

# The Semantics of Shared Variable Concurrency in Relaxed Memory Models

Colloque Anniversaire

en l'honneur de

G rard Berry & Jean-Jacques L vy

## Claim.

☞ relaxed memory = a **truly concurrent** setting for program execution

## Claim.

☞ relaxed memory = a **truly concurrent** setting for program execution

☞ true concurrency = **Berry & Lévy's** equivalence by permutation of transitions (POPL 1977)

## Claim.

☞ relaxed memory = a **truly concurrent** setting for program execution


☞ true concurrency = **Berry & Lévy's** equivalence by permutation of transitions (POPL 1977)

(I came to know about that at the weekly Thursday Seminar organized by Maurice Nivat, who managed to attract GG & JJ – and some others – out of the pure light of their Bat 8 Sanctuary, down to some obscure room at Paris 7 University – thanks Maurice.)

# SHARED MEMORY CONCURRENCY

---

- ▶ sequential programs (threads) running in parallel, sharing a global memory



- ▶ semantics: small step transitions, [interleaved](#), i.e. one thread at a time, non deterministically chosen, for one step

Hennesy & Plotkin 1979


precursors: Bekič 1971, Milner 1973, Cadiou & [Lévy](#) 1973

- ▶ aka [sequential consistency](#) (SC), Lamport 1979

# SYNCHRONIZATION

(1/2)

[Basic idea for] **mutual exclusion**, in a ML-like syntax:



The computation must start with updating  $\text{flag}_0$  or  $\text{flag}_1$


↳ **at most one thread**, but not both, can update  $x$

# SYNCHRONIZATION

---

(2/2)

A **producer/consumer** scenario, or safe publication:



To update  $x$  the consumer must wait for **flag** being updated

↳  $x$  **cannot get** the **old** value


# UNFORTUNATELY...

(1/2)

The standard, input/output equivalence of sequential programs is **not compatible** with parallel composition

$$\begin{array}{l} \text{flag}_0 := ff; \\ \text{if } !\text{flag}_1 \text{ then } x := !x + 1 \end{array} \approx \begin{array}{l} \text{if } !\text{flag}_1 \text{ then } x := !x + 1; \\ \text{flag}_0 := ff \end{array}$$

but cannot replace each other in a concurrent context:



may end up with  $x = 2$



# UNFORTUNATELY...

---

(2/2)

- ▶ standard equivalence of sequential program is the basis for **optimizing compilers**, and **hardware**, especially **relaxed memory models**
- ▶ correct – and efficient! – for sequential programs, but
- ↳ do **not** implement SC for concurrent threads

# RELAXED MEMORY MODELS

---

memory model = reordering, i.e. **relaxation** of program order, for  
memory accesses (write/read)  
+ write visibility  
+ means to maintain program order in some cases  
(fences, memory barriers, acquire/release...)

Some examples:

- ▶ TSO: relaxing write/read (on different memory locations)  
+ read own writes
- ▶ PSO: TSO + write/write (idem)
- ▶ RMO: PSO + read/write + read/read

# QUESTION

---

How to formalize the **semantics of concurrent programs** in the context of a relaxed memory model?

Two approaches:

- ▶▶ **axiomatic**: relations between memory events, satisfying some properties, see Francesco & Luc & Jade
- ▶▶ **operational**: transitions between configurations, as with programming languages (abstract machines)

# An Example: WRITE BUFFERING

---

(1/2)

Boudol & Petri POPL 2009, Sewell & al. CACM 2010

- configurations:

$$C = (S, (B_1, P_1), \dots, (B_n, P_n))$$

$S$  = shared memory

$B_i$  = write buffers

= fifo sequences of writes  $\langle x := v \rangle$

$P_i$  = processors = multisets of threads

- labelled transitions: action from a component  $i$
- actions: read/write memory locations  $\mathbf{rd}_{x,v}$  and  $\mathbf{wr}_{y,w}$  (and more)

# An Example: WRITE BUFFERING

(2/2)

