



COLLÈGE
DE FRANCE
—1530—

Program logics for weakly-consistent memory

Xavier Leroy

Collège de France, chair of software sciences

`xavier.leroy@college-de-france.fr`

**Sequential consistency:
an idealized model
for concurrent programming**


Sequential consistency

A multiprocessor system is sequentially consistent if 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
(L. Lamport, 1978)

For a shared-memory system: the state of the shared memory is the result of an **interleaving** of the memory operations of the processes.


Two sequentially-consistent (SC) implementations

1. Time-sharing on a monoprocessor.
2. Multiprocessor system with a single “port” to access memory.




Two sequentially-consistent (SC) implementations

1. Time-sharing on a monoprocessor.
2. Multiprocessor system with a single “port” to access memory.



Two sequentially-consistent (SC) implementations

1. Time-sharing on a monoprocessor.
2. Multiprocessor system with a single “port” to access memory.



A straightforward formal semantics

Reminder

Semantics for the $c_1 \parallel c_2$ construct:

an **interleaving** of the reductions of c_1 and c_2 .

$(a_1 \parallel a_2)/h \rightarrow 0/h$ (or any combination of a_1 and a_2)

$(c_1 \parallel c_2)/h \rightarrow (c'_1 \parallel c_2)/h'$ if $c_1/h \rightarrow c'_1/h'$

$(c_1 \parallel c_2)/h \rightarrow (c_1 \parallel c'_2)/h'$ if $c_2/h \rightarrow c'_2/h'$

$(c_1 \parallel c_2)/h \rightarrow \text{err}$ if $c_1/h \rightarrow \text{err}$ or $c_2/h \rightarrow \text{err}$

Our semantics for parallelism in PTR was already SC!

The SC model defines precisely the semantics of concurrent programs **even when they contain race conditions**.

This opens the way to concurrent algorithms that use race conditions in a well-controlled manner.

These algorithms are useful when the atomic instructions of the processor or the critical sections of the language are lacking or too expensive.

Peterson's mutual exclusion algorithm

```
flag : bool[2] = {false, false}; turn : {0, 1};
```

```
flag[0] := true;  
turn := 1;  
while flag[1] ∧ turn = 1  
do skip done;  
// enter critical section  
...  
// leave critical section  
flag[0] := false
```

```
flag[1] := true;  
turn := 0;  
while flag[0] ∧ turn = 0  
do skip done;  
// enter critical section  
...  
// leave critical section  
flag[1] := false
```

Peterson's mutual exclusion algorithm


```
flag : bool[2] = {false, false}; turn : {0, 1};

flag[0] := true;
turn := 1;
while flag[1] ∧ turn = 1
do skip done;
// enter critical section
...
// leave critical section
flag[0] := false

flag[1] := true;
turn := 0;
while flag[0] ∧ turn = 0
do skip done;
// enter critical section
...
// leave critical section
flag[1] := false
```

With an enumeration of SC interleavings, we can show that both processes are simultaneously in their critical sections only if $flag[0] = flag[1] = true \wedge turn = 0 \wedge turn = 1$, which is impossible.

The ticket lock algorithm



A process that tries to enter the critical section

- takes the next ticket (atomic increment)
- waits for the number on its ticket to be displayed.

When it leaves the critical section, it ensures the next number is displayed.

The ticket lock algorithm

Two global variables: `next` and `now_serving`, initially 0.

```
lock() {  
    int t = fetch_and_add(&next, 1);    // atomic increment  
    while (now_serving != t) pause();  // nonatomic read  
}
```

```
unlock() {  
    now_serving = now_serving + 1;    // nonatomic increment  
}
```

The increment of `next` must be atomic
(otherwise two processes could get the same ticket).

Accesses to `now_serving` need not be atomic.

**The real world:
weakly-consistent
memory models**

A litmus test

A fragment of Dekker's algorithm for mutual exclusion:

$$\begin{array}{l} \text{set}(X, 1); \\ \text{let } a = \text{get}(Y) \end{array} \parallel \parallel \begin{array}{l} \text{set}(Y, 1); \\ \text{let } b = \text{get}(X) \end{array}$$

Initially, $X = Y = 0$.

In the SC model:

- either $\text{set}(X, 1)$ executes first, and then $b = 1$ at the end;
- either $\text{set}(Y, 1)$ executes first, and then $a = 1$ at the end.

Therefore, the final state $a = b = 0$ is impossible.

Experimental refutation

We write the test in x86 assembly (to get full control on the code):

```
X86_64 Dekker
{ want0=0; want1=0; }
  P0          | P1 ;
  movl $1,(want0) | movl $1,(want1) ;
  movl (want1),%eax | movl (want0),%eax ;
exists (0:rax=0 /\ 1:rax=0)
```


We execute the test with the litmus7 tool:

```
Test Dekker Allowed
Histogram (4 states)
178      *>0:rax=0; 1:rax=0;
1999870:>0:rax=1; 1:rax=0;
1999881:>0:rax=0; 1:rax=1;
71       :>0:rax=1; 1:rax=1;
```

The sequential consistency model is respected

- neither by the modern **hardware architectures** (in order to provide faster memory subsystems).
- nor by **optimizing compilers** (in order to increase performance of generated code).

