

# Effective Lock Handling in Stateless Model Checking

---

*Michalis Kokologiannakis*   Azalea Raad   Viktor Vafeiadis

October 25, 2019

Max Planck Institute for Software Systems (MPI-SWS)

# Why Locks?

# Why Locks?

(Coarse-grained) locking is the most widespread form of synchronization

- Programming languages provide syntactic support for it (e.g., Java)
- Modern hardware provides lock elision support

# Why Locks?

(Coarse-grained) locking is the most widespread form of synchronization

- Programming languages provide syntactic support for it (e.g., Java)
- Modern hardware provides lock elision support

Verifying programs with locks is hard!

# Verifying Programs With Locks Is Hard

## Verifying Programs With Locks Is Hard

Consider a **fine-grained** hashtable

Let us try to verify it using a state-of-the-art **stateless model checking** (SMC) tool

# Verifying Programs With Locks Is Hard

Consider a **fine-grained** hashtable

Let us try to verify it using a state-of-the-art **stateless model checking** (SMC) tool




#Threads	Fine-grained HT
5	0.05 s
6	0.05 s
7	0.06 s
8	0.08 s

Verification time using GENMC

# Verifying Programs With Locks Is Hard

Consider a **fine-grained** hashtable

Let us try to verify it using a state-of-the-art **stateless model checking** (SMC) tool

#Threads	Fine-grained HT	Coarse-grained HT
5	0.05 s	104 s
6	0.05 s	
7	0.06 s	
8	0.08 s	

Verification time using GENMC



# Why Is Verifying Programs With Locks Hard?

## Why Is Verifying Programs With Locks Hard?

SMC enumerates all the interleavings of a program

↪ it is usually combined with **Partial Order Reduction (POR)**

# Why Is Verifying Programs With Locks Hard?

SMC enumerates all the interleavings of a program

↪ it is usually combined with **Partial Order Reduction (POR)**

$[x = 0]$

$a := x; \parallel y := 42;$

Executions:

# Why Is Verifying Programs With Locks Hard?

SMC enumerates all the interleavings of a program

↪ it is usually combined with **Partial Order Reduction (POR)**

$[x = 0]$

$a := x; \parallel y := 42;$

**Executions:**

SMC-naive : 2

# Why Is Verifying Programs With Locks Hard?

SMC enumerates all the interleavings of a program

↪ it is usually combined with **Partial Order Reduction (POR)**

$[x = 0]$

$a := x; \parallel y := 42;$

**Executions:**

SMC-naive : 2

SMC+POR : 1

# Why Is Verifying Programs With Locks Hard?

SMC enumerates all the interleavings of a program

↪ it is usually combined with **Partial Order Reduction (POR)**

$[x = 0]$

$a := x; \parallel y := 42;$

**Executions:**

SMC-naive : 2

SMC+POR : 1

$[l = 0]$

$lock(l); \parallel lock(l);$

**Executions:**

# Why Is Verifying Programs With Locks Hard?

SMC enumerates all the interleavings of a program

↪ it is usually combined with **Partial Order Reduction (POR)**

$[x = 0]$

$a := x; \parallel y := 42;$

**Executions:**

SMC-naive : 2

SMC+POR : 1

$[l = 0]$

$lock(l); \parallel lock(l);$

**Executions:**

SMC-naive : 2

# Why Is Verifying Programs With Locks Hard?

SMC enumerates all the interleavings of a program

↪ it is usually combined with **Partial Order Reduction (POR)**

$[x = 0]$

$a := x; \parallel y := 42;$

**Executions:**

SMC-naive : 2

SMC+POR : 1

$[l = 0]$

$lock(l); \parallel lock(l);$

**Executions:**

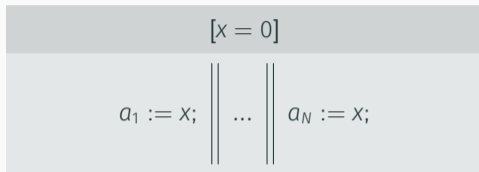
SMC-naive : 2

SMC+POR : 2



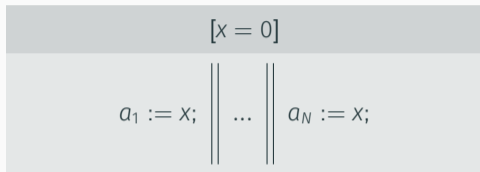
# Why Is Verifying Programs With Locks Hard?

## Why Is Verifying Programs With Locks Hard?



Executions:

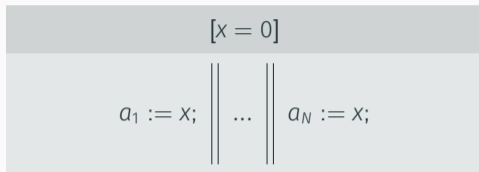
# Why Is Verifying Programs With Locks Hard?



**Executions:**

SMC-naive :  $N!$

# Why Is Verifying Programs With Locks Hard?

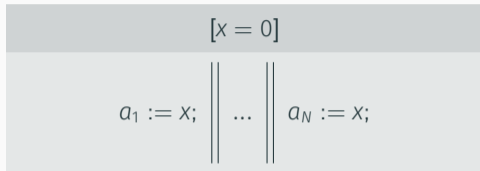


## Executions:

SMC-naive :  $N!$

SMC+POR : 1

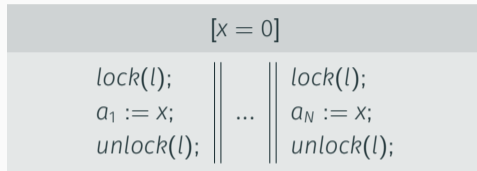
# Why Is Verifying Programs With Locks Hard?



**Executions:**

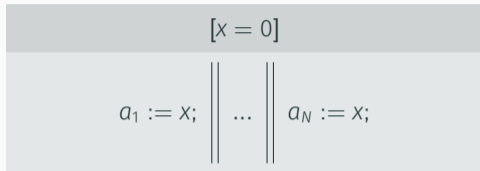
SMC-naive :  $N!$

SMC+POR : 1



**Executions:**

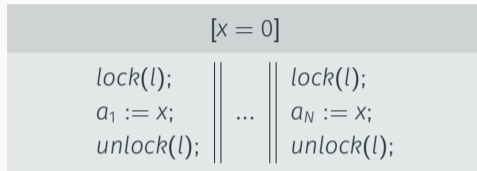
# Why Is Verifying Programs With Locks Hard?



