

# Optimal Reads-From Consistency Checking for C11-Style Memory Models

---

Hünkar Can Tunç

Parosh Aziz Abdulla

Soham Chakraborty

Shankaranarayanan Krishna

Umang Mathur

Andreas Pavlogiannis



## Scenario 1

$x := 0, y := 0$

**Thread 1**

$x := 1$


$a := y$

**Thread 2**

$y := 1$

$b := x$

## Scenario 1




### Under sequential consistency:

- Is *Scenario 1* possible?

$\hookrightarrow a = b = 0$

## Scenario 1




**Under sequential consistency:**


- Is *Scenario 1* possible?

$\hookrightarrow a = b = 0$  ✗

## Scenario I




## Scenario II




### Under sequential consistency:

- Is *Scenario I* possible?  
 $\hookrightarrow a = b = 0$  ✗
- Is *Scenario II* possible?  
 $\hookrightarrow a = b = 1$

## Scenario I



## Scenario II



### Under sequential consistency:

- Is *Scenario I* possible?  
 $\hookrightarrow a = b = 0$  ✗
- Is *Scenario II* possible?  
 $\hookrightarrow a = b = 1$  ✓

**Reads-From consistency checking:**

## Reads-From consistency checking:

- **Input:** Partial execution  $X$  in a memory model  $\mathcal{M}$
- **Task:** Check if  $X$  can be consistent in  $\mathcal{M}$



## Reads-From consistency checking:

- **Input:** Partial execution  $X$  in a memory model  $\mathcal{M}$
- **Task:** Check if  $X$  can be consistent in  $\mathcal{M}$
- Practical applications:
  - ↔ Model checking, testing

## Reads-From consistency checking:

- **Input:** Partial execution  $X$  in a memory model  $\mathcal{M}$
- **Task:** Check if  $X$  can be consistent in  $\mathcal{M}$
- Practical applications:
  - ↪ Model checking, testing
- Well understood for traditional memory models
  - ↪ Sequential consistency, x86-TSO

## Reads-From consistency checking:

- **Input:** Partial execution  $X$  in a memory model  $\mathcal{M}$
- **Task:** Check if  $X$  can be consistent in  $\mathcal{M}$
- Practical applications:
  - ↔ Model checking, testing
- Well understood for traditional memory models
  - ↔ Sequential consistency, x86-TSO
- Little is known about the variants of C11 memory model

## Reads-From consistency checking:

- **Input:** Partial execution  $X$  in a memory model  $\mathcal{M}$
- **Task:** Check if  $X$  can be consistent in  $\mathcal{M}$
- Practical applications:
  - ↔ Model checking, testing
- Well understood for traditional memory models
  - ↔ Sequential consistency, x86-TSO
- Little is known about the variants of C11 memory model
  - ↔ **This work fills this gap!**

- We study reads-from consistency checking in various variants of C11
- **Main results:**
  - Efficient algorithms
    - ↪ Optimal or nearly-optimal
  - Complexity characterization
    - ↪ Fine-grained optimality or  $\mathcal{NP}$ -hardness results
  - Empirical evaluation
    - ↪ Shows the impact of new algorithms in practice

# C11 Memory Model

- Introduced by the ISO C/C++ 2011 standards



# C11 Memory Model

- Introduced by the ISO C/C++ 2011 standards
- Support for low-level atomic operations
  - ↔ Load, Store, RMW
  - ↔ Used for communication between threads



# C11 Memory Model

- Introduced by the ISO C/C++ 2011 standards
- Support for low-level atomic operations
  - ↔ Load, Store, RMW
  - ↔ Used for communication between threads
- Memory accesses levels:
  - ↔ Synchronization guarantees
  - ↔ Implementation cost

```
atomic<int> x (0);  
x.store(1, memory_order_relaxed);  
x.load(memory_order_acquire);
```



Sequentially consistent






## C11 Memory Model

- Semantics of a program is defined as a set of consistent executions
- Each execution is a graph
  - ↪ Nodes are instructions in the program
  - ↪ Edges represent certain relations among the instructions

Thread 1	Thread 2
$x := 0;$	$y := 0;$
$a := x;$	$b := y;$
$y := 1;$	$x := 1;$



# C11 Memory Model

- Standard relations:

↪ Program order (po): Precedence among the same thread events

**Thread 1**      **Thread 2**

$x := 0;$

$y := 0;$

$a := x;$

$b := y;$

$y := 1;$

$x := 1;$

**Thread 1**

$W(x, 0)$

po ↓

$R(x, ?)$

po ↓

$W(y, 1)$

**Thread 2**

$W(y, 0)$

po ↓

$R(y, ?)$

po ↓

$W(x, 1)$

# C11 Memory Model

- Standard relations:

↪ Program order (po): Precedence among the same thread events

↪ Reads-from (rf): Relates the writes to the loads which read their value

**Thread 1**      **Thread 2**

$x := 0;$

$y := 0;$

$a := x;$

$b := y;$

$y := 1;$

$x := 1;$

**Thread 1**

**Thread 2**

$W(x, 0)$

$W(y, 0)$

$R(x, 1)$

$R(y, 1)$

$W(y, 1)$

$W(x, 1)$

*rf*


*rf*

# C11 Memory Model

- Standard relations:


- ↪ Program order (po): Precedence among the same thread events
- ↪ Reads-from (rf): Relates the writes to the loads which read their value
- ↪ Modification order (mo): A total order of the writes on a given location

Thread 1	Thread 2
$x := 0;$	$y := 0;$
$a := x;$	$b := y;$
$y := 1;$	$x := 1;$




## C11-Style Memory Models

- A memory model restricts which executions are consistent



## C11-Style Memory Models


- A memory model restricts which executions are consistent



Memory model  $\mathcal{M}1$  :  $\text{acyclic}(\text{po} \cup \text{mo})$

## C11-Style Memory Models


- A memory model restricts which executions are consistent



Memory model  $\mathcal{M}1$  :  $\text{acyclic}(\text{po} \cup \text{mo})$

## C11-Style Memory Models

- A memory model restricts which executions are consistent




Memory model  $\mathcal{M}1$  :  $\text{acyclic}(\text{po} \cup \text{mo})$

Memory model  $\mathcal{M}2$  :  $\text{acyclic}(\text{po} \cup \text{mo}) \wedge \text{acyclic}(\text{po} \cup \text{rf})$



## C11-Style Memory Models

- A memory model restricts which executions are consistent



Memory model  $\mathcal{M}_1$  :  $\text{acyclic}(\text{po} \cup \text{mo})$

Memory model  $\mathcal{M}_2$  :  $\text{acyclic}(\text{po} \cup \text{mo}) \wedge \text{acyclic}(\text{po} \cup \text{rf})$

## Reads-From Consistency Checking

- **Input:** A partial execution  $X$  in a memory model  $\mathcal{M}$ 
  - $\hookrightarrow X$  contains  $po$  and  $rf$
  - $\hookrightarrow X$  lacks  $mo$

## Reads-From Consistency Checking

- **Input:** A partial execution  $X$  in a memory model  $\mathcal{M}$ 
  - $\hookrightarrow X$  contains  $po$  and  $rf$
  - $\hookrightarrow X$  lacks  $mo$
- **Task:** Check if  $X$  can be extended to a complete execution consistent in  $\mathcal{M}$ 
  - $\hookrightarrow$  Find an  $mo$  that turns  $X$  consistent

## RC20 Memory Model

- Captures a rich fragment of C11
  - ↪ Contains Release, Acquire and Relaxed accesses
  - ↪ Lacks Sequentially Consistent accesses

- Captures a rich fragment of C11
  - ↪ Contains Release, Acquire and Relaxed accesses
  - ↪ Lacks Sequentially Consistent accesses
- Reads-from consistency checking is a bottleneck
- **Previous works:**  $O(n^3 \cdot k)$ ,  $O(n^2 \cdot k)$ 
  - ↪ For  $n$  events and  $k$  threads

- Captures a rich fragment of C11
  - ↪ Contains Release, Acquire and Relaxed accesses
  - ↪ Lacks Sequentially Consistent accesses
- Reads-from consistency checking is a bottleneck
- **Previous works:**  $O(n^3 \cdot k)$ ,  $O(n^2 \cdot k)$ 
  - ↪ For  $n$  events and  $k$  threads
- **Our result:**  $O(n \cdot k)$ 
  - ↪ Key idea: *minimal coherence* witness relation

## Witness relation

- ↪ Serves as a witness for consistency
- ↪ Construct a partially ordered **mo**
  - ↪ Include *necessary* orderings enforced by the memory model
  - ↪ It should be extendable to a total **mo**

## Witness relation

- ↪ Serves as a witness for consistency
- ↪ Construct a partially ordered **mo**
  - ↪ Include *necessary* orderings enforced by the memory model
  - ↪ It should be extendable to a total **mo**


Recall →

### Reads-from consistency checking:


- ↪ Task is to find an **mo**
- ↪ **mo** is a total order on the same location writes



# Witness Relation




## Witness Relation




- $W_2(x)$  must be **mo** ordered before  $W_1(x)$

## Witness Relation



- $W_2(x)$  must be **mo** ordered before  $W_1(x)$
- $W_3(x)$  is not relevant
  - ↪ It can be left unordered