Hardware: write buffers, store buffers



(Higham, Jackson, Kawash, 2007)

Each processor puts its writes in a buffer while they are transmitted to the shared main memory.

Writes performed by a processor are immediately visible by this processor, but not immediately by the other processors.

A non-SC execution

```
set(X, 1);           || set(Y, 1);  
let a = get(Y)      || let b = get(X)
```

	Processor 1	Processor 2
Time $t = 0$	puts $X \leftarrow 1$ in its buffer	puts $Y \leftarrow 1$ in its buffer
Time $t = 1$	reads $Y = 0$ from main memory	
Time $t = 2$		reads $X = 0$ from main memory
Time $t = 3$	sends $X \leftarrow 1$ to main memory	
Time $t = 4$		sends $Y \leftarrow 1$ to main memory

Hardware: out-of-order, speculative execution

The processor can reorder instructions on the fly, so as to start long-running instructions earlier (e.g. memory reads).

`write X; . . . ; read Y` \rightarrow `read Y; write X; . . .`

This out-of-order execution is often speculative: if the processor realizes that $X = Y$, it cancels the anticipated read from Y , or satisfies it with the value written to X (forwarding).

Another non-SC execution

Machine code:

<code>set(X, 1);</code>		<code>set(Y, 1);</code>
<code>let a = get(Y)</code>		<code>let b = get(X)</code>

Code actually executed by the processor after on-the-fly reordering:

<code>let a = get(Y) in</code>		<code>let b = get(X) in</code>
<code>set(X, 1)</code>		<code>set(Y, 1)</code>

The reordered code can obviously terminate with $a = b = 0$.

Hardware: split memory accesses

A misaligned datum can span two cache lines, requiring two memory accesses per access to the datum.

```
set(X, 0x12345678) || let a = get(X)
```

can be executed like

```
set(X1, 0x1234); || let a1 = get(X1) in  
set(X2, 0x5678); || let a2 = get(X2) in  
|| let a = a1 << 16 | a2
```

A non-SC with values “out of thin air”

```
set(X, 0x12345678) || let a = get(X)
```

Starting from $X = 0$, we have two SC executions:

$a = 0$ or $a = 0x12345678$.

```
set(X1, 0x1234); || let a1 = get(X1) in  
set(X2, 0x5678); || let a2 = get(X2) in  
|| let a = a1 << 16 | a2
```

After splitting the memory accesses, a third result is possible:

$a = 0x12340000$, coming from $a_1 = 0x1234$ and $a_2 = 0$.

Note: the value $0x12340000$ appears nowhere in the initial code.

It appears **out of thin air!**

Possible remedies

Use **barrier instructions** that prevent the processor from reordering certain memory accesses:

- “Strong” barrier: preserves ordering between accesses before the barrier and accesses after the barrier.
- “Weak” barrier: preserves ordering between reads before the barrier and accesses after the barrier.

Other instructions with special memory behaviors:

- “locked” instructions (the x86 `lock` prefix);
- *load-acquire* and *store-release* (Itanium, ARM);
- etc.

Compiler optimizations: instruction reordering

The compiler can reorder independent reads and writes (at addresses X , Y that are guaranteed to be different).

Typically, reads are anticipated while writes are delayed.

`write X; . . . ; read Y` \rightarrow `read Y; write X; . . .`

\rightarrow same non-SC behaviors as dynamic reordering by the processor.

Compiler optimization: factoring redundant reads

A special case of common subexpression elimination (CSE).

let $a = \text{get}(X)$ in		let $a = \text{get}(X)$ in
...		...
(no writes to X)	\rightsquigarrow	...
...		...
let $b = \text{get}(X)$ in		let $b = a$ in
...		...

A non-SC execution

```
let a = get(X) in   || set(X, 1);  
let b = get(Y) in   || set(Y, 1);  
let c = get(X)
```

With $X = Y = 0$ initially, no SC execution terminates with $(a, b, c) = (0, 1, 0)$.

After factoring of `get(X)`, we have `let c = a` and the result $(a, b, c) = (0, 1, 0)$ is possible.

Compiler optimizations: loop-invariant code motion

A computation performed repeatedly at each loop iteration can be performed once before the loop:

		<code>t := j × 10;</code>
<code>for i = 0 to 99 do</code>	\rightsquigarrow	<code>for i = 0 to 99 do</code>
<code>A[i] := i + j × 10</code>		<code>A[i] := i + t</code>
<code>done</code>		<code>done</code>

This can break codes based on busy waiting:

		<code>t := get(X);</code>
<code>do</code>	\rightsquigarrow	<code>do</code>
<code>t := get(X)</code>		<code>skip</code>
<code>while t = 0</code>		<code>while t = 0</code>

Turn optimizations off? Never!

Inform the compiler of which memory accesses implement inter-process communications, so as to compile them specially:

- the `volatile` modifier (C, C++, Java)
- a library of low-level atomic operations (C/C++ 2011)
- etc.

