

# Reasoning over Relaxed Shared Memory Models: A Tutorial

Brijesh Dongol  
University of Surrey, UK

with thanks to

Lara Bargmann, Mark Batty, Sadegh Dalvandi, Simon Doherty, Ori Lahav, Azalea Raad,  
Mikhail Semenyuk, Eleni Vafeiadi Bila, Heike Wehrheim, John Wickerson, Daniel Wright

# Tutorial Plan

## 1 Background

- Sequential consistency
- Owicki-Gries reasoning
- Isabelle/HOL
- Examples: Message passing, Lamport buffer

## 2 Relaxed Memory

- Motivation
- RC11 Memory Model (Relaxed Accesses)

## 3 Verification

- Owicki-Gries reasoning, revisited
- Examples, revisited
- Isabelle/HOL

## 4 Formal Model

## 5 Release/Acquire Accesses

- Lamport buffer, revisited again
- Message passing
- Isabelle/HOL


# Background

# Sequential Consistency

- “... the result of **any execution** is the same as if
- **the operations of all the processors** were executed **in some sequential order**, and
  - **the operations of each individual processor appear** in this sequence **in the order specified by its program.**”

— *[Lamport, 1979]*

# Sequentially Consistent Architecture




Concurrent execution  $T_1 || T_2 || \dots || T_n$

$\equiv$

Interleaving of atomic statements in

$T_1, T_2, \dots, T_n$

# Sequentially Consistent Architecture



Concurrent execution  $T_1 \parallel T_2 \parallel \dots \parallel T_n$

$\equiv$

Interleaving of atomic statements in

$T_1, T_2, \dots, T_n$

## Example

I: $y = 0$	
$a_1: x := 1;$	$b_1: r := y;$
$a_2: y := 1$	$b_2: s := x$

## Possible executions

- I;  $a_1$ ;  $a_2$ ;  $b_1$ ;  $b_2$
- I;  $a_1$ ;  $b_1$ ;  $a_2$ ;  $b_2$
- I;  $a_1$ ;  $b_1$ ;  $b_2$ ;  $a_2$
- + symmetric cases starting with  $b_1$

# Sequentially Consistent Architecture

*... makes a memory operation appear to execute atomically (instantaneously) with respect to other memory operations. — [Adve and Gharachorloo, 1996]*

# Sequentially Consistent Architecture

*... makes a memory operation appear to execute atomically (instantaneously) with respect to other memory operations. — [Adve and Gharachorloo, 1996]*

# Sequentially Consistent Architecture

*... makes a memory operation appear to execute atomically (instantaneously) with respect to other memory operations. — [Adve and Gharachorloo, 1996]*

Convenient state model:  $\Sigma \hat{=} Loc \rightarrow Val$

# Sequentially Consistent Architecture

... makes a memory operation appear to execute atomically (instantaneously) with respect to other memory operations. — [Adve and Gharachorloo, 1996]

Convenient state model:  $\Sigma \hat{=} Loc \rightarrow Val$

**Example.**

$$\begin{array}{c} \text{I: } y = 0 \\ \mathbf{a}_1: x := 1; \quad \parallel \quad \mathbf{b}_1: r := y; \\ \mathbf{a}_2: y := 1 \quad \parallel \quad \mathbf{b}_2: s := x \end{array}$$

Traces are of the form:

$$\begin{array}{c} [x \mapsto ?, \\ y \mapsto 0, \\ r \mapsto ?, \\ s \mapsto ?] \end{array} \xrightarrow{x:=1} \begin{array}{c} [x \mapsto 1, \\ y \mapsto 0, \\ r \mapsto ?, \\ s \mapsto ?] \end{array} \xrightarrow{r:=y} \begin{array}{c} [x \mapsto 1, \\ y \mapsto 0, \\ r \mapsto 0, \\ s \mapsto ?] \end{array} \xrightarrow{y:=1} \begin{array}{c} [x \mapsto 1, \\ y \mapsto 1, \\ r \mapsto 0, \\ s \mapsto ?] \end{array} \xrightarrow{s:=x} \begin{array}{c} [x \mapsto 1, \\ y \mapsto 1, \\ r \mapsto 0, \\ s \mapsto 1] \end{array}$$

# Verification

- Dozens of techniques developed since the 1970s [de Roever et al., 2001]:
  - Owicki-Gries
  - Rely-guarantee
  - TLA
  - Action systems
  - Input/output automata
  - Concurrent separation logic
  - ...
- Let's focus on one: Owicki-Gries [Owicki, 1975]
- Extension of Hoare Logic with rules for “interference freedom”

## Owicki-Gries proof outline

**Example (Message passing).** For *any* interleaving of threads, if thread  $b$  reads  $y = 1$ , then it must read  $x = 1$ .

$$\begin{array}{l|l} \{true\} & \{y = 0\} \\ x := 1; & \{y = 1 \Rightarrow x = 1\} \\ \{x = 1\} & r := y; \\ y := 1 & \{r = 1 \Rightarrow x = 1\} \\ \{true\} & s := x \\ & \{r = 1 \Rightarrow s = 1\} \\ & \{r = 1 \Rightarrow s = 1\} \end{array}$$

## Owicki-Gries proof outline

**Example (Message passing).** For *any* interleaving of threads, if thread  $b$  reads  $y = 1$ , then it must read  $x = 1$ .

$$\boxed{\begin{array}{l|l} \{true\} & \{y = 0\} \\ x := 1; & \{y = 1 \Rightarrow x = 1\} \\ \{x = 1\} & r := y; \\ y := 1 & \{r = 1 \Rightarrow x = 1\} \\ \{true\} & s := x \\ & \{r = 1 \Rightarrow s = 1\} \\ & \{r = 1 \Rightarrow s = 1\} \end{array}}$$

### 1 Initialisation

1. Initialisation proof:

$$y = 0 \Rightarrow true \wedge (y = 1 \Rightarrow x = 1)$$

# Owicki-Gries proof outline

**Example (Message passing).** For *any* interleaving of threads, if thread  $b$  reads  $y = 1$ , then it must read  $x = 1$ .

$$\begin{array}{l|l} \{true\} & \{y = 0\} \\ x := 1; & \{y = 1 \Rightarrow x = 1\} \\ \{x = 1\} & r := y; \\ y := 1 & \{r = 1 \Rightarrow x = 1\} \\ \{true\} & s := x \\ & \{r = 1 \Rightarrow s = 1\} \\ & \{r = 1 \Rightarrow s = 1\} \end{array}$$

- 1 Initialisation
- 2 Local correctness

2. Local correctness proofs:

$$\begin{array}{l} \{true\} x := 1 \{x = 1\} \\ \{x = 1\} y := 1 \{true\} \end{array}$$

$$\begin{array}{l} \{y = 1 \Rightarrow x = 1\} r := y \{r = 1 \Rightarrow x = 1\} \\ \{r = 1 \Rightarrow x = 1\} s := x \{r = 1 \Rightarrow s = 1\} \end{array}$$

## Owicki-Gries proof outline

**Example (Message passing).** For *any* interleaving of threads, if thread  $b$  reads  $y = 1$ , then it must read  $x = 1$ .

$\{true\}$	$\{y = 0\}$
$x := 1;$	$\{y = 1 \Rightarrow x = 1\}$
$\{x = 1\}$	$r := y;$
$y := 1$	$\{r = 1 \Rightarrow x = 1\}$
$\{true\}$	$s := x$
	$\{r = 1 \Rightarrow s = 1\}$

- 1 Initialisation
- 2 Local correctness
- 3 Global correctness

### 3. Global correctness proofs:

For  $P \in \{true, x = 1\}$ ,

$\{P \wedge (y = 1 \Rightarrow x = 1)\} r := y \{P\}$

$\{P \wedge (r = 1 \Rightarrow x = 1)\} s := x \{P\}$


For  $Q \in \left\{ \begin{array}{l} y = 1 \Rightarrow x = 1, \quad r = 1 \Rightarrow x = 1, \\ r = 1 \Rightarrow s = 1 \end{array} \right\}$ ,

$\{Q \wedge true\} x := 1 \{Q\}$

$\{Q \wedge x = 1\} y := 1 \{Q\}$

## Owicki-Gries proof outline

**Example (Message passing).** For *any* interleaving of threads, if thread  $b$  reads  $y = 1$ , then it must read  $x = 1$ .



