

# Model Checking for Weakly Consistent Libraries

---

*Michalis Kokologiannakis*   Azalea Raad   Viktor Vafeiadis

June 24, 2019

Max Planck Institute for Software Systems (MPI-SWS)

# How Do We Verify Concurrent Programs?

# How Do We Verify Concurrent Programs?

**Stateless Model Checking (SMC):** enumerates all executions

- **without** explicitly *storing* the visited states

**Challenges:**

- State space explosion
- Weak memory

# Challenge #1: State Space Explosion

## Challenge #1: State Space Explosion

$[x = y = 0]$

$x := 1 \parallel x := 2 \parallel y := 42$

Executions:

## Challenge #1: State Space Explosion

$[x = y = 0]$

$x := 1 \parallel x := 2 \parallel y := 42$

Executions:

SMC : 6

## Challenge #1: State Space Explosion

$[x = y = 0]$

$x := 1 \parallel x := 2 \parallel y := 42$

### Executions:

SMC : 6

SMC+POR<sup>mo</sup> : 2

## Challenge #1: State Space Explosion

$[x = y = 0]$

$x := 1 \parallel x := 2 \parallel y := 42$

### Executions:

SMC : 6

SMC+POR<sup>mo</sup> : 2

SMC+POR<sup>porf</sup> : 1



## Challenge #2: Weak Memory Models

## Challenge #2: Weak Memory Models

All current techniques are **memory-model specific**  
⇒ with the exception of **herd**

## Challenge #2: Weak Memory Models

All current techniques are **memory-model specific**  
⇒ with the exception of **herd**

What memory model **properties** are **sufficient** for efficient SMC?

# Our contribution

## Our contribution

- We present sufficient properties for efficient SMC

# Our contribution

- We present sufficient properties for efficient SMC
- **GENMC**: an SMC procedure
  - **parametric** in the choice of the memory model

# Our contribution

- We present sufficient properties for efficient SMC
- **GENMC**: an SMC procedure
  - **parametric** in the choice of the memory model
  - sound, complete, optimal, and efficient

# Our contribution

- We present sufficient properties for efficient SMC
- **GENMC**: an SMC procedure
  - **parametric** in the choice of the memory model (+ libraries!)
  - sound, complete, optimal, and efficient



# Generic Model Checking

---

**Goal:** Enumerate all consistent execution graphs of  $P$  for *any* memory model.


**Goal:** Enumerate all consistent execution graphs of  $P$  for *any* memory model.

$[x = 0]$

$x := 1 \parallel a := x \parallel x := 2$

# GENMC: Generic Model Checking

**Goal:** Enumerate all consistent execution graphs of  $P$  for *any* memory model.



# Systematically Enumerate All Graphs

# Systematically Enumerate All Graphs

 $[x = y = 0]$ 
$$\begin{array}{l|l} a := y & b := x \\ \hline x := a & y := b \end{array}$$


# Systematically Enumerate All Graphs

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

Can the reads of this program read 42?

# Systematically Enumerate All Graphs

|                           |
|---------------------------|
| $[x = y = 0]$             |
| $a := y \parallel b := x$ |
| $x := a \parallel y := b$ |




Can the reads of this program read 42?



# Systematically Enumerate All Graphs

|                           |
|---------------------------|
| $[x = y = 0]$             |
| $a := y \parallel b := x$ |
| $x := a \parallel y := b$ |



Can the reads of this program read 42?

The number of executions may be infinite!

**Goal:** Enumerate all consistent execution graphs of  $P$  for *any* memory model

**Goal:** Enumerate all consistent execution graphs of  $P$  for *any* memory model ,  
where

- $po \cup rf$  is acyclic

## Checking Consistency At Each Step

# Checking Consistency At Each Step

|  |
|--|
| $[x = 0]$                                  |
| $x := 1 \parallel a := x \parallel x := 2$ |



# Checking Consistency At Each Step



|  |
|--|
| $[x = 0]$                                  |
| $x := 1 \parallel a := x \parallel x := 2$ |



②  $[init]$



# Checking Consistency At Each Step

|  |
|--|
| $[x = 0]$                                  |
| $x := 1 \parallel a := x \parallel x := 2$ |



# Checking Consistency At Each Step


|  |
|--|
| $[x = 0]$                                  |
| $x := 1 \parallel a := x \parallel x := 2$ |





# Checking Consistency At Each Step

|  |
|--|
| $[x = 0]$                                  |
| $x := 1 \parallel a := x \parallel x := 2$ |



**Goal:** Enumerate all consistent execution graphs of  $P$  for *any* memory model,  
where

- $po \cup rf$  is acyclic

**Goal:** Enumerate all consistent execution graphs of  $P$  for *any* memory model,  
where

- $po \cup rf$  is acyclic
- Each execution can be obtained from some **linear extension** of  $po \cup rf$

