ==>  
  inf(B,N) | tc(B,N).
```

```
r(I,J) ==>  
  inf(I,J) | tc(I,J).
```



## Detection of incoherent execution

Reminder : incoherent execution = cycle in the perserved relations

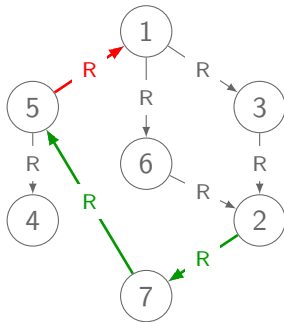
```
:- chr_constraint r/2, tc/2, cycle/1.
```

```
tc(B,E) \ tc(B,E) <=>  
true.
```

```
tc(B,E), r(E,B) <=>  
cycle(B).
```

```
tc(B,E), r(E,N) ==>  
inf(B,N) | tc(B,N).
```

```
r(I,J) ==>  
inf(I,J) | tc(I,J).
```



## Detection of incoherent execution

Reminder : incoherent execution = cycle in the perserved relations

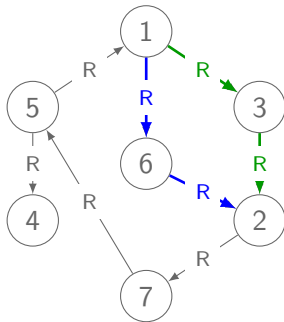
```
:- chr_constraint r/2, tc/2, cycle/1.
```

```
tc(B,E) \ tc(B,E) <=>  
true.
```

```
tc(B,E), r(E,B) <=>  
cycle(B).
```

```
tc(B,E), r(E,N) ==>  
inf(B,N) | tc(B,N).
```

```
r(I,J) ==>  
inf(I,J) | tc(I,J).
```



## Detection of incoherent execution

Reminder : incoherent execution = cycle in the perserved relations

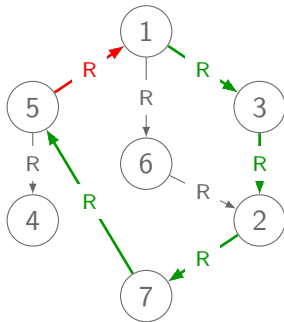
```
:- chr_constraint r/2, tc/2, cycle/1.
```

```
tc(B,E) \ tc(B,E) <=>  
  true.
```

```
tc(B,E), r(E,B) <=>  
  cycle(B).
```

```
tc(B,E), r(E,N) ==>  
  inf(B,N) | tc(B,N).
```

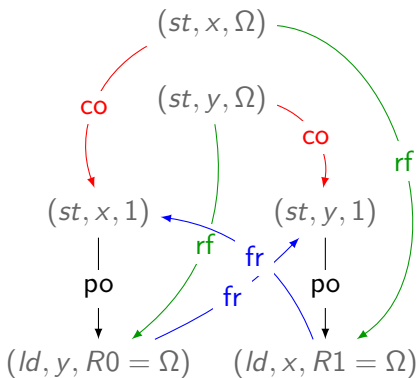
```
r(I,J) ==>  
  inf(I,J) | tc(I,J).
```





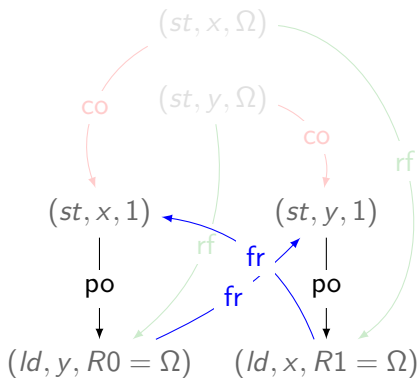
## Definition of a model : SC

$\text{co}(I, J) \implies \text{sc}(I, J).$   
 $\text{rf}(I, J) \implies \text{sc}(I, J).$   
 $\text{fr}(I, J) \implies \text{sc}(I, J).$   
 $\text{po}(I, J) \implies \text{sc}(I, J).$



## Definition of a model : SC

$\text{co}(I, J) \implies \text{sc}(I, J).$   
 $\text{rf}(I, J) \implies \text{sc}(I, J).$   
 $\text{fr}(I, J) \implies \text{sc}(I, J).$   
 $\text{po}(I, J) \implies \text{sc}(I, J).$



## Definition of a model: TSO (simplified)

$$\text{ppo}((st, \_, \_), (ld, \_, \_)) \Leftrightarrow \text{true}.$$

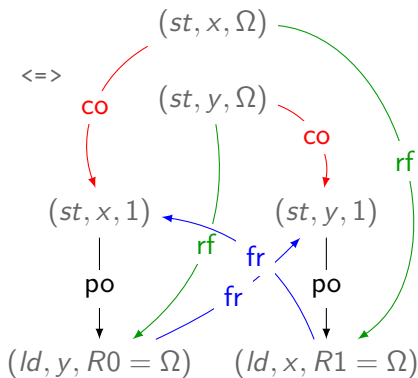
$$\text{ipo}(I, J) \Rightarrow \text{ppo}(I, J).$$

$$\text{bar}(I, J) \Rightarrow \text{tso}(I, J).$$

$$\text{ppo}(I, J) \Rightarrow \text{tso}(I, J).$$

$$\text{rf}(I, J) \Rightarrow \text{tso}(I, J).$$

$$\text{co}(I, J) \Rightarrow \text{tso}(I, J).$$

$$\text{fr}(I, J) \Rightarrow \text{tso}(I, J).$$


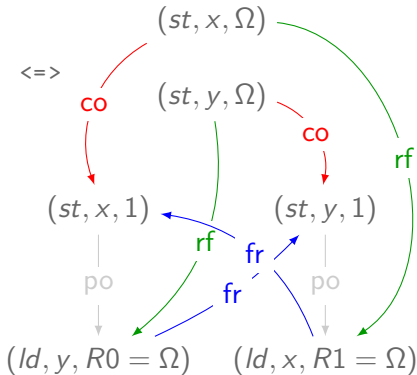
## Definition of a model: TSO (simplified)