- 1 Initialisation
- 2 Local correctness
- 3 Global correctness
- 4 Finalisation



4. Finalisation proof:

$$true \wedge (r = 1 \Rightarrow s = 1) \Rightarrow r = 1 \Rightarrow s = 1$$

# Message Passing under SC

Isabelle/HOL Demo

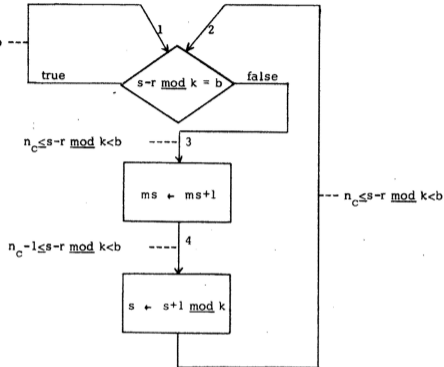
# Lamport Buffer



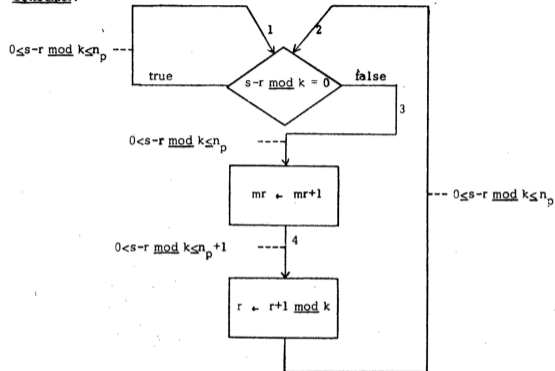
- Synchronisation via shared variables  $s$  and  $r$
- Auxiliary variables
  - $m_s = \#$ messages sent
  - $m_r = \#$ messages received
- Safety invariant  $0 \leq m_s - m_r \leq b$

# Lamport's Proof

**Producer:**



**Consumer:**




Let  $n = m_s - m_r$

$$n_c = \begin{cases} n + 1 & \text{if the consumer's token is on arc 4} \\ n & \text{otherwise} \end{cases}$$

$$n_p = \begin{cases} n - 1 & \text{if the producer's token is on arc 4} \\ n & \text{otherwise} \end{cases}$$

# Owicki-Gries Proof of Lamport's Buffer




- Define  $I \hat{=} 0 \leq n \leq b$
- Highlighted assertions can be omitted

Isabelle/HOL Demo

# Relaxed Memory

# (Modern?) Multiprocessors and Languages


## Hardware memory models



- Store buffer:
  - Intel (FIFO)
  - Arm, Power (Weakly ordered)
- Shared memory:
  - Intel, Armv8/v9 (Multi-copy atomic)
  - Armv7, Power (Not multi-copy atomic)

# (Modern?) Multiprocessors and Languages

## Hardware memory models



- Store buffer:  
Intel (FIFO)  
Arm, Power (Weakly ordered)
- Shared memory:  
Intel, Armv8/v9 (Multi-copy atomic)  
Armv7, Power (Not multi-copy atomic)

## Language memory models

C11, C20, RC11, OCaml, Java, LLVM, WASM ...

## Kernel memory models


Linux, OpenBSD, ...

## Heterogenous models

CUDA, PTX, RDMA, P<sub>x</sub>86, ...

# (Modern?) Multiprocessors and Languages

## Hardware memory models



- Store buffer:  
Intel (FIFO)  
Arm, Power (Weakly ordered)
- Shared memory:  
Intel, Armv8/v9 (Multi-copy atomic)  
Armv7, Power (Not multi-copy atomic)

## Language memory models

C11, C20, RC11, OCaml, Java, LLVM, WASM ...

## Kernel memory models

Linux, OpenBSD, ...

## Heterogenous models

CUDA, PTX, RDMA, P<sub>x</sub>86, ...

## Message Passing Under Relaxed Memory

$$\begin{array}{c} y = 0 \\ x := 1; \parallel r := y; \\ y := 1 \parallel s := x \end{array}$$

Is the final outcome  $r = 1 \wedge s \neq 1$  allowed?

- Arm ✓
- Intel-TSO ✗
- C11 - it depends

## Message Passing Under Relaxed Memory

$$\begin{array}{c} y = 0 \\ x := 1; \parallel r := y; \\ y := 1 \parallel s := x \end{array}$$

Is the final outcome  $r = 1 \wedge s \neq 1$  allowed?

- Arm ✓
- Intel-TSO ✗
- C11 - it depends

**Conclusion:** Relaxed memory makes concurrency *much* more complicated

# Message Passing Under Relaxed Memory

$$\begin{array}{c} y = 0 \\ x := 1; \parallel r := y; \\ y := 1 \parallel s := x \end{array}$$


Is the final outcome  $r = 1 \wedge s \neq 1$  allowed?

- Arm ✓
- Intel-TSO ✗
- C11 - it depends

**Conclusion:** Relaxed memory makes concurrency *much* more complicated


... but this is not a new phenomenon!

# Relaxed Memory (Historically)



Memory consistency hierarchy  
[Steinke and Nutt, 2004]

# Relaxed Memory (Historically)



Memory consistency hierarchy  
[Steinke and Nutt, 2004]

1996

Theme Feature



## Shared Memory Consistency Models: A Tutorial

**T**he shared memory programming model has several advantages over the message passing model. In particular, it simplifies data partitioning and dynamic load distribution. Shared memory systems are therefore gaining wide acceptance for both technical and commercial computing.

To write correct and efficient shared memory programs, programmers need a precise notion of shared memory semantics. For example, in the program in Figure 1 (a fragment from a program in the Splash application suite), processor P1 repeatedly updates a data field in a new task record and then inserts the record into a task queue. When no tasks are left, P1 updates a pointer, `Head`, to point to the first record in the task queue. Meanwhile, the other processors wait for `Head` to have a non-null value, dequeue the task pointed to by `Head` in a critical section, and read the data in the dequeued task. To ensure correct execution, a programmer expects that the data value read should be the same as that written by P1. However, in many commercial shared memory systems, the processors may observe an older value, causing unexpected behavior.

Sarita V. Adve  
Rice University

Kourosh Gharachorloo  
Digital Equipment Corporation

[Adve and Gharachorloo, 1996]

# Relaxed Memory (Historically)

1993

## Memory Consistency Models

David Mosberger

Department of Computer Science  
University of Arizona  
Tucson, AZ 85721  
davidm@cs.arizona.edu

### Abstract

This paper discusses memory consistency models and their influence on software in the context of parallel machines. In the first part we review previous work on memory consistency models. The second part discusses the issues that arise due to weakening memory consistency. We are especially interested in the influence that weakened consistency models have on language, compiler, and runtime system design. We conclude that tighter interaction between those parts and the memory system might improve performance considerably.

### 1 Introduction

Traditionally, memory consistency models were of interest only to computer architects designing parallel machines. The goal was to present a model as close as possible to the model exhibited by sequential machines. The model of choice was sequential consistency (SC). Sequential consistency guarantees that the result of any execution of  $n$  processors is the same as if the operations of all processors were executed in some sequential

as the consistency model becomes weaker. That is, an architecture can employ a weaker memory model only if the software using it is prepared to deal with the new programming model. Consequently, memory consistency models are now of concern to operating system and language designers too.

