

Designing a Programming Language Shared-Memory Concurrency Semantics

Ori Lahav

<http://www.cs.tau.ac.il/~orilahav/>



ceClub -The Technion Computer Engineering Club

February 17, 2021

Example: Dekker's mutual exclusion

Initially, $x = y = 0$.

```
x := 1;
```

```
a := y;
```

```
if (a = 0) then
```

```
  /* critical section */
```

```
y := 1;
```

```
b := x;
```

```
if (b = 0) then
```

```
  /* critical section */
```

Example: Dekker's mutual exclusion

Initially, $x = y = 0$.

```
x := 1;
```

```
a := y;
```

```
if (a = 0) then
```

```
  /* critical section */
```

```
||
```

```
y := 1;
```

```
b := x;
```

```
if (b = 0) then
```

```
  /* critical section */
```

Is it safe?


Example: Dekker's mutual exclusion

Initially, $x = y = 0$.

<code>x := 1;</code>		<code>y := 1;</code>
<code>a := y; // 0</code>		<code>b := x; // 0</code>
<code>if (a = 0) then</code>		<code>if (b = 0) then</code>
<code> /* critical section */</code>		<code> /* critical section */</code>

Is it safe?

Yes, if we assume sequential consistency (SC):




Example: Dekker's mutual exclusion

Initially, $x = y = 0$.

<code>x := 1;</code>		<code>y := 1;</code>
<code>a := y; // 0</code>		<code>b := x; // 0</code>
<code>if (a = 0) then</code>		<code>if (b = 0) then</code>
<code> /* critical section */</code>		<code> /* critical section */</code>

Is it safe?

Yes, if we assume sequential consistency (SC):



No existing hardware implements SC!

- ▶ SC is very expensive (memory ~ 100 times slower than CPU).

Example: Shared-memory concurrency in C++

```
int X, Y, a, b;

void thread1() {
    X = 1;
    a = Y;
}

void thread2() {
    Y = 1;
    b = X;
}
```

```
int main () {
    int cnt = 0;

    do {
        X = 0; Y = 0;

        thread first(thread1);
        thread second(thread2);

        first.join();
        second.join();
        cnt++;

    } while (a != 0 || b != 0);

    printf("%d\n", cnt);
    return 0;
}
```


Example: Shared-memory concurrency in C++

```
int X, Y, a, b;  
  
void thread1() {  
    X = 1;  
    a = Y;  
}  
  
void thread2() {  
    Y = 1;  
    b = X;  
}
```

If Dekker's mutual exclusion is safe, this program will not terminate

```
int main () {  
    int cnt = 0;  
  
    do {  
        X = 0; Y = 0;  
  
        thread first(thread1);  
        thread second(thread2);  
  
        first.join();  
        second.join();  
        cnt++;  
  
    } while (a != 0 || b != 0);  
  
    printf("%d\n", cnt);  
    return 0;  
}
```

Weak memory models




We look for a *substitute for SC*

- ▶ What are the possible outcomes of a multi-threaded program in a high-level language?

Typically called a *weak memory model (WMM)*

- ▶ Allows more behaviors than SC.

Weak memory models



We look for a *substitute for SC*

- ▶ What are the possible outcomes of a multi-threaded program in a high-level language?

Typically called a *weak memory model (WMM)*

- ▶ Allows more behaviors than SC.

But it is not easy to get right

- ▶ The Java memory model (JMM) the the C/C++11 model are both flawed...

The Problem of Programming Language Concurrency Semantics

Mark Batty, Kayvan Memarian, Kyndylan Nienhuis, Jean Pichon-Pharabod,
and Peter Sewell

University of Cambridge

“Disturbingly, 40+ years after the first relaxed-memory hardware was introduced (the IBM 370/158MP), the field still *does not have a credible proposal for the concurrency semantics* of any general-purpose high-level language that includes high performance shared-memory concurrency primitives. This is a *major open problem* for programming language semantics.”

European Symposium on Programming (ESOP) 2015


Plan for rest of the talk

1. Challenges for programming language memory models
2. The C/C++11 memory model as a prototype
3. The “out-of-thin-air” problem
4. The “promising semantics” solution


Plan for rest of the talk

1. Challenges for programming language memory models
2. The C/C++11 memory model as a prototype
3. The “out-of-thin-air” problem
4. The “promising semantics” solution


Challenge 1: Various target models




x86-TSO  
(2010)



POWER 
(2011)



ARMv8 
(2016)



Store buffering in x86-TSO




Initially, $x = y = 0$.

$x := 1;$

$y := 1;$

$a := y; // 0$

$b := x; // 0$



Store buffering in x86-TSO




Initially, $x = y = 0$.

► $x := 1;$

► $y := 1;$

$a := y; // 0$

$b := x; // 0$




Store buffering in x86-TSO



Initially, $x = y = 0$.

<code>x := 1;</code>		<code>▶ y := 1;</code>
<code>▶ a := y; // 0</code>		<code>b := x; // 0</code>




Store buffering in x86-TSO



Initially, $x = y = 0$.

```
x := 1;           ||           y := 1;  
▶ a := y; // 0    ||    ▶ b := x; // 0
```




Store buffering in x86-TSO



Initially, $x = y = 0$.

```
x := 1;  
fence;  
a := y; // 0
```

```
y := 1;  
fence;  
b := x; // 0
```




Load buffering in ARM



Initially, $x = y = 0$.

```
a := x; // 1    ||    b := y; // 1
y := 1;         ||    x := b;
```




Load buffering in ARM



Initially, $x = y = 0$.

```
a := x; // 1    ||    b := y; // 1
y := 1;         ||    x := b;
```




Load buffering in ARM



Initially, $x = y = 0$.

```
a := x; // 1    ||    b := y; // 1
y := 1;        ||    x := b;
```




Load buffering in ARM



Initially, $x = y = 0$.

```
a := x; // 1    ||    b := y; // 1  
y := 1;        ||    x := b;
```

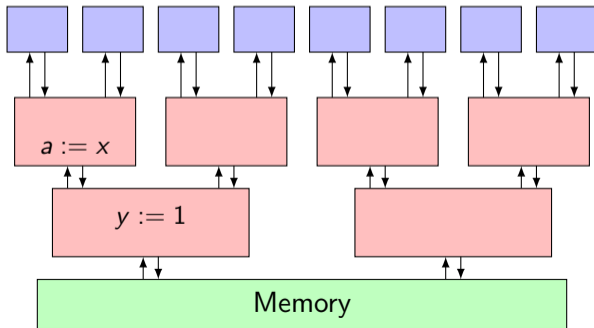


Load buffering in ARM



Initially, $x = y = 0$.

```
    a := x; // 1    ||    b := y; // 1  
    y := 1;        ||    x := b;
```



Challenge 2: Compilers stir the pot


Initially, $x = y = 0$.

```
x := 1; || a := x;  
y := 1; || b := y; // 1  
        || c := x; // 0
```

X forbidden under SC

Challenge 2: Compilers stir the pot

Initially, $x = y = 0$.



