

# Data races & Data Race Freedom (DRF)

- Two operations are conflicting if they both access the same location and at least one is a write
- Two operations are concurrent if they are not ordered by happens-before.
- A data race (or simply, a race) occurs whenever we have two concurrent conflicting operations, such that at least one is non-atomic.

$$\text{race}(a, b) \stackrel{\text{def}}{=} a \neq b \wedge \neg \text{hb}(a, b) \wedge \neg \text{hb}(b, a) \\ \wedge (\text{isWrite}(a) \vee \text{isWrite}(b)) \\ \wedge (\text{isNA}(a) \vee \text{isNA}(b)) \\ \wedge (\exists l. \text{loc}(a) = \text{loc}(b) = l)$$

- A program is DRF (data race free) if every consistent execution of the program has no races.
- The semantics of C11 programs is defined to be:
  - either the set of all its consistent executions if they are DRF,
  - or "undefined" if the program has a consistent execution with a data race.
- This style of defining weak memory models is also known as the "DRF" style.
- Motivation for giving undefined semantics to racy programs comes from optimising compilers.  
A program transformation  $S \rightarrow T$  is deemed correct if  $\text{Behaviours}(T) \subseteq \text{Behaviours}(S)$ .

# Examples:

## Sequential:

(1) `int x;`  $\rightsquigarrow$  `print(0)`  
`print(x)`

(can print arbitrary value)

(2) `*NULL = 3`  $\rightsquigarrow$  skip [MSVC does this]  
(undefined/crash)

## SC

(3) `x = 1 || y = 2`  $\rightsquigarrow$  `x = 1; y = 2`

(4)

Common  
Subexpression  
elimination  
(CSE)

`x = 0; p = &x;`  
`a = x;`  
`b = *p;`  
`c = x;`  
`print(a, b, c);`  $\parallel$  `x = 1`

cannot print 0, 1, 0  
(Coherence)

`x = 0; p = &x`  
`a = x;`  
`b = *p;`  
`c = a;`  
`print(a, b, c);`  $\parallel$  `x = 1`

can print 0, 1, 0.

## C11

(4) The LHS is undefined. The compiler can do anything it likes!

$\rightarrow$  But race freedom is a global property, and makes it impossible to reason locally against a malicious client. Consider the program:

`secret = 7 || Client`

where the Client doesn't know the secret (and does not mention the "secret" variable). How can we ensure that it never finds it out?

[We cannot. If the client has a race, the compiler is allowed to leak the secret.]

## The DRF theorem

- ⇒ If under SC, a program has no races, then:
- $$\text{Behaviours}_{\text{weakMem}}(\text{Program}) = \text{Behaviours}_{\text{SC}}(\text{Program})$$
- ⇒ Considered as a sanity check for memory models.
- ⇒ Allows reasoning under SC for DRF programs both for (a) showing that the program is DRF, and (b) proving some property of the reachable program states.
- 

Q: What's a race under SC?

A: Two concurrent, conflicting memory accesses where at least one is non-atomic. (as before)

w.r.t. the happens-before order induced by the program order & synchronization (i.e. whenever we have a  $\xrightarrow{rf}$  edge)

again, we assume some operations are marked as "atomic"/"volatile" and are used to synchronize between threads.


the others are "non-atomic" and concurrent accesses to them result in a race.

Q: What are  $\text{Behaviours}_{\text{weakMem}}$  /  $\text{Behaviours}_{\text{SC}}$ ?

A: The set of consistent executions of the program according to the appropriate memory model. (And in case of a DRF model such as C11,  $\text{Behaviours}_{\text{DRFmodel}}$  is the universe set if the program has a consistent racy execution.)

# Does the DRF theorem hold for C11?

→ Not as stated. Example:




→ OK. So how about the following:

"If under SC the program  $P$  has no concurrent, conflicting accesses such that at least one of them is a non-SC access, then

$$\text{Behaviours}_{C11}(P) = \text{Behaviours}_{SC}(P). "$$

→ Again, doesn't hold because of the dependency cycles:



→ What if we remove dependency cycles either by

(a) Strengthening the model by adding the

axiom:  $\text{acyclic}(hb_{USC})$

or (b) assume that there are no relaxed accesses in the program? (in this case  $rf \subseteq hb$ .)