We can also turn the coin around. A compiler normally considers memory accesses to be expensive and therefore tries to replace them by accesses to registers. In terms of a memory consistency model, this means that certain accesses suddenly are not observable any more. In effect, compilers implicitly generate weak memory consistency. This is possible because a compiler knows exactly (or estimates conservatively) the points where memory has to be consistent. For example, compilers typically write back register values before a function call, thus ensuring consistency. It is only natural to attempt to make this implicit weakening explicit in order to let the memory system take advantage too. In fact, it is anticipated that software could gain from a weak model to a much higher degree than hardware [GGH91] by enabling optimizations such as code scheduling or delaying updates that are not legal under SC.

1993

## Memory Consistency Models

David Mosberger

Department of Computer Science  
University of Arizona  
Tucson, AZ 85721  
davidm@cs.arizona.edu

### Abstract

This paper discusses memory consistency models and their influence on software in the context of parallel machines. In the first part we review previous work on memory consistency models. The second part discusses the issues that arise due to weakening memory consistency. We are especially interested in the influence that weakened consistency models have on language, compiler, and runtime system design. We conclude that tighter interaction between those parts and the memory system might improve performance considerably.

### 1 Introduction

Traditionally, memory consistency models were of interest only to computer architects designing parallel machines. The goal was to present a model as close as possible to the model exhibited by sequential machines. The model of choice was sequential consistency (SC). Sequential consistency guarantees that the result of any execution of  $n$  processors is the same as if the operations of all processors were executed in some sequential

as the consistency model becomes weaker. That is, an architecture can employ a weaker memory model only if the software using it is prepared to deal with the new programming model. Consequently, memory consistency models are now of concern to operating system and language designers too.

We can also turn the coin around. A compiler normally considers memory accesses to be expensive and therefore tries to replace them by accesses to registers. In terms of a memory consistency model, this means that certain accesses suddenly are not observable any more. In effect, compilers implicitly generate weak memory consistency. This is possible because a compiler knows exactly (or estimates conservatively) the points where memory has to be consistent. For example, compilers typically write back register values before a function call, thus ensuring consistency. It is only natural to attempt to make this implicit weakening explicit in order to let the memory system take advantage too. In fact, it is anticipated that software could gain from a weak model to a much higher degree than hardware [GGH91] by enabling optimizations such as code scheduling or delaying updates that are not legal under SC.

1988

## Efficient and Correct Execution of Parallel Programs that Share Memory

DENNIS SHASHA

Courant Institute of Mathematical Sciences, New York University  
and

MARC SNIR

IBM T. J. Watson Research Center

---

In this paper we consider an optimization problem that arises in the execution of parallel programs on shared-memory multiple-instruction-stream, multiple-data-stream (MIMD) computers. A program on such machines consists of many sequential program segments, each executed by a single processor. These segments interact as they access shared variables. Access to memory is asynchronous, and memory accesses are not necessarily executed in the order they were issued. An execution is correct if it is sequentially consistent: It should seem as if all the instructions were executed sequentially, in an order obtained by interleaving the instruction streams of the processors. Sequential consistency can be enforced by delaying each access to shared memory until the previous access of the same processor has terminated. For performance reasons, however, we want to allow several accesses by the same processor to proceed concurrently. Our analysis finds a minimal set of delays that enforces sequential consistency. The analysis extends to interprocessor synchronization constraints and to code where blocks of operations have to execute atomically. We use a conflict graph similar to that used to schedule transactions in distributed databases. Our graph incorporates the order on operations given by the program text, enabling us to do without locks even when database conflict graphs would suggest that locks are necessary. Our work has implications for the design of multiprocessors; it offers new compiler optimization techniques for parallel languages that support shared variables.

# Relaxed Memory (Historically)

... achieving sequential consistency may not be worth the price of slowing down the processors.

1979

## How to Make a Multiprocessor Computer That Correctly Executes Multiprocess Programs

LESLIE LAMPORT

*Abstract*—Many large sequential computers execute operations in a different order than is specified by the program. A correct execution is achieved if the results produced are the same as would be produced by executing the program steps in order. For a multiprocessor computer, such a correct execution by each processor does not guarantee the correct execution of the entire program. Additional conditions are given which do guarantee that a computer correctly executes multiprocess programs.

# Relaxed Memory (Historically)

1979

## How to Make a Multiprocessor Computer That Correctly Executes Multiprocess Programs

LESLIE LAMPORT

*Abstract*—Many large sequential computers execute operations in a different order than is specified by the program. A correct execution is achieved if the results produced are the same as would be produced by executing the program steps in order. For a multiprocessor computer, such a correct execution by each processor does not guarantee the correct execution of the entire program. Additional conditions are given which do guarantee that a computer correctly executes multiprocess programs.

... achieving sequential consistency may not be worth the price of slowing down the processors. *In this case, one must be aware that conventional methods for designing multiprocess algorithms cannot be relied upon to produce correctly executing programs.*

# Relaxed Memory (Historically)

1979

## How to Make a Multiprocessor Computer That Correctly Executes Multiprocess Programs

LESLIE LAMPORT

*Abstract*—Many large sequential computers execute operations in a different order than is specified by the program. A correct execution is achieved if the results produced are the same as would be produced by executing the program steps in order. For a multiprocessor computer, such a correct execution by each processor does not guarantee the correct execution of the entire program. Additional conditions are given which do guarantee that a computer correctly executes multiprocess programs.

... achieving sequential consistency may not be worth the price of slowing down the processors. *In this case, one must be aware that conventional methods for designing multiprocess algorithms cannot be relied upon to produce correctly executing programs. Protocols for synchronizing the processors must be designed at the lowest level of the machine instruction code, and verifying their correctness becomes a monumental task.*

# Relaxed Memory (Recently)

2010s/2020s explosion of formal models and results

- [Batty et al., 2011] presented the first formalisation of the C11 memory model
- [Alglave et al., 2014] defined the axiomatic CAT language for formalising consistency models
- [Lahav et al., 2017] discovered inconsistencies in the C11 specification, and proposed a fix called RC11

For this tutorial, we focus on **RC11** —

- Program verification focusses on language-level relaxed-memory
- Rely on compiler to map to hardware memory model

## Repaired C11 (RC11)

# Repairing C11

[Lahav et al., 2017] discovered three key problems with the C11 memory model

- 1 Compilation from C11 to PowerPC is unsound
- 2 Fences (used to enforce instruction order) are too weak
- 3 Model allows out-of-thin-air reads

# Repairing C11


[Lahav et al., 2017] discovered three key problems with the C11 memory model

- 1 Compilation from C11 to PowerPC is unsound
- 2 Fences (used to enforce instruction order) are too weak
- 3 Model allows out-of-thin-air reads


Proposed RC11 memory model as a fix

- Introduces an axiom that **maintains program order**
- Allows us to view weak memory programs as an interleaving of threads
- Relaxed behaviour induces **delayed propagation** of writes

# RC11 Memory Accesses



# RC11 Memory Accesses



**We will focus on relaxed, release and acquire**


## RC11 Models

$$\begin{array}{c} y = 0 \\ [x] : \overset{x}{=} 1; \quad \parallel \quad r : \overset{x}{=} [y]; \\ [y] : \overset{x}{=} 1 \quad \parallel \quad s : \overset{x}{=} [x] \end{array}$$

## RC11 Models

$$\boxed{\begin{array}{c} y = 0 \\ [x] : \stackrel{x}{=} 1; \quad \parallel \quad r : \stackrel{x}{=} [y]; \\ [y] : \stackrel{x}{=} 1 \quad \parallel \quad s : \stackrel{x}{=} [x] \end{array}}$$


Axiomatic (aka declarative) semantics



## RC11 Models

$$\boxed{\begin{array}{c} y = 0 \\ [x] : \overset{x}{=} 1; \quad \parallel \quad r : \overset{x}{=} [y]; \\ [y] : \overset{x}{=} 1 \quad \parallel \quad s : \overset{x}{=} [x] \end{array}}$$


Axiomatic (aka declarative) semantics



## RC11 Models

$$\begin{array}{c} y = 0 \\ [x] : \overset{x}{=} 1; \quad \parallel \quad r : \overset{x}{=} [y]; \\ [y] : \overset{x}{=} 1 \quad \parallel \quad s : \overset{x}{=} [x] \end{array}$$