Some rules:

$$\frac{P_i \xrightarrow{\text{wr}_{x,v}} P'_i}{(S, \dots, (B_i, P_i), \dots) \xrightarrow{i,\tau} (S, \dots, (B_i \cdot \langle x := v \rangle, P'_i), \dots)}$$

$$\frac{P_i \xrightarrow{\text{rd}_{x,v}} P'_i \quad B_i(x) = v}{(S, \dots, (B_i, P_i), \dots) \xrightarrow{i,\tau} (S, \dots, (B_i, P'_i), \dots)}$$

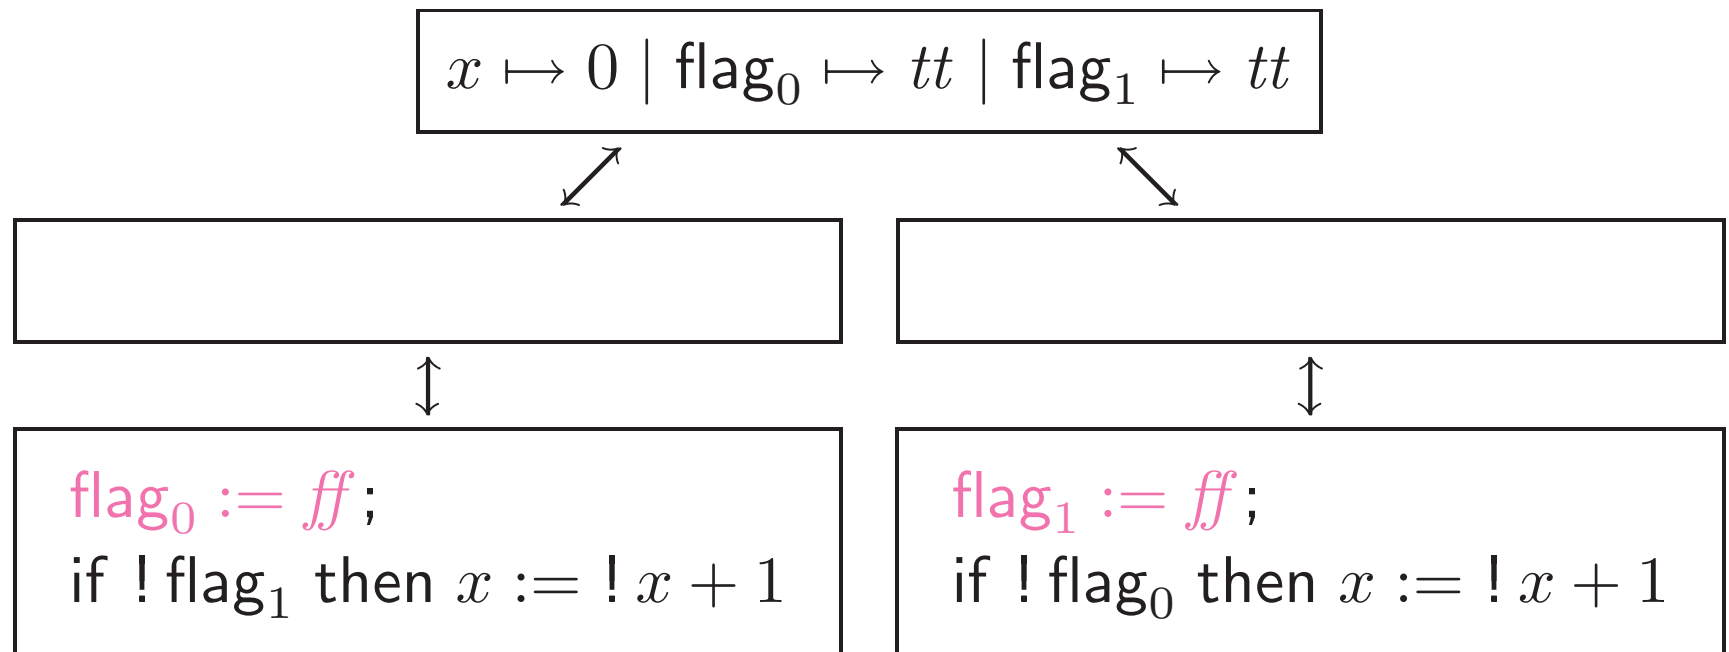
$$\frac{P_i \xrightarrow{\text{rd}_{x,v}} P'_i \quad S(x) = v \quad x \notin \text{dom}(B_i)}{(S, \dots, (B_i, P_i), \dots) \xrightarrow{i,\text{rd}_{x,v}} (S, \dots, (B_i, P'_i), \dots)}$$

$$\frac{B_i \equiv_{\text{TSO/PSO}} \langle x := v \rangle \cdot B'_i}{(S, \dots, (B_i, P_i), \dots) \xrightarrow{i,\text{wr}_{x,v}} (S[x := v], \dots, (B'_i, P_i), \dots)}$$