The *low-level atomics* in C/C++ 2011

Various atomic operations: read, write, fetch-and-add, compare-and-swap,

Each operation is annotated with the consistency model expected by the programmer:

<code>memory_order_seq_cst</code>		sequential consistency
<code>memory_order_acq_rel</code>	}	just enough for message passing
<code>memory_order_acquire</code>		
<code>memory_order_release</code>		
<code>memory_order_consume</code>		
<code>memory_order_relaxed</code>		no guarantees beyond atomicity

DRF + CSL = ♥

**Concurrent separation logic
and the DRF guarantee**

The Data Race Free (DRF) guarantee

A property of a relaxed memory model:

*If a program executes in the SC model without race conditions,
then it executes in the relaxed model exactly like in the SC model.*

In other words: for a program free of race conditions, the relaxations of the memory model do not add more behaviors beyond those permitted by SC.

This “DRF guarantee” seems to hold / is claimed to hold for all the known memory models (hardware models + language models).

The extended DRF guarantee

If a program comprising critical sections, atomic operations, and other synchronization devices executes in the SC model without race conditions,
then it executes in the relaxed model exactly like in the SC model,
provided the synchronization devices are correctly implemented.

“Correctly implemented” = with enough memory barriers and special instructions to rule out non-SC behaviors.

Examples: implementing locks

x86: lock

```
    movl  $1, %edx
.L2: movl  %edx, %eax
    xchgb (%rdi), %al
    testb %al, %al
    jne   .L2
```

unlock

```
    movb  $0, (%rdi)
```

Power: lock

```
.L2: lbarx  9,0,3
     stbcx. 10,0,3
     bne    0, .L2
     isync
     andi.  9,9,0xff
     bne    0, .L2
```

unlock

```
     lwsync
     li    9,0
     stb   9,0(3)
```

A subtle point of the x86 implementation

Before 1999, the Linux kernel implemented `unlock` with an atomic instruction

```
lock; btr $0, (...)
```

instead of a nonatomic write

```
movb $0, (...)
```

A long discussion concluded that a nonatomic write is enough, because the x86 memory model is TSO.

The write of 0 in the lock is not immediately visible by other processes waiting for the lock. But when it becomes visible, all preceding writes are already visible, and all preceding reads have been performed.

A great many compiler optimizations are sound for programs that contain no race conditions.

(Sound = all behaviors of the optimized program are possible behaviors of the original program. Optimization did not introduce additional behaviors.)

Compatibility with compiler optimizations

Transformation	SC	DRF guarantee	JMM
Trace-preserving transformations	✓	✓	✓
Reordering normal memory accesses	×	✓	×
Redundant read after read elimination	✓	✓	×
Redundant read after write elimination	✓	✓	✓
Irrelevant read elimination	✓	✓	✓
Irrelevant read introduction	✓	?	×
Redundant write before write elimination	✓	✓	✓
Redundant write after read elimination	✓	✓	×
Roach-motel reordering	×(✓ for locks)	✓	×
External action reordering	×	✓	×

J. Ševčík, *Program Transformations in Weak Memory Models*, PhD, 2008.

How to prove the absence of race conditions?

The assumption “the program has no race conditions in the SC model” is strong! How to establish it?

- *Ad-hoc* proof.
- Type system (\approx Rust).
- Static analysis (Infer, etc).
- Deductive verification in concurrent separation logic!

Reminder (lecture #4): if $J \vdash \{ P \} c \{ Q \}$, then c executes without race conditions in an interleaving semantics that is equivalent to the SC model.

The CSL guarantee (Concurrent Separation Logic)

*If a program is **provable in concurrent separation logic** (including critical sections, atomic sections, etc),
then it executes correctly in a relaxed memory model that respects the DRF guarantee,
provided that critical sections and atomic sections are correctly implemented.*

A direct proof of the CSL guarantee for a TSO model

We now show semantic soundness for concurrent separation logic in a variant of our PTR language extended with write buffers (TSO model):

Write buffers: $s ::= \varepsilon \mid (\ell, v) \cdot s$

We consider whole-program configurations

$$((c_1/s_1) \parallel \dots \parallel (c_n/s_n)) / h$$

composed of n processes $c_1 \dots c_n$, each with its own buffer s_i , plus a global heap h .

We also consider local, per-process configurations

$$c/s/h$$

At any time a process can perform the oldest write from its buffer:

$$c/s \cdot (\ell, v)/h \rightarrow c/s/h[\ell \leftarrow v]$$

The base language constructs have their usual semantics:

$$(\text{let } x = a \text{ in } c)/s/h \rightarrow c[x \leftarrow \llbracket a \rrbracket]/s/h$$

$$(\text{let } x = c_1 \text{ in } c_2)/s/h \rightarrow (\text{let } x = c'_1 \text{ in } c_2)/s'/h' \\ \text{if } c_1/s/h \rightarrow c'_1/s'/h'$$

$$(\text{let } x = c_1 \text{ in } c_2)/s/h \rightarrow \text{err} \quad \text{if } c_1/s/h \rightarrow \text{err}$$