Axiomatic (aka declarative) semantics



## RC11 Models

$$\begin{array}{c} y = 0 \\ [x] : \stackrel{x}{=} 1; \quad \parallel \quad r : \stackrel{x}{=} [y]; \\ [y] : \stackrel{x}{=} 1 \quad \parallel \quad s : \stackrel{x}{=} [x] \end{array}$$


Axiomatic (aka declarative) semantics



# RC11 Models

$$\begin{array}{c}
 y = 0 \\
 [x] : \stackrel{x}{=} 1; \quad \parallel \quad r : \stackrel{x}{=} [y]; \\
 [y] : \stackrel{x}{=} 1 \quad \parallel \quad s : \stackrel{x}{=} [x]
 \end{array}$$

Axiomatic (aka declarative) semantics




**Allowed**

# RC11 Models

$$\begin{array}{c}
 y = 0 \\
 [x] : \overset{x}{=} 1; \quad \parallel \quad r : \overset{x}{=} [y]; \\
 [y] : \overset{x}{=} 1 \quad \parallel \quad s : \overset{x}{=} [x]
 \end{array}$$

Axiomatic (aka declarative) semantics




**Allowed**

# RC11 Models

$$\begin{array}{c}
 y = 0 \\
 [x] : \stackrel{x}{=} 1; \quad \parallel \quad r : \stackrel{x}{=} [y]; \\
 [y] : \stackrel{x}{=} 1 \quad \parallel \quad s : \stackrel{x}{=} [x]
 \end{array}$$

Axiomatic (aka declarative) semantics




**Allowed**

# RC11 Models

$$\begin{array}{c}
 y = 0 \\
 [x] : \stackrel{x}{=} 1; \quad \parallel \quad r : \stackrel{x}{=} [y]; \\
 [y] : \stackrel{x}{=} 1 \quad \parallel \quad s : \stackrel{x}{=} [x]
 \end{array}$$

Axiomatic (aka declarative) semantics




**Allowed**


# RC11 Models

$$\begin{array}{c}
 y = 0 \\
 [x] : \stackrel{x}{=} 1; \quad \parallel \quad r : \stackrel{x}{=} [y]; \\
 [y] : \stackrel{x}{=} 1 \quad \parallel \quad s : \stackrel{x}{=} [x]
 \end{array}$$

Axiomatic (aka declarative) semantics



**Allowed**



**Disallowed**

# Axiomatic Models

- Often the “authoritative” formal model, particularly in hardware E.g.,  
<https://github.com/herd/herdtools7/blob/master/herd/libdir/aarch64.cat>

```
(** External visibility requirement **)
irreflexive ob as external

(** Internal visibility requirements **)
irreflexive [Exp & R]; ((po & same-loc) | rmw); [Exp & W];
    rfi; [Exp & R] as coRW1-Exp
irreflexive [Imp & Tag & R]; (po & same-loc); [Exp & Tag & W];
    rfi; [Imp & Tag & R] as coRW1-MTE
irreflexive [Exp & W]; (po & same-loc); [Exp & W];
    (ca & int); [Exp & W] as coWW-Exp
irreflexive [Exp & W]; (po & same-loc); [Exp & R];
    (ca & int); [Exp & W] as coWR-Exp
irreflexive [Exp & Tag & W]; (po & same-loc); [Imp & Tag & R];
    (ca & int) as coWR-MTE
```

## Axiomatic Models

- Often the “authoritative” formal model, particularly in hardware E.g.,  
<https://github.com/herd/herdtools7/blob/master/herd/libdir/aarch64.cat>

```
(** External visibility requirement **)
irreflexive ob as external

(** Internal visibility requirements **)
irreflexive [Exp & R]; ((po & same-loc) | rmw); [Exp & W];
    rfi; [Exp & R] as coRW1-Exp
irreflexive [Imp & Tag & R]; (po & same-loc); [Exp & Tag & W];
    rfi; [Imp & Tag & R] as coRW1-MTE
irreflexive [Exp & W]; (po & same-loc); [Exp & W];
    (ca & int); [Exp & W] as coWW-Exp
irreflexive [Exp & W]; (po & same-loc); [Exp & R];
    (ca & int); [Exp & W] as coWR-Exp
irreflexive [Exp & Tag & W]; (po & same-loc); [Imp & Tag & R];
    (ca & int) as coWR-MTE
```

- ... but cannot directly be connected to traditional operational proof techniques

# Axiomatic Models

- Often the “authoritative” formal model, particularly in hardware E.g.,  
<https://github.com/herd/herdtools7/blob/master/herd/libdir/aarch64.cat>


```
(** External visibility requirement **)
irreflexive ob as external

(** Internal visibility requirements **)
irreflexive [Exp & R]; ((po & same-loc) | rmw); [Exp & W];
    rfi; [Exp & R] as coRW1-Exp
irreflexive [Imp & Tag & R]; (po & same-loc); [Exp & Tag & W];
    rfi; [Imp & Tag & R] as coRW1-MTE
irreflexive [Exp & W]; (po & same-loc); [Exp & W];
    (ca & int); [Exp & W] as coWW-Exp
irreflexive [Exp & W]; (po & same-loc); [Exp & R];
    (ca & int); [Exp & W] as coWR-Exp
irreflexive [Exp & Tag & W]; (po & same-loc); [Imp & Tag & R];
    (ca & int) as coWR-MTE
```

- ... but cannot directly be connected to traditional operational proof techniques
- ... however every axiomatic model gives rise to 1-1 operational model  
— check the axioms of the memory model at every step

# RC11 Models II

## Example 1.



Operational (view-based) models [Kang et al., 2017, Kaiser et al., 2017]


$\left[ \begin{array}{l} r \mapsto ?, \\ s \mapsto ? \end{array} \right]$

0  
 $x \bigcirc$   
 $T_1, T_2$


0  
 $y \bigcirc$   
 $T_1, T_2$

# RC11 Models II

## Example 1.




Operational (view-based) models [Kang et al., 2017, Kaiser et al., 2017]




# RC11 Models II

## Example 1.




Operational (view-based) models [Kang et al., 2017, Kaiser et al., 2017]


$$\left[ \begin{array}{l} r \mapsto?, \\ s \mapsto? \end{array} \right]$$


# RC11 Models II


## Example 1.



Operational (view-based) models [Kang et al., 2017, Kaiser et al., 2017]




$$\left[ \begin{array}{l} r \mapsto 1, \\ s \mapsto ? \end{array} \right]$$





# RC11 Models II

## Example 1.




Operational (view-based) models [Kang et al., 2017, Kaiser et al., 2017]

$$\left[ \begin{array}{l} r \mapsto 1, \\ s \mapsto 0 \end{array} \right]$$



## RC11 Models II

### Example 2.



Operational (view-based) models [Kang et al., 2017, Kaiser et al., 2017]


$$\left[ \begin{array}{l} r \mapsto ?, \\ s \mapsto ? \end{array} \right]$$

$$\begin{array}{c} 0 \\ x \bigcirc \\ T_1, T_2 \end{array}$$

$$\begin{array}{c} 0 \\ y \bigcirc \\ T_1, T_2 \end{array}$$

## RC11 Models II

### Example 2.




Operational (view-based) models [Kang et al., 2017, Kaiser et al., 2017]


$$\left[ \begin{array}{l} r \mapsto ?, \\ s \mapsto 0 \end{array} \right]$$
$$\begin{array}{c} 0 \\ x \bigcirc \\ \mathbf{T_1, T_2} \end{array}$$
$$\begin{array}{c} 0 \\ y \bigcirc \\ \mathbf{T_1, T_2} \end{array}$$

# RC11 Models II

## Example 2.




Operational (view-based) models [Kang et al., 2017, Kaiser et al., 2017]




## RC11 Models II

### Example 2.




Operational (view-based) models [Kang et al., 2017, Kaiser et al., 2017]