# For Instance – TSO

---

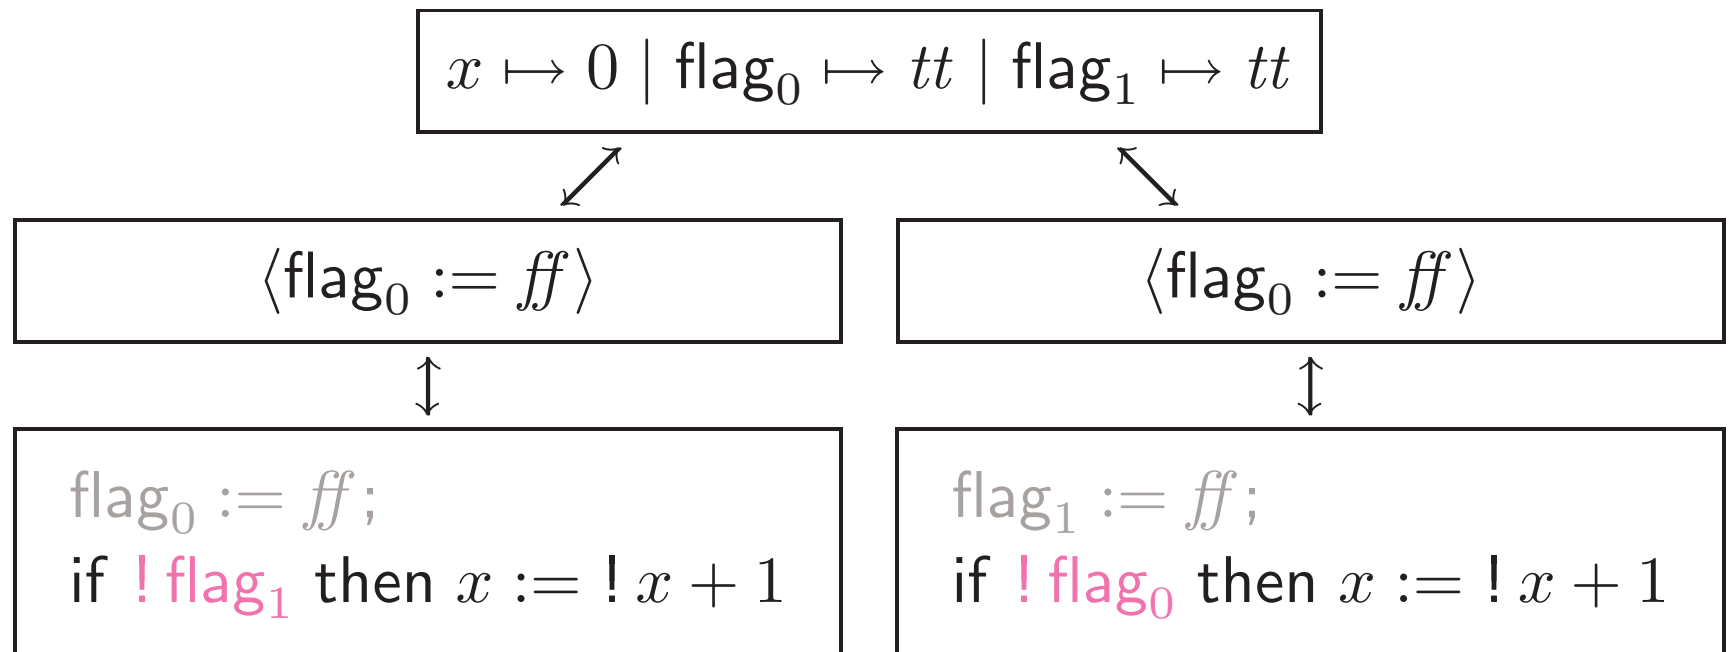
The “mutual exclusion” example:



# For Instance – TSO

---

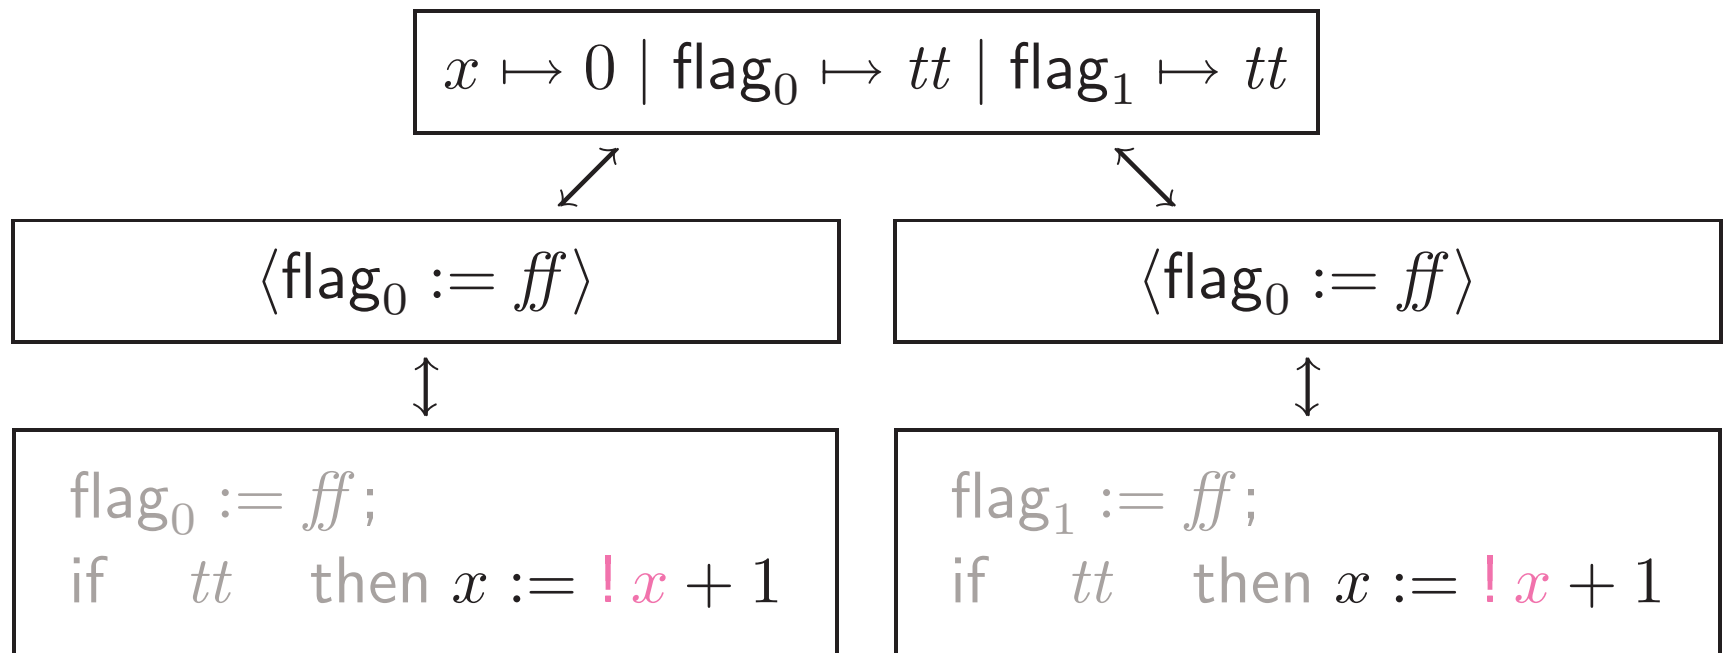
The writes are **issued** and put **in the buffers**:



