

Program logics for relaxed consistency

UPMARC Summer School 2014

Viktor Vafeiadis

Max Planck Institute for Software Systems (MPI-SWS)

2nd Lecture, 29 July 2014

Topics covered yesterday:

- ▶ The C11 memory model
- ▶ Separation logic
- ▶ Relaxed separation logic

Today:

- ▶ Compare and swap
- ▶ GPS
- ▶ Advanced features

Recap: Rules for release/acquire accesses

Ownership transfer by rel-acq synchronizations.

- ▶ Atomic allocation \rightsquigarrow pick loc. invariant Q .

$$\{Q(v)\} x = \text{alloc}(v); \{W_Q(x) * R_Q(x)\}$$

- ▶ Release write \rightsquigarrow give away permissions.

$$\{W_Q(x) * Q(v)\} x.\text{store}(v, \text{rel}); \{W_Q(x)\}$$

- ▶ Acquire read \rightsquigarrow gain permissions.

$$\{R_Q(x)\} t = x.\text{load}(\text{acq}); \{Q(t) * R_{Q[t:=\text{emp}]}(x)\}$$

Basically, disallow ownership transfer.

- ▶ Relaxed reads:

$$\{ \mathbf{R}_Q(x) \} \quad t = x.load(rlx) \quad \left\{ \begin{array}{l} \mathbf{R}_Q(x) \wedge \\ (Q(t) \neq \text{false}) \end{array} \right\}$$

- ▶ Relaxed writes:

$$\frac{Q(v) = \text{emp}}{\{ \mathbf{W}_Q(x) \} \quad x.store(v, rlx) \quad \{ \mathbf{W}_Q(x) \}}$$

Compare and swap (CAS)

A standard primitive for implementing concurrent algorithms

```
x.CAS( $v$ ,  $v'$ ,  $M$ )  $\stackrel{\text{def}}{=}$   
  atomic {  
    if ( $x.\text{load}(M) == v$ ) {  
       $x.\text{store}(v', M)$ ;  
      return true;  
    }  
    return false;  
  }
```

Reasoning about CAS in RSL

- ▶ New assertion form, $P := \dots \mid \mathbf{C}_Q(x)$.
- ▶ “Permission to do a CAS”
- ▶ Duplicable:

$$\mathbf{C}_Q(x) \iff \mathbf{C}_Q(x) * \mathbf{C}_Q(x)$$

- ▶ Also allows writing:

$$\mathbf{C}_Q(x) \iff \mathbf{C}_Q(x) * \mathbf{W}_Q(x)$$

- ▶ And reading without ownership transfer:

$$\mathbf{C}_Q(x) \iff \mathbf{C}_Q(x) * \mathbf{R}_{\text{emp}}(x)$$

Allocation rule:

$$\{Q(v)\} x = \text{alloc}(v); \{C_Q(x)\}$$

CAS rule:

$$\begin{array}{l} t \wedge P * Q(v) \Rightarrow Q(v') * R \\ \neg t \wedge P \Rightarrow R \\ X \in \{rel, rlx\} \Rightarrow Q(v) \equiv \text{emp} \\ X \in \{acq, rlx\} \Rightarrow Q(v') \equiv \text{emp} \\ \hline \{C_Q(x) * P\} t = x.\text{CAS}(v, v', X) \{R\} \end{array}$$

Mutual exclusion locks

Attach a 'resource invariant' at each lock:

$$Lock(x, J) \iff Lock(x, J) * Lock(x, J)$$

Specifications for mutex operations:

$$\{J\} x = new-lock() \{Lock(x, J)\}$$

$$\{Lock(x, J)\} \quad lock(x) \quad \{Lock(x, J) * J\}$$

$$\{J * Lock(x, J)\} \quad unlock(x) \quad \{Lock(x, J)\}$$

Mutual exclusion locks

Let $Q_J(v) \stackrel{\text{def}}{=} (v = 0 \wedge \text{emp}) \vee (v = 1 \wedge J)$
 $Lock(x, J) \stackrel{\text{def}}{=} \mathbf{C}_{Q_J}(x)$

$new_lock() \stackrel{\text{def}}{=}$
 $\{J\}$
 $res = \mathbf{alloc}(1)$
 $\{Lock(res, J)\}$