**Executions:**

SMC-naive : N!

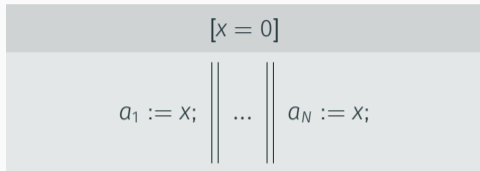
SMC+POR : 1



**Executions:**

SMC-naive : N!

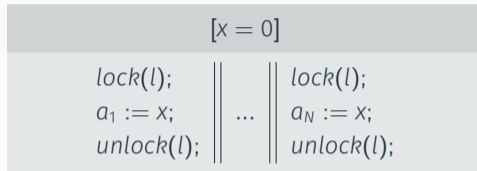
# Why Is Verifying Programs With Locks Hard?



**Executions:**

SMC-naive : N!

SMC+POR : 1

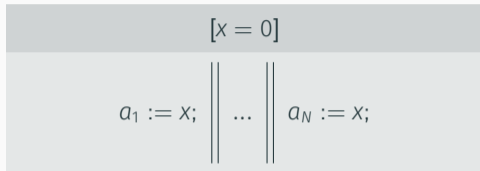


**Executions:**

SMC-naive : N!

SMC+POR : N!

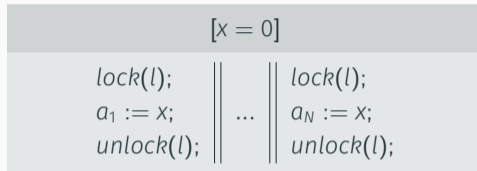
# Why Is Verifying Programs With Locks Hard?



**Executions:**

SMC-naive : N!

SMC+POR : 1



**Executions:**

SMC-naive : N!

SMC+POR : N!

Locks **inhibit** Partial Order Reduction (POR)





- LAPOR: a **Lock-Aware Partial Order Reduction** algorithm
  - independence at both instruction and **critical-section** level

# Our Contribution

- LAPOR: a **Lock-Aware Partial Order Reduction** algorithm
  - independence at both instruction and **critical-section** level
- We prove that LAPOR is sound, complete, optimal

# Our Contribution

- LAPOR: a **Lock-Aware Partial Order Reduction** algorithm
  - independence at both instruction and **critical-section** level
- We prove that LAPOR is sound, complete, optimal
- Built on top of the GENMC framework
  - Works for a **variety** of memory models

# Our Contribution

- LAPOR: a **Lock-Aware Partial Order Reduction** algorithm
  - independence at both instruction and **critical-section** level
- We prove that LAPOR is sound, complete, optimal
- Built on top of the GENMC framework
  - Works for a **variety** of memory models
- Exponentially fewer executions than state-of-the-art

# Partial Order Reduction

---

Equivalent interleavings are represented by graphs:

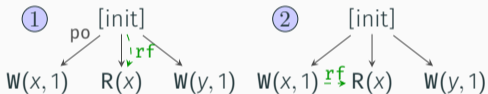
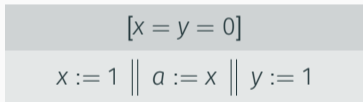
Equivalent interleavings are represented by graphs:

$$[x = y = 0]$$
$$x := 1 \parallel a := x \parallel y := 1$$



# SMC + POR: Background

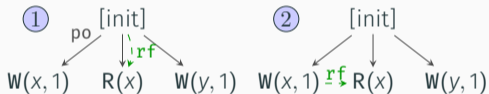
Equivalent interleavings are represented by graphs:



# SMC + POR: Background

Equivalent interleavings are represented by graphs:

$[x = y = 0]$
$x := 1 \parallel a := x \parallel y := 1$



**Goal:** Enumerate all consistent execution graphs of a program

## SMC+POR: An Example

## SMC+POR: An Example

$[x = y = 0]$

$x := 1 \parallel a := x \parallel y := 1$

## SMC+POR: An Example

$[x = y = 0]$

$x := 1 \parallel a := x \parallel y := 1$

[init]

## SMC+POR: An Example

$[x = y = 0]$

$x := 1 \parallel a := x \parallel y := 1$

$W(x, 1)$  ← [init]

# SMC+POR: An Example

$[x = y = 0]$

$x := 1 \parallel a := x \parallel y := 1$

$[init]$   
 $\swarrow$   
 $W(x, 1) \xrightarrow{rf} R(x)$   
 $\downarrow$

# SMC+POR: An Example

$[x = y = 0]$

$x := 1 \parallel a := x \parallel y := 1$

$[init]$   
↙   ↓  
 $W(x, 1) \xrightarrow{rf} R(x)$

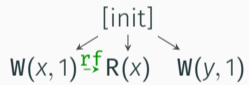
$[init]$



# SMC+POR: An Example

$[x = y = 0]$

$x := 1 \parallel a := x \parallel y := 1$

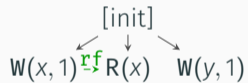
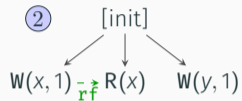


$[init]$

# SMC+POR: An Example

$[x = y = 0]$

$x := 1 \parallel a := x \parallel y := 1$

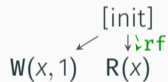
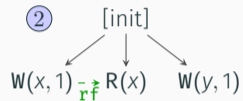


[init]

# SMC+POR: An Example

$[x = y = 0]$

$x := 1 \parallel a := x \parallel y := 1$

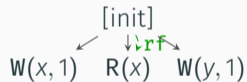
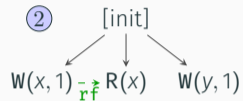


[init] ✓

# SMC+POR: An Example

$[x = y = 0]$

$x := 1 \parallel a := x \parallel y := 1$

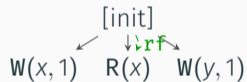
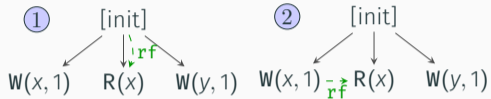