# For Instance – TSO

---

The flags are read ( $tt$ ) from the shared memory, then branches selected:






# For Instance – TSO

---

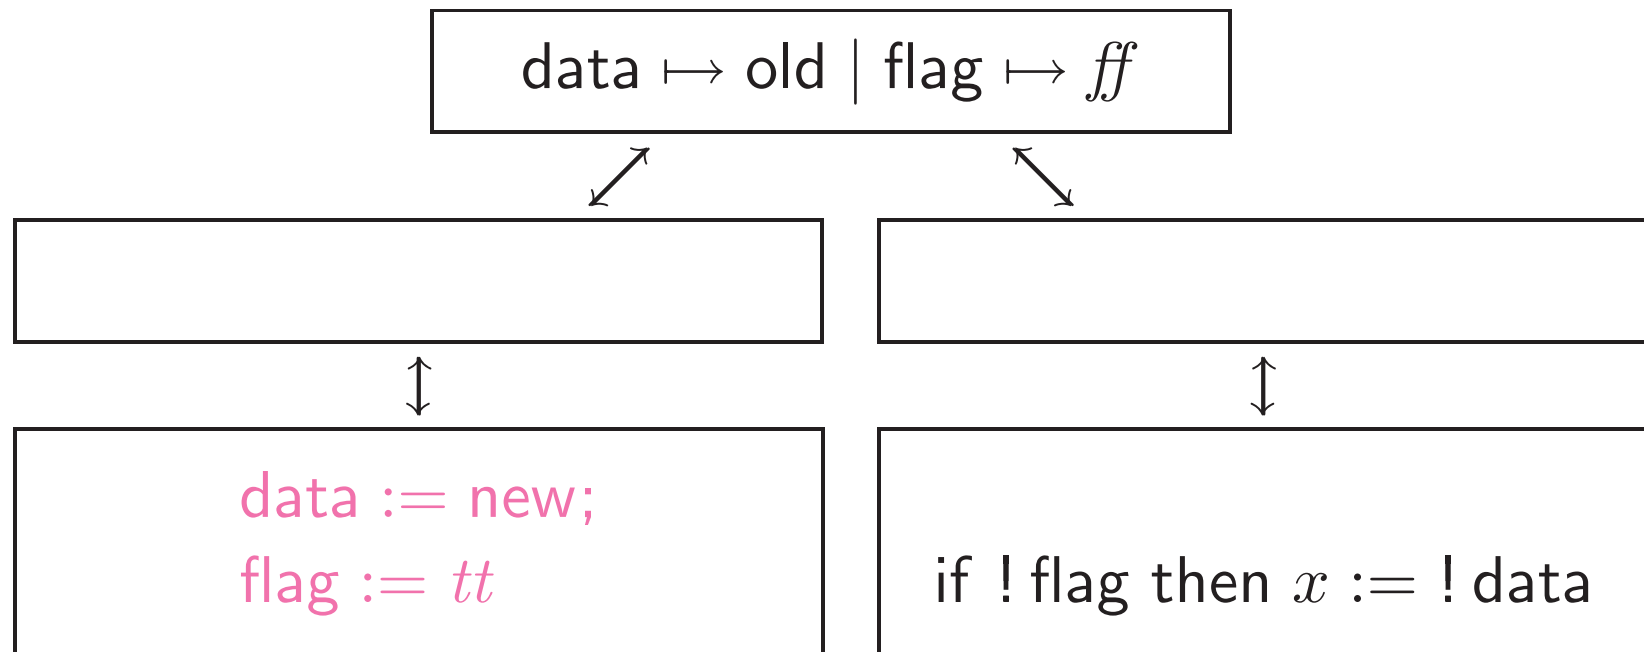
Value of  $x$  taken [from the shared memory](#) – not in the buffers:



# For Instance – PSO

---

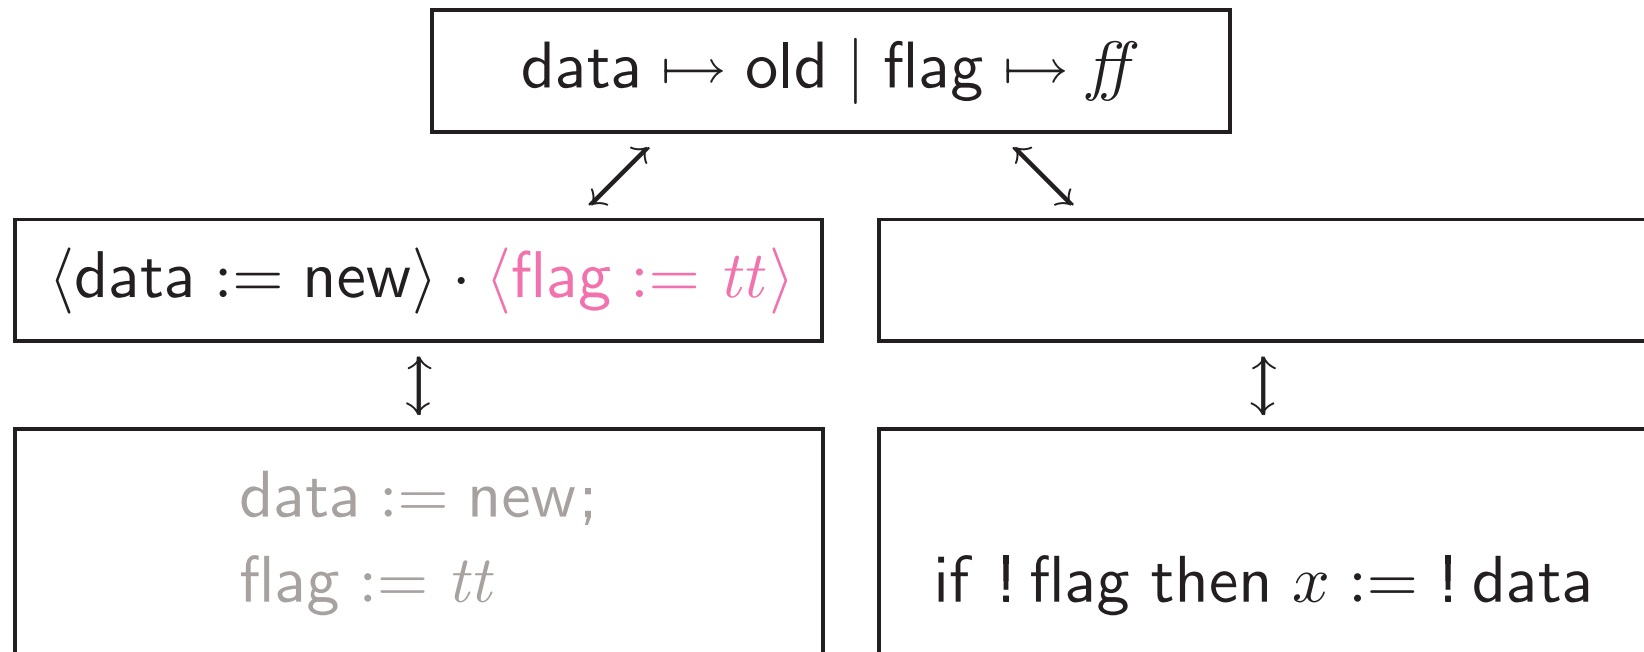
The producer/consumer example