$unlock(x) \stackrel{\text{def}}{=}$
 $\{J * Lock(x, J)\}$
 $x.store(1, rel)$
 $\{Lock(x, J)\}$


$lock(x) \stackrel{\text{def}}{=}$
 $\{Lock(x, J)\}$
repeat
 $\{Lock(x, J)\}$
 $t = x.CAS(1, 0, acq)$
 $\{Lock(x, J) * (t \Rightarrow J)\}$
until t
 $\{J * Lock(x, J)\}$

GPS: Towards a better logic for C11

- ▶ Protocols
- ▶ Ghosts & escrows

GPS: A better logic for release-acquire

Three key features:

- ▶ Location ~~invariants~~ **protocols**
- ▶ Ghost state/tokens 
- ▶ Escrows for ownership transfer

Example (Racy message passing)

Initially, $x = y = 0$.

$$\begin{array}{l} x.\text{store}(1, \text{rel}); \parallel x.\text{store}(1, \text{rel}); \parallel t = y.\text{load}(\text{acq}); \\ y.\text{store}(1, \text{rel}); \parallel y.\text{store}(1, \text{rel}); \parallel t' = x.\text{load}(\text{acq}); \end{array}$$

Cannot get $t = 1 \wedge t' = 0$.

Racy message passing in GPS

Protocol for x : $\mathbf{A: } x = 0 \longrightarrow \mathbf{B: } x = 1$

Protocol for y : $\mathbf{C: } y = 0 \longrightarrow \mathbf{D: } y = 1 \wedge x.st \geq \mathbf{B}$

Acquire reads gain knowledge, not ownership.

$\{x.st \geq \mathbf{A} \wedge y.st \geq \mathbf{C}\}$		$\{x.st \geq \mathbf{A} \wedge y.st \geq \mathbf{C}\}$
$x.store(1, rel);$		$t = y.load(acq);$
$\{x.st \geq \mathbf{B} \wedge y.st \geq \mathbf{C}\}$		$\left\{ \begin{array}{l} t = 0 \wedge x.st \geq \mathbf{A} \\ \vee t = 1 \wedge x.st \geq \mathbf{B} \end{array} \right\}$
$y.store(1, rel);$		$t' = x.load(acq);$
$\{x.st \geq \mathbf{B} \wedge y.st \geq \mathbf{D}\}$		$\{t = 0 \vee (t = 1 \wedge t' = 1)\}$

Rules for reads and writes

Read rule:

$$\forall s' \geq_{\tau} s. \mathbf{inv}_{\tau}(s', t) * P \Rightarrow Q$$
$$Q \Leftrightarrow Q * Q$$

$$\left\{ \begin{array}{l} x.st \geq_{\tau} s \\ * P \end{array} \right\} t = x.load(acq); \left\{ \begin{array}{l} \exists s'. x.st \geq_{\tau} s' \\ * P * Q \end{array} \right\}$$

Write rule:

$$P \Rightarrow \mathbf{inv}_{\tau}(s'', v) * Q$$
$$\forall s' \geq_{\tau} s. \mathbf{inv}_{\tau}(s', _) * P \Rightarrow s'' \geq_{\tau} s'$$

$$\left\{ x.st \geq_{\tau} s * P \right\} x.store(v, rel); \left\{ x.st \geq_{\tau} s'' * Q \right\}$$

We can create ghost unduplicable tokens:

$$\frac{K \text{ is fresh}}{P \Rightarrow P * K} \quad \frac{}{K * K \Rightarrow \text{false}}$$

We can also create escrows:

$$\frac{P * P \Rightarrow \text{false}}{Q \Rightarrow \mathbf{Esc}(P, Q)} \quad \frac{}{\mathbf{Esc}(P, Q) * P \Rightarrow Q}$$

Escrows are duplicable:

$$\mathbf{Esc}(P, Q) \iff \mathbf{Esc}(P, Q) * \mathbf{Esc}(P, Q)$$

but only one component can ‘unlock’ them.