Local reductions: imperative constructs

Imperative constructs write to s (the buffer) and read from $s \triangleright h$, the heap h updated as described by s :

$$\varepsilon \triangleright h = h \quad ((\ell, v) \cdot s) \triangleright h = (s \triangleright h)[\ell \leftarrow v]$$

$$\text{get}(a)/s/h \rightarrow (s \triangleright h)(\llbracket a \rrbracket)/s/h \quad \text{if } \llbracket a \rrbracket \in \text{Dom}(s \triangleright h)$$

$$\text{set}(a, a')/s/h \rightarrow 0/(\llbracket a \rrbracket, \llbracket a' \rrbracket) \cdot s/h \quad \text{if } \llbracket a \rrbracket \in \text{Dom}(s \triangleright h)$$

$$\text{get}(a)/s/h \rightarrow \text{err} \quad \text{if } \llbracket a \rrbracket \notin \text{Dom}(s \triangleright h)$$

$$\text{set}(a, a')/s/h \rightarrow \text{err} \quad \text{if } \llbracket a \rrbracket \notin \text{Dom}(s \triangleright h)$$

Local reductions: atomic sections

As in PTR, atomic sections execute in a single “big step”:

$$\begin{array}{ll} \text{atomic}(c)/s/h \rightarrow a/\varepsilon/h' & \text{if } c/s/h \xrightarrow{*} a/\varepsilon/h' \\ \text{atomic}(c)/s/h \rightarrow \text{err} & \text{if } c/s/h \xrightarrow{*} \text{err} \end{array}$$

However, the buffer must be empty at the end of the atomic section (\approx there is a write barrier at the end).

When we create a resource invariant, the buffer must also be empty, hence the `mkinv(c)` construct:

$$\text{mkinv}(c)/\varepsilon/h \rightarrow c/\varepsilon/h$$

At each step we locally reduce one of the processes c_i/s_i from the parallel composition; the other processes are unchanged.

$$\frac{c_i/s_i/h \rightarrow c'/s'/h'}{(\dots \parallel (c_i/s_i) \parallel \dots) / h \rightarrow (\dots \parallel (c'/s') \parallel \dots) / h'}$$

$$\frac{c_i/s_i/h \rightarrow \text{err}}{(\dots \parallel (c_i/s_i) \parallel \dots) / h \rightarrow \text{err}}$$

Semantic soundness of concurrent separation logic

For PTR in the SC model, we decomposed the current heap h in three disjoint parts:

$$h = h_1 \uplus h_j \uplus h_f$$

h_1 is the private memory for c .

h_j is the shared memory accessible to atomic sections.

h_f is the “frame” memory, including the private memories of the processes that execute in parallel with c .

Semantic soundness of concurrent separation logic

For PTR in the TSO model, we decompose the main heap h and the buffer s for the current process c as follows:

$$h = h_u \uplus h_j \quad s \triangleright h_u = h_1 \uplus h_f$$

The main heap h decomposes into shared memory h_j and unshared memory h_u .

The unshared memory h_u , updated according to the buffer s , decomposes into h_1 , the private memory for c , and h_f , the frame.

Equivalent presentation:

$$s \triangleright h = h_1 \uplus h_j \uplus h_f \quad \text{with} \quad \text{Dom}(s) \cap \text{Dom}(h_j) = \emptyset$$

The semantic triple

We define the semantic triple $J \models \{\{P\}\} c \{\{Q\}\}$ as follows:

$$J \models \{\{P\}\} c \{\{Q\}\} \stackrel{\text{def}}{=} \forall n, h, P h \Rightarrow \text{Safe}^n c h Q J$$

As in lecture #4, we have:

$$\text{Safe}^0 c h Q J \quad \frac{Q \llbracket a \rrbracket h}{\text{Safe}^{n+1} a h Q J} \quad \frac{(\forall a, c \neq a) \quad \dots}{\text{Safe}^{n+1} c h Q J}$$

The semantic triple

The recursive case: one reduction step for $c/s/h$.

$$\forall a, c \neq a$$

$$\forall s, h, h_j, h_f, s \triangleright h = h_1 \uplus h_j \uplus h_f \wedge \text{Dom}(s) \cap \text{Dom}(h_j) = \emptyset \wedge J h_j \Rightarrow \\ c/s/h \not\rightarrow \text{err}$$

$$\forall s, h, h_j, h_f, c', s', h',$$

$$s \triangleright h = h_1 \uplus h_j \uplus h_f \wedge \text{Dom}(s) \cap \text{Dom}(h_j) = \emptyset$$

$$\wedge J h_j \wedge c/s/h \rightarrow c'/s'/h'$$

$$\Rightarrow \exists h'_1, h'_j, s' \triangleright h' = h'_1 \uplus h'_j \uplus h_f \wedge \text{Dom}(s') \cap \text{Dom}(h'_j) = \emptyset$$

$$\wedge J h'_j \wedge \text{Safe}^n c' h'_1 Q$$

$$\text{Safe}^{n+1} c h_1 Q$$

Validating the rules of the logic

It remains to show that this semantic triple $J \models \{P\} c \{Q\}$ validates the rules of concurrent separation logic. The two interesting cases deal with atomic sections.