[init] ✓

# SMC+POR: An Example

$[x = y = 0]$

$x := 1 \parallel a := x \parallel y := 1$



[init] ✓



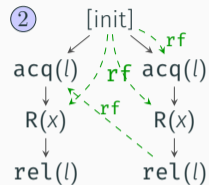
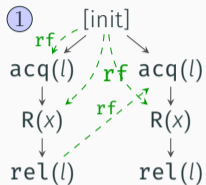
## SMC+POR: Handling Locks

$[init(l)]$

$lock(l);$	$\parallel$	$lock(l);$
$a_1 := x;$		$a_2 := x;$
$unlock(l);$		$unlock(l);$

# SMC+POR: Handling Locks

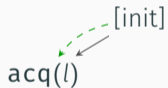
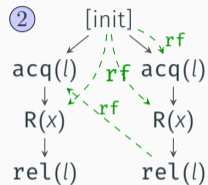
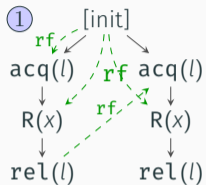
[init(l)]	
<i>lock(l);</i>	<i>lock(l);</i>
<i>a<sub>1</sub> := x;</i>	<i>a<sub>2</sub> := x;</i>
<i>unlock(l);</i>	<i>unlock(l);</i>





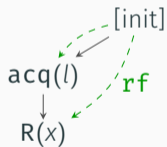
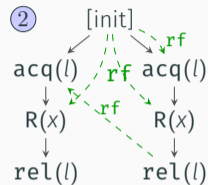
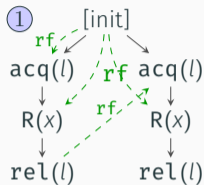
# SMC+POR: Handling Locks

$[init(l)]$	
$lock(l);$	$lock(l);$
$a_1 := x;$	$a_2 := x;$
$unlock(l);$	$unlock(l);$

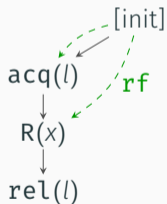
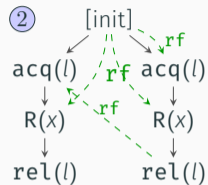
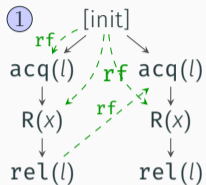
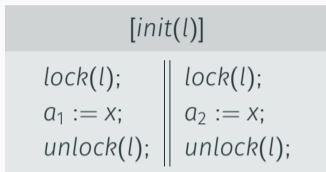


# SMC+POR: Handling Locks

$[init(l)]$	
$lock(l);$	$lock(l);$
$a_1 := x;$	$a_2 := x;$
$unlock(l);$	$unlock(l);$

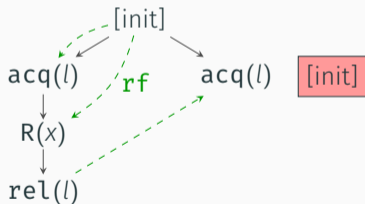
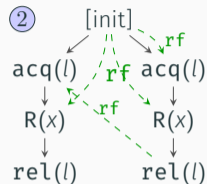
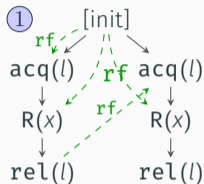


# SMC+POR: Handling Locks



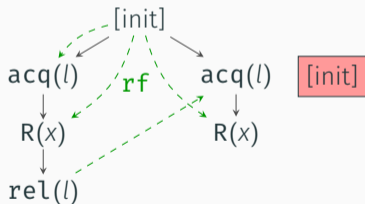
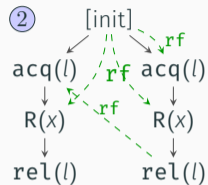
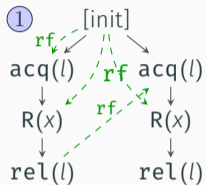
# SMC+POR: Handling Locks

$[init(l)]$	
$lock(l);$	$lock(l);$
$a_1 := x;$	$a_2 := x;$
$unlock(l);$	$unlock(l);$



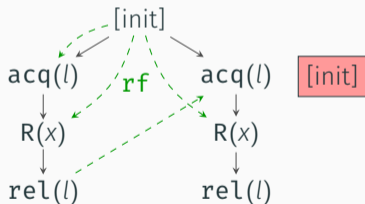
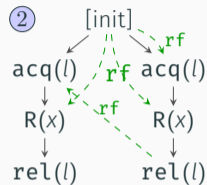
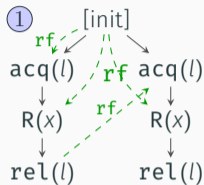
# SMC+POR: Handling Locks

$[init(l)]$	
$lock(l);$	$lock(l);$
$a_1 := x;$	$a_2 := x;$
$unlock(l);$	$unlock(l);$

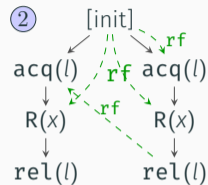
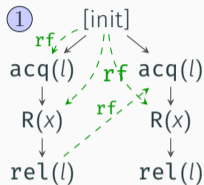
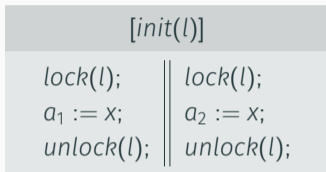


# SMC+POR: Handling Locks

$[init(l)]$	
$lock(l);$	$lock(l);$
$a_1 := x;$	$a_2 := x;$
$unlock(l);$	$unlock(l);$



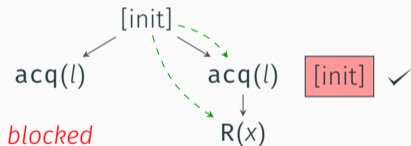
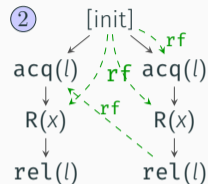
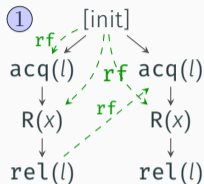
# SMC+POR: Handling Locks