# For Instance – PSO

---

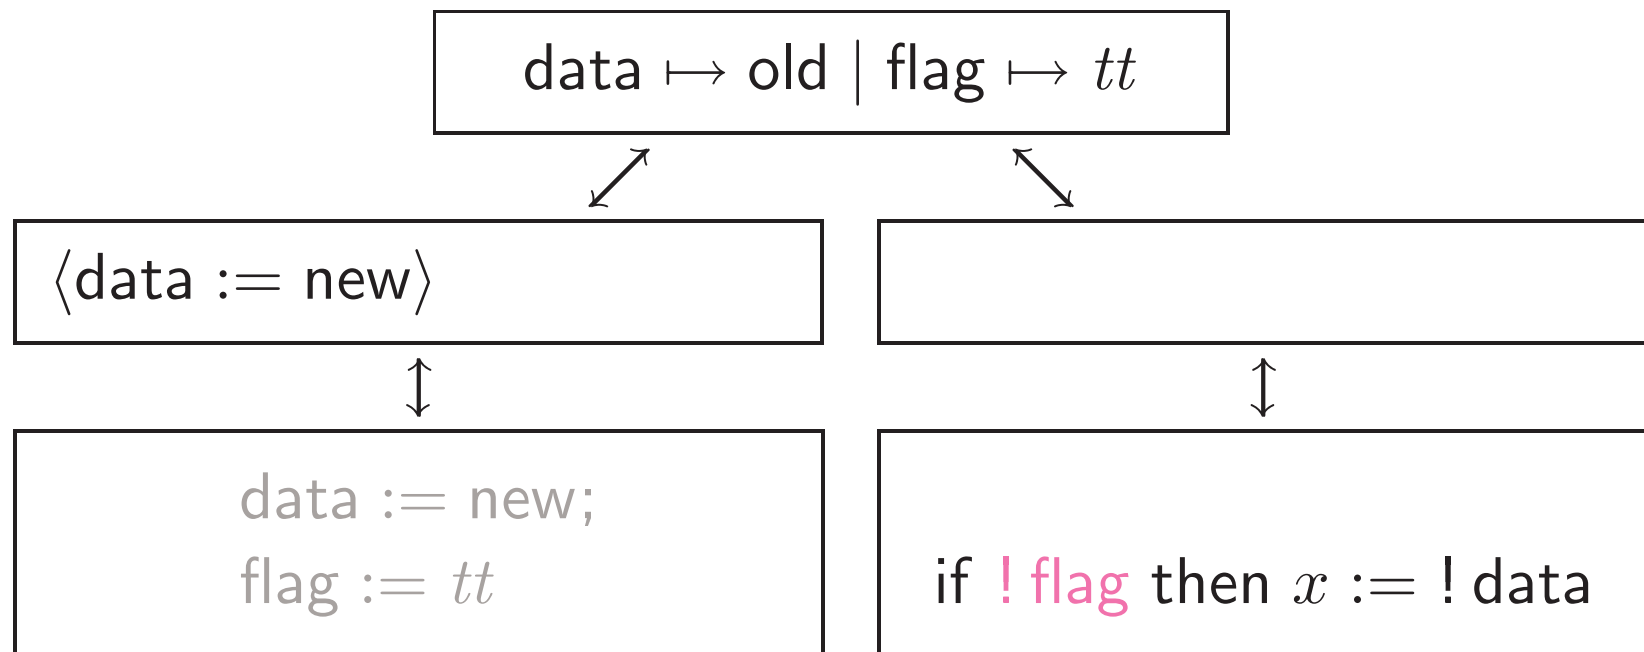
The writes are issued:



# For Instance – PSO

---

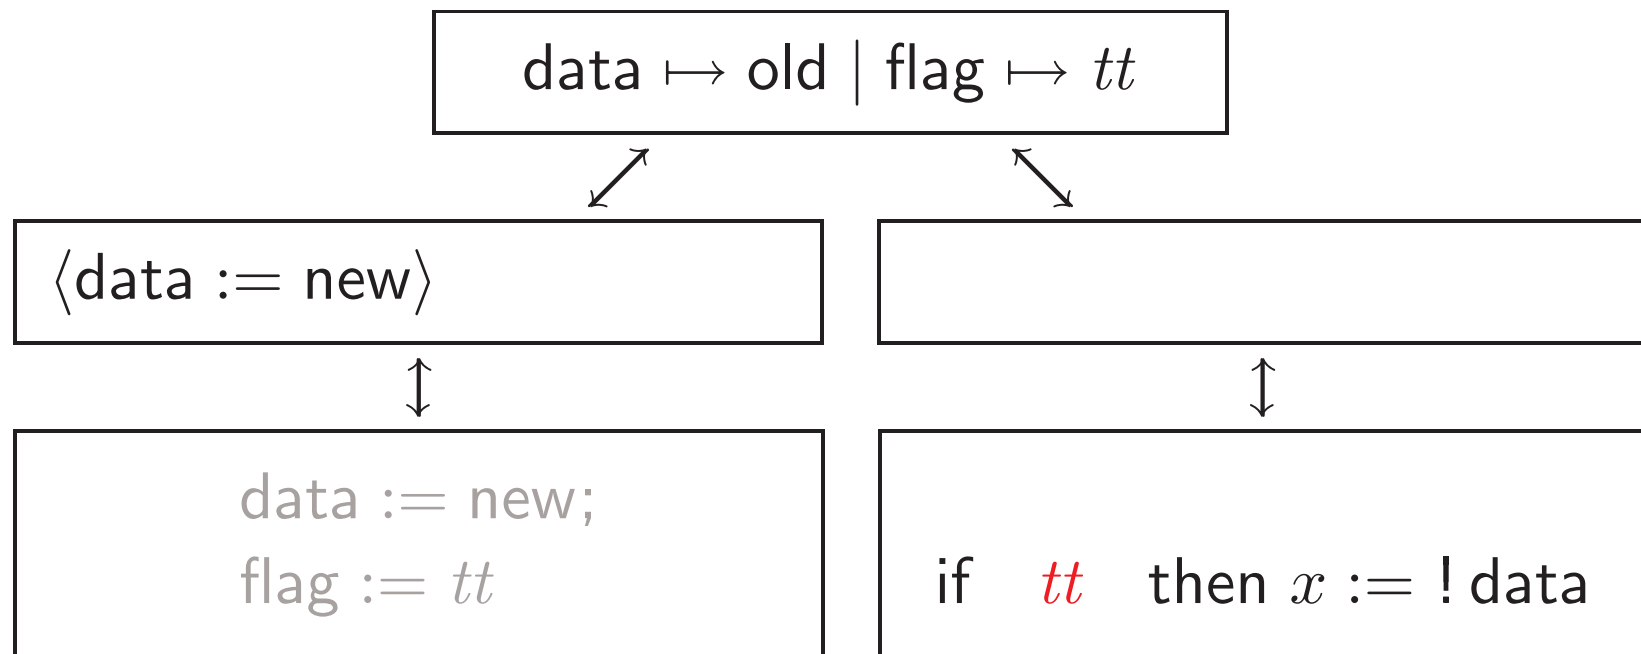
The memory is partially updated:



# For Instance – PSO

---

The **flag** is read from the memory:



then  $x$  may take the **old** value


# A METATHEORETICAL QUESTION

---

How to prove that a program appears to have a sequentially consistent semantics?

↳ an ideal tool: **GGJJ**'s equivalence by **permutation of independent steps**

Why? Under certain conditions (e.g. DRF) one can turn a **weak** computation into an **interleaving** computation, by permuting independent steps:



# Some RESULTS

---

[Boudol & Petri 2009]: a proof of the “DRF guarantee” for the TSO/PSO weak memory models using [GGJ's equivalence](#):

- DRF programs are “well-synchronized”
- DRF  $\Rightarrow$  (weak) computations are equivalent (by permutations) to SC (i.e. interleaving) computations

# SPECULATIVE COMPUTATIONS

---

(1/2)

Relaxation not allowed with write buffering: read/write order (RMO, PowerPC...)

↳ a more general approach: **speculative computations**

= computing (almost) anywhere in the code, introducing **speculation contexts**, e.g.

- ▶  $e; []$ , or more generally  $(\text{let } x = e \text{ in } []) = (\lambda x [] e)$ : **out of order evaluation**
- ▶  $(\text{if } e \text{ then } [] \text{ else } e_1)$  and  $(\text{if } e \text{ then } e_0 \text{ else } [])$ : **branch prediction**



# SPECULATIVE COMPUTATIONS

---

(2/2)

Speculations can go wrong:

- ▶ violation of **data dependency** – read/write on the same location
- ▶ violation of **control dependency** – computing in the wrong branch
- ↳ **validity** criterion: a speculation is valid iff it is, up to the reorderings allowed in a given memory model, a **reordering** – à la **GGJJ** – of a **normal computation**

A result: write buffering coincides with speculations where write/read (TSO) and write/write (PSO) relaxations are allowed.

More in the PhD Thesis of **Gustavo Petri** (November 2010)

thanks again, GG & JJ

and thanks to Pierre-Louis and Paul André