$$\frac{\text{emp} \vdash \{P \star J\} c \{\lambda v. Q v \star J\}}{J \vdash \{P\} \text{atomic} c \{Q\}}$$

At the end of the execution of c we have a decomposition $\varepsilon \triangleright h' = (h'_1 \uplus h'_j) \uplus \emptyset \uplus h_f$ that we rewrite as $\varepsilon \triangleright h' = h'_1 \uplus h'_j \uplus h_f$. It is crucial that the final buffer s' is ε , otherwise the constraint $\text{Dom}(s') \cap \text{Dom}(h'_j) = \emptyset$ could not be satisfied.

Validating the rules of the logic

Consider adding an invariant J' to J :

$$\frac{J \star J' \vdash \{P\} c \{Q\}}{J \vdash \{P \star J'\} \text{mkinv} c \{\lambda v. Q v \star J'\}}$$

At the beginning of the execution, we have a decomposition

$s \triangleright h = (h_1 \uplus h'_j) \uplus h_j \uplus h_f$ that we rewrite as

$s \triangleright h = h_1 \uplus (h_j \uplus h'_j) \uplus h_f.$

Here too we must enforce $s = \varepsilon$ to satisfy

$\text{Dom}(s) \cap \text{Dom}(h_j \uplus h'_j) = \emptyset.$

The `mkinv` construct forces the buffer to be empty.

A limitation and a strength

A limitation: our formalization fails to justify the implementation of `unlock(ℓ) = atomic(set(ℓ , 0))` by a normal write, without flushing the store buffer.

This is an aspect of TSO that our formalization does not capture.

A strength: this makes our proof reusable for memory models that are more relaxed than TSO, in particular PSO (*Partial Store Ordering*), where this optimization of `unlock` is invalid.

In the PSO model, writes to different locations can be reordered, and therefore “leave” the write buffer in a different order than execution order:

$$c/s_1 \cdot (\ell, v) \cdot s_2/h \rightarrow c/s_1 \cdot s_2/h[\ell \leftarrow v] \quad \text{if } \ell \notin \text{Dom}(s_2)$$

Here, the write to ℓ “overtakes” the writes in s_2 .

This makes no difference for the soundness proof of the logic, since

$$(s_1 \cdot (\ell, v) \cdot s_2) \triangleright h = (s_1 \cdot s_2) \triangleright (h[\ell \leftarrow v]) \quad \text{if } \ell \notin \text{Dom}(s_2)$$

A logic for release-acquire

The release-acquire model

A write marked “release” guarantees that all preceding reads and writes have been performed.

(No reordering of $X; W_{rel}$ into $W_{rel}; X$.)

A read marked “acquire” guarantees that all following reads and writes have not started yet.

(No reordering of $R_{acq}; X$ into $X; R_{acq}$.)

Use for message passing

```
// preparing the message
// nonatomic (na) writes
setna(msg, ...);
setna(msg + 1, ...);
setna(msg + 2, ...);
// sending the message
setrel(ready, 1)
```

```
// waiting for the message
// read ready until ≠ 0
while getacq(ready) = 0 do skip;
// accessing the message
// nonatomic reads
let x = getna(msg) in
...
```

A lightweight form of synchronization and resource transfer, without mutual exclusion.

Use for implementing locks

Unlocking a lock is just one release write:

$$\text{unlock}(\ell) = \text{set}_{rel}(\ell, 0)$$

Locking a lock requires an atomic instruction such as Compare And Swap, marked “acquire”:

$$\text{lock}(\ell) = \text{while } (\text{CAS}_{acq}(\ell, 0, 1) \neq 0) \text{ do skip}$$

We can improve performance with a busy-wait loop using relaxed reads:

$$\text{spin}(\ell) = \text{while } (\text{get}_{rlx}(\ell) \neq 0) \text{ do skip}$$
$$\text{lock}(\ell) = \text{do } \text{spin}(\ell) \text{ while } (\text{CAS}_{acq}(\ell, 0, 1) \neq 0)$$

Implementing the release/acquire model

For TSO architectures like x86: nothing to do!

- Ordinary writes have release semantics.
- Ordinary loads have acquire semantics.

For more relaxed architectures like Power and ARM:

- Memory barriers that are less costly than the barriers needed to guarantee SC.

Relaxed Separation Logic (Vafeiadis & Narayan, 2013)

A concurrent separation logic for a fragment of the low-level atomics from C/C++ 2011.