*blocked*

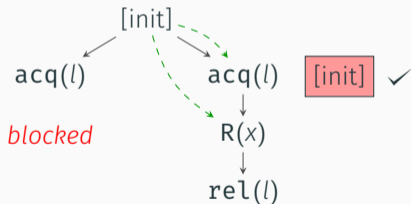
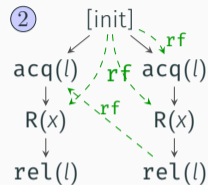
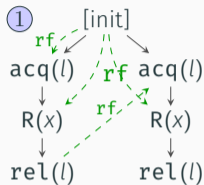
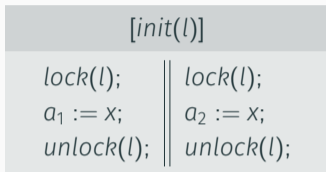
# SMC+POR: Handling Locks

[init(l)]	
<code>lock(l);</code>	<code>lock(l);</code>
<code>a<sub>1</sub> := x;</code>	<code>a<sub>2</sub> := x;</code>
<code>unlock(l);</code>	<code>unlock(l);</code>



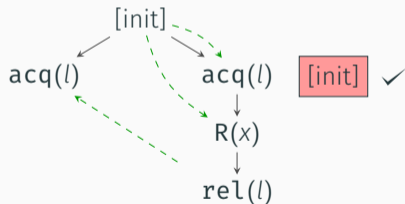
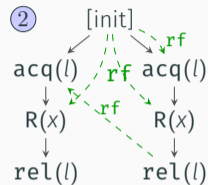
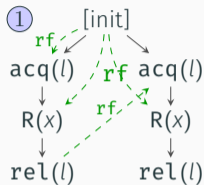


# SMC+POR: Handling Locks



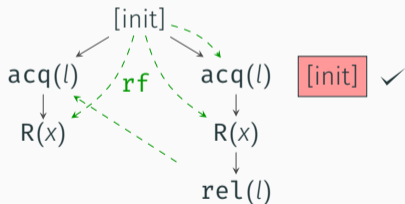
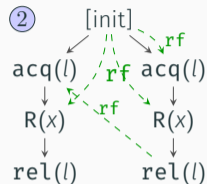
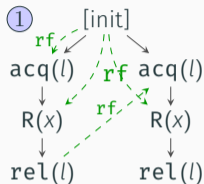
# SMC+POR: Handling Locks

$[init(l)]$	
$lock(l);$	$lock(l);$
$a_1 := x;$	$a_2 := x;$
$unlock(l);$	$unlock(l);$



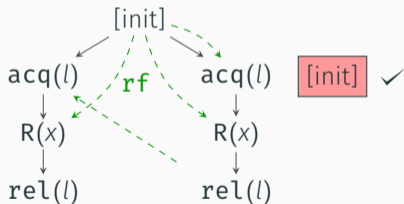
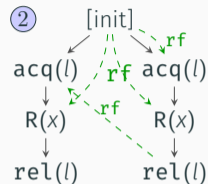
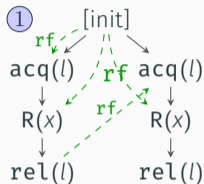
# SMC+POR: Handling Locks

$[init(l)]$	
$lock(l);$	$lock(l);$
$a_1 := x;$	$a_2 := x;$
$unlock(l);$	$unlock(l);$



# SMC+POR: Handling Locks

$[init(l)]$	
$lock(l);$	$lock(l);$
$a_1 := x;$	$a_2 := x;$
$unlock(l);$	$unlock(l);$



# Lock-Aware Partial Order Reduction

---

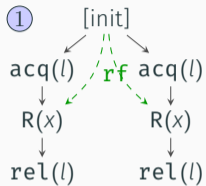
## LAPOR: Keeping Locks Unordered

$[init(l)]$

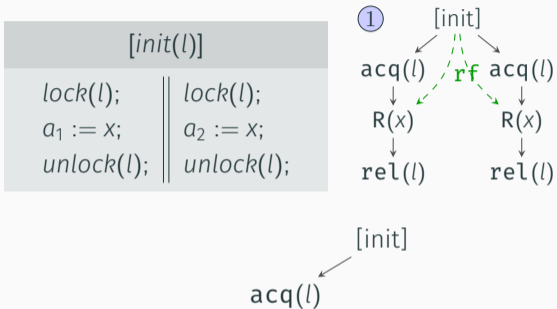
$lock(l);$	$\parallel$	$lock(l);$
$a_1 := x;$	$\parallel$	$a_2 := x;$
$unlock(l);$	$\parallel$	$unlock(l);$

# LAPOR: Keeping Locks Unordered

[init(l)]	
<i>lock(l);</i>	<i>lock(l);</i>
<i>a<sub>1</sub> := x;</i>	<i>a<sub>2</sub> := x;</i>
<i>unlock(l);</i>	<i>unlock(l);</i>



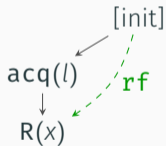
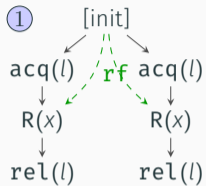
# LAPOR: Keeping Locks Unordered





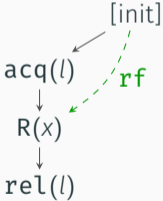
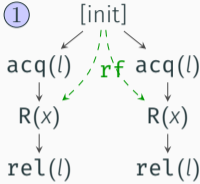
# LAPOR: Keeping Locks Unordered

[init(l)]	
<i>lock(l);</i>	<i>lock(l);</i>
<i>a<sub>1</sub> := x;</i>	<i>a<sub>2</sub> := x;</i>
<i>unlock(l);</i>	<i>unlock(l);</i>

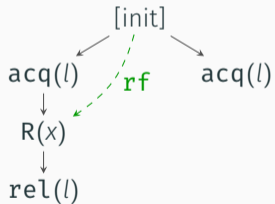
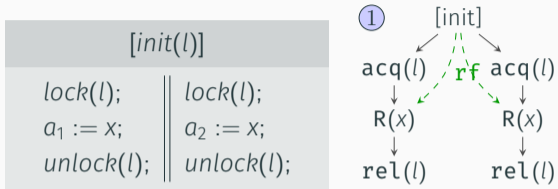


# LAPOR: Keeping Locks Unordered

[init(l)]	
lock(l);	lock(l);
a <sub>1</sub> := x;	a <sub>2</sub> := x;
unlock(l);	unlock(l);

