


Formalizing the Concurrency Semantics of an LLVM Fragment

Soham Chakraborty, Viktor Vafeiadis


Max Planck Institute for Software Systems (MPI-SWS)

CGO 2017


LLVM Compilation



LLVM Concurrency Compilation




LLVM Concurrency Compilation




formalized


formalized



Correctness of the transformations is unclear



Valid opt is removed by over-restriction in bug fix



formalized

informal

formalized

Formalized fragment of LLVM concurrency
(except monotonic/relaxed accesses and fences)

Proved correctness of transformations

Informal text in *Language Reference Manual*

Frequent references to C11 concurrency

- *"This model is inspired by the C++0x memory model."*
- *"These semantics are borrowed from Java and C++0x, but are somewhat more colloquial."*
- *This is intended to match shared variables in C/C++ ..."*
- ...

Why not adopt C11 concurrency?

Subtle differences

- A program has write-read race on non-atomics
 - C11: the behavior of the program is *undefined*
 - LLVM: *defined* behavior;
 - * racy read returns **undef(u)**

```
X = 1; | if(X)
        |   t = 4;
        | else
        |   t = 4;
```

$t \neq 4$? : C11 ✓ LLVM ✗


- Set of allowed optimizations are different

Formalization by Event Structure

- Program

```
int X = 0, Y = 0;  
a = X; || b = Y;  
Y = 1; || X = 1;
```

- Event Structure



Event Structure Construction

```
int X = 0, Y = 0;
```

```
a = X; || b = Y;
```

```
Y = 1; || X = 1;
```

WX0

↓ program-order


WY0

Event Structure Construction

```
int X = 0, Y = 0;
```

```
a = X; || b = Y;
```

```
Y = 1; || X = 1;
```




Event Structure Construction

```
int X = 0, Y = 0;
```


```
a = X; || b = Y;
```

```
Y = 1; || X = 1;
```




Event Structure Construction

```
int X = 0, Y = 0;  
a = X; || b = Y;  
Y = 1; || X = 1;
```



Event Structure Construction

```
int X = 0, Y = 0;  
a = X; || b = Y;  
Y = 1; || X = 1;
```




Event Structure Construction

```
int X = 0, Y = 0;
```

```
a = X; || b = Y;
```

```
Y = 1; || X = 1;
```




Event Structure Construction

```
int X = 0, Y = 0;
```

```
a = X; || b = Y;
```

```
Y = 1; || X = 1;
```




Event Structure Construction

```
int X = 0, Y = 0;
```

```
a = X; || b = Y;
```

```
Y = 1; || X = 1;
```



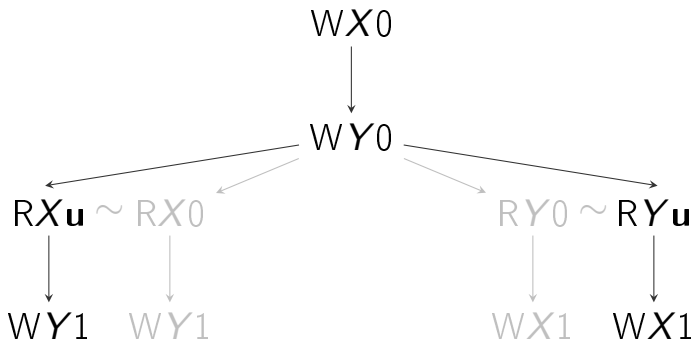
```
int X = 0, Y = 0;  
a = X; || b = Y;  
Y = 1; || X = 1;  
a = b = 1 ? ✓
```

```
int X = 0, Y = 0;  
a = X; || b = Y;  
Y = 1; || X = 1;  
a = b = 1 ? ✓
```

```
int X = 0, Y = 0;  
( a = X; || b = Y; ) ~> int X = 0, Y = 0;  
  Y = 1; || X = 1; )  
a = X; || b = Y;
```


Execution from Event Structure

```
int X = 0, Y = 0;  
a = X; || b = Y;  
Y = 1; || X = 1;
```




Execution from Event Structure

```
int X = 0, Y = 0;  
a = X; || b = Y;  
Y = 1; || X = 1;
```



Execution from Event Structure


```
int X = 0, Y = 0;  
a = X; || b = Y;  
Y = 1; || X = 1;
```



$a = u = 1, b = u = 1$

- Proposed formalization handles
 - Memory operations: load, store, CAS
 - Memory orders: non-atomic, acquire, release, acquire_release, sequentially consistent (SC)
- Preserves *consistency* at each construction step
- Multiple consistent event structures per program


Transformation Correctness



$$\text{Behavior}(P_{tgt}) \subseteq \text{Behavior}(P_{src})$$

Behavior. final values observed in each location

Transformation Correctness



$$\text{Behavior}(P_{tgt}) \subseteq \text{Behavior}(P_{src})$$

Behavior. final values observed in each location



$$\text{Behavior}(G_{tgt}) \subseteq \text{Behavior}(G_{src})$$

LLVM performs these eliminations

Adjacent read after read/write elimination

- $a = X_o; b = X_{na}; \rightsquigarrow a = X_o; b = a;$
- $X_o = v; b = X_{na}; \rightsquigarrow X_o = v; b = v;$

Adjacent overwritten write elimination

- $X_{na} = v'; X_{na} = v; \rightsquigarrow X_{na} = v;$

Non-adjacent overwritten write elimination

- $X_{na} = v'; C; X_{na} = v; \rightsquigarrow C; X_{na} = v;$
where $\text{rel-acq-pair} \notin C$

LLVM does NOT perform these eliminations

Adjacent read after read/write elimination

- $a = X_{\text{acq}}; b = X_{\text{acq}}; \rightsquigarrow a = X_{\text{acq}}; b = a;$
- $a = X_{\text{sc}}; b = X_{(\text{acq}|\text{sc})}; \rightsquigarrow a = X_{\text{sc}}; b = a;$
- $X_{\text{rel}} = v; b = X_{\text{acq}}; \rightsquigarrow X_{\text{rel}} = v; b = v;$
- $X_{\text{sc}} = v; b = X_{(\text{acq}|\text{sc})}; \rightsquigarrow X_{\text{sc}} = v; b = v;$

Adjacent overwritten write elimination

- $X_{\text{rel}} = v'; X_{\text{rel}} = v; \rightsquigarrow X_{\text{rel}} = v;$
- $X_{(\text{rel}|\text{sc})} = v'; X_{\text{sc}} = v; \rightsquigarrow X_{\text{sc}} = v;$

LLVM does NOT perform these eliminations

Adjacent read after read/write elimination

- $a = X_{acq}; b = X_{acq}; \rightsquigarrow a = X_{acq}; b = a;$
- $a = X_{sc}; b = X_{(acq|sc)}; \rightsquigarrow a = X_{sc}; b = a;$
- $X_{rel} = v; b = X_{acq}; \rightsquigarrow X_{rel} = v; b = v;$
- $X_{sc} = v; b = X_{(acq|sc)}; \rightsquigarrow X_{sc} = v; b = v;$

Adjacent overwritten write elimination

- $X_{rel} = v'; X_{rel} = v; \rightsquigarrow X_{rel} = v;$
- $X_{(rel|sc)} = v'; X_{sc} = v; \rightsquigarrow X_{sc} = v;$

Non-adjacent read after write elimination

- $X_{na} = v; C; a = X_{na}; \rightsquigarrow X_{na} = v; C; a = v;$
where rel-acq-pair $\notin C$

LLVM performs(✓) these reorderings

 $a; b \rightsquigarrow b; a$

$\downarrow a \setminus b \rightarrow$	$(\text{St} \text{Ld})_{na}$	St_{rel}	Ld_{acq}	Ld_{sc}	$U_{(\text{acq_rel} \text{sc})}$
$(\text{St} \text{Ld})_{na}$	✓	-	✓	✓	-
St_{rel}	✓	-	-	-	-
St_{sc}	✓	-	-	-	-
Ld_{acq}	-	-	-	-	-
$U_{(\text{acq_rel} \text{sc})}$	-	-	-	-	-

$$X_{rel} = v; Y_{na} = v'; \rightsquigarrow Y_{na} = v'; X_{rel} = v; \quad \checkmark$$

LLVM restricts(×) these reorderings

 $a; b \rightsquigarrow b; a$

$\downarrow a \setminus b \rightarrow$	$(\text{St} \text{Ld})_{\text{na}}$	St_{rel}	Ld_{acq}	Ld_{sc}	$\text{U}_{(\text{acq_rel} \text{sc})}$
$(\text{St} \text{Ld})_{\text{na}}$	✓	×	✓	✓	×
St_{rel}	✓	×	-	-	×
St_{sc}	✓	×	-	×	×
Ld_{acq}	×	×	×	×	×
$\text{U}_{(\text{acq_rel} \text{sc})}$	×	×	×	×	×

$$Y_{\text{na}} = v'; X_{\text{rel}} = v; \rightsquigarrow X_{\text{rel}} = v; Y_{\text{na}} = v'; \quad \times$$


LLVM does NOT perform these reorderings

 $a; b \rightsquigarrow b; a$

$\downarrow a \setminus b \rightarrow$	$(St Ld)_{na}$	St_{rel}	Ld_{acq}	Ld_{sc}	$U_{(acq_rel sc)}$
$(St Ld)_{na}$	✓	✗	✓	✓	✗
St_{rel}	✓	✗	✓	✓	✗
St_{sc}	✓	✗	✓	✗	✗
Ld_{acq}	✗	✗	✗	✗	✗
$U_{(acq_rel sc)}$	✗	✗	✗	✗	✗


$$X_{rel} = v; t = Y_{acq}; \rightsquigarrow t = Y_{acq}; X_{rel} = v; \quad \checkmark$$

C11 to LLVM Mapping Correctness



- LLVM has operations (Ld/St/CAS) and memory orders (na/rel/acq/acq_rel/SC) similar to C11.
- LLVM model is stronger than C11.