## RC11 Models II

### Example 2.



Operational (view-based) models [Kang et al., 2017, Kaiser et al., 2017]



# Verification

## RC11: Towards a Reasoning Framework

- View model for RC11 supports **interleaved execution**
- Allows relaxed behaviours — reads may **return stale values**
- Difficulty comes from the state model
  - $Loc \rightarrow Val$  states only valid for thread-local variables
  - View-based states required to model shared variables

**Key idea:** develop high-level assertions directly over view states


## RC11: View assertions

Only a handful of assertions needed to handle most examples

- *Last value predicate*  $x \ll \tau$   
thread  $\tau$  is viewing the last-written value for shared variable  $x$
- *Values view function*  $|\tau \rangle x$   
the set of values that thread  $\tau$  can see for shared variable  $x$
- *Definite value predicate*  $x \stackrel{\tau}{=} xv$   
 $x \ll \tau \wedge |\tau \rangle x = \{xv\}$

## View Assertions (Example)

Consider the state  $\sigma$ :



The following properties hold for  $\sigma$ :

$$\neg x \ll T_1 \wedge y \ll T_1$$

$$x \ll T_2 \wedge \neg y \ll T_2$$

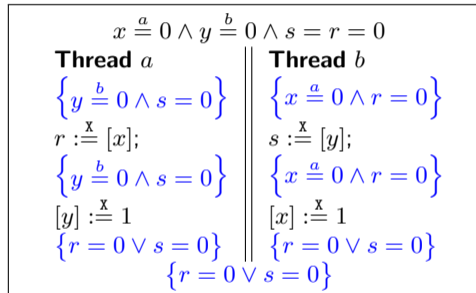

$$|T_1\rangle x = \{0, 1\}$$

$$|T_2\rangle x = \{1\}$$

$$y \stackrel{T_1}{=} 1$$

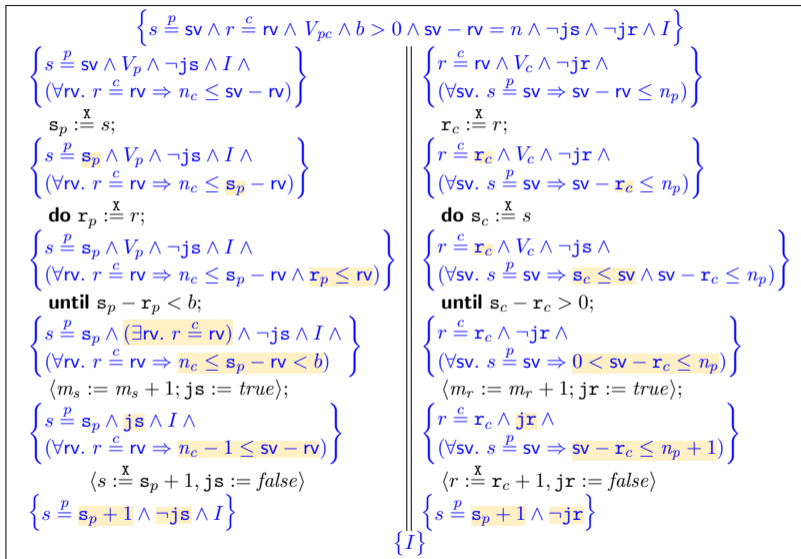
$$x \stackrel{T_2}{=} 1$$

# Verifying Load Buffering



Isabelle/HOL Demo

# Lamport Buffer



Isabelle/HOL Demo

# Lamport Buffer

```

{m_s = m_r = s = r = 0 ∧ 0 < b ∧ ¬js ∧ ¬jr}
{n_c ≤ s - r ∧ ¬js ∧ I}
  while true
    inv n_c ≤ s - r ∧ ¬js ∧ I
      {n_c ≤ s - r ∧ ¬js ∧ I}
        wait s - r < b end;
      {n_c ≤ s - r < b ∧ ¬js ∧ I}
        ⟨m_s := m_s + 1; js := true⟩;
      {n_c - 1 ≤ s - r ∧ js ∧ I}
        ⟨s := s + 1; js := false⟩
    {I ∧ ¬js}
  {I}
  
```

```

{...}
{s ≐ sv ∧ V_p ∧ ¬js ∧ I ∧
 (∀rv. r ≐ rv ⇒ n_c ≤ sv - rv)}
  s_p ≐ s;
{s ≐ s_p ∧ V_p ∧ ¬js ∧ I ∧
 (∀rv. r ≐ rv ⇒ n_c ≤ s_p - rv)}
  do r_p ≐ r;
{s ≐ s_p ∧ V_p ∧ ¬js ∧ I ∧
 (∀rv. r ≐ rv ⇒ n_c ≤ s_p - rv ∧ r_p ≤ rv)}
  until s_p - r_p < b;
{s ≐ s_p ∧ (∃rv. r ≐ rv) ∧ ¬js ∧ I ∧
 (∀rv. r ≐ rv ⇒ n_c ≤ s_p - rv < b)}
  ⟨m_s := m_s + 1; js := true⟩;
{s ≐ s_p ∧ js ∧ I ∧
 (∀rv. r ≐ rv ⇒ n_c - 1 ≤ sv - rv)}
  ⟨s ≐ s_p + 1, js := false⟩
{s ≐ s_p + 1 ∧ ¬js ∧ I}
  {I}
  
```

Isabelle/HOL Demo

# Formal Model

# Syntax

$$\begin{aligned} \text{Exp} &::= v \mid \mathbf{r} \mid \ominus e \mid e_1 \oplus e_2 & \text{Aux} &::= v \mid \mathbf{r} \mid \mathbf{a} \mid \ominus f \mid f_1 \oplus f_2 \\ \text{Asgn} &::= \mathbf{r} := e \mid \mathbf{r} \stackrel{\mathbf{M}}{:=} x \mid x \stackrel{\mathbf{M}}{:=} e \\ \text{Com} &::= a \mid \langle a; \bar{\mathbf{a}} := \bar{f} \rangle \mid c_1; c_2 \mid \mathbf{if } b \mathbf{ then } c_1 \mathbf{ else } c_2 \mid \mathbf{do } c \mathbf{ until } b \\ \text{Prog} &::= \lambda \tau \in \text{Tid}. c \end{aligned}$$

# Operational Semantics (Uninterpreted)

RD

$$\frac{\gamma \xrightarrow{\langle x, \text{read}(v), \mathbf{M} \rangle} \gamma'}{\langle \sigma, \gamma, \mathbf{r} :=^{\mathbf{M}} x \rangle \rightarrow_{\tau} \langle \sigma[\mathbf{r} \mapsto v], \gamma', \perp \rangle}$$

WR

$$\frac{\gamma \xrightarrow{\langle x, \text{write}(\llbracket e \rrbracket_{\sigma}), \mathbf{M} \rangle} \gamma'}{\langle \sigma, \gamma, x :=^{\mathbf{M}} e \rangle \rightarrow_{\tau} \langle \sigma, \gamma', \perp \rangle}$$

PROG

$$\frac{\langle \sigma, \gamma, \Pi(\tau) \rangle \rightarrow_{\tau} \langle \sigma', \gamma', c' \rangle}{\langle \sigma, \gamma, \Pi \rangle \rightarrow \langle \sigma', \gamma', \Pi[\tau \mapsto c'] \rangle}$$

# RC11 State

- Use encoding by [Castañeda et al., 2024]
- Only need two components  $\langle m, \mathcal{V} \rangle$ 
  - $m$  = set of timestamped **messages** (writes)
  - $\mathcal{V}$  = **view** of each thread
- Each message  $\mu = \langle x, v, t, V, M \rangle$  where:
  - $x \in Var$  is the variable of the message
  - $v \in Val$  is the written value
  - $t \in \mathbb{Q}$  is the timestamp
  - $V \in View \hat{=} Var \rightarrow \mathbb{Q}$  is the view of the message (ignored for now)
  - $M \in \{X, R\}$  is the memory ordering
- Thread view:  $\mathcal{V} \in Tid \rightarrow View$