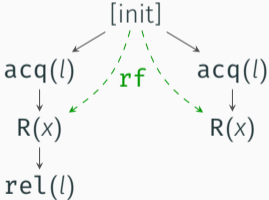
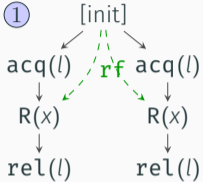


# LAPOR: Keeping Locks Unordered



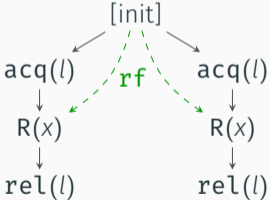
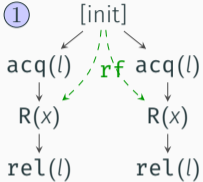
# LAPOR: Keeping Locks Unordered

[init(l)]	
<i>lock(l);</i>	<i>lock(l);</i>
<i>a<sub>1</sub> := x;</i>	<i>a<sub>2</sub> := x;</i>
<i>unlock(l);</i>	<i>unlock(l);</i>

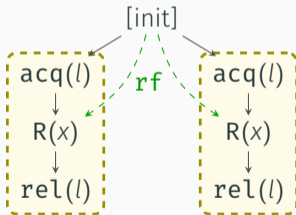
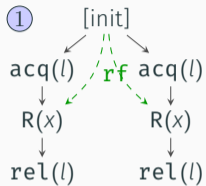
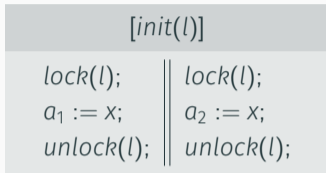


# LAPOR: Keeping Locks Unordered

[init(l)]	
<i>lock(l);</i>	<i>lock(l);</i>
<i>a<sub>1</sub> := x;</i>	<i>a<sub>2</sub> := x;</i>
<i>unlock(l);</i>	<i>unlock(l);</i>



# LAPOR: Keeping Locks Unordered



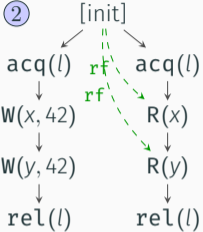
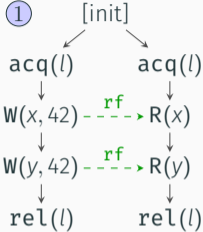
# LAPOR: Keeping Locks Unordered

$[init(l)]$

<i>lock(l);</i>		<i>lock(l);</i>
<i>x := 42;</i>		<i>a := x;</i>
<i>y := 42;</i>		<i>b := y;</i>
<i>unlock(l);</i>		<i>unlock(l);</i>

# LAPOR: Keeping Locks Unordered

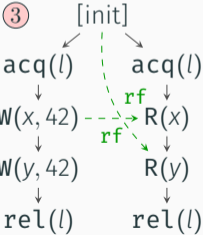
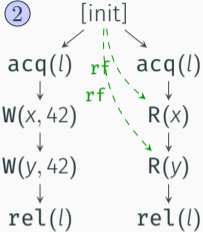
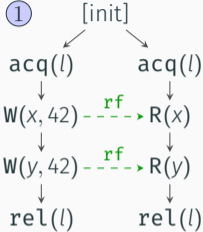
[init(l)]	
lock(l);	lock(l);
x := 42;	a := x;
y := 42;	b := y;
unlock(l);	unlock(l);





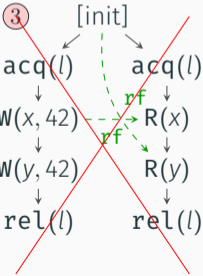
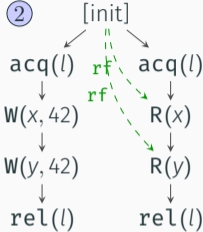
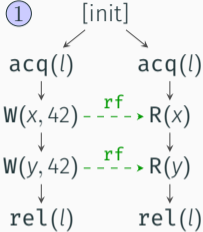
# LAPOR: Keeping Locks Unordered

[init(l)]	
lock(l);	lock(l);
x := 42;	a := x;
y := 42;	b := y;
unlock(l);	unlock(l);



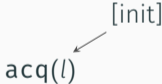
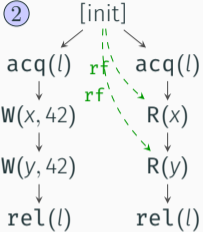
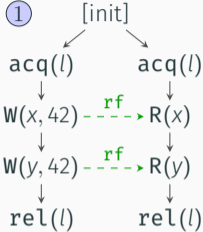
# LAPOR: Keeping Locks Unordered

[init(l)]	
lock(l);	lock(l);
x := 42;	a := x;
y := 42;	b := y;
unlock(l);	unlock(l);



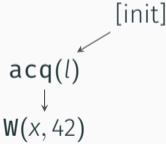
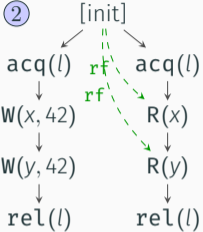
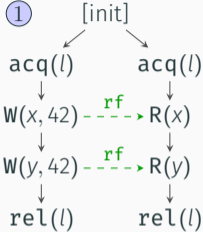
# LAPOR: Keeping Locks Unordered

[init(l)]	
<i>lock(l);</i>	<i>lock(l);</i>
<i>x := 42;</i>	<i>a := x;</i>
<i>y := 42;</i>	<i>b := y;</i>
<i>unlock(l);</i>	<i>unlock(l);</i>



# LAPOR: Keeping Locks Unordered

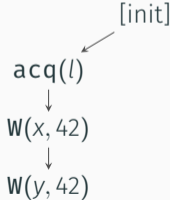
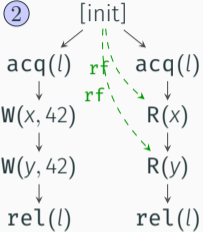
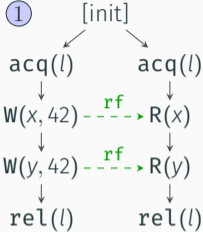
[init(l)]	
lock(l);	lock(l);
x := 42;	a := x;
y := 42;	b := y;
unlock(l);	unlock(l);