To gain ownership, we use ghost state & escrows.

$$\frac{P * P \Rightarrow \text{false}}{Q \Rightarrow \mathbf{Esc}(P, Q)} \quad \frac{}{\mathbf{Esc}(P, Q) * P \Rightarrow Q}$$

Example (Message passing using escrows)

Invariant for x : $x = 0 \vee \mathbf{Esc}(K, \&a \mapsto 7)$.

<pre> {&a ↦ 0} a = 7; {&a ↦ 7} {Esc(K, &a ↦ 7)} x.store(1, rel); </pre>	<pre> {K} if (x.load(acq) ≠ 0) {K * Esc(K, &a ↦ 7)} {&a ↦ 7} print(a); </pre>
---	---

With a successful CAS we can gain not only knowledge, but also ownership:

$$\forall s'' \geq_{\tau} s. \mathbf{inv}_{\tau}(s'', v) * P \Rightarrow \mathbf{inv}_{\tau}(s', v') * Q \wedge s' \geq_{\tau} s''$$

$$\forall s'' \geq_{\tau} s. \forall v'' \neq v. \mathbf{inv}_{\tau}(s'', v'') * P \Rightarrow R$$

$$R \Leftrightarrow R * R$$

$$\left\{ \begin{array}{l} x.st \geq_{\tau} s \\ * P \end{array} \right\} t = x.CAS \quad (v, v', rel-acq); \left\{ \begin{array}{l} (t \wedge x.st \geq_{\tau} s' * Q) \\ \vee \neg t \wedge P * R \end{array} \right\}$$


Reasoning about advanced C11 features

(Work in progress)

- ▶ Fences
- ▶ Consume reads


Message passing

```
int a; atomic_int x = 0;  
( a = 7;                               || if (x.load(acq) ≠ 0){  
  x.store(1, rel);                       ||   print(a); } )
```




Incorrect message passing

```
int a; atomic_int x = 0;
( a = 7;                               || if (x.load(rlx) ≠ 0){
  x.store(1, rlx);                       print(a); } )
```



Message passing with C11 memory fences

```
int a; atomic_int x = 0;
( a = 7;                               | if (x.load(rlx) ≠ 0){
  fence(release);                       |   fence(acq);
  x.store(1, rlx);                       |   print(a); }
)
```



Introduce two ‘modalities’ in the logic.

$$\{P\} \text{ fence}(\text{release}) \{\Delta P\}$$

$$\{\nabla P\} \text{ fence}(\text{acq}) \{P\}$$

$$\{\mathbf{R}_Q(x)\} t := x.\text{load}(r/x) \{\mathbf{R}_{Q[t:=\text{emp}]}(x) * \nabla Q(t)\}$$

$$\{\mathbf{W}_Q(x) * \Delta Q(v)\} x.\text{store}(v, r/x) \{\mathbf{W}_Q(x)\}$$

Reasoning about fences

Let $Q(v) \stackrel{\text{def}}{=} v = 0 \vee \&a \mapsto 5$.

$$\left(\begin{array}{l|l} \{ \&a \mapsto 0 * \mathbf{W}_Q(x) * \mathbf{R}_Q(x) \} & \\ \{ \&a \mapsto 0 * \mathbf{W}_Q(x) \} & t = x.\text{load}(r/x); \\ a = 5; & \{ \nabla(t = 0 \vee \&a \mapsto 5) \} \\ \{ \&a \mapsto 5 * \mathbf{W}_Q(x) \} & \text{if } (t \neq 0) \\ \text{fence}(\text{release}); & \text{fence}(\text{acq}); \\ \{ \Delta(\&a \mapsto 5) * \mathbf{W}_Q(x) \} & \{ \&a \mapsto 5 \} \\ x.\text{store}(1, r/x); & \mathbf{print}(a); \\ \{ \text{true} \} & \{ \text{true} \} \end{array} \right)$$

Why two modalities?

Consider the program, where initially $x = y = 0$:

$a = 5;$	$t = x.\text{load}(rlx);$	$t' = y.\text{load}(rlx);$
$\text{fence}(\text{release});$	if ($t \neq 0$)	if ($t' \neq 0$) {
$x.\text{store}(1, rlx);$	$y.\text{store}(1, rlx);$	$\text{fence}(\text{acq});$
		print (a);
		}

If $\nabla P \Rightarrow \Delta P$, we can 'verify' this program.