## Isabelle/HOL encoding

```
record ('Var, 'Val) Msg =  
  var  :: 'Var  
  val  :: 'Val  
  ts   :: TS  
  view :: "'Var View"  
  mord :: Mord
```

```
record ('Tid, 'Var, 'Val) C11State =  
  msgs :: "('Var, 'Val) Msg set"  
  TView :: "'Tid  $\Rightarrow$  'Var View"
```

# RC11 Relaxed Reads / Writes

$$\frac{\text{READ-RLX} \quad e = \langle x, \text{read}(v), X \rangle \quad \langle x, v, t, -, - \rangle \in m \quad \mathcal{V}(\tau)(x) \leq t \quad V' = \mathcal{V}(\tau)[x \mapsto t]}{\langle m, \mathcal{V} \rangle \xrightarrow{e}_{\tau} \langle m, \mathcal{V}[\tau \mapsto V'] \rangle}$$

$$\frac{\text{WRITE-RLX-REL} \quad e = \langle x, \text{write}(v), M \rangle \quad V' = \mathcal{V}(\tau)[x \mapsto t] \quad \mathcal{V}(\tau)(x) < t \quad t \notin \text{ts}(m|_x) \quad \mu = \langle x, v, t, V', M \rangle}{\langle m, \mathcal{V} \rangle \xrightarrow{e}_{\tau} \langle m \cup \{\mu\}, \mathcal{V}[\tau \mapsto V'] \rangle}$$

- READ-RLX reads from the message with timestamp  $t$
- WRITE-RLX-REL introduces a new write message with *fresh* timestamp  $t$

```
definition "readRlxTrans  $\tau$  x  $\mu$   $\sigma$   $\equiv$ 
  let  $V' = (\text{TView } \sigma \ \tau)(x := \text{ts } \mu)$  in
   $\sigma$  ;; updateTView  $\tau$   $V'$ "
```

```
definition "readRlx  $\tau$  x v  $\sigma$   $\sigma'$   $\equiv$ 
   $\exists \mu \in (\text{obs } \tau \ \sigma)|_x$  .
  val  $\mu = v \wedge \sigma' = \text{readRlxTrans } \tau \ x \ \mu \ \sigma$ "
```

```
definition "writeTransExp b  $\tau$  x v t  $\sigma$   $\equiv$ 
  let  $V' = (\text{TView } \sigma \ \tau)(x := t)$ ;
   $\mu = (\text{var}=x, \text{val}=v, \text{ts}=t,$ 
     $\text{view}=V', \text{mord}=b)$ 
  in  $\sigma$  ;; updateMsgs  $\mu$  ;; updateTView  $\tau$   $V'$ "
```

```
definition "writeTrans b  $\tau$  x v  $\sigma$   $\sigma'$   $\equiv$ 
   $\exists t$ . validFreshTS  $\sigma \ \tau \ x \ t \wedge$ 
   $\sigma' = \text{writeTransExp } b \ \tau \ x \ v \ t \ \sigma$ "
```

## Formalising View Assertions

- *Last value predicate*: Thread  $\tau$  is viewing the last-written value for  $x$

$$x \ll \tau \hat{=} \lambda \langle m, \mathcal{V} \rangle. \mathcal{V}(\tau)(x) = \max(\text{ts}(m|_x))$$

```
definition lastView :: "... " ("_<<_ _" ...)
  where "lastView x \tau \sigma \equiv \forall \text{ts} \in \text{tss} ((\text{msgs } \sigma)|_x). \text{ts} \leq \text{TView } \sigma \tau x "
```

- *Values view function*: The set of values that thread  $\tau$  can see for  $x$

$$\begin{aligned} \text{obs}(\tau) &\hat{=} \lambda \langle m, \mathcal{V} \rangle. \{ \mu \in m \mid \text{ts}(\mu) \geq \mathcal{V}(\tau)(\text{var}(\mu)) \} \\ |\tau\rangle x &\hat{=} \lambda \sigma. \text{val}(\text{obs}(\tau)(\sigma)|_x) \end{aligned}$$

```
definition valsView :: ... ("|_)_ _" ...)
  where "valsView \tau x \sigma \equiv \text{vals} ((\text{obs } \tau \sigma)|_x) "
```

- *Definite value predicate*

$$x \stackrel{\tau}{=} \text{xv} \hat{=} x \ll \tau \wedge |\tau\rangle x = \{\text{xv}\}$$

# Hoare Triples

## Definition

Given assertions  $P, Q : \Sigma \times \Gamma \rightarrow \mathbb{B}$ , thread  $\tau \in Tid$ , command  $c \in Com$  and program  $\Pi \in Prog$ :

$$\{P\}_{\tau} \triangleright c \{Q\} \hat{=} \forall \sigma, \gamma. P(\langle \sigma, \gamma \rangle) \wedge \langle \sigma, \gamma, c \rangle \rightarrow_{\tau}^* \langle \sigma', \gamma', \perp \rangle \Rightarrow Q(\langle \sigma', \gamma' \rangle)$$

$$\{P\}_{\Pi} \{Q\} \hat{=} \forall \sigma, \gamma. P(\langle \sigma, \gamma \rangle) \wedge \langle \sigma, \gamma, \Pi \rangle \rightarrow^* \langle \sigma', \gamma', \lambda\tau. \perp \rangle \Rightarrow Q(\langle \sigma', \gamma' \rangle)$$

# RC11 Axioms

Examples:

RD-STABLE

$$\frac{}{\left\{x \stackrel{\tau}{=} xv\right\} \tau' \triangleright \mathbf{r} \stackrel{\mathbf{M}}{=} y \left\{x \stackrel{\tau}{=} xv\right\}}$$

WR-LATEST

$$\frac{}{\left\{x \ll \tau\right\} \tau \triangleright x \stackrel{\mathbf{M}}{=} e \left\{x \stackrel{\tau}{=} e\right\}}$$

Example soundness proof for WR-LATEST:

```
lemma wrAtLast_stable_sameTh: "x⟨⟨τ σ ⇒ writeStep τ y v σ σ' ⇒ x⟨⟨τ σ'⟩⟩"
  by (auto simp add: unfold_defs, auto)
```

```
lemma lastView_Wr: "x⟨⟨τ σ ⇒ writeStep τ x v σ σ' ⇒ |τ⟩x σ' = {v}"
  apply (auto simp add: unfold_defs, auto)
  apply ((metis Msg.select_convs(2))+)[2]
  by (metis Msg.select_convs(1,2,3) dual_order.refl)+
```

```
lemma DVal_Update_Wr:
  "x⟨⟨τ σ ⇒ writeStep τ x v σ σ' ⇒ x =τ v σ'"
  by (meson defVal_def lastView_Wr wrAtLast_stable_sameTh)
```

# Release-Acquire

# Lamport Buffer

<pre><b>while</b> true   <math>s_p :=^X s</math>;   <b>do</b> <math>r_p :=^X r</math>;   <b>until</b> <math>s_p - r_p &lt; b</math>;   <math>\langle m_s := m_s + 1; js := true \rangle</math>;   <math>\langle s :=^X s_p + 1, js := false \rangle</math></pre>	<pre><b>while</b> true   <math>r_c :=^X r</math>;   <b>do</b> <math>s_c :=^X s</math>   <b>until</b> <math>s_c - r_c &gt; 0</math>;   <math>\langle m_r := m_r + 1; jr := true \rangle</math>;   <math>\langle r :=^X r_c + 1, jr := false \rangle</math></pre>
--	---

- Program guarantees the invariant  $I = 0 \leq m_s - m_r \leq b$  under RC11
- Question: What is the problem?
- Answer: Messages sent by the producer may not arrive at the consumer in time
- This problem is due to weak memory effects