Common sub-expression elimination is **unsound** under SC

Challenge 3: Transformations do not suffice

Program transformations fail short to explain some weak behaviors.

- ▶ In C/C++, no reordering is allowed in the following program:

Message passing (MP)

```
x := 1;      || a := yacq; // 1  
y :=rel 1; || b := x; // 0
```

Challenge 3: Transformations do not suffice

Program transformations fail short to explain some weak behaviors.

- ▶ In C/C++, no reordering is allowed in the following program:


Message passing (MP)

```
x := 1;      || a := yacq; // 1
y :=rel 1; || b := x; // 0
```

- ▶ And yet, since C/C++ is intended to be compiled to a *non-multi-copy-atomic* architectures:

Independent reads of independent writes (IRIW)

```
a := xacq; // 1 || x :=rel 1; || y :=rel 1; || c := yacq; // 1
b := yacq; // 0 ||           ||           || d := xacq; // 0
```



WMM desiderata

1. Formal and comprehensive
2. Not too weak (good for programmers)
3. Not too strong (good for hardware)
4. Admits optimizations (good for compilers)

Implementability vs. **Programmability**

DRF-SC: A fundamental programmability guarantee

DRF-SC guarantee

no data races under SC \implies only SC behaviors

In most cases, programmers can avoid data races by using provided *synchronization mechanisms* (e.g., locks), and need not understand the full semantics.

DRF-SC: A fundamental programmability guarantee

DRF-SC guarantee

no data races under SC \implies only SC behaviors

In most cases, programmers can avoid data races by using provided *synchronization mechanisms* (e.g., locks), and need not understand the full semantics.

Establishing more refined programmability guarantees is an active research area:

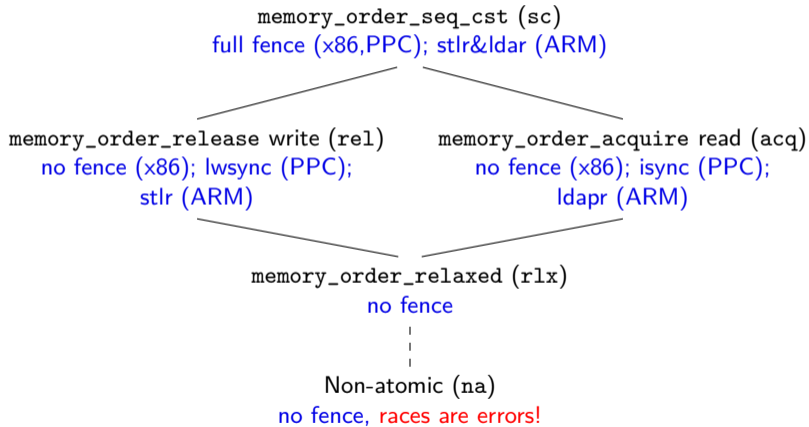
- ▶ *Local DRF* for OCaml memory model [Dolan, Sivaramakrishnan, Madhavapeddy PLDI'18]
- ▶ DRF wrt fragments *weaker than SC* [Kang, Hur, L, Vafeiadis, Dreyer POPL'17]

The C11 memory model

- ▶ Introduced by the ISO C/C++ 2011 standards.
- ▶ Serves as a solid basis for:
 - ▶ LLVM
 - ▶ WebAssembly memory model [Watt et al. OOPSLA 19]
 - ▶ JavaScript memory model [Watt et al. PLDI 20]
 - ▶ Java 9 [Bender & Palsberg OOPSLA 19]
 - ▶ Rust

A spectrum of access modes

non-atomic □ relaxed □ release/acquire □ sc



+ Explicit primitives for language level fences

C11: a declarative memory model





Declarative semantics abstracts away from implementation details.

- ▶ Became the “standard” in weak memory models
- ▶ Mature formalisms and tools (e.g., Herd [Alglave, Maranget, Tautschnig. TOPLAS'14])

1. a program \rightsquigarrow a set of directed **graphs**.
2. The model defines what graphs are *consistent*.

Execution graphs

Store buffering (SB)

$$\begin{array}{l} x = y = 0 \\ x :=_{rlx} 1; \parallel y :=_{rlx} 1; \\ a := y_{rlx}; \parallel b := x_{rlx}; \end{array}$$