But the program is racy.


Release-consume synchronization

Initially $a = x = 0$.

```
 $a = 5;$  | |  $t = x.load(consume);$   
 $x.store(release, \&a);$  | | if ( $t \neq 0$ )  $print(*t);$ 
```

This program cannot crash nor print 0.

Justification:



Release-consume synchronization

Initially $a = x = 0$. Let $J(t) \stackrel{\text{def}}{=} t = 0 \vee t \mapsto 5$.

$\{\&a \mapsto 0 * \mathbf{W}_J(x)\}$		$\{\mathbf{R}_J(x)\}$
$a = 5;$		$t = x.load(\text{consume});$
$\{\&a \mapsto 5 * \mathbf{W}_J(x)\}$		$\{\nabla_t(t = 0 \vee t \mapsto 5)\}$
$x.store(\text{release}, \&a);$		if $(t \neq 0)$ $print(*t);$

This program cannot crash nor print 0.

Index the ∇ with program variable t .
 t data dependence \implies locally open ∇_t .

Proposed rules for consume accesses

$$\{\mathbf{R}_Q(x)\} t := x.\text{load}(\text{cons}) \{ \mathbf{R}_{Q[t:=\text{emp}]}(x) * \nabla_t Q(t) \}$$

$$\frac{\begin{array}{c} \{P\} C \{Q\} \\ C \text{ is basic command mentioning } t \end{array}}{\{\nabla_t P\} C \{\nabla_t Q\}}$$

Question: Is the following valid?

$$\{\mathbf{W}_Q(x) * \nabla_t Q(v)\} x.\text{store}(v, \text{rel}); \{\mathbf{W}_Q(x)\}$$

Release-acquire too weak in the presence of consume

Initially $x = y = 0$.

```
a = 1;
x.store(1, release);
while (x.read(consume) ≠ 1);
y.store(1, release);
(* while (y.load(acquire) ≠ 1);
(* a = 2;
```

C11 deems this program racy.

- ▶ Only different thread rel-acq synchronize.

What goes wrong in PL:

On ownership transfers, we must prove that we don't read from the same thread.

Release-acquire too weak in the presence of consume

Initially $x = y = 0$.

```
a = 1;
x.store(1, release);
while (x.read(consume) ≠ 1);
y.store(1, release);
(* while (y.load(acquire) ≠ 1);
(* a = 2;
```

C11 deems this program racy. But, it is not racy:

- ▶ On x86-TSO, Power, ARM, and Itanium.
- ▶ Or if we move the (*) lines to a new thread.

So, drop the “different thread” restriction.

We know how to reason about:

- ▶ Release-acquire
- ▶ Consume reads
- ▶ C11 memory fences

We found a number of bugs in the model:


- ▶ Dependency cycles (also in [Batty et al. '03])
- ▶ Same thread rel-acq don't synchronize
- ▶ Semantics of SC accesses odd and too weak...
...when mixed with non-SC accesses
- ▶ Release sequences too strong

Soundness proof challenges

- ▶ Assertions in heaps
 - ⇒ Store syntactic assertions (modulo *-ACI)
- ▶ No (global) notions of state and time
 - ⇒ Define a *logical* local notion of state
 - ⇒ Annotate hb edges with logical state
- ▶ No operational semantics
 - ⇒ Use the axiomatic semantics
 - ⇒ Induct over max hb-path distance from top



- ▶ Annotate hb edges of executions with heaps.



- ▶ Local annot. validity: $\sum ins + \text{node-effect} = \sum outs$.
- ▶ Configuration safety: can extend a valid annotation for n further events.

Definition (Pairwise independence)

\mathcal{T} is pairwise independent iff $\forall (a, a'), (b, b') \in \mathcal{T}$,
 $(a', b) \notin hb^*$.

Lemma (Independent heap compatibility)

If $hmap$ is a valid annotation, and $\mathcal{T} \subseteq hb$ is pairwise independent, then $\bigoplus_{x \in \mathcal{T}} hmap(x)$ is defined.

Formal reasoning about weak memory
is possible & not too difficult.

We're not quite there yet; there's still a lot to do:

Liveness, refinement, tool support, ...

Relaxed program logics
are a useful tool for
understanding
weak memory models