# LAPOR: Keeping Locks Unordered

```

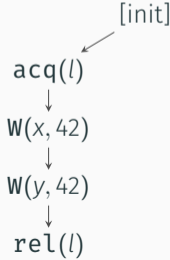
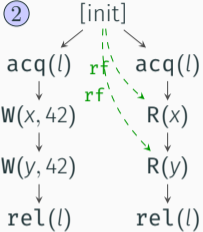
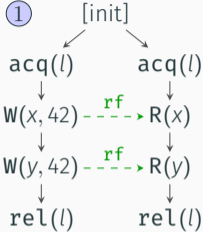
[init(l)]
lock(l);   || lock(l);
x := 42;   || a := x;
y := 42;   || b := y;
unlock(l); || unlock(l);
    
```



# LAPOR: Keeping Locks Unordered

```

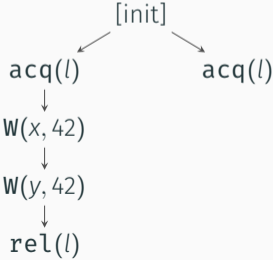
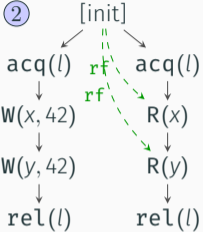
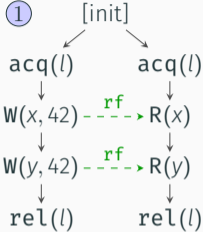
[init(l)]
lock(l);   || lock(l);
x := 42;   || a := x;
y := 42;   || b := y;
unlock(l); || unlock(l);
    
```



# LAPOR: Keeping Locks Unordered

```

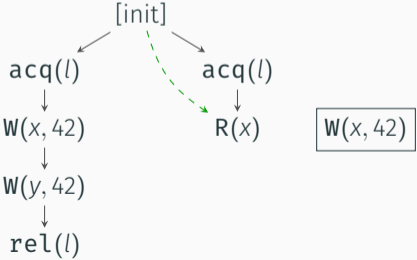
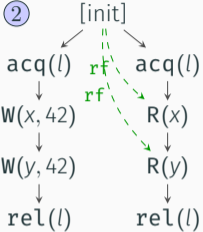
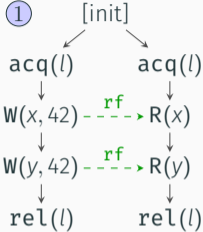
[init(l)]
lock(l);   || lock(l);
x := 42;   || a := x;
y := 42;   || b := y;
unlock(l); || unlock(l);
    
```



# LAPOR: Keeping Locks Unordered

```

[init(l)]
lock(l);   || lock(l);
x := 42;   || a := x;
y := 42;   || b := y;
unlock(l); || unlock(l);
    
```

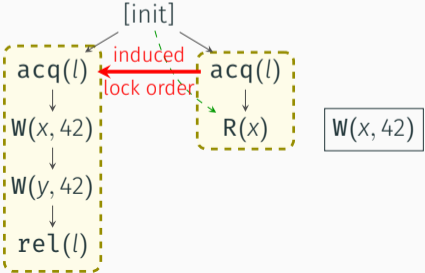
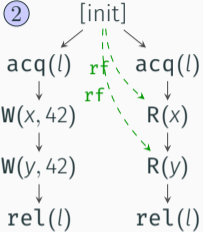
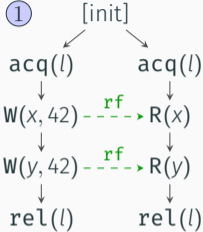




# LAPOR: Keeping Locks Unordered

```

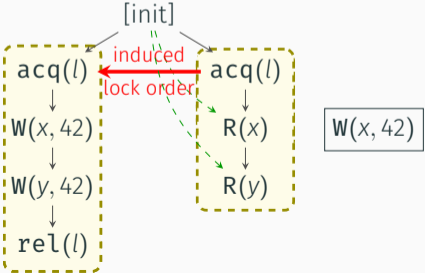
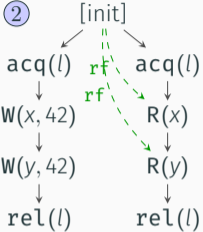
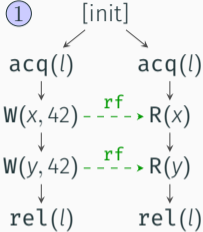
[init(l)]
lock(l);   || lock(l);
x := 42;   || a := x;
y := 42;   || b := y;
unlock(l); || unlock(l);
    
```



# LAPOR: Keeping Locks Unordered

```

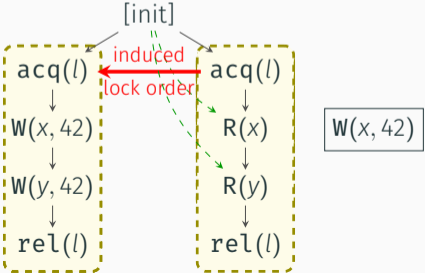
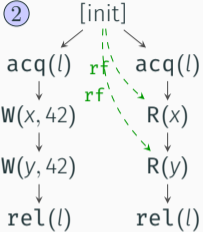
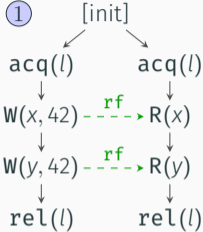
[init(l)]
lock(l);   || lock(l);
x := 42;   || a := x;
y := 42;   || b := y;
unlock(l); || unlock(l);
    
```



# LAPOR: Keeping Locks Unordered

```

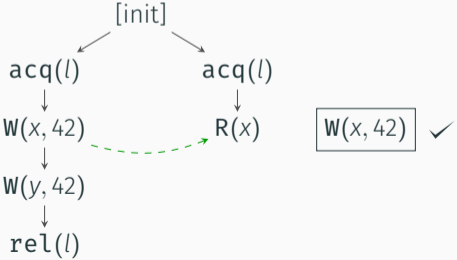
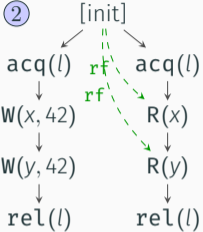
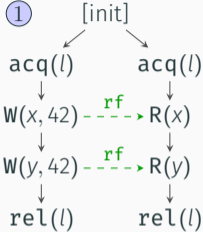
[init(l)]
lock(l);   || lock(l);
x := 42;   || a := x;
y := 42;   || b := y;
unlock(l); || unlock(l);
    
```