Relations

- ▶ Program order, po
- ▶ Reads-from, rf

$$\begin{aligned}
 [-] &: \text{CExp} \rightarrow \mathbb{P}(\text{res} : \text{Val} \cup \{\perp\}, \mathcal{A} : \mathbb{P}(\text{AName}), \text{lab} : \mathcal{A} \rightarrow \text{Act}, \text{sb} : \mathbb{P}(\mathcal{A} \times \mathcal{A}), \text{fst} : \mathcal{A}, \text{lst} : \mathcal{A}) \\
 [v] &\stackrel{\text{def}}{=} \{(v, (a), \text{lab}, \theta, a, a) \mid a \in \text{AName} \wedge \text{lab}(a) = \text{skip}\} \\
 [\text{alloc}()] &\stackrel{\text{def}}{=} \{(\ell, (a), \text{lab}, \theta, a, a) \mid a \in \text{AName} \wedge \ell \in \text{Loc} \wedge \text{lab}(a) = \Lambda(\ell)\} \\
 [[v]_x := v'] &\stackrel{\text{def}}{=} \{(v', (a), \text{lab}, \theta, a, a) \mid a \in \text{AName} \wedge \text{lab}(a) = W_x(v, v')\} \\
 [[v]_x] &\stackrel{\text{def}}{=} \{(v', (a), \text{lab}, \theta, a, a) \mid a \in \text{AName} \wedge v' \in \text{Val} \wedge \text{lab}(a) = R_x(v, v')\} \\
 [\text{CAS}_{X,Y}(v, v_0, v_0)] &\stackrel{\text{def}}{=} \{(v', (a), \text{lab}, \theta, a, a) \mid a \in \text{AName} \wedge v' \in \text{Val} \wedge v' \neq v_0 \wedge \text{lab}(a) = R_Y(v, v')\} \\
 \cup \{(v_0, (a), \text{lab}, \theta, a, a) \mid a \in \text{AName} \wedge \text{lab}(a) = \text{RMW}_X(v, v_0, v_0)\} \\
 [\text{let } x = E_1 \text{ in } E_2] &\stackrel{\text{def}}{=} \{(\perp, \mathcal{A}_1, \text{lab}_1, \text{sb}_1, \text{fst}_1, \text{lst}_1) \mid (\perp, \mathcal{A}_1, \text{lab}_1, \text{sb}_1, \text{fst}_1, \text{lst}_1) \in [E_1]\} \\
 \cup \{(\text{res}_2, \mathcal{A}_1 \uplus \mathcal{A}_2, \text{lab}_1 \cup \text{lab}_2, \text{sb}_1 \cup \text{sb}_2 \cup \{(\text{lst}_1, \text{fst}_2 \}, \text{fst}_2, \text{lst}_2) \mid \\
 (v_1, \mathcal{A}_1, \text{lab}_1, \text{sb}_1, \text{fst}_1, \text{lst}_1) \in [E_1] \wedge (\text{res}_2, \mathcal{A}_2, \text{lab}_2, \text{sb}_2, \text{fst}_2, \text{lst}_2) \in [E_2][v_1/x]\} \\
 [\text{repeat } E \text{ end}] &\stackrel{\text{def}}{=} \{(\text{res}_N, \bigcup_{i \in [1..N]} \mathcal{A}_i, \bigcup_{i \in [1..N]} \text{lab}_i, \bigcup_{i \in [1..N]} \text{sb}_i \cup \{(\text{lst}_1, \text{fst}_2, \dots, (\text{lst}_{N-1}, \text{fst}_N) \}, \text{fst}_1, \text{lst}_N) \mid \\
 \forall i. (\text{res}_i, \mathcal{A}_i, \text{lab}_i, \text{sb}_i, \text{fst}_i, \text{lst}_i) \in [E] \wedge (i \neq N \implies \text{res}_i = 0) \wedge \text{res}_N \neq 0\} \\
 [E_1][E_2] &\stackrel{\text{def}}{=} \{(\text{combine}(\text{res}_1, \text{res}_2), \mathcal{A}_1 \uplus \mathcal{A}_2 \uplus \{ \text{at}_{\text{lock}}, \text{a}_{\text{join}} \}, \text{lab}_1 \cup \text{lab}_2 \cup \{ \text{at}_{\text{lock}} \mapsto \text{skip}, \text{a}_{\text{join}} \mapsto \text{skip} \}, \\
 \text{sb}_1 \cup \text{sb}_2 \cup \{ \text{at}_{\text{lock}}, \text{fst}_1 \}, \{ \text{at}_{\text{lock}}, \text{fst}_2 \}, \{ \text{lst}_1, \text{a}_{\text{join}} \}, \{ \text{lst}_2, \text{a}_{\text{join}} \}, \text{at}_{\text{lock}}, \text{a}_{\text{join}} \} \mid \\
 (\text{res}_1, \mathcal{A}_1, \text{sb}_1, \text{fst}_1, \text{lst}_1) \in [E_1] \wedge (\text{res}_2, \mathcal{A}_2, \text{sb}_2, \text{fst}_2, \text{lst}_2) \in [E_2] \wedge \text{at}_{\text{lock}}, \text{a}_{\text{join}} \in \text{AName}\}
 \end{aligned}$$

Figure 2. Semantics of closed program expressions.

$$\begin{aligned}
 \exists x. \text{hb}(x, x) & \quad (\text{IrreflexiveHB}) \\
 \forall \ell. \text{totalorder}(\{a \in \mathcal{A} \mid \text{iswrite}_\ell(a)\}, \text{mo}) \wedge \text{hb}_\ell \subseteq \text{mo} & \quad (\text{ConsistentMO}) \\
 \text{totalorder}(\{a \in \mathcal{A} \mid \text{isSeqCst}(a)\}, \text{sc}) \wedge \text{hb}_{\text{SeqCst}} \subseteq \text{sc} \wedge \text{mo}_{\text{SeqCst}} \subseteq \text{sc} & \quad (\text{ConsistentSC}) \\
 \forall b. \text{rf}(b) \neq \perp \iff \exists \ell. a. \text{iswrite}_\ell(a) \wedge \text{isread}_\ell(b) \wedge \text{hb}(a, b) & \quad (\text{ConsistentRFdom}) \\
 \forall a, b. \text{rf}(b) = a \implies \exists \ell. v. \text{iswrite}_{\ell, v}(a) \wedge \text{isread}_{\ell, v}(b) \wedge \neg \text{hb}(a, b) & \quad (\text{ConsistentRF}) \\
 \forall a, b. \text{rf}(b) = a \wedge (\text{mode}(a) = \text{ra} \vee \text{mode}(b) = \text{ra}) \implies \text{hb}(a, b) & \quad (\text{ConsistentRFna}) \\
 \forall a, b. \text{rf}(b) = a \wedge \text{isSeqCst}(b) \implies \text{isc}(a, b) \vee \neg \text{isSeqCst}(a) \wedge (\forall x. \text{isc}(x, b) \implies \neg \text{hb}(a, x)) & \quad (\text{RestrSCReads}) \\
 \exists a, b. \text{hb}(a, b) \wedge \text{mo}(\text{rf}(b), \text{rf}(a)) \wedge \text{locs}(a) = \text{locs}(b) & \quad (\text{CoherentRR}) \\
 \exists a, b. \text{hb}(a, b) \wedge \text{mo}(\text{rf}(b), a) \wedge \text{iswrite}(a) \wedge \text{locs}(a) = \text{locs}(b) & \quad (\text{CoherentWR}) \\
 \exists a, b. \text{hb}(a, b) \wedge \text{mo}(b, \text{rf}(a)) \wedge \text{iswrite}(b) \wedge \text{locs}(a) = \text{locs}(b) & \quad (\text{CoherentRW}) \\
 \forall a. \text{isrmw}(a) \wedge \text{rf}(a) \neq \perp \implies \text{mo}(\text{rf}(a), a) \wedge \exists c. \text{mo}(\text{rf}(a), c) \wedge \text{mo}(c, a) & \quad (\text{AtomicRMW}) \\
 \forall a, b, \ell. \text{lab}(a) = \text{lab}(b) = \Lambda(\ell) \implies a = b & \quad (\text{ConsistentAlloc})
 \end{aligned}$$