$c ::= a$	pure expression
$\text{let } x = c \text{ in } c'$	sequencing and binding
$\text{if } a \text{ then } c_1 \text{ else } c_2$	conditional
$\text{repeat } c$	repeat until not 0
$c_1 \parallel c_2$	parallel execution
$\text{alloc}()$	allocation
$\text{get}_x(a)$	memory read
$\text{set}_y(a, a')$	memory write
$\text{CAS}_{z,x}(a, a', a'')$	Compare And Swap
$X ::= \text{sc} \mid \text{acq} \mid \text{rlx} \mid \text{na}$	type of read
$Y ::= \text{sc} \mid \text{rel} \mid \text{rlx} \mid \text{na}$	type of store
$Z ::= \text{sc} \mid \text{rel_acq} \mid \text{acq} \mid \text{rel}$	type of CAS

Assertions, preconditions:

$P ::= \langle A \rangle \mid \text{true} \mid P \star P'$

| $\ell \stackrel{\pi}{\mapsto} v$ location ℓ contains v with permission π

| $\text{Acq}(\ell, \Phi) \mid \text{Rel}(\ell, \Phi)$ resource invariants

| $\text{RMWAcq}(\ell, \Phi)$ resource invariant

| $\text{Init}(\ell)$ invariant was established

| $\text{Uninit}(\ell)$ location ℓ is freshly allocated

Postconditions, resource invariants:

$Q, \Phi ::= \lambda v. P$

Nonatomic accesses

Nonatomic reads and writes follow standard separation logic:

$$\begin{array}{l} \{ \text{emp} \} \quad \text{alloc}() \quad \{ \lambda \ell. \text{Uninit}(\ell) \} \\ \{ \ell \stackrel{\pi}{\mapsto} \mathbf{v} \} \quad \text{get}_{na}(\ell) \quad \{ \lambda x. \langle x = \mathbf{v} \rangle \star \ell \stackrel{\pi}{\mapsto} \mathbf{v} \} \\ \{ \text{Uninit}(\ell) \vee \ell \stackrel{1}{\mapsto} - \} \quad \text{set}_{na}(\ell, \mathbf{v}) \quad \{ \lambda -. \ell \stackrel{1}{\mapsto} \mathbf{v} \} \end{array}$$

(The role of *Uninit* is to prevent us from reading from a freshly allocated memory location that has not been initialized yet.)

Release writes, acquire reads

$Rel(\ell, \Phi)$ grants permission to write to location ℓ a value v provided we have the resource Φv .

$Acq(\ell, \Phi)$, in conjunction with $Init(\ell)$, grants permission to read from location ℓ , obtaining a value v and the resource Φv .

$$\begin{array}{l} \{ \text{emp} \} \quad \text{alloc}() \quad \{ \lambda \ell. Rel(\ell, \Phi) \star Acq(\ell, \Phi) \} \\ \{ Rel(\ell, \Phi) \star \Phi v \} \quad \text{set}_{rel}(\ell, v) \quad \{ Rel(\ell, \Phi) \star Init(\ell) \} \\ \{ Acq(\ell, \Phi) \star Init(\ell) \} \quad \text{get}_{acq}(\ell) \quad \{ \lambda v. \Phi v \star Acq(\ell, \Phi[v \leftarrow \text{emp}]) \} \end{array}$$

We can read the same value multiple times, but the second and subsequent reads transfer no resources:

$$\Phi[v \leftarrow \text{emp}] \stackrel{def}{=} \lambda v'. \text{if } v' = v \text{ then emp else } \Phi v'.$$

Example of resource transfer

We have a (nonatomic) buffer b and an (atomic) flag x .

We take $\Phi = \lambda v. \text{if } v = 0 \text{ then emp else } b \xrightarrow{1} 53$.

```
let  $x = \text{alloc}()$  in let  $b = \text{alloc}()$  in  $\text{set}_{rel}(x, 0)$ ;  
  {  $\text{Uinit}(b) * \text{Rel}(x, \Phi) * \text{Init}(x) * \text{Acq}(x, \Phi)$  }
```

```
{  $\text{Uinit}(b) * \text{Rel}(x, \Phi)$  }  
   $\text{set}_{na}(b, 53)$ ;  
  {  $b \xrightarrow{1} 53 * \text{Rel}(x, \Phi)$  }  
 $\Rightarrow$  {  $\Phi 1 * \text{Rel}(x, \Phi)$  }  
   $\text{set}_{rel}(x, 1)$ 
```

```
{  $\text{Init}(x) * \text{Acq}(x, \Phi)$  }  
  repeat  $\text{get}_{acq}(x)$ ;  
  {  $\exists v \neq 0, \Phi v$  }  $\Rightarrow$  {  $b \xrightarrow{1} 53$  }  
  let  $n = \text{get}_{na}(b)$  in  
  {  $b \xrightarrow{1} 53 * \langle n = 53 \rangle$  }
```

Multiple readers, multiple writers

Write permissions can be duplicated:

$$\text{Init}(\ell) = \text{Init}(\ell) \star \text{Init}(\ell) \quad \text{Rel}(\ell, \Phi) = \text{Rel}(\ell, \Phi) \star \text{Rel}(\ell, \Phi)$$

Read permissions can be split:

$$\text{Acq}(\ell, \lambda v. \Phi_1 v \star \Phi_2 v) = \text{Acq}(\ell, \Phi_1) \star \text{Acq}(\ell, \Phi_2)$$

Example: one writer, two readers.

