The DRF theorem

Ori Lahav   Viktor Vafeiadis

31 August 2017

Motivation for the DRF theorem

WMC is complicated:

- Most programmers “do not understand” WMC.
- Leads to subtle bugs $\rightsquigarrow$ hard to debug and fix.

Define programming disciplines that:

- Avoid weak behaviors.
- Can be understood without referring to the WMM.

The DRF discipline:

- Do not have any data races.
- Just use locks for synchronization.

DRF: A minimal requirement for memory models

Definition (DRF property)

A memory model X satisfies the *DRF property* if
for every program that is race-free under SC semantics,
its allowed outcomes under X are the same as under SC.

- A programming discipline to avoid weak behavior.
- The premise requires us to establish race-freedom *under* SC.
- So a defensive programmer does not need to understand WMM.

For specific memory models, one can establish more permissive programming disciplines that ensure the absence of weak behaviors.

What models satisfy the DRF property?

Among the models we saw so far, which satisfy the DRF property?

- COH
- StrongCOH
- RA
- C11
- TSO

The DRF property can be also taken as a definition of a "catch fire" crude model:

- If the program is race-free under SC, then the allowed outcomes are the same as under SC
- Otherwise, "undefined behavior" (*i.e.*, any outcome is allowed!)

What constitutes a race under SC? (operationally)

Definition (racy program under SC (operationally))

P is called *racy under* SC if there exist P', S', M' such that the following hold:

- $P, S_0, M_0 \Rightarrow^* P', S', M'$
- $P', S' \xrightarrow{i_1 : l_1} \_$ and $P', S' \xrightarrow{i_2 : l_2} \_$ for some $i_1 \neq i_2$, and labels l_1 and l_2, such that $\texttt{loc}(l_1) = \texttt{loc}(l_2)$, and $\{\texttt{typ}(l_1), \texttt{typ}(l_2)\} \cap \{\texttt{W}, \texttt{RMW}\} \neq \emptyset$

What constitutes a race under SC? (declaratively)

Definition (race)

Given an execution graph G and a relation $R \subseteq G.\mathtt{E} \times G.\mathtt{E}$, we say that two events a, b *R-race* in G if the following hold:

- $a \neq b$
- $\mathtt{loc}(a) = \mathtt{loc}(b)$
- $\{\mathtt{typ}(a), \mathtt{typ}(b)\} \cap \{\mathtt{W}, \mathtt{RMW}\} \neq \emptyset$
- $\langle a, b\rangle \notin R^+$ and $\langle b, a\rangle \notin R^+$

Definition (racy execution)

An execution graph G is called *R-racy* if there are two events that R-race in G.

Definition (racy program under SC (declaratively))

P is called *racy under* SC if there exists an execution graph G such that the following hold:

- G is a $(\mathrm{po} \cup \mathtt{rf})^+$-prefix of an execution of P
- G is SC-consistent
- G is $(\mathrm{po} \cup \mathtt{rf})$-racy

What constitutes a race under SC?

- The two definitions differ for programs with RMW's:

$$
\begin{array}{l} x := 1; \\ a := \textbf{FAI}(y) \quad /\!\!/ 0 \end{array}
\;\Bigg\|\;
\begin{array}{l} b := \textbf{FAI}(y); \quad /\!\!/ 1 \\ \textbf{if } b \textbf{ then} \\ \quad c := x \quad /\!\!/ 0 \end{array}
$$

  (**FAI**(y) is an atomic fetch-and-increment)

  - Operational definition: the program is racy under SC
  - Declarative definition: the program is not racy under SC

- Declaratively racy under SC $\Rightarrow$ operationally racy under SC

- For programs without RMW's, the definitions coincide.

- Next, for simplicity, we assume the declarative definition (and restrict RMW's when needed).

Among the models we saw so far, which satisfy the DRF property?

- ✗ COH: The out-of-thin-air (OOTA) problem:

$a := x;$ *// 1*	$b := y;$ *// 1*
if a **then**	**if** b **then**
$\quad y := 1$	$\quad x := 1$

- ✓ StrongCOH
- ✓ RA
- ✗ C11: same reason as for COH (using **rlx** accesses)
- ✓ RC11 (C11 with $(\mathtt{po} \cup \mathtt{rf})$ acyclicity)
- ✓ TSO

Proving DRF for RA

To prove that RA satisfies the DRF property, we have:

1. The easy part of the proof:

Lemma

If an RA*-consistent execution graph* G *contains no* $(\texttt{po} \cup \texttt{rf})$*-races, then it is also* SC*-consistent.*

2. The more difficult part:

Lemma

If P *has an* RA*-consistent* $(\texttt{po} \cup \texttt{rf})$*-racy execution graph, then* P *is racy under* SC*.*

We prove the latter by considering the "first" race of the execution.

Proof outline (1)

- Let G be a the an RA-consistent $(\mathrm{po} \cup \mathtt{rf})$-racy execution graph of P.
- Let G' be a minimal $(\mathrm{po} \cup \mathtt{rf})$-prefix of G that is $(\mathrm{po} \cup \mathtt{rf})$-racy.
- NB: This prefix might not be unique (*e.g.*, SB).
- Let a, b be two events that $(\mathrm{po} \cup \mathtt{rf})$-race in G'.
- Let $x = \mathtt{loc}(a) = \mathtt{loc}(b)$.
- G' is RA-consistent. (why?)
- Let $G'' \triangleq G' \setminus \{a, b\}$.
  - G'' is RA-consistent.
  - G'' is not $(\mathrm{po} \cup \mathtt{rf})$-racy.

  Therefore, G'' is SC-consistent.

Proof outline (2)

- Possible cases:
  - $\texttt{typ}(a) = \texttt{W}$ and $\texttt{typ}(b) = \texttt{W}$
  - $\texttt{typ}(a) \in \{\texttt{R}, \texttt{RMW}\}$ and $\texttt{typ}(b) = \texttt{W}$
  - $\texttt{typ}(a) = \texttt{W}$ and $\texttt{typ}(b) \in \{\texttt{R}, \texttt{RMW}\}$ (symmetric)
  - $\texttt{typ}(a) = \texttt{R}$ and $\texttt{typ}(b) = \texttt{RMW}$
  - $\texttt{typ}(a) = \texttt{RMW}$ and $\texttt{typ}(b) = \texttt{R}$ (symmetric)

- We cannot have $\texttt{typ}(a) = \texttt{RMW}$ and $\texttt{typ}(b) = \texttt{RMW}$. (why?)

Proof outline (3)

CASE 1: $\texttt{typ}(a) = \texttt{W}$ and $\texttt{typ}(b) = \texttt{W}$

- G' is SC-consistent.
  (Take an sc-order for G'' and add a and b at the end)

Proof outline (4)

CASE 2: $\mathtt{typ}(a) \in \{\mathtt{R}, \mathtt{RMW}\}$ and $\mathtt{typ}(b) = \mathtt{W}$

- There exists $a' \sim a$ (a and a' are identical except for the read value, and a' may be a read if a is an RMW) such that some $G_a \in \mathsf{Add}(G'', a')$ is SC-consistent.
  (read from the last write to x in the sc-order for G'')
- Let $G_{ab} \in \mathsf{Add}(G_a, b)$.
- G_{ab} is SC-consistent and $(\mathtt{po} \cup \mathtt{rf})$-racy.

CASE 3: $\mathtt{typ}(a) = \mathtt{R}$ and $\mathtt{typ}(b) = \mathtt{RMW}$

- Let $G_b \triangleq G' \setminus \{a\}$.
- G_b is SC-consistent. (why?)
- b is the $(\mathrm{po} \cup \mathtt{rf})^+$-maximal write to x in G_b.
- There exists $a' \sim a$ (a and a' are identical except for the read value) such that some $G_{ba} \in \mathsf{Add}(G_b, a')$ is SC-consistent and $\langle b, a' \rangle \notin G_{ba}.\mathtt{rf}$.
  (read from the $(\mathrm{po} \cup \mathtt{rf})^+$-maximal write to x in G'')
- G_{ba} is $(\mathrm{po} \cup \mathtt{rf})$-racy.

What properties did we use?

- $(\mathtt{po} \cup \mathtt{rf})$-acyclicity
- RA-consistency is $(\mathtt{po} \cup \mathtt{rf})$-prefix closed
- Receptiveness (changing the value of a final read)

Can we actually write useful programs that are not racy under SC?

- Not really...

lock(l) :	**unlock**(l) :
$r := 0$	$l := 0$
while $\neg r$ **do** $r :=$ **CAS**$(l, 0, 1)$	

- Formally, a lock induces races between the failed lock acquisition attempts and the RMW's/writes to the lock location.
- However, it suffices to consider only executions of the program in which lock acquisitions never fail (why?).
- All successful lock acquisitions and lock releases are totally ordered by $(\mathtt{po} \cup \mathtt{rf})^+$.
- In some models (*e.g.*, full C11), locks are also primitives.

Other theorems

- Triangular race freedom for TSO. (Owens, ECOOP 2010)
- SC fences between every two racy accesses.

Suppose that we change the definition of an R-race and require also that $\mathtt{R} \in \{\mathtt{typ}(a), \mathtt{typ}(b)\}$ (that is, R-concurrent writes are not considered racy).

- Does RA satisfy the corresponding DRF-property?
- Does TSO satisfy the corresponding DRF-property?

Exercise: DRF property for RC11 via DRF-RA

Let RC11 be the simplified C11 model strengthened with $(\mathtt{po} \cup \mathtt{rf})$ acyclicity.

Let P be a program without RMW's. Suppose that in every RA-consistent execution graph, which is a $(\mathtt{po} \cup \mathtt{rf})^{+}$-prefix of an execution graph of P, there are no two events a, b that $(\mathtt{po} \cup \mathtt{rf})$-race and satisfy $\mathbf{rlx} \in \{\mathtt{mod}(a), \mathtt{mod}(b)\}$.

- Show that the outcomes of P under RC11 are the same as under RA.
- Conclude that RC11 satisfies the DRF-property.
- What happens if P contains RMW's?