$$\begin{aligned}
 \text{where } \text{iswrite}_{\ell, v}(a) &\stackrel{\text{def}}{=} \exists X, v_{\text{old}}. \text{lab}(a) \in \{W_X(\ell, v), \text{RMW}_X(\ell, v_{\text{old}}, v)\} \quad \text{iswrite}_\ell(a) \stackrel{\text{def}}{=} \exists v. \text{iswrite}_{\ell, v}(a) \\
 \text{isread}_{\ell, v}(a) &\stackrel{\text{def}}{=} \exists X, v_{\text{new}}. \text{lab}(a) \in \{R_X(\ell, v), \text{RMW}_X(\ell, v, v_{\text{new}})\} \quad \text{etc.} \\
 \text{rsElem}(a, b) &\stackrel{\text{def}}{=} \text{sameThread}(a, b) \vee \text{isrmw}(b) \\
 \text{rseq}(a) &\stackrel{\text{def}}{=} \{a\} \cup \{b \mid \text{rsElem}(a, b) \wedge \text{mo}(a, b) \wedge (\forall c. \text{mo}(a, c) \wedge \text{mo}(c, b) \implies \text{rsElem}(a, c))\} \\
 \text{sw} &\stackrel{\text{def}}{=} \{(a, b) \mid \text{mode}(a) \in \{\text{rel}, \text{rel_acq}, \text{sc}\} \wedge \text{mode}(b) \in \{\text{acq}, \text{rel_acq}, \text{sc}\} \wedge \text{rf}(b) \in \text{rseq}(a)\} \\
 \text{hb} &\stackrel{\text{def}}{=} (\text{sb} \cup \text{sw})^+ \\
 \text{hb}_\ell &\stackrel{\text{def}}{=} \{(a, b) \in \text{hb} \mid \text{iswrite}_\ell(a) \wedge \text{iswrite}_\ell(b)\} \\
 X_{\text{SeqCst}} &\stackrel{\text{def}}{=} \{(a, b) \in X \mid \text{isSeqCst}(a) \wedge \text{isSeqCst}(b)\} \\
 \text{isc}(a, b) &\stackrel{\text{def}}{=} \text{iswrite}_{\text{lock}(b)}(a) \wedge \text{sc}(a, b) \wedge \exists c. \text{sc}(a, c) \wedge \text{sc}(c, b) \wedge \text{iswrite}_{\text{lock}(b)}(c)
 \end{aligned}$$

Figure 3. Axioms satisfied by consistent C11 executions, Consistent(A, lab, sb, rf, mo, sc).




Figure 4. Sample executions violating coherency conditions (Batty et al. 2011).


Require the existence of several orders that satisfy certain constraints:

- ▶ SC-per-location (a.k.a. coherence)
- ▶ Release/acquire synchronization
- ▶ Global conditions on SC accesses

Example: flag-based synchronization


Message passing (MP)

```
y :=rlx 42; || a := xrlx; // 1  
x :=rlx 1; || b := yrlx; // 0
```



Message passing (MP)

```
y :=rlx 42; || a := xacq; // 1  
x :=rel 1; || b := yrlx; // 0
```



- ▶ The semantics of SC accesses is the *most complicated* part of the model.

- ▶ The semantics of SC accesses is the *most complicated* part of the model.
- ▶ C/C++11 provides *too strong* semantics (a correctness problem!)

$$\begin{array}{l}
 a := x_{\text{acq}}; // 1 \\
 b := y_{\text{sc}}; // 0
 \end{array}
 \parallel
 \begin{array}{l}
 x :=_{\text{sc}} 1; \\
 y :=_{\text{sc}} 1;
 \end{array}
 \parallel
 \begin{array}{l}
 c := y_{\text{acq}}; // 1 \\
 d := x_{\text{sc}}; // 0
 \end{array}$$

- ▶ In addition, its semantics for SC fences is *too weak*.

$$\begin{array}{l}
 a := x_{\text{acq}}; // 1 \\
 \mathbf{fence}_{\text{sc}}; \\
 b := y_{\text{acq}}; // 0
 \end{array}
 \parallel
 \begin{array}{l}
 x :=_{\text{rel}} 1; \\
 y :=_{\text{rel}} 1;
 \end{array}
 \parallel
 \begin{array}{l}
 c := y_{\text{acq}}; // 1 \\
 \mathbf{fence}_{\text{sc}}; \\
 d := x_{\text{acq}}; // 0
 \end{array}$$

- ▶ The semantics of SC accesses is the *most complicated* part of the model.
- ▶ C/C++11 provides *too strong* semantics (a correctness problem!)

$$\begin{array}{l}
 a := x_{\text{acq}}; // 1 \\
 b := y_{\text{sc}}; // 0
 \end{array}
 \parallel
 \begin{array}{l}
 x :=_{\text{sc}} 1; \\
 y :=_{\text{sc}} 1;
 \end{array}
 \parallel
 \begin{array}{l}
 c := y_{\text{acq}}; // 1 \\
 d := x_{\text{sc}}; // 0
 \end{array}$$

- ▶ In addition, its semantics for SC fences is *too weak*.

$$\begin{array}{l}
 a := x_{\text{acq}}; // 1 \\
 \mathbf{fence}_{\text{sc}}; \\
 b := y_{\text{acq}}; // 0
 \end{array}
 \parallel
 \begin{array}{l}
 x :=_{\text{rel}} 1; \\
 y :=_{\text{rel}} 1;
 \end{array}
 \parallel
 \begin{array}{l}
 c := y_{\text{acq}}; // 1 \\
 \mathbf{fence}_{\text{sc}}; \\
 d := x_{\text{acq}}; // 0
 \end{array}$$

- ▶ The standard committee fixed the specification to solve these problems in C++20.