# LAPOR: Keeping Locks Unordered

```

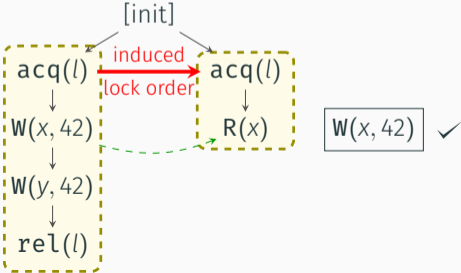
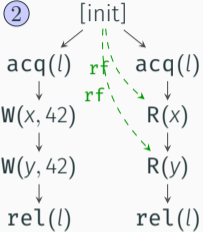
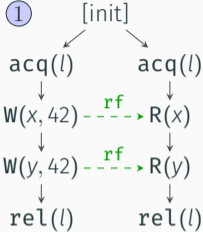
[init(l)]
lock(l);   || lock(l);
x := 42;   || a := x;
y := 42;   || b := y;
unlock(l); || unlock(l);
    
```



# LAPOR: Keeping Locks Unordered

```

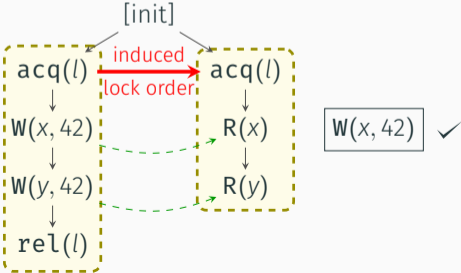
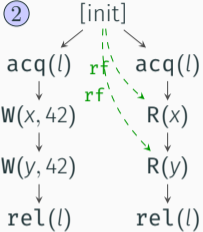
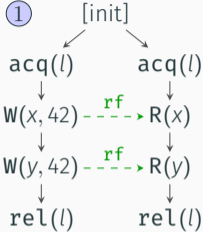
[init(l)]
lock(l);   || lock(l);
x := 42;   || a := x;
y := 42;   || b := y;
unlock(l); || unlock(l);
    
```



# LAPOR: Keeping Locks Unordered

```

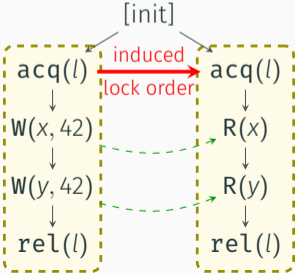
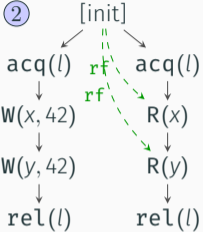
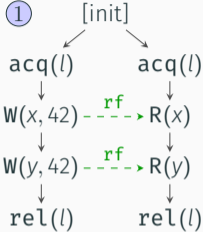
[init(l)]
lock(l);   || lock(l);
x := 42;   || a := x;
y := 42;   || b := y;
unlock(l); || unlock(l);
    
```



# LAPOR: Keeping Locks Unordered

```

[init(l)]
lock(l);   || lock(l);
x := 42;   || a := x;
y := 42;   || b := y;
unlock(l); || unlock(l);
    
```



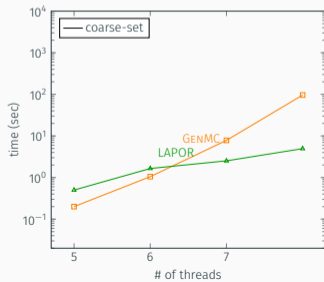
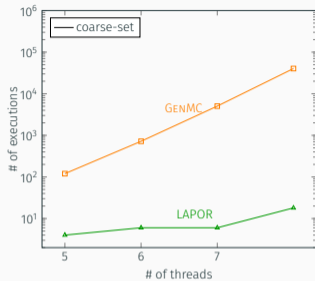
W(x, 42) ✓

## Results

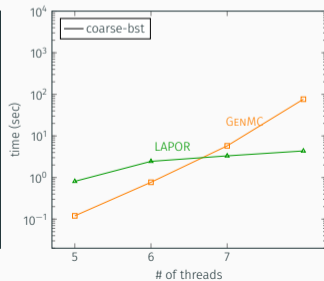
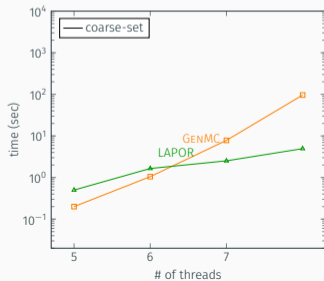
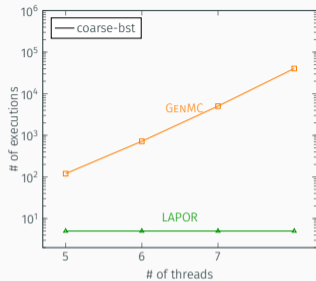
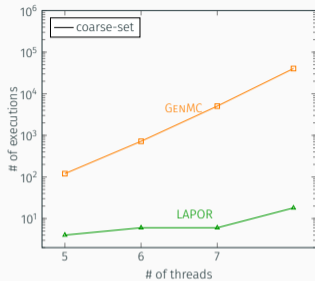
---



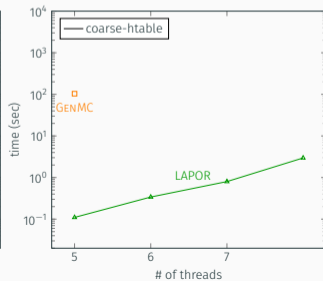
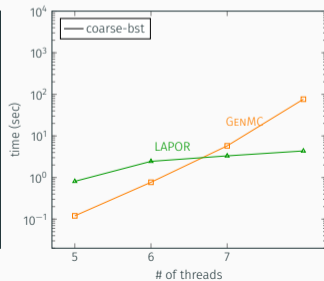
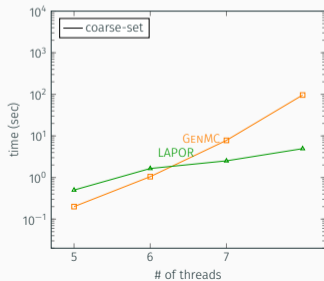
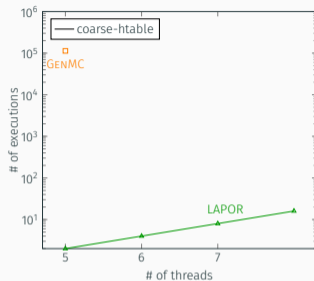
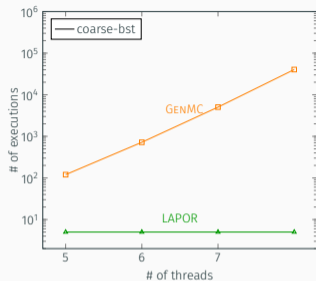
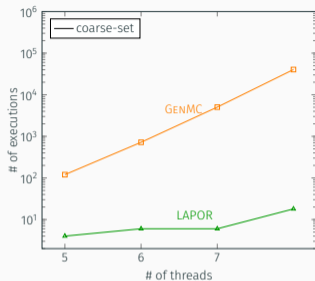
# Coarse-Grained Data-Structures