<code>set_{na}(a, 13);</code>		<code>repeat get_{acq}(x);</code>		<code>repeat get_{acq}(x);</code>
<code>set_{na}(b, 17);</code>		<code>{ a $\xrightarrow{1}$ 13 }</code>		<code>{ b $\xrightarrow{1}$ 17 }</code>
<code>set_{rel}(x, 1);</code>				

Compare And Swap

$$\{ \text{emp} \} \text{alloc}() \{ \lambda \ell. \text{Rel}(\ell, \Phi) \star \text{RMWAcq}(\ell, \Phi) \}$$

RMWAcq permissions can be duplicated:

$$\text{RMWAcq}(\ell, \Phi) \star \text{RMWAcq}(\ell, \Phi) = \text{RMWAcq}(\ell, \Phi)$$

The rule for $\text{CAS}_{X,rlx}$:

$$P \Rightarrow \text{Init}(\ell) \star \text{RMWAcq}(\ell, \Phi) \star \text{true}$$

$$P \star \Phi v \Rightarrow \text{Rel}(\ell, \Psi) \star \Psi v' \star R 1$$

$$P \Rightarrow R 0$$

$$X \in \{ \text{rel}, \text{rlx} \} \Rightarrow \Phi v = \text{emp} \quad X \in \{ \text{acq}, \text{rlx} \} \Rightarrow \Psi v' = \text{emp}$$

$$\{ P \} \text{CAS}_{X,rlx}(\ell, v, v') \{ R \}$$

A logic for relaxed accesses

Relaxed reads and writes

Intuition: like a release write or an acquire load, but without resource transfer.

$$\{ \text{Acq}(\ell, \Phi) \} \text{get}_{rlx}(\ell) \{ \lambda v. \langle \Phi v \neq \text{false} \rangle \}$$

$$\Phi v = \text{emp} \text{ (i.e. } \Phi v \text{ is pure and true)}$$

$$\{ \text{Rel}(\ell, \Phi) \} \text{set}_{rlx}(\ell, v) \{ \text{Rel}(\ell, \Phi) \}$$

A modest application: control the set of all possible values for location ℓ , for instance $\Phi = \lambda v. \langle 0 \leq v \leq 10 \rangle$.

Relaxed accesses and values out of thin air

Vafeiadis & Narayan point out that these rules are incorrect for C/C++ 2011, since relaxed accesses are allowed to produce values out of thin air.

$$\begin{array}{l} \text{let } a = \text{get}_{rlx}(X) \text{ in} \\ \text{set}_{rlx}(Y, a) \end{array} \quad \parallel \quad \begin{array}{l} \text{let } b = \text{get}_{rlx}(Y) \text{ in} \\ \text{set}_{rlx}(X, b) \end{array}$$

Starting with $X = Y = 0$, we can (according to the C/C++ 2011 standard) end with $X = Y = 1$.

According to Vafeiadis & Narayan's rules, $\Phi = \lambda v. \langle v = 0 \rangle$ is a correct invariant for X and for Y .

The problem with values out of thin air

A major risk: they break type safety and open security holes.

A theoretical risk: no known architecture or compiler exhibits “out of thin air” behaviors.

A specification problem: axiomatic definitions (using event structures) of C11-style memory models have a hard time distinguishing between

- behaviors involving values out of thin air;
- speculative behaviors (of the “load buffering” kind) that are correct.

A promising semantics

Kang et al (2017) describe an operational semantics for C/C++ 2011 atomics, of the (simplified) shape below.

1- Shared memory $M =$ a set of write messages

A message is $\langle \ell : v @ t \rangle$, representing the write of value v at location ℓ at timestamp t .

There is at most one message $\langle \ell : v @ t \rangle$ for a given ℓ and a given t .

A promising semantics

Kang et al (2017) describe an operational semantics for C/C++ 2011 atomics, of the (simplified) shape below.

1- Shared memory M = a set of write messages

2- A process = a view V of the shared memory ...

A view = a function location \rightarrow time at which the contents of this location was observed most recently.

Reading from location ℓ = observing a message $\langle \ell : v @ t \rangle$ with $t \geq V(\ell)$.

Writing v to location ℓ = sending a message $\langle \ell : v @ t \rangle$ with $t > V(\ell)$ a fresh timestamp.

In both cases, V is updated to $V[\ell \leftarrow t]$.

A promising semantics

Kang et al (2017) describe an operational semantics for C/C++ 2011 atomics, of the (simplified) shape below.

1- Shared memory M = a set of write messages

2- A process = a view V of the shared memory ...

3- ... plus a set P of promises.

A promise = a speculative write = a message already in the shared memory, but which still needs to be realized by an actual write later in the process execution.

A promising semantics


Kang et al (2017) describe an operational semantics for C/C++ 2011 atomics, of the (simplified) shape below.

1- Shared memory M = a set of write messages

2- A process = a view V of the shared memory ...

3- ... plus a set P of promises.

Invariant enforced at each reduction step: it is always possible to reduce so as to realize all promises. This prevents out-of-thin-air behaviors.



SLR: a separation logic for the promising semantics

(Svendsen *et al*, 2018.)

