



Verification of weak memory models

Elli Anastasiadi

OPCT 2023, Bertinoro, Italy

26-30 June



- 1 Weak memory
 - Why - how - examples

- 2 Verification
 - Basic problems
 - Basic principles

Starting point: distributed programs & architectures.



Starting point: distributed programs & architectures.
→ Need for verification.



Starting point: distributed programs & architectures.
→ Need for verification. ...so far so good.



Starting point: distributed programs & architectures.

→ Need for verification. ...so far so good.

Necessary: a model for distributed execution.



Starting point: distributed programs & architectures.

→ Need for verification. ...so far so good.

Necessary: a model for distributed execution.

→ Before weak memory: interleaving.



Starting point: distributed programs & architectures.

→ Need for verification. ...so far so good.

Necessary: a model for distributed execution.

→ Before weak memory: interleaving.

An example (Dekker protocol):

$$P_1 :=$$
$$x := 1$$
$$a := y$$
$$\text{if } a = 0 \text{ then}$$
$$CS_1$$
$$P_2 :=$$
$$y := 1$$
$$b := x$$
$$\text{if } b = 0 \text{ then}$$
$$CS_2$$

Assertion: $\text{not}(CS_1 \text{ and } CS_2)$



Starting point: distributed programs & architectures.

→ Need for verification. ...so far so good.

Necessary: a model for distributed execution.

→ Before weak memory: interleaving.

An example (Dekker protocol):

$P_1 :=$

$x := 1$

$a := y$

if $a = 0$ then

CS_1


$P_2 :=$

$y := 1$

$b := x$

if $b = 0$ then

CS_2



Assertion: $\text{not}(CS_1 \text{ and } CS_2)$

Starting point: distributed programs & architectures.

→ Need for verification. ...so far so good.

Necessary: a model for distributed execution.

→ Before weak memory: interleaving.

An example (Dekker protocol):

$P_1 :=$

$x := 1$

$\mathbf{a:=y}$

if $a = 0$ then

CS_1


$P_2 :=$

$y := 1$

$b := x$

if $b = 0$ then

CS_2



Assertion: $\text{not}(CS_1 \text{ and } CS_2)$

Starting point: distributed programs & architectures.

→ Need for verification. ...so far so good.

Necessary: a model for distributed execution.

→ Before weak memory: interleaving.

An example (Dekker protocol):

$P_1 :=$

$x := 1$

$a := y$

if $a = 0$ **then**

CS_1


$P_2 :=$

$y := 1$

$b := x$

if $b = 0$ **then**

CS_2



Assertion: $\text{not}(CS_1 \text{ and } CS_2)$

Starting point: distributed programs & architectures.

→ Need for verification. ...so far so good.

Necessary: a model for distributed execution.

→ Before weak memory: interleaving.

An example (Dekker protocol):

$P_1 :=$

$x := 1$

$a := y$

if $a = 0$ then

CS_1


$P_2 :=$

$y := 1$

$b := x$

if $b = 0$ then

CS_2



Assertion: $\text{not}(CS_1 \text{ and } CS_2)$

Starting point: distributed programs & architectures.

→ Need for verification. ...so far so good.

Necessary: a model for distributed execution.

→ Before weak memory: interleaving.

An example (Dekker protocol):

$P_1 :=$

$x := 1$

$a := y$

if $a = 0$ then

CS_1


$P_2 :=$

$y := 1$

$\mathbf{b} := \mathbf{x}$

if $b = 0$ then

CS_2



Assertion: $\text{not}(CS_1 \text{ and } CS_2)$

Starting point: distributed programs & architectures.

→ Need for verification. ...so far so good.

Necessary: a model for distributed execution.

→ Before weak memory: interleaving.

An example (Dekker protocol):

$P_1 :=$

$x := 1$

$a := y$

if $a = 0$ then

CS_1


$P_2 :=$

$y := 1$

$b := x$

if $b = 0$ then

CS_2




Assertion: $\text{not}(CS_1 \text{ and } CS_2)$ ✓

Interleaving: implementation



Interleaving: implementation

- atomic writes,
- read-from-memory



Interleaving: formally



Interleaving: formally

Sequential consistency - SC

On a global trace, the events of a process will occur in the order stated locally for the process.

Or: only the events of different processes can be shuffled.



So what is weak memory?



So what is weak memory?

Rule of thumb: Anything “below” interleaving.





So what is weak memory?

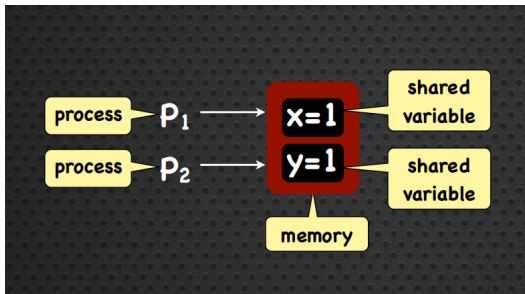
Rule of thumb: Anything “below” interleaving.

...extra reorderings




How?


How?



How?



How?