The “out-of-thin-air” problem

C/C++11 is too weak

non-atomic □ relaxed □ release/acquire □ sc

C/C++11 is too weak

non-atomic relaxed release/acquire sc

Load-buffering

```
a := x; // 1    ||    b := y; // 1
y := 1;         ||    x := b;
```

C/C++11 allows this behavior
because **POWER & ARM allow it!**


C/C++11 is too weak

non-atomic \sqsubset relaxed \sqsubset release/acquire \sqsubset sc

Load-buffering

```
a := x; // 1    ||    b := y; // 1
y := 1;         ||    x := b;
```

C/C++11 allows this behavior
because **POWER & ARM allow it!**




C/C++11 is too weak

non-atomic \square relaxed \square release/acquire \square sc

Load-buffering

```
a := x; // 1    ||    b := y; // 1
y := 1;         ||    x := b;
```

C/C++11 allows this behavior
because **POWER & ARM allow it!**




C/C++11 is too weak

non-atomic \square relaxed \square release/acquire \square sc

Load-buffering

```
a := x; // 1    ||    b := y; // 1
y := 1;         ||    x := b;
```

C/C++11 allows this behavior
because **POWER & ARM allow it!**




C/C++11 is too weak

non-atomic \square relaxed \square release/acquire \square sc

Load-buffering

```
a := x; // 1    ||    b := y; // 1
y := 1;        ||    x := b;
```

C/C++11 allows this behavior
because **POWER & ARM allow it!**



C/C++11 is too weak

non-atomic □ **relaxed** □ release/acquire □ sc


Load-buffering

```
a := x; // 1     ||     b := y; // 1
y := 1;           ||     x := b;
```

C/C++11 allows this behavior
because **POWER & ARM allow it!**

Load-buffering + data dependency

```
a := x; // 1     ||     b := y; // 1
y := a;           ||     x := b;
```



C/C++11 is too weak

non-atomic □ relaxed □ release/acquire □ sc

Load-buffering


```
a := x; // 1     ||     b := y; // 1
y := 1;           ||     x := b;
```

C/C++11 allows this behavior
because **POWER & ARM allow it!**

Load-buffering + data dependency

```
a := x; // 1     ||     b := y; // 1
y := a;           ||     x := b;
```

C/C++11 allows this behavior



C/C++11 is too weak

non-atomic \square relaxed \square release/acquire \square sc

Load-buffering


```
a := x; // 1    ||    b := y; // 1
y := 1;         ||    x := b;
```

C/C++11 allows this behavior
because **POWER & ARM allow it!**

Load-buffering + data dependency

```
a := x; // 1    ||    b := y; // 1
y := a;         ||    x := b;
```

C/C++11 allows this behavior
Values appear out-of-thin-air!
(no hardware/compiler exhibit this behavior)




C/C++11 is too weak

non-atomic □ relaxed □ release/acquire □ sc

Load-buffering + control dependency

```
a := x; // 1      ||      b := y; // 1
if (a = 1)        ||      if (b = 1)
  y := 1;         ||      x := 1;
```




C/C++11 is too weak

non-atomic \square relaxed \square release/acquire \square sc

Load-buffering + control dependency

```
a := x; // 1      ||      b := y; // 1
if (a = 1)        ||      if (b = 1)
  y := 1;         ||      x := 1;
```

C/C++11 allows this behavior



C/C++11 is too weak


non-atomic □ relaxed □ release/acquire □ sc

Load-buffering + control dependency

```
a := x; // 1      ||      b := y; // 1
if (a = 1)        ||      if (b = 1)
  y := 1;         ||      x := 1;
```

C/C++11 allows this behavior

The DRF guarantee is broken!



C/C++11 is too weak

non-atomic relaxed release/acquire sc


Load-buffering + control dependency

```
a := x; // 1     ||     b := y; // 1  
if (a = 1)       ||     if (b = 1)
```

**The three examples have
the same execution graph!**

C/C++11 shows this behavior

The DRF guarantee is broken!



The hardware solution


Keep track of **syntactic dependencies** and forbid **dependency cycles**.

Load-buffering

<code>a := x; // 1</code>		<code>b := y; // 1</code>
<code>y := 1;</code>		<code>x := b;</code>

Load-buffering + data dependency

<code>a := x; // 1</code>		<code>b := y; // 1</code>
<code>y := a;</code>		<code>x := b;</code>



The hardware solution

Keep track of **syntactic dependencies** and forbid **dependency cycles**.

Load-buffering


```
a := x; // 1    ||    b := y; // 1
y := 1;         ||    x := b;
```

Load-buffering + data dependency

```
a := x; // 1    ||    b := y; // 1
y := a;         ||    x := b;
```

Load-buffering + fake dependency

```
a := x; // 1    ||    b := y; // 1
y := a + 1 - a; ||    x := b;
```



The hardware solution

Keep track of **syntactic dependencies** and forbid **dependency cycles**.

Load-buffering


```
a := x; // 1    ||    b := y; // 1
y := 1;         ||    x := b;
```

Load-buffering + data dependency

```
a := x; // 1    ||    b := y; // 1
y := a;         ||    x := b;
```

Load-buffering + fake dependency

```
a := x; // 1    ||    b := y; // 1
y := a + 1 - a; ||    x := b;
```



Unsuitable for PL: **Compilers do not preserve syntactic dependencies.**

The “out-of-thin-air” problem

C/C++11 is too weak

- ▶ Values might appear *out-of-thin-air*.
- ▶ The *DRF guarantee* is broken.

The C++14 standard states:

“Implementations should ensure that no “out-of-thin-air” values are computed that circularly depend on their own computation.”

A straightforward solution

- ▶ Disallow `po U rf` cycles!
- ▶ On weak hardware it carries a certain *implementation cost*.

[Ou & Demsky. Towards understanding the costs of avoiding out-of-thin-air results. OOPSLA'18]
Slowdown on ARMv8 is 3.1% on average and 17.6% max (on some benchmarks...)

A straightforward solution

- ▶ Disallow `po U rf` cycles!
- ▶ On weak hardware it carries a certain *implementation cost*.
[Ou & Demsky. Towards understanding the costs of avoiding out-of-thin-air results. OOPSLA'18]
Slowdown on ARMv8 is 3.1% on average and 17.6% max (on some benchmarks...)