**Goal:** Enumerate all consistent execution graphs of  $P$  for *any* memory model,  
where

- $po \cup rf$  is acyclic
- Each execution can be obtained from ~~some~~ every linear extension of  $po \cup rf$

**Goal:** Enumerate all consistent execution graphs of  $P$  for *any* memory model,  
where

- $po \cup rf$  is acyclic
- Consistency is prefix-closed

# Fixing The Construction Order

## Fixing The Construction Order

|  |
|--|
| $[x = 0]$                                  |
| $x := 1 \parallel a := x \parallel x := 2$ |

## Fixing The Construction Order

|  |
|--|
| $[x = 0]$                                  |
| $x := 1 \parallel a := x \parallel x := 2$ |

[init]



## Fixing The Construction Order

|  |
|--|
| $[x = 0]$                                  |
| $x := 1 \parallel a := x \parallel x := 2$ |

$w(x, 1)$  ← [init]


# Fixing The Construction Order

|  |
|--|
| $[x = 0]$                                  |
| $x := 1 \parallel a := x \parallel x := 2$ |

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


# Fixing The Construction Order

$[x = 0]$   
 $x := 1 \parallel a := x \parallel x := 2$




# Fixing The Construction Order

$[x = 0]$   
 $x := 1 \parallel a := x \parallel x := 2$






# Fixing The Construction Order

[x = 0]  
x := 1 || a := x || x := 2





# Fixing The Construction Order

$[x = 0]$   
 $x := 1 \parallel a := x \parallel x := 2$





# Fixing The Construction Order

[x = 0]  
x := 1 || a := x || x := 2



# Fixing The Construction Order



$[x = 0]$   
 $x := 1 \parallel a := x \parallel x := 2$





# Fixing The Construction Order

$[x = 0]$   
 $x := 1 \parallel a := x \parallel x := 2$



**Goal:** Enumerate all consistent execution graphs of  $P$  for *any* memory model,  
where

- $po \cup rf$  is acyclic
- Consistency is prefix-closed

**Goal:** Enumerate all consistent execution graphs of  $P$  for *any* memory model, *where*

- $po \cup rf$  is acyclic
- Consistency is *prefix-closed*
- The memory model is *extensible*

**Goal:** Enumerate all consistent execution graphs of  $P$  for *any* memory model, *where*

- $po \cup rf$  is acyclic
- Consistency is *prefix-closed*
- The memory model is *extensible*

These are fulfilled by SC, TSO, PSO, RC11

**Goal:** Enumerate all consistent execution graphs of  $P$  for *any* memory model, *where*

- $po \cup rf$  is acyclic
- Consistency is *prefix-closed*
- The memory model is *extensible*

These are fulfilled by SC, TSO, PSO, RC11  
but *not* by POWER and ARM






# Handling Locks

*[init(l)]*

|                  |  |                  |
|------------------|--|------------------|
| <i>lock(l)</i>   |  | <i>lock(l)</i>   |
| <i>unlock(l)</i> |  | <i>unlock(l)</i> |

# Handling Locks





## Results

---

## An Interesting Example

|            | NIDHUGG |                 | RCMC               | GENMC              |                      |
|------------|---------|-----------------|--------------------|--------------------|----------------------|
|            | SC      | SC <sup>o</sup> | RC11 <sup>mo</sup> | RC11 <sup>mo</sup> | RC11 <sup>porf</sup> |
| lamport(2) | .13     | .10             | .04                | .03                | .03                  |
| lamport(3) | 7.53    | 4.49            | 5.40               | 6.87               | 1.36                 |
| lamport(4) | ⏸       | ⏸               | ⏸                  | ⏸                  | ⏸                    |

⏸ = tool did not finish within 2 days

All times are in seconds





## An Interesting Example

|            | NIDHUGG |                 | RCMC               | GENMC              |                      |     |
|------------|---------|-----------------|--------------------|--------------------|----------------------|-----|
|            | SC      | SC <sup>o</sup> | RC11 <sup>mo</sup> | RC11 <sup>mo</sup> | RC11 <sup>porf</sup> | LIB |
| lamport(2) | .13     | .10             | .04                | .03                | .03                  | .09 |
| lamport(3) | 7.53    | 4.49            | 5.40               | 6.87               | 1.36                 | .09 |
| lamport(4) | ⌚       | ⌚               | ⌚                  | ⌚                  | ⌚                    | .09 |


⌚ = tool did not finish within 2 days

All times are in seconds

# Equivalence Partitionings



# Equivalence Partitionings (202 benchmarks)



## More in the paper

- Detailed description of the **algorithm**
- **Formalization** of memory model assumptions
- More **benchmarks** and **evaluation**