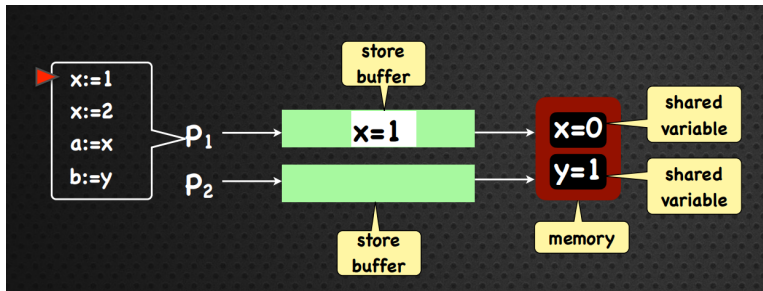


An example: total store order (TSO)



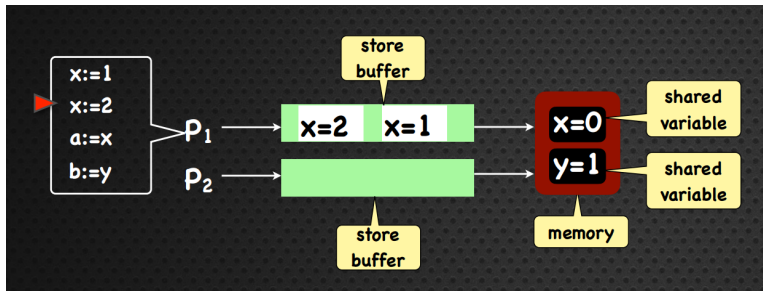
An example: total store order (TSO)

- non-atomic writes
- read locally or from memory



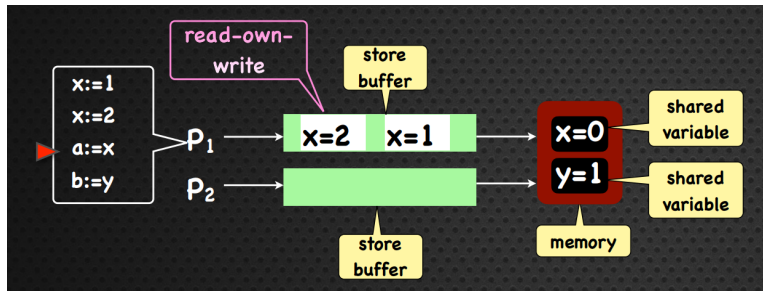
An example: total store order (TSO)

- non-atomic writes
- read locally or from memory



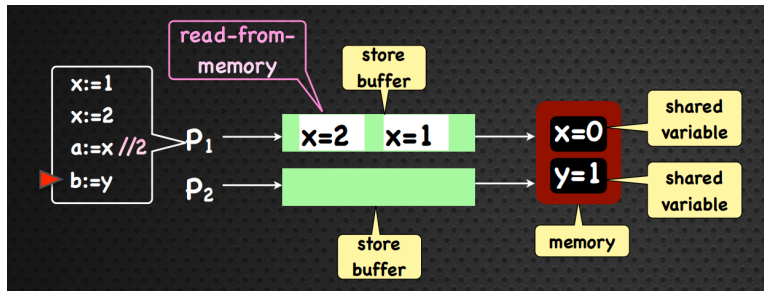
An example: total store order (TSO)

- non-atomic writes
- read locally or from memory




An example: total store order (TSO)

- non-atomic writes
- read locally or from memory



A *specification* of how an implementation will tackle parallel memory access is called a *memory model*.



A *specification* of how an implementation will tackle parallel memory access is called a *memory model*.

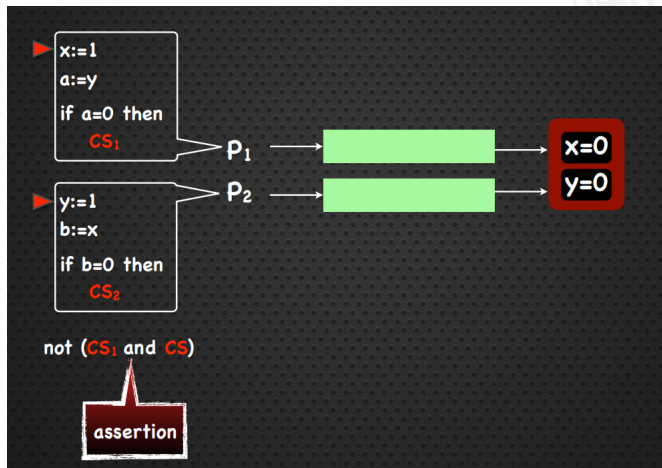
A **weak** memory model is one that allows extra behaviors.

A *specification* of how an implementation will tackle parallel memory access is called a *memory model*.

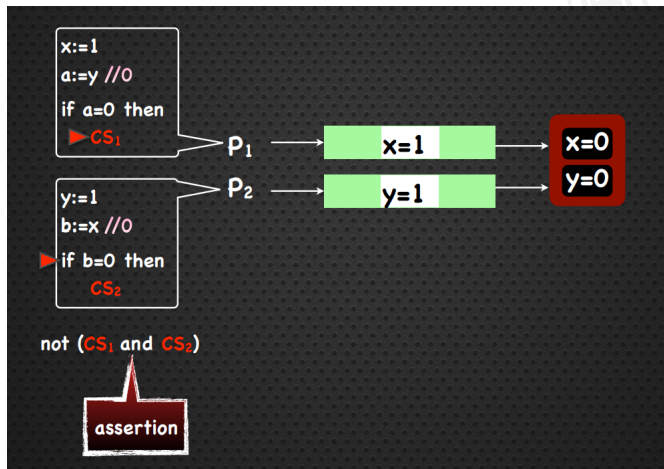
A **weak** memory model is one that allows extra behaviors.

potentially bad behaviors


Dekker under TSO



Dekker under TSO



Dekker under TSO




The dichotomy

Sequential consistency

- intuitive, well-researched 😊
- slow, unrealistic 😞

Weak memory

- **not** intuitive 😞
- fast, realistic 😊



The dichotomy

Sequential consistency

- intuitive, well-researched 😊
- slow, unrealistic 😞

Weak memory

- **not** intuitive 😞
- fast, realistic 😊

What about verification?

The dichotomy

Sequential consistency

- intuitive, well-researched 😊
- slow, unrealistic 😞

Weak memory

- **not** intuitive 😞
- fast, realistic 😊

What about verification?

Road-map: Semantics → Complexity → Techniques

- New software: only works when the architecture below satisfies **at least** a specific weak memory model.

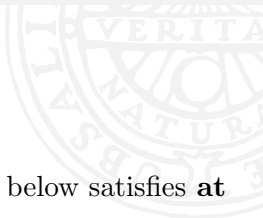
We need: algorithm for detecting “illegal” memory accesses.

- New software: only works when the architecture below satisfies **at least** a specific weak memory model.

We need: algorithm for detecting “illegal” memory accesses.

- New architecture: claims it satisfies some weak memory model.

We need: guarantee all runs of the new architecture are safe.



Scenario 1





Potentially we don't know the architecture.



Scenario 1

Potentially we don't know the architecture.

We look at the memory accesses of different processes:



Scenario 1

Potentially we don't know the architecture.

We look at the memory accesses of different processes:



- P_1 : write(x,0), write(y,1), read(x,1)
 P_2 : read(y,1), write(x,1), write(x,0)

Scenario 1

Potentially we don't know the architecture.


We look at the memory accesses of different processes:



- P_1 : write(x,0), write(y,1), read(x,1)
 P_2 : read(y,1), write(x,1), write(x,0)
- P_1 : write(x,0), write(x,1), write(y,1)
 P_2 : read(y,1), read(x,1), read(x,0)

Some verification primitives: Shasha-Snir traces

P_1 : read(x,1), write(y,1)
 P_2 : read(y,1), write(x,1)

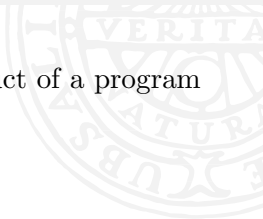


Example use: for SC the execution graph must be acyclic.

Sad result: Given the traces, to solve this is NP-complete

Some verification primitives: reachability

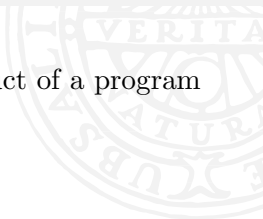
Execution graphs are in away the composition/product of a program with the semantics of a memory model.



Some verification primitives: reachability

Execution graphs are in away the composition/product of a program with the semantics of a memory model.

Why would reachability be hard?




Some verification primitives: reachability

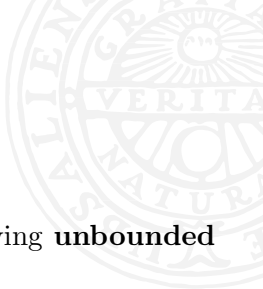
Execution graphs are in away the composition/product of a program with the semantics of a memory model.

Why would reachability be hard?

Reminder:



Weak memory (almost always) has semantics involving **unbounded data structures**.



Weak memory (almost always) has semantics involving **unbounded data structures**.

..sad realization

Some cool things

- Bad patterns
- Well quasi-orderings & monotonicity
- Well-structured systems



Bad patterns

Event sequences that are known to be violating a weak memory model.

When we are lucky: a finite set of bad patterns characterizes a given weak memory model.

Use: Only test these and we get answers for **any** sequence of events.

Some cool things

Well quasi orderings

Orderings between states of the (unavoidable) unbounded data structure associated to the semantics of the model.

When we are lucky: ordering between configurations means monotonicity in satisfaction of memory model.

Use: help us prune the infinitely large tree of configurations that we explore for reachability.

Well-structured systems

Semantic models for weak memory that have:

- Unbounded but FIFO components.
- Monotonicity

When we are lucky: we manage to transform the semantics of a weak memory model to a form that is well-structured.

Use: (Theorem) Reachability is always decidable.



Summary:

- basic idea of how weak memory shows up
- why is it problematic
- what we (usually) do about it.

Future:

- develop good algorithms for specific memory models
- develop hardness results
- (maybe) unification results.



Thank you for your
attention!





Questions?