An extension of RSL with extra assertions:

$O(\ell, v, t)$ (generalizes $Init(\ell)$)

I observed value v in location ℓ at time t .

$W^\pi(\ell, X)$ (generalizes $\ell \xrightarrow{\pi} v$)

I have permission π on location ℓ .

$X = \{(v_1, t_1), \dots, (v_n, t_n)\}$ is a set of timestamped writes to this location.

If $\pi = 1$ (exclusive permission), X contains all the writes to ℓ ever performed.

If $\pi < 1$, X is a subset of these writes.

Some properties of assertion W

The assertion can be split:

$$W^{\pi_1 + \pi_2}(\ell, X_1 \cup X_2) = W^{\pi_1}(\ell, X_1) \star W^{\pi_2}(\ell, X_2)$$

Writes are consistent (unique value for a given timestamp):

$$W^\pi(\ell, X) \star \langle (v, t) \in X \wedge (v', t') \in X \wedge v \neq v' \rangle \Rightarrow W^\pi(\ell, X) \star \langle t \neq t' \rangle$$

All writes are observed:

$$W^\pi(\ell, X) \star \langle (v, t) \in X \rangle \Rightarrow W^\pi(\ell, X) \star O(\ell, v, t)$$

The converse is true if the permission is exclusive:

$$W^1(\ell, X) \star O(\ell, v, t) \Rightarrow W^1(\ell, X) \star O(\ell, v, t) \star \langle (v, t) \in X \rangle$$

Reasoning about relaxed atomic accesses

$\Phi v = \text{emp}$ (i.e. Φv is pure and true)

$$\left\{ \begin{array}{l} W^\pi(\ell, X) \\ \star \text{Rel}(\ell, \Phi) \\ \star O(\ell, -, t) \end{array} \right\} \text{set}_{rlx}(\ell, v) \left\{ \begin{array}{l} \lambda \dots \exists t' > t, \\ W^\pi(\ell, \{(v, t')\} \cup X) \end{array} \right\}$$

As in RSL, a relaxed write transfers no resources ($\Phi v = \text{emp}$).

The write is reflected in assertion $W^\pi(\ell, X)$, which does not need to be exclusive. ($\pi < 1$ is allowed!)

$O(\ell, -, t)$ proves that ℓ is initialized and gives a lower bound for the new timestamp t' .

Reasoning about relaxed atomic accesses

$$\left\{ \begin{array}{l} \text{Acq}(\ell, \Phi) \\ \star O(\ell, -, t) \end{array} \right\} \text{get}_{rlx}(\ell) \left\{ \begin{array}{l} \lambda v. \exists t' \geq t, \\ \text{Acq}(\ell, \Phi) \star O(\ell, v, t') \star \nabla(\Phi v) \end{array} \right\}$$

A relaxed read of value v gives access to the pure part $\nabla(\Phi v)$ of the resource invariant Φv , and to a new observation $O(\ell, v, t')$.

If we own the full permission on ℓ , the value read is determined by the most recent write.

$$\left\{ \begin{array}{l} \text{Acq}(\ell, \Phi) \\ \star W^1(\ell, X) \end{array} \right\} \text{get}_{rlx}(\ell) \left\{ \begin{array}{l} \lambda v. \exists t, \langle (v, t) = \max(X) \rangle \\ \star \text{Acq}(\ell, \Phi) \star W^1(\ell, X) \star \nabla(\Phi v) \end{array} \right\}$$

Summary

Summary

A realization: relaxed memory models such as those of Java or C/C++ are complicated and not fully understood yet.

There are two ways of constructing a software design. One way is to make it so simple that there are obviously no deficiencies. And the other way is to make it so complicated that there are no obvious deficiencies

(C. A. R. Hoare)

Summary

A realization: relaxed memory models such as those of Java or C/C++ are complicated and not fully understood yet.

A hope: it affects a handful of libraries only;
the bulk of parallel computation codes are still written using conventional synchronization primitives.

Summary

A realization: relaxed memory models such as those of Java or C/C++ are complicated and not fully understood yet.

A hope: it affects a handful of libraries only; the bulk of parallel computation codes are still written using conventional synchronization primitives.

A most necessary tool: program logics!

- To abstract over some of the complexity of the memory model (cf. RSL, SLR).
- To combine reasoning in standard separation logic with reasoning specific to a given memory model.

References

An introduction to weakly-consistent memory models:

- S. V. Adve, H. J. Boehm, *Memory models: a case for rethinking parallel languages and hardware*, Comm. ACM, 2010.

The RSL and SLR program logics:

- V. Vafeiadis, C. Narayan, *Relaxed separation logic: a program logic for C11 concurrency*, OOPSLA 2013.
- K. Svendsen, J. Pichon-Pharabod, M. Doko, O. Lahav, V. Vafeiadis, *A Separation Logic for a Promising Semantics*, ESOP 2018.

The promising semantics for C/C++ 2011:

- J. Kang, C.-K. Hur, O. Lahav, V. Vafeiadis, D. Dreyer, *A promising semantics for relaxed-memory concurrency*, POPL 2017.