## Summary

- **Sound, complete, and optimal** SMC procedure for memory models that are:
  - $po \cup rf$ -acyclic
  - prefix-closed
  - extensible
- GENMC can be **exponentially faster** than existing tools
- GENMC is **available** at [github.com/MPI-SWS/genmc](https://github.com/MPI-SWS/genmc)

## Summary

- **Sound, complete, and optimal** SMC procedure for memory models that are:
  - $po \cup rf$ -acyclic
  - prefix-closed
  - extensible
- GENMC can be **exponentially faster** than existing tools
- GENMC is **available** at [github.com/MPI-SWS/genmc](https://github.com/MPI-SWS/genmc)

## Future work

- Can we relax the memory-model assumptions?

Thank You!





## Why Extensibility Is Necessary

$$[x = y = 0]$$
$$a := x \parallel b := y \parallel x := 42$$

## Why Extensibility Is Necessary

$$[x = y = 0]$$
$$a := x \parallel b := y \parallel x := 42$$


Under a memory model that dictates the following:

*“If a read of  $y$  reads 0, then there cannot be a read of  $x$  that also reads 0”*

# Why Extensibility Is Necessary

$[x = y = 0]$

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




Under a memory model that dictates the following:

*"If a read of y reads 0, then there cannot be a read of x that also reads 0"*

# Why Extensibility Is Necessary

$[x = y = 0]$

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



Under a memory model that dictates the following:


*“If a read of y reads 0, then there cannot be a read of x that also reads 0”*

① [init]

# Why Extensibility Is Necessary


$[x = y = 0]$

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



Under a memory model that dictates the following:


*“If a read of  $y$  reads 0, then there cannot be a read of  $x$  that also reads 0”*



# Why Extensibility Is Necessary


$[x = y = 0]$

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



Under a memory model that dictates the following:


*“If a read of  $y$  reads 0, then there cannot be a read of  $x$  that also reads 0”*



# Why Extensibility Is Necessary


$[x = y = 0]$

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



Under a memory model that dictates the following:


*“If a read of  $y$  reads 0, then there cannot be a read of  $x$  that also reads 0”*







# Handling Locks

| [init(l)]        |                  |
|------------------|------------------|
| $a : lock(l)$    | $b : lock(l)$    |
| $a' : unlock(l)$ | $b' : unlock(l)$ |



# Handling Locks


| [init(l)]      |                |
|----------------|----------------|
| a : lock(l)    | b : lock(l)    |
| a' : unlock(l) | b' : unlock(l) |



W[b]  
[init]




# Handling Locks

| [init(l)]        |                  |
|------------------|------------------|
| $a : lock(l)$    | $b : lock(l)$    |
| $a' : unlock(l)$ | $b' : unlock(l)$ |



# Handling Locks

| $[init(l)]$      |                  |
|------------------|------------------|
| $a : lock(l)$    | $b : lock(l)$    |
| $a' : unlock(l)$ | $b' : unlock(l)$ |



# Handling Locks

|                |                |
|----------------|----------------|
| [init(l)]      |                |
| a : lock(l)    | b : lock(l)    |
| a' : unlock(l) | b' : unlock(l) |



# Linux-Kernel Benchmarks

|                 | <i>Nidhugg<br/>SC</i> | <i>Nidhugg<br/>TSO</i> | <i>Nidhugg<br/>PSO</i> | <i>RCMC<br/>RC17</i> | <i>RCMC<br/>WRC17</i> | <i>GENMC<br/>MO</i> | <i>GENMC<br/>WB</i> |
|-----------------|-----------------------|------------------------|------------------------|----------------------|-----------------------|---------------------|---------------------|
| mcs_spinlock(2) | .12                   | .09                    | .10                    | .05                  | .05                   | .05                 | .05                 |
| mcs_spinlock(3) | 2.98                  | 6.84                   | 12.54                  | .84                  | .67                   | .89                 | .78                 |
| mcs_spinlock(4) | 0.68h                 | 1.51h                  | 3.32h                  | 0.16h                | 0.15h                 | 0.42h               | 0.26h               |
| qspinlock(2)    | .17                   | .11                    | .11                    | .04                  | .04                   | .04                 | .04                 |
| qspinlock(3)    | 10.93                 | 18.20                  | 23.43                  | 2.13                 | 2.08                  | 1.10                | 1.12                |
| seqlock(2)      | .10                   | .09                    | .10                    | .04                  | .04                   | .04                 | .04                 |
| seqlock(3)      | 1.64                  | 3.07                   | 11.00                  | .49                  | .51                   | .37                 | .37                 |