

Compare and Swap (CAS)

(Also known as compare-exchange, CMPXCHG on x86)

- CAS is an atomic instruction that performs a read and a write if the read gets a certain value.

```
x.CAS (old, new)  $\stackrel{\text{def}}{=}$ 
    atomic {
        if (x.load() == old) {
            x.store(new);
            return true;
        } else {
            return false;
        }
    }
```

- It is widely used for synchronisation, e.g. implementing locks:

```
lock()  $\stackrel{\text{def}}{=}$  while (!x.CAS(0, 1));
unlock()  $\stackrel{\text{def}}{=}$  x.store(0);
```

or an atomic update:

```
< x := f(x) >  $\rightsquigarrow$  do {
    t = x.load();
    n = f(t);
} while (!x.CAS(t, n));
```

- In C11, it also gets an access-type argument:

RLX, REL, ACQ, REL-ACQ, SC


These intuitively mean that the load & the store of the CAS get the appropriate annotation:

load: RLX, REL, ACQ, REL-ACQ, SC

store: RLX, REL, ACQ, REL-ACQ, SC.

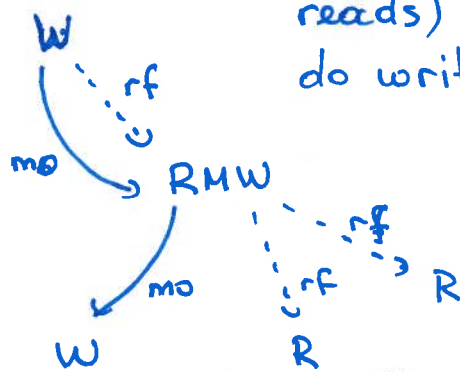
Consistency Axioms for CAS

- In C11 executions, successful CAS operations (i.e., ones where the store happens) are modeled by a new kind of event:



Failed CAS operations (i.e., where the store does not happen) are modeled as normal read events.

- In the C11 axioms, RMW events are treated both as reads and writes. They have incoming rf edges (as do reads) and outgoing rf edges (as do writes). Moreover, they are included in the memory order ("coherence order"), mo .



- There is one additional axiom asserting that RMWs are atomic:

$$\forall a, b. \quad rf(b) = a \wedge isRMW(b) \Rightarrow mo(a, b) \wedge \nexists c. mo(a, c) \wedge mo(c, b).$$

That is, RMWs read from their immediate mo -predecessor. In essence, this axiom rules out the following pictures



Proof Rules for CAS

- We introduce a new permission for performing CAS:

$$P ::= \dots \mid C_J(x)$$

- This permission is duplicable:

$$C_J(x) \Leftrightarrow C_J(x) * C_J(x)$$

& generated at allocation:

$$\{ J(v) \} \quad x = \text{alloc}(v); \quad \{ \underbrace{C_J(x) * W_J(x)}_{\text{-----}} \}$$

Note: No $R_J(x)$!!!

At allocation, choose whether it's a location used for reads or CAS.

- It allows doing a read without gaining ownership (cf. the relaxed-read rule):

$$\{ C_J(x) \} \quad t = x.\text{load}(typ) \quad \{ C_J(x) \wedge \text{"J(t) satisfiable"} \}$$

- And importantly, it allows one to do a CAS:

$$\begin{array}{l} t \wedge P * J(v) \Rightarrow Q * J(v') \quad \leftarrow \text{successful case} \\ \neg t \wedge P \Rightarrow Q \quad \leftarrow \text{unsuccessful CAS} \\ typ \in \{rel, rlx\} \Rightarrow J(v) = emp \quad \leftarrow \text{not ACQ} \\ typ \in \{acq, rlx\} \Rightarrow J(v') = emp \quad \leftarrow \text{not REL} \\ \hline \{ C_J(x) * P \} \quad t = x.\text{CAS}_{typ}(v, v') \quad \{ C_J(x) * Q \} \end{array}$$

- A successful CAS gains the ownership of $J(v)$ and loses ownership of $J(v')$.
- A failed CAS does no ownership transfer.
- If the CAS is not at least of ACQ kind, it cannot acquire any ownership as it doesn't synchronize (so $J(v) = emp$)
- Symmetrically for release. - 3 -

Verifying the CAS-based lock

Specs:

There exists a lock predicate $\text{Lock}(l, P)$ such that:

$$\text{Lock}(l, P) \Leftrightarrow \text{Lock}(l, P) * \text{Lock}(l, P)$$

$$\{ P \} \quad l = \text{newlock}() \quad \{ \text{Lock}(l, P) \}$$

$$\{ \text{Lock}(l, P) \} \quad \text{lock}(l) \quad \{ \text{Lock}(l, P) * P \}$$

$$\{ \text{Lock}(l, P) * P \} \quad \text{unlock}(l) \quad \{ \text{Lock}(l, P) \}$$

Implementation / proof:

$$\text{Let } J(v) \stackrel{\text{def}}{=} (v = 0 \wedge P) \vee (v = 1 \wedge \text{emp})$$

$$\text{Let } \text{Lock}(l, P) \stackrel{\text{def}}{=} C_J(l) * W_J(l)$$

- Duplicability of Lock follows from that of $C \& W$.

- newlock:

$$\{ P \} \quad l = \text{alloc}(0) \quad \{ C_J(l) * W_J(l) \}$$

by the allocation rule

- unlock:

$$\{ \text{Lock}(l, P) * P \}$$

$$\{ W_J(l) * J(0) * C_J(l) \}$$

$l.\text{store}(0, \text{rel});$

$$\{ W_J(l) * C_J(l) \}$$

$$\{ \text{Lock}(l, P) \}$$

W-REL & FRAME rules

Verifying the CAS-based lock (ctd.)

- lock:

{ Lock(l, P) }

do

{ Lock(l, P) }

t = l.CAS_{acq}(0, 1);

{ Lock(l, P) * (t ∧ P ∨ ¬t ∧ emp) }

FRAME away
W_J(l)
/ & CAS rule

while (¬t);

{ Lock(l, P) * P }

Conditions of the CAS rule:

- $t \wedge J(1) \Rightarrow (t \wedge P \vee \neg t \wedge emp) * J(0)$
- $\neg t \wedge emp \Rightarrow (t \wedge P \vee \neg t \wedge emp)$
- $J(0) = emp$.

Using this implementation of locks, we can now implement CSL atomic blocks as

atomic(c) \rightsquigarrow lock(l); c; unlock(l)

where l is a global lock, allocated at the start of the program. Moreover, we can encode (& derive) the CSL rules as follows:

$J \vdash_{CSL} \{P\} c \{Q\} \rightsquigarrow$

$\vdash_{RSL} \{Lock(l, J) * P\} c \{Lock(l, J) * Q\}$