# Message Passing

Relaxed accesses

$$\begin{array}{c} y = 0 \\ [x] \stackrel{\text{X}}{:=} 1; \quad \parallel \quad r \stackrel{\text{X}}{:=} [y]; \\ [y] \stackrel{\text{X}}{:=} 1 \quad \parallel \quad s \stackrel{\text{X}}{:=} [x] \end{array}$$

# Message Passing

## Relaxed accesses

$$y = 0$$

$$[x] :=^X 1; \parallel r :=^X [y];$$


$$[y] :=^X 1 \parallel s :=^X [x]$$

## Release-acquire

$$y = 0$$

$$[x] :=^X 1; \parallel r :=^A [y];$$


$$[y] :=^R 1 \parallel s :=^X [x]$$





Allowed

# Message Passing


## Relaxed accesses




## Release-acquire



**Allowed**



**Disallowed**



**Allowed**

# RA Synchronisation (Operationally)

**Recall.** Write rule models both relaxed and releasing writes

$$\begin{array}{c} \text{WRITE-RLX-REL} \\ e = \langle x, \text{write}(v), \mathbf{M} \rangle \quad V' = \mathcal{V}(\tau)[x \mapsto t] \\ \mathcal{V}(\tau)(x) < t \quad t \notin \text{ts}(m|_x) \\ \mu = \langle x, v, t, V', \mathbf{M} \rangle \\ \hline \langle m, \mathcal{V} \rangle \xrightarrow{e} \langle m \cup \{\mu\}, \mathcal{V}[\tau \mapsto V'] \rangle \end{array}$$

## Key ideas.

- Each message records the *view of the executing thread* at the time of writing
- The view indicates the thread's *happens-before knowledge* (at the time of writing)
- Later if RA synchronisation occurs,  
view transfer  $\Leftrightarrow$  happens-before transfer

# RA Synchronisation

Write transition: Recall  $M \in \{X, R\}$

$$\begin{array}{c}
 \text{WRITE-RLX-REL} \\
 e = \langle x, \text{write}(v), M \rangle \quad V' = \mathcal{V}(\tau)[x \mapsto t] \\
 \mathcal{V}(\tau)(x) < t \quad t \notin \text{ts}(m|_x) \\
 \mu = \langle x, v, t, V', M \rangle \\
 \hline
 \langle m, \mathcal{V} \rangle \xrightarrow{e}_\tau \langle m \cup \{\mu\}, \mathcal{V}[\tau \mapsto V'] \rangle
 \end{array}$$

Acquiring reads:

$$\begin{array}{c}
 \text{ACQ-RF-RLX} \\
 e = \langle x, \text{read}(v), A \rangle \quad \langle x, v, t, -, X \rangle \in m \\
 \mathcal{V}(\tau)(x) \leq t \quad V' = \mathcal{V}(\tau)[x \mapsto t] \\
 \hline
 \langle m, \mathcal{V} \rangle \xrightarrow{e}_\tau \langle m, \mathcal{V}[\tau \mapsto V'] \rangle
 \end{array}$$

$$\begin{array}{c}
 \text{ACQ-RF-REL} \\
 e = \langle x, \text{read}(v), A \rangle \quad \langle x, v, t, V, R \rangle \in m \\
 \mathcal{V}(\tau)(x) \leq t \quad V' = \mathcal{V}(\tau) \sqcup V \\
 \hline
 \langle m, \mathcal{V} \rangle \xrightarrow{e}_\tau \langle m, \mathcal{V}[\tau \mapsto V'] \rangle
 \end{array}$$

where  $V_1 \sqcup V_2 \hat{=} \lambda x \in \text{Var}. \max\{V_1(x), V_2(x)\}$

# RA in Isabelle

ACQ-RF-RLX

$$\frac{e = \langle x, \text{read}(v), \mathbf{A} \rangle \quad \langle x, v, t, -, \mathbf{X} \rangle \in m \quad \mathcal{V}(\tau)(x) \leq t \quad V' = \mathcal{V}(\tau)[x \mapsto t]}{\langle m, \mathcal{V} \rangle \xrightarrow{e}_{\tau} \langle m, \mathcal{V}[\tau \mapsto V'] \rangle}$$

ACQ-RF-REL

$$\frac{e = \langle x, \text{read}(v), \mathbf{A} \rangle \quad \langle x, v, t, \mathbf{V}, \mathbf{R} \rangle \in m \quad \mathcal{V}(\tau)(x) \leq t \quad V' = \mathcal{V}(\tau) \sqcup V}{\langle m, \mathcal{V} \rangle \xrightarrow{e}_{\tau} \langle m, \mathcal{V}[\tau \mapsto V'] \rangle}$$





where  $V_1 \sqcup V_2 \hat{=} \lambda x \in \text{Var}. \max\{V_1(x), V_2(x)\}$

```
definition maxView:: "... w" ("_  $\sqcup$  _" [100, 100]) where  
"maxView V1 V2  $\equiv$   $\lambda$  x. max (V1 x) (V2 x)"
```





definition

```
"readAcqTrans  $\tau$  x  $\mu$   $\sigma$   $\equiv$   
let V' = if mord  $\mu$  = rlx then (TView  $\sigma$   $\tau$ )(x := ts  $\mu$ ) else (view  $\mu$ )  $\sqcup$  (TView  $\sigma$   $\tau$ ) in  
 $\sigma$  ;; updateTView  $\tau$  V'"
```





# Recap: Relaxed Message-Passing Execution







# Recap: Relaxed Message-Passing Execution







# Recap: Relaxed Message-Passing Execution



# Recap: Relaxed Message-Passing Execution




# Recap: Relaxed Message-Passing Execution



## View-Transfer Assertion

Use dynamic **modal** predicate  $\pi \overset{\tau}{\rightsquigarrow} \varphi$


- $\pi$  is predicate over messages (a read filter)
- all messages observable to  $\tau$  satisfying  $\pi$  are releasing
- performing an acquiring read when the read trigger holds guarantees  $\varphi$  after the read



## View-Transfer Assertion

Use dynamic **modal** predicate  $\pi \overset{\tau}{\rightsquigarrow} \varphi$


- $\pi$  is predicate over messages (a read filter)
- all messages observable to  $\tau$  satisfying  $\pi$  are releasing
- performing an acquiring read when the read trigger holds guarantees  $\varphi$  after the read



## View-Transfer Assertion

Use dynamic **modal** predicate  $\pi \overset{\tau}{\rightsquigarrow} \varphi$


- $\pi$  is predicate over messages (a read filter)
- all messages observable to  $\tau$  satisfying  $\pi$  are releasing
- performing an acquiring read when the read trigger holds guarantees  $\varphi$  after the read



## View-Transfer Assertion

Use dynamic **modal** predicate  $\pi \overset{\tau}{\rightsquigarrow} \varphi$


- $\pi$  is predicate over messages (a read filter)
- all messages observable to  $\tau$  satisfying  $\pi$  are releasing
- performing an acquiring read when the read trigger holds guarantees  $\varphi$  after the read




## View-Transfer Assertion

Use dynamic **modal** predicate  $\pi \overset{\tau}{\rightsquigarrow} \varphi$

- $\pi$  is predicate over messages (a read filter)
- all messages observable to  $\tau$  satisfying  $\pi$  are releasing
- performing an acquiring read when the read trigger holds guarantees  $\varphi$  after the read



# Formalising the View-Transfer Assertion



$$\begin{array}{c}
 \text{ACQ-RF-REL} \\
 e = \langle x, \text{read}(v), \mathbf{A} \rangle \quad \langle x, v, t, \mathbf{V}, \mathbf{R} \rangle \in m \\
 \mathcal{V}(\tau)(x) \leq t \quad \mathbf{V}' = \mathcal{V}(\tau) \sqcup \mathbf{V} \\
 \hline
 \langle m, \mathcal{V} \rangle \xrightarrow{e}_{\tau} \langle m, \mathcal{V}[\tau \mapsto \mathbf{V}'] \rangle
 \end{array}$$