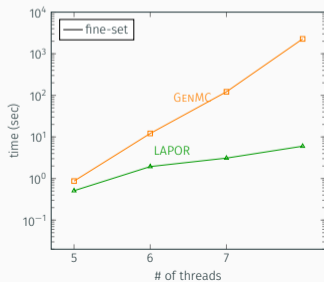
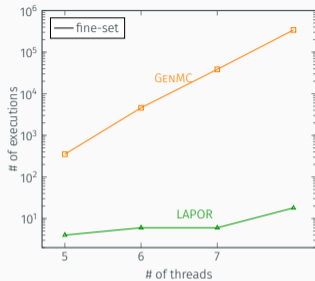
# Coarse-Grained Data-Structures



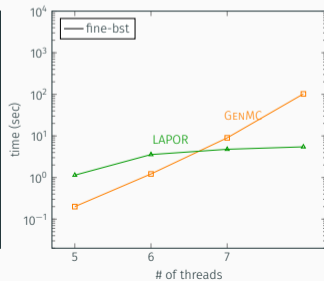
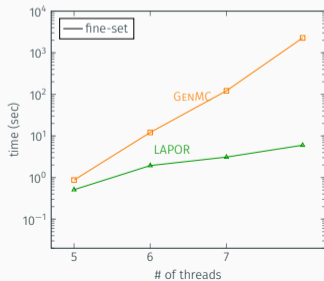
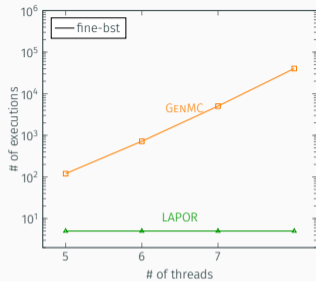
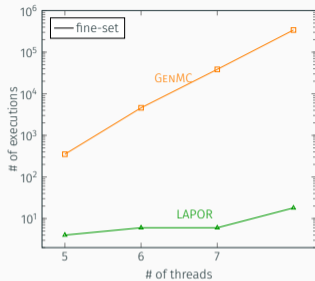
# Coarse-Grained Data-Structures



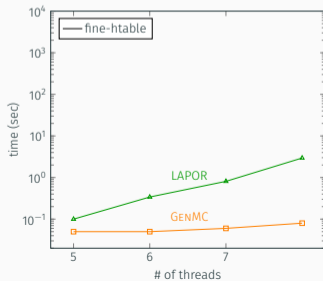
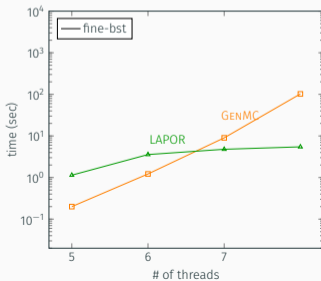
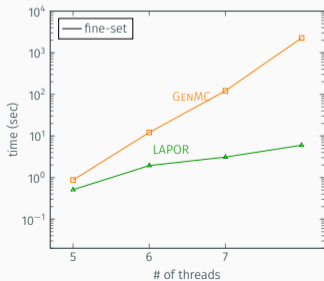
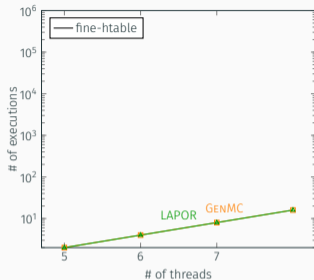
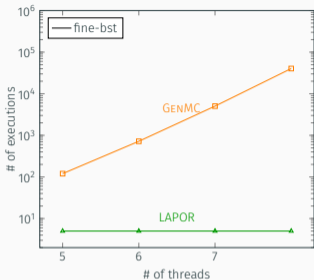
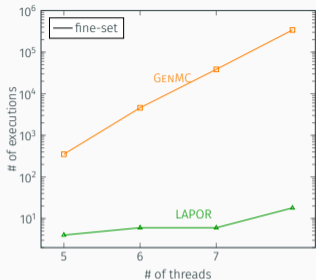
# Fine-Grained Data-Structures



# Fine-Grained Data-Structures



# Fine-Grained Data-Structures



- Detailed description of the **algorithm**
- **Formalization** of the induced ordering's calculation
- More **benchmarks** and **evaluation**

## Summary

- A **Lock-Aware** Partial Order Reduction Algorithm
  - independence at both instruction and critical section levels
  - sound, complete, and optimal
  - works for a variety of memory models
- LAPOR can be **exponentially faster** in programs with locks
- LAPOR is **available** at [github.com/MPI-SWS/genmc](https://github.com/MPI-SWS/genmc)



## Summary

- A **Lock-Aware** Partial Order Reduction Algorithm
  - independence at both instruction and critical section levels
  - sound, complete, and optimal
  - works for a variety of memory models
- LAPOR can be **exponentially faster** in programs with locks
- LAPOR is **available** at [github.com/MPI-SWS/genmc](https://github.com/MPI-SWS/genmc)

## Future work

- Extend the same idea to transactions and other synchronization idioms

**Thank You!**



## Treating Critical Sections As Atomic Blocks

$[l = x = 0]$

<i>lock</i> ( <i>l</i> );		<i>a</i> := <i>x</i>
<i>x</i> := 1;		
<i>x</i> := 2;		
<i>unlock</i> ( <i>l</i> )		