```

ppo((st,_,_), (ld,_,_)) <=>
  true.
ipo(I,J) ==> ppo(I,J).

bar(I,J) ==> tso(I,J).
ppo(I,J) ==> tso(I,J).
rf(I,J) ==> tso(I,J).
co(I,J) ==> tso(I,J).
fr(I,J) ==> tso(I,J).

```



## Definition of a model: TSO (simplified)

```
ppo((st,_,_), (ld,_,_)) <=>
  true.
```

```
ipo(I,J) ==> ppo(I,J).
```

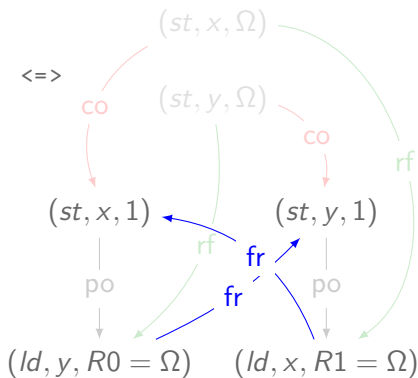
```
bar(I,J) ==> tso(I,J).
```

```
ppo(I,J) ==> tso(I,J).
```

```
rf(I,J) ==> tso(I,J).
```

```
co(I,J) ==> tso(I,J).
```

```
fr(I,J) ==> tso(I,J).
```



## Definition of a model: TSO (simplified)

```
ppo((st,_,_), (ld,_,_)) <=>
  true.
```

```
ipo(I,J) ==> ppo(I,J).
```

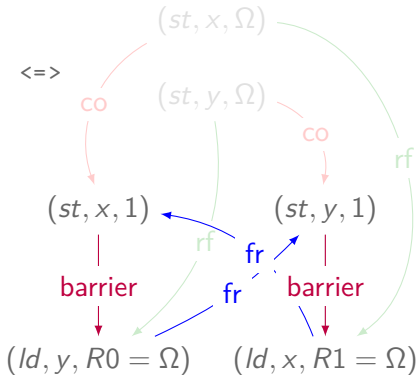
```
bar(I,J) ==> tso(I,J).
```

```
ppo(I,J) ==> tso(I,J).
```

```
rf(I,J) ==> tso(I,J).
```

```
co(I,J) ==> tso(I,J).
```

```
fr(I,J) ==> tso(I,J).
```



# Implemented Memory Models

- SC
- TSO
- PSO
- ARM (Almost finished)

## Program samples from Herd

We tested the implementation on 18 examples of litmus programs.

- message passing,
- message passing by address passing,
- basic uniproc relations,
- store buffering,
- ...

We observe the same results found by Herd.

<http://virginia.cs.ucl.ac.uk/herd/>



## Multiple message passing

Example	Model	#exec	Herd	CHR Solver
MP3	Generic	147 436	1.2s	3.3s
	PSO	2 258	3.8s	6.4s
	TSO	800	4.1s	3.2s
	SC	678	5.5s	3.3s
MP4	Generic	255 000 000	1405s	> 1h
	PSO	516 030	> 1h	2796s
	TSO	96 498	> 1h	752s
	SC	81 882	> 1h	747s

Early pruning makes the solver efficient for “long” programs.

# Conclusion

## A CHR based solver for weak memory behaviors

- currently for SC, TSO and PSO
- easy to use and to extend for new memory models
- it was not so hard to implement
- despite the (awful) complexity, execution time is not so bad

## Future work

- Features :
  - finalize ARM
  - add read-modify-write operations and branching
- Evaluate on further benchmarks

# Thank you ! Questions ?

Commissariat à l'énergie atomique et aux énergies alternatives  
CEA Tech List  
Centre de Saclay — 91191 Gif-sur-Yvette Cedex  
[www-list.cea.fr](http://www-list.cea.fr)

Etablissement public à caractère industriel et commercial — RCS Paris B 775 685 019