Construct the intersection:

$$\text{filt}(\tau, \pi) \hat{=} \lambda \alpha. \{ \mu \in \text{obs}(\tau)(\alpha) \mid \pi(\mu) \}$$

Define modal view transfer:

$$\begin{array}{l}
 \pi \xrightarrow{\tau} \varphi \hat{=} \lambda \alpha. \mathbf{let} \ \alpha = \langle \sigma, \langle m, \mathcal{V} \rangle \rangle \mathbf{in} \\
 \quad \forall \mu \in \text{filt}(\tau, \pi)(\alpha). \text{ord}(\mu) = \mathbf{R} \wedge \varphi(\langle \sigma, \langle m, \mathcal{V}[\tau \mapsto \mathcal{V}(\tau) \sqcup \text{view}(\mu)] \rangle \rangle)
 \end{array}$$

# Formalising the View-Transfer Assertion

Special case:

$$(x = xv) \rightsquigarrow^{\tau} \varphi \hat{=} (\lambda\mu. \text{var}(\mu) = x \wedge \text{val}(\mu) = xv) \rightsquigarrow^{\tau} \varphi$$

Whenever  $\tau$  acquire-reads  $xv$  for  $x$ ,  $\varphi$  is guaranteed in the post state

Using the special case of view-transfer is easy

ACQ-VIEW-TRANSFER

$$\frac{\{x = xv \rightsquigarrow^{\tau} \varphi\}}{\{x = xv \rightsquigarrow^{\tau} \varphi\} \tau \triangleright \mathbf{r} :=^A x \{ \mathbf{r} = xv \Rightarrow \varphi \}}$$

Establishing view-transfer is not

REL-DEF-VIEW-SET

$$\forall i \in S. x \neq y_i$$

$$\frac{\{xv \notin |\tau'|x \wedge (\forall i \in S. y_i \stackrel{\tau}{=} yv_i)\}}{\{x = xv \rightsquigarrow^{\tau'} (\forall i \in S. y_i \stackrel{\tau'}{=} yv_i)\}} \tau \triangleright x :=^R xv$$

lemma

$$"(x = xv) \rightsquigarrow^{\tau} \varphi \sigma$$

$$\Rightarrow \sigma \tau \triangleright \mathbf{r} :=^A [\mathbf{x}] \sigma'$$

$$\Rightarrow \mathbf{r} = xv \rightarrow \varphi \sigma''"$$

by (auto simp add: unfold\_defs)

lemma

$$"wfs \sigma \Rightarrow x \neq y \Rightarrow v \notin |\tau'|x \sigma \Rightarrow y =^{\tau} u \sigma$$


$$\Rightarrow \sigma \tau \triangleright [\mathbf{x}] :=^R v \sigma'$$

$$\Rightarrow (x = v) \rightsquigarrow^{\tau'} (\lambda \sigma. y =^{\tau'} u \sigma) \sigma''"$$

proof

(approx 30 lines)

# View-Transfer in Action




$\left[ \begin{array}{l} r \mapsto?, \\ s \mapsto? \end{array} \right]$



0  
 $x \bigcirc$   
 $T_1, T_2$

0  
 $y \bigcirc$   
 $T_1, T_2$


# View-Transfer in Action




$\left[ \begin{array}{l} r \mapsto?, \\ s \mapsto? \end{array} \right]$



# View-Transfer in Action




$$\begin{bmatrix} r \mapsto?, \\ s \mapsto? \end{bmatrix}$$




$T_1$  establishes

$$y = 1 \overset{T_2}{\rightsquigarrow} (x \overset{T_2}{=} 1)$$

# View-Transfer in Action



$$\left[ \begin{array}{l} r \mapsto 1, \\ s \mapsto ? \end{array} \right]$$



$T_2$  uses

$y = 1 \stackrel{T_2}{\rightsquigarrow} (x \stackrel{T_2}{=} 1)$  to establish

$x \stackrel{T_2}{=} 1$

# View-Transfer in Action

$$\begin{array}{c}
 \{x \ll \mathbf{T}_1 \wedge 1 \notin \mathbf{T}_2\} y \\
 \text{Thread } \mathbf{T}_1 \\
 \{1 \notin \mathbf{T}_2\} y \wedge x \ll \mathbf{T}_1 \\
 [x] :=^{\mathbf{X}} 1; \\
 \{1 \notin \mathbf{T}_2\} y \wedge x \stackrel{\mathbf{T}_1}{=} 1 \\
 [y] :=^{\mathbf{R}} 1 \\
 \{true\}
 \end{array}
 \parallel
 \begin{array}{c}
 \text{Thread } \mathbf{T}_2 \\
 \{y = 1 \stackrel{\mathbf{T}_2}{\rightsquigarrow} (x \stackrel{\mathbf{T}_2}{=} 1)\} \\
 r :=^{\mathbf{A}} [y]; \\
 \{r = 1 \Rightarrow x \stackrel{\mathbf{T}_2}{=} 1\} \\
 s :=^{\mathbf{X}} [x] \\
 \{r = 1 \Rightarrow s = 1\}
 \end{array}$$

$$\{r = 1 \Rightarrow s = 1\}$$

Isabelle/HOL Demo

# Questions?

## 1 Background

- Sequential consistency
- Owicki-Gries reasoning
- Isabelle/HOL
- Examples: Message passing, Lamport buffer

## 2 Relaxed Memory

- Motivation
- RC11 Memory Model (Relaxed Accesses)

## 3 Verification






- Owicki-Gries reasoning, revisited
- Examples, revisited
- Isabelle/HOL

## 4 Formal Model







## 5 Release/Acquire Accesses

- Lamport buffer, revisited again
- Message passing
- Isabelle/HOL

# References I

-  Adve, S. V. and Gharachorloo, K. (1996).  
Shared memory consistency models: A tutorial.  
*IEEE Computer*, 29(12):66–76.
-  Alglave, J., Maranget, L., and Tautschnig, M. (2014).  
Herdng cats: Modelling, simulation, testing, and data mining for weak memory.  
*ACM Trans. Program. Lang. Syst.*, 36(2):7:1–7:74.
-  Batty, M., Owens, S., Sarkar, S., Sewell, P., and Weber, T. (2011).  
Mathematizing C++ concurrency.  
In Ball, T. and Sagiv, M., editors, *POPL*, pages 55–66. ACM.
-  Castañeda, A., Chockler, G., Dongol, B., and Lahav, O. (2024).  
What Cannot Be Implemented on Weak Memory?  
In Alistarh, D., editor, *DISC*, volume 319 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 11:1–11:22, Dagstuhl, Germany. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
-  de Roever, W. P., de Boer, F. S., Hannemann, U., Hooman, J., Lakhnech, Y., Poel, M., and Zwiers, J. (2001).  
*Concurrency Verification: Introduction to Compositional and Noncompositional Methods*, volume 54 of *Cambridge Tracts in Theoretical Computer Science*.  
Cambridge University Press.

## References II

-  Kaiser, J., Dang, H., Dreyer, D., Lahav, O., and Vafeiadis, V. (2017). Strong logic for weak memory: Reasoning about release-acquire consistency in iris. In *ECOOP*, pages 17:1–17:29.
-  Kang, J., Hur, C., Lahav, O., Vafeiadis, V., and Dreyer, D. (2017). A promising semantics for relaxed-memory concurrency. In *POPL*, pages 175–189. ACM.
-  Lahav, O., Vafeiadis, V., Kang, J., Hur, C., and Dreyer, D. (2017). Repairing sequential consistency in C/C++11. In Cohen, A. and Vechev, M. T., editors, *PLDI*, pages 618–632. ACM.
-  Lamport, L. (1979). How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Trans. Computers*, 28(9):690–691.
-  Owicki, S. S. (1975). *Axiomatic Proof Techniques for Parallel Programs*. Outstanding Dissertations in the Computer Sciences. Garland Publishing, New York.
-  Steinke, R. C. and Nutt, G. J. (2004). A unified theory of shared memory consistency. *J. ACM*, 51(5):800–849.