

# Promising-ARM/RISC-V

~~a simpler and faster operational concurrency model~~

an equivalent simpler and faster presentation of ARMv8/RISC-V

---

Christopher Pulte<sup>1</sup> Jean Pichon-Pharabod<sup>1</sup> Jeehoon Kang<sup>2</sup> Sung-Hwan Lee<sup>3</sup> Chung-Kil Hur<sup>3</sup>

24 June 2019

<sup>1</sup>University of Cambridge

<sup>2</sup>Korea Advanced Institute of Science and Technology

<sup>3</sup>Seoul National University

## Concurrency programming

```
void produce(int v) {  
    D = v;  
    F = 1;  
}  
  
int consume() {  
    while (1) {  
        if (F) break;  
    }  
    return D;  
}
```

```
produce(42) || d = consume()
```

**Initially:** D = F = 0

**Finally:** d = 42

# Concurrency programming

```
void produce(int v) {  
    D = v; ← compiler re-ordering,  
    F = 1;     hardware out-of-order  
}
```

```
int consume() {  
    while (1) {  
        if (F) break; ← re-ordering past  
    }           control dependency  
    return D;  
}
```

**Initially:** D = F = 0

**Finally:** d = 42

**or:** d = 0

```
produce(42) || d = consume()
```

## Concurrency programming

```
void produce(int v) {
    D = v;
    F[rel] = 1;
}

int consume() {
    while (1) {
        if (F[acq]) break;
    }
    return D;
}
```

```
produce(42) || d = consume()
```

**Initially:**  $D = F = 0$

**Finally:**  $d = 42$

**or:**  $d = 0$

```
void produce(int v) {  
    D = v;  
    F[rel] = 1;  
}  
  
int consume() {  
    while (1) {  
        if (F[acq]) break;  
    }  
    return D;  
}
```

```
produce(42) || d = consume()
```

**Need precise, simple semantics  
and tool support.**

### **Axiomatic**

- official reference model
- + abstract, concise
- not incremental: global axioms

### **Flat operational**

- proved equivalent
- + incremental
- + ISA, mixed-size support
- + closer relation to H/W
- complex

### Axiomatic

- official reference model
- + abstract, concise
- not incremental: global axioms

### Flat operational

- proved equivalent
- + incremental
- + ISA, mixed-size support
- + closer relation to H/W
- complex

### Promising-ARM/RISC-V

- + simple, incremental
- + Coq equivalence proof with Axiomatic (excl. ISA model)
- + ISA support
- no mixed-size support yet
- + fast enough for checking data structure examples

inspired by Promising C11 [Kang et al]

## **Model overview**

---



**Idea 1: out-of-order read execution  
by reading from message history**

**Out-of-order reads**

**T1**

```
store D := 42  
store F := 1
```

**T2**

```
r1 = load F      //reads 1  
if (r1 == 1) then  
  r2 = load D    //reads 0  
else  
  ..
```

## Idea 1: out-of-order read execution by reading from message history

## Out-of-order reads

→ **T1**  
store D := 42  
store F := 1

→ **T2**  
r1 = load F //reads 1  
if (r1 == 1) then  
  r2 = load D //reads 0  
else  
  ..

**T1.regs:** ...  
**T1.promises:** ...  
...

**T2.regs:** r1 ↦ 0@0, r2 ↦ 0@0 ...  
**T2.promises:** ...  
...

**memory:** (init)@0

## Idea 1: out-of-order read execution by reading from message history

## Out-of