LLVM to Architecture Mapping Correctness



$(\text{LLVM} \rightsquigarrow \text{x86/Power}) = (\text{C11} \rightsquigarrow \text{x86/Power})$

Proved correctness of these mappings

- LLVM to SC
- LLVM to SPower

Ensure correctness of $\text{LLVM} \rightsquigarrow \text{x86/Power}$
(results from Lahav & Vafeiadis. FM'16)

Event structure construction rules

$$\frac{e \in V \quad \forall e'' \in V. e'.id \neq e''.id \quad e.code \xrightarrow{e'.lab} e'.code \quad \forall e''. po(e, e'') \implies e'.lab \neq e''.lab}{\langle V, po, rf \rangle \xrightarrow{e'} \langle V \cup \{e'\}, (po \cup \{(e, e')\})^+, rf \rangle} \text{ (BASIC)}$$

$$\frac{WWrace(G)}{G \rightsquigarrow G'} \text{ (WW-RACE)} \quad \frac{G \xrightarrow{e'} G' \quad e' \notin \mathcal{R}}{G \rightsquigarrow G'} \text{ (NON-READ)}$$

$$\frac{G \xrightarrow{e'} G' \quad e' \in \mathcal{R} \quad e'.rval = \mathbf{u} \quad \neg G'.hbW(e')}{G \rightsquigarrow G'} \text{ (R-UNINIT)}$$

$$\frac{G \xrightarrow{e'} G' \quad e' \in \mathcal{R} \quad e'.rval = \mathbf{u} \quad G'.hbW(e') \quad \exists w. G'.race(w, e')}{G \rightsquigarrow G'} \text{ (R-RACE)}$$

$$\frac{G \xrightarrow{e'} G' \quad e' \in \mathcal{R} \quad e'.rval = w.wval \quad G'.hbW(e') \quad \neg G'.race(w, e') \quad G'' = \text{AddRF}(G', w, e') \quad \text{isCons}(G'')}{G \rightsquigarrow G''} \text{ (R-NORACE)}$$

Event structure construction rules

Consistency constraints

$$\begin{aligned} \text{isCons}(G) \triangleq & \text{irreflexive}(\text{wb}) \wedge \text{irreflexive}(\text{cf}; \text{hb}) \\ & \wedge \text{irreflexive}(\text{rf}; \text{hb}^{-1}; \text{rf}^{-1}; \text{cf}) \\ & \wedge \text{acyclic}((\text{hb} \cup \text{wb} \cup \text{fr}); [\text{SC}]) \end{aligned}$$

Event structure construction rules

Consistency constraints

Data race freedom (DRF) theorems

Event structure construction rules

Consistency constraints

Data race freedom (DRF) theorems

More transformations

- Speculative load
- Strengthening memory order of accesses

Event structure construction rules

Consistency constraints

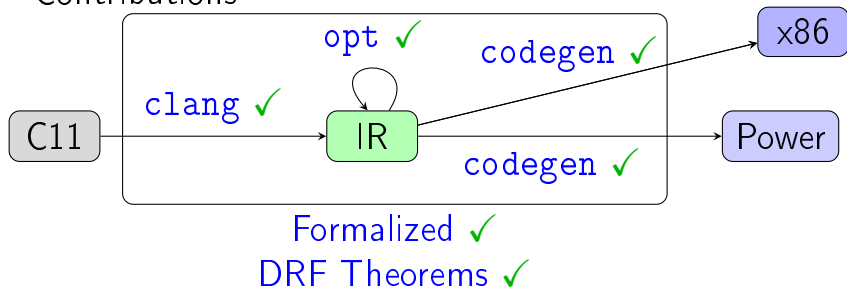
Data race freedom (DRF) theorems

More transformations

- Speculative load
- Strengthening memory order of accesses

Proofs: <http://plv.mpi-sws.org/llvmcs/>

- Contributions



- Future: extend the LLVM concurrency model
 - With relaxed accesses and fences
 - Prove/disprove more optimizations
 - Mechanize the formalization

Thank You !

Examples



```
int X = 0, Y = 0;  
a = X; || b = Y;  
Y = 1; || X = 1;  
a = b = 1 ✓
```

(LB)


```
int X = 0, Y = 0;  
a = X; || b = Y;  
if(a == 1) || if(b == 1)  
Y = 1; || X = 1;  
a = b = 1 ✗
```

(CYC)

LLVM



C11



LLVM performs speculative load

$$X = 1; \left\| \begin{array}{l} \text{if}(flag)\{ \\ \quad a = X; \\ \} \end{array} \right. \rightsquigarrow X = 1; \left\| \begin{array}{l} t = X; // \text{undef} \\ \text{if}(flag)\{ \\ \quad a = t; \\ \} \end{array} \right.$$