## Witness Relation




- $W_2(x)$  must be **mo** ordered before  $W_1(x)$
- $W_3(x)$  is not relevant
  - ↪ It can be left unordered
- Witness should always be extendable to a total **mo**
  - ↪  $W_3(x) \overset{\text{mo}}{\dashrightarrow} W_2(x)$
  - ↪  $W_1(x) \overset{\text{mo}}{\dashrightarrow} W_3(x)$

## Minimal coherence


- ↪ Serves as a witness for consistency
- ↪ Weaker than prior witness relations
- ↪ Allows efficient consistency checking algorithm for RC20

# Minimal Coherence for RC20

## Minimal coherence




## Prior witness relations




# Minimal Coherence for RC20

## Minimal coherence




## Prior witness relations




# Minimal Coherence for RC20

## Minimal coherence




## Prior witness relations






## Minimal coherence



## Prior witness relations



- Minimal coherence is weaker!  
     $\hookrightarrow$  More efficient to compute

# Experimental Results



- Focused on the RC20/Release-Acquire (RA) fragments
  - ↪ RA is a fragment of RC20
- Performed an evaluation in two scenarios
  - ↪ Model checking
  - ↪ Testing
- Modified only the consistency checking components

# Experimental Results - Model Checking

- Implemented minimal coherence inside GenMC<sup>1</sup>
- Compared with the original GenMC
  - ↪ 25 standard benchmarks
  - ↪ 2 hour timeout



	<b>GenMC</b>	<b>Our Algorithm</b>
Average time per execution	14.5 sec	0.26 sec
Total number of executions	356K	4.6M

<sup>1</sup>Michalis Kokologiannakis, Viktor Vafeiadis. GenMC: A Model Checker for Weak Memory Models. CAV'21

# Experimental Results - Testing

- Implemented minimal coherence inside C11Tester<sup>1</sup>
  - ↪ Online version
  - ↪  $O(n \cdot k)$  bound does not apply
- Compared with the original C11Tester
  - ↪ 32 standard benchmarks



	<b>C11Tester</b>	<b>Our Algorithm</b>
Total analysis time	286 sec	170 sec

<sup>1</sup>Weiyu Luo, Brian Demsky. C11Tester: a race detector for C/C++ atomics. ASPLOS'21

## Summary of Established Bounds

- **NP-hard:**  $O(k \cdot n^{k+1})$

↪ Strong Release-Acquire (SRA)

$n \rightarrow$  number of events

$k \rightarrow$  number of threads

## Summary of Established Bounds

- **NP-hard:**  $O(k \cdot n^{k+1})$ 
  - ↪ Strong Release-Acquire (SRA)
- **Super-linear:**  $O(n \cdot k)$ 
  - ↪ RC20
  - ↪ Weak Release-Acquire (WRA)
  - ↪ RMW-free SRA

$n \rightarrow$  number of events

$k \rightarrow$  number of threads

## Summary of Established Bounds

- **NP-hard:**  $O(k \cdot n^{k+1})$ 
  - ↪ Strong Release-Acquire (SRA)
- **Super-linear:**  $O(n \cdot k)$ 
  - ↪ RC20
  - ↪ Weak Release-Acquire (WRA)
  - ↪ RMW-free SRA
- **Linear:**  $O(n)$ 
  - ↪ Relaxed

$n \rightarrow$  number of events

$k \rightarrow$  number of threads

## Summary of Established Bounds

- **NP-hard:**  $O(k \cdot n^{k+1})$ 
    - ↪ Strong Release-Acquire (SRA)
  - **Super-linear:**  $O(n \cdot k)$ 
    - ↪ RC20
    - ↪ Weak Release-Acquire (WRA)
    - ↪ RMW-free SRA
  - **Linear:**  $O(n)$ 
    - ↪ Relaxed
  - **Super-linear lower bound:**
    - ↪ RMW-free RA, SRA, WRA
    - ↪ Improving  $O(n \cdot k)$  bounds would be non-trivial
- $n \rightarrow$  number of events  
 $k \rightarrow$  number of threads



- Addressed the reads-from consistency checking problem in variants of C11

## Conclusion

- Addressed the reads-from consistency checking problem in variants of C11
- Collection of optimal or nearly-optimal algorithms for different variants

## Conclusion

- Addressed the reads-from consistency checking problem in variants of C11
- Collection of optimal or nearly-optimal algorithms for different variants
- Established fine-grained complexity results

## Conclusion

- Addressed the reads-from consistency checking problem in variants of C11
- Collection of optimal or nearly-optimal algorithms for different variants
- Established fine-grained complexity results
- Experimental evaluation confirms the impact of the new algorithms

- Addressed the reads-from consistency checking problem in variants of C11
- Collection of optimal or nearly-optimal algorithms for different variants
- Established fine-grained complexity results
- Experimental evaluation confirms the impact of the new algorithms

**Thank you!**