RC11 (Repaired C11) model

[L, Vafeiadis, Kang, Hur, Dreyer. PLDI'17]

- ▶ (Modified) compilation schemes are correct.
- ▶ DRF holds and no OOTA-values.
- ▶ Model checking tools [Kokologiannakis, L, Sagonas, Vafeiadis. POPL'18]
<http://plv.mpi-sws.org/rcmc/>

A straightforward solution

- ▶ Disallow `po U rf` cycles!
- ▶ On weak hardware it carries a certain *implementation cost*.
[Ou & Demsky. Towards understanding the costs of avoiding out-of-thin-air results. OOPSLA'18]
Slowdown on ARMv8 is 3.1% on average and 17.6% max (on some benchmarks...)

RC11 (Repaired C11) model

[L, Vafeiadis, Kang, Hur, Dreyer. PLDI'17]

- ▶ (Modified) compilation schemes are correct.
- ▶ DRF holds and no OOTA-values.
- ▶ Model checking tools [Kokologiannakis, L, Sagonas, Vafeiadis. POPL'18]
<http://plv.mpi-sws.org/rcmc/>

- ▶ Solving the problem without changing the compilation schemes will require a *major revision* of the standard.

A 'promising' solution to OOTA

[Kang, Hur, L, Vafeiadis, Dreyer. POPL'17]

A 'promising' solution to OOTA

[Kang, Hur, L, Vafeiadis, Dreyer. POPL'17]

Key idea: Start with an operational interleaving semantics, but allow threads to **promise** to write in the future.

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
      x = y = 0
x = 1;  ||  y = 1;
a = y; //0 || b = x; //0
```

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
x = y = 0
┆ x = 1;      ┆ y = 1;
┆ a = y; // 0 ┆ b = x; // 0
```

Memory

$\langle x : 0@0 \rangle$
 $\langle y : 0@0 \rangle$

T_1 's view

x	y
0	0

T_2 's view

x	y
0	0

- ▶ Global memory is a pool of messages of the form

$\langle location : value@timestamp \rangle$

- ▶ Each thread maintains a *thread-local view* recording the last observed timestamp for every location

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
x = y = 0
  x = 1;      |      y = 1;
  ▶ a = y; // 0 |      ▶ b = x; // 0
```

Memory

```
⟨x : 0@0⟩
⟨y : 0@0⟩
⟨x : 1@5⟩
```

T₁'s view

x	y
0	0
5	

T₂'s view

x	y
0	0

- ▶ Global memory is a pool of messages of the form

$\langle location : value @ timestamp \rangle$

- ▶ Each thread maintains a *thread-local view* recording the last observed timestamp for every location

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
x = y = 0
x = 1;      y = 1;
▶ a = y; // 0  ||  ▶ b = x; // 0
```

Memory

```
⟨x : 0@0⟩
⟨y : 0@0⟩
⟨x : 1@5⟩
⟨y : 1@5⟩
```

T₁'s view

x	y
0	0
5	

T₂'s view

x	y
0	0
	5

- ▶ Global memory is a pool of messages of the form

⟨location : value@timestamp⟩

- ▶ Each thread maintains a *thread-local view* recording the last observed timestamp for every location

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
x = y = 0
x = 1;      ||      y = 1;
a = y; // 0  ||      b = x; // 0
```

Memory

```
<x : 0@0>
<y : 0@0>
<x : 1@5>
<y : 1@5>
```

T₁'s view

x	y
0	0
5	

T₂'s view

x	y
0	0
	5

- ▶ Global memory is a pool of messages of the form

$\langle location : value @ timestamp \rangle$

- ▶ Each thread maintains a *thread-local view* recording the last observed timestamp for every location

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
x = y = 0
x = 1;      ||      y = 1;
a = y; // 0 ||      b = x; // 0
```

Memory

```
<x : 0@0>
<y : 0@0>
<x : 1@5>
<y : 1@5>
```

T₁'s view

x	y
0	0
5	

T₂'s view

x	y
0	0
	5

- ▶ Global memory is a pool of messages of the form

$\langle location : value @ timestamp \rangle$

- ▶ Each thread maintains a *thread-local view* recording the last observed timestamp for every location

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
x = y = 0
x = 1;      ||      y = 1;
a = y; // 0 ||      b = x; // 0
```

Memory

```
<x: 0@0>
<y: 0@0>
<x: 1@5>
<y: 1@5>
```

T_1 's view

x	y
0	0
5	

T_2 's view

x	y
0	0
	5

Coherence Test

```
x = 0
x := 1;  ||  x := 2;
a = x; // 2 ||  b = x; // 1
```

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
x = y = 0
x = 1;      |      y = 1;
a = y; // 0 |      b = x; // 0
```

Memory

```
<x : 0@0>
<y : 0@0>
<x : 1@5>
<y : 1@5>
```

T₁'s view

x	y
0	0
5	

T₂'s view

x	y
0	0
	5

Coherence Test

```
x = 0
x := 1; | x := 2;
a = x; // 2 | b = x; // 1
```

Memory

```
<x : 0@0>
```

T₁'s view

x
0

T₂'s view

x
0

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
x = y = 0
x = 1;      |      y = 1;
a = y; // 0 |      b = x; // 0
```

Memory

```
<x: 0@0>
<y: 0@0>
<x: 1@5>
<y: 1@5>
```

T₁'s view

x	y
0	0
5	

T₂'s view

x	y
0	0
	5

Coherence Test

```
x = 0
x := 1;    |      x := 2;
a = x; // 2 |      b = x; // 1
```

Memory

```
<x: 0@0>
<x: 1@5>
```

T₁'s view

x
0
5

T₂'s view

x
0

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
      x = y = 0
x = 1;  ||  y = 1;
a = y; //0 || b = x; //0
```

Memory

```
<x: 0@0>
<y: 0@0>
<x: 1@5>
<y: 1@5>
```

T₁'s view

x	y
0	0
5	

T₂'s view

x	y
0	0
	5

Coherence Test

```
      x = 0
x := 1;  ||  x := 2;
a = x; //2 || b = x; //1
```