→ Claim: the theorem now holds.

[Batty et al. POPL'12] prove a slightly simplified version with no REL/ACQ.

For simplicity, we will deal with case (b)

To show:

( $\supseteq$ ) easy  $\otimes$   $\text{Consistent}_{sc}(A, lab, po, rf, sc) \Rightarrow$

$\exists mo', sc'. \text{Consistent}_{c11}(A, lab, po, rf, mo', sc')$

[Pick  $mo' := \{ (a,b) \mid sc(a,b) \wedge \exists l. isWrite_l(a) \wedge isWrite_l(b) \}$   
 $sc' := \{ (a,b) \mid sc(a,b) \wedge isSC(a) \wedge isSC(b) \}$ ]

$\otimes$   $\text{DRF}_{sc}(A, lab, po, rf, sc) \Rightarrow$

$\text{DRF}_{c11}(A, lab, po, rf, mo', sc')$

[the happens-before relation has not changed.]

( $\subseteq$ )

Lemma:  $\text{DRF}_{c11}(A, lab, po, rf, mo) \wedge \text{Consistent}_{c11}(A, lab, po, rf, mo, sc)$

[first attempt]  $\Rightarrow \exists sc'. \text{Consistent}_{sc}(A, lab, po, rf, sc')$

Recall  $\text{DRF}_{c11}$  in this context: "all conflicting concurrent accesses must be SC-accesses"

in  $(mo \cup rf \cup fr)^+$  not in  $hb \cup hb^{-1} \cup (=)$

So:  $\left( \begin{array}{l} \text{Non-SC} \times \text{Non-SC} \\ \cup \text{SC} \times \text{Non-SC} \\ \cup \text{Non-SC} \times \text{SC} \end{array} \right) \cap (mo \cup rf \cup fr)^+ \subseteq hb \cup hb^{-1} \cup (=)$

SC-accesses are ordered by the sc order, and we know that  $hb|_{sc} \subseteq sc$ . (Consistent SCHb)

ruled out by coherence.

Therefore, acyclic  $(hb \cup mo \cup rf \cup fr)$ .

Pick  $sc'$  to be any total irreflexive order extending  $(hb \cup mo \cup rf \cup fr)$ .

Oops, we proved a weaker result, i.e. that race-free c11 executions correspond to SC ones. What about racy c11 executions?

- Well, if acyclic  $(hb \cup \underbrace{mo \cup rf \cup fr}_{com})$ , then we can find an SC execution.



- So what if there is a  $hb \cup com$  cycle?

# Dealing with (hb ucom) cycles

Observe that:



Given a program  $P$  and a consistent execution  $X$  of  $P$ , we can remove any  $(hb_{usc} \cup rf)$ -maximal event from  $X$  and get another consistent execution of  $P$ .

- Now consider an execution with a (minimal)  $(hb \cup com)$  cycle.
- Remove a maximal  $(hb_{usc} \cup rf)$  event. Either the cycle remains, or it vanishes. In the former case, we continue removing events. (As the graph is finite, we will eventually reach the second case.)
- In the latter case, there were two possibilities:



(The other two  $\searrow$   $hb$  and  $\searrow$   $rf$  are not  $(hb_{usc} \cup rf)$ -maximal.)

- In both cases, we can construct a cycle-free execution that preserves the memory accesses of the execution with the cycle.



- Since the execution is  $(hb_{usc} \cup rf)$ -acyclic, we can pick  $sc'$  to be any strict total order extension of  $(hb_{usc} \cup rf)$ , thereby ensuring that the execution is also consistent under  $sc'$ .
- Finally, since we haven't changed the events of the cycle, nor the  $hb$  relation (as in subcase (a), the read is a non-SC access), the race persists in the consistent SC execution, contradicting our assumption that all consistent SC executions of  $P$  are race-free. — 6 —