Memory

```
<x: 0@0>
<x: 1@5>
<x: 2@7>
```

T₁'s view

x
0
5

T₂'s view

x
0
7

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
      x = y = 0
x = 1;  ||  y = 1;
a = y; //0 || b = x; //0
```

Memory

```
<x : 0@0>
<y : 0@0>
<x : 1@5>
<y : 1@5>
```

T₁'s view

x	y
0	0
5	

T₂'s view

x	y
0	0
	5

Coherence Test

```
      x = 0
x := 1;  ||  x := 2;
a = x; //2  ▶ b = x; //1
```

Memory

```
<x : 0@0>
<x : 1@5>
<x : 2@7>
```

T₁'s view

x
0
1
7

T₂'s view

x
0
7

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
      x = y = 0
x = 1;  ||  y = 1;
a = y; //0 || b = x; //0
```

Memory

```
<x : 0@0>
<y : 0@0>
<x : 1@5>
<y : 1@5>
```

T₁'s view

x	y
0	0
5	

T₂'s view

x	y
0	0
	5

Coherence Test

```
      x = 0
x := 1;  ||  x := 2;
a = x; //2 || b = x; //1
```

Memory

```
<x : 0@0>
<x : 1@5>
<x : 2@7>
```

T₁'s view

x
0
1
7

T₂'s view

x
0
7

Load-buffering

```
a := x; // 1  
y := 1;  |||  x := y;
```

- ▶ To model load-store reordering, we allow **“promises”**.
- ▶ At any point, a thread may promise to write a message in the future, allowing other threads to read from the promised message.

Promises

Load-buffering

```
▶ a := x; // 1 || ▶ x := y;  
  y := 1;
```

Memory

```
<x : 0@0>  
<y : 0@0>
```

T_1 's view

x	y
0	0

T_2 's view

x	y
0	0

- ▶ To model load-store reordering, we allow **“promises”**.
- ▶ At any point, a thread may promise to write a message in the future, allowing other threads to read from the promised message.

Promises

Load-buffering

```
▶ a := x; // 1 || ▶ x := y;  
  y := 1;
```

Memory

$\langle x : 0@0 \rangle$

$\langle y : 0@0 \rangle$

$\langle y : 1@5 \rangle$

T_1 's view

x	y
0	0

T_2 's view

x	y
0	0

- ▶ To model load-store reordering, we allow **“promises”**.
- ▶ At any point, a thread may promise to write a message in the future, allowing other threads to read from the promised message.

Promises

Load-buffering

```
▶ a := x; // 1 || ▶ x := y;  
  y := 1;
```

Memory

⟨x : 0@0⟩

⟨y : 0@0⟩

⟨y : 1@5⟩

T₁'s view

x	y
0	0

T₂'s view

x	y
0	0 5

- ▶ To model load-store reordering, we allow **“promises”**.
- ▶ At any point, a thread may promise to write a message in the future, allowing other threads to read from the promised message.

Promises

Load-buffering

```
▶ a := x; // 1 || x := y;  
y := 1; ▶
```

Memory

```
⟨x : 0@0⟩  
⟨y : 0@0⟩  
⟨y : 1@5⟩  
⟨x : 1@5⟩
```

T_1 's view

x	y
0	0

T_2 's view

x	y
0	0
5	5

- ▶ To model load-store reordering, we allow **“promises”**.
- ▶ At any point, a thread may promise to write a message in the future, allowing other threads to read from the promised message.

Promises

Load-buffering

```
a := x; // 1
▶ y := 1;  ||  x := y;
```

Memory

```
<x : 0@0>
<y : 0@0>
<y : 1@5>
<x : 1@5>
```

T_1 's view

x	y
0	0
5	

T_2 's view

x	y
0	0
5	5

- ▶ To model load-store reordering, we allow **“promises”**.
- ▶ At any point, a thread may promise to write a message in the future, allowing other threads to read from the promised message.

Promises

Load-buffering

```
a := x; // 1  
y := 1;  || x := y;
```

Memory

```
<x : 0@0>  
<y : 0@0>  
<y : 1@5>  
<x : 1@5>
```

T_1 's view

x	y
0	0
5	5

T_2 's view

x	y
0	0
5	5

- ▶ To model load-store reordering, we allow **“promises”**.
- ▶ At any point, a thread may promise to write a message in the future, allowing other threads to read from the promised message.

Promises

Load-buffering

```
a := x; // 1 || x := y;  
y := 1;      ▶
```

Load-buffering + dependency

```
a := x; // 1 || x := y;  
y := a;      ||
```

Memory

```
⟨x : 0@0⟩  
⟨y : 0@0⟩  
⟨y : 1@5⟩  
⟨x : 1@5⟩
```

T_1 's view

x	y
0	0
5	5

T_2 's view

x	y
0	0
5	5

Must not admit the same execution!

Certified promises

Thread-local certification

A thread can promise to write a message, if it can *thread-locally certify* that its promise will be fulfilled.

Load-buffering

$$\begin{array}{l} a := x; \text{ // } 1 \\ y := 1; \end{array} \parallel x := y;$$

T_1 **may promise** $y := 1$, since it is able to write $y := 1$ by itself.

Load buffering + fake dependency

$$\begin{array}{l} a := x; \text{ // } 1 \\ y := a + 1 - a; \end{array} \parallel x := y;$$

T_1 **may NOT promise** $y := 1$, since it is not able to write $y := 1$ by itself.

Load buffering + dependency

$$\begin{array}{l} a := x; \text{ // } 1 \\ y := a; \end{array} \parallel x := y;$$

Is this behavior possible?

```
a := x; // 42
y := a;
|
|
| b := y;
| if (b = 42)
|   c := 1;
| else
|   c := 2;
|   b := 42;
| x := b;
| print (c); // prints 1
```

Is this behavior possible?

```
a := x; // 42
y := a;
b := y;
if (b = 42)
    c := 1;
else
    c := 2;
    b := 42;
x := b;
print (c); // prints 1
```

Yes. And it can be obtained by standard compiler optimizations!

The full model

We have extended this basic idea to handle:

- ▶ Atomic Read-Modify-Writes (e.g., CAS, fetch-and-add)
- ▶ Release/acquire accesses and fences
- ▶ SC fences
- ▶ Plain accesses (C11's non-atomics & Java's normal accesses)

Results

- ▶ No “out-of-thin-air” values
- ▶ DRF guarantees
- ▶ Compiler optimizations (incl. reorderings, eliminations)
- ▶ Efficient h/w mappings (x86-TSO, Power, ARM)

The full model

We have extended this basic idea to handle:

- ▶ Atomic Read-Modify-Writes (e.g., CAS, fetch-and-add)
- ▶ Release/acquire accesses and fences
- ▶ SC fences
- ▶ Plain accesses (C11's non-atomics & Java's normal accesses)



The **Coq**
proof assistant



Results

- ▶ No “out-of-thin-air” values
- ▶ DRF guarantees
- ▶ Compiler optimizations (incl. reorderings, eliminations)
- ▶ Efficient h/w mappings (x86-TSO, Power, ARM)

The full model

We have extended this basic idea to handle:

- ▶ Atomic Read-Modify-Writes (e.g., CAS, fetch-and-add)
- ▶ Release/acquire accesses and fences
- ▶ SC fences
- ▶ Plain accesses (C11's non-atomics & Java's normal accesses)

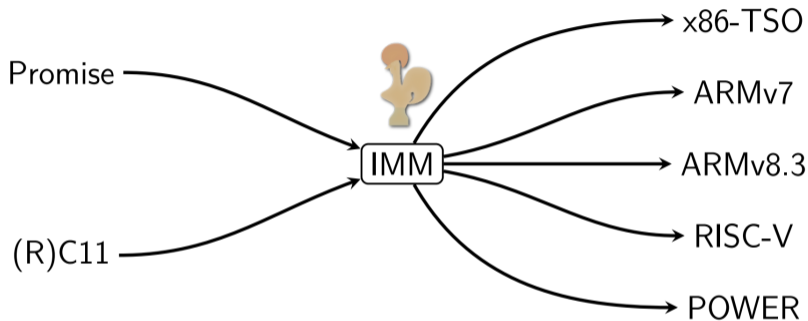


The **Coq**
proof assistant



Results

- ▶ No “out-of-thin-air” values
- ▶ DRF guarantees
- ▶ Compiler optimizations (incl. reorderings, eliminations)
- ▶ Efficient h/w mappings (x86-TSO, Power, ARM)



- ▶ A *common denominator* of existing models
- ▶ Formulated in the *declarative style*
- ▶ Simplifies *compilation correctness* proofs

Certification from current memory is not enough!

<pre>a := FADD(x, 1, acq-rel) // 0 if a = 0 then y := 1</pre>	<pre>b := FADD(x, 1, acq-rel) // 0 if b = 0 then c := y // 1 if c = 1 then x := 0</pre>
--------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------

- ▶ The only race is on an *acquire-release RMW*.
- ▶ The DRF-RA guarantee entails the annotated behavior should be **disallowed**.
- ▶ Thus, the behavior must be forbidden by the promising model.
- ▶ We forbid it by requiring a certification from *any extension* of the current memory.

Complex compilation issues... (1/2)

$$\begin{array}{l} a := y // 1 \\ z := a \end{array} \parallel \begin{array}{l} b := z // 1 \\ c := \mathbf{FADD}(x, 1) // 0 \\ y := c + 1 \end{array}$$

- ▶ The promising model **forbids** this behavior.
- ▶ But, it is **allowed** when compiling to ARMv8.
- ▶ **Register promotion** is unsound.

Complex compilation issues... (2/2)

```
a := CAS(x, 0, 1) // 1 || x := 42
if a < 10 then         || b := y // 1
    y := 1             || x := b
```

- ▶ The promising model **forbids** this behavior.
- ▶ But, it can be obtained by local compiler optimization + *global value-range analysis*.

Complex compilation issues... (2/2)

```
a := CAS(x, 0, 1) // 1 || x := 42
if a < 10 then         || b := y // 1
    y := 1              || x := b
```

- ▶ The promising model **forbids** this behavior.
- ▶ But, it can be obtained by local compiler optimization + *global value-range analysis*.

Promising 2.0

[Lee, Cho, Podkopaev, Chakraborty, Hur, L, Vafeiadis PLDI'20]

These issues were fixed by a better “forall future memory” certification requirement.

An ongoing challenge: A local DRF guarantee

Existing programmability guarantees are non-modular!

<code>a := pop(S)</code>		<code>b := pop(S)</code>
<code>lock()</code>		<code>lock()</code>
process <i>a</i> accessing <i>x, y</i>		process <i>b</i> accessing <i>x, y</i>
<code>unlock()</code>		<code>unlock()</code>

- ▶ We want to assume SC semantics for the accesses to *x* and *y*.
- ▶ The stack implementation may have (benign) *races*.
- ▶ The (global) DRF-SC guarantee is **inapplicable**.

An ongoing challenge: A local DRF guarantee

Existing programmability guarantees are non-modular!

<code>a := pop(S)</code>		<code>b := pop(S)</code>
<code>lock()</code>		<code>lock()</code>
process <i>a</i> accessing <i>x, y</i>		process <i>b</i> accessing <i>x, y</i>
<code>unlock()</code>		<code>unlock()</code>


- ▶ We want to assume SC semantics for the accesses to *x* and *y*.
- ▶ The stack implementation may have (benign) *races*.
- ▶ The (global) DRF-SC guarantee is **inapplicable**.

A bad surprise...

[Lee, Cho, Hur, L submitted]


Standard compiler optimizations are inconsistent with local DRF guarantees.

Summary




- ▶ The challenges in designing a WMM.
- ▶ The C/C++11 model.
- ▶ C/C++11 is **broken**:
 - ▶ Most problems are **locally fixable**.
 - ▶ But ruling out **OOTA** requires an entirely different approach.
- ▶ The **promising model** may be the solution.

Summary



- ▶ The challenges in designing a WMM.
- ▶ The C/C++11 model.
- ▶ C/C++11 is **broken**:
 - ▶ Most problems are **locally fixable**.
 - ▶ But ruling out **OOA** requires an entirely different approach.
- ▶ The **promising model** may be the solution.

Summary



- ▶ The challenges in designing a WMM.
- ▶ The C/C++11 model.
- ▶ C/C++11 is **broken**:
 - ▶ Most problems are **locally fixable**.
 - ▶ But ruling out **OOA** requires an entirely different approach.
- ▶ The **promising model** may be the solution.

Thank you!

<http://www.cs.tau.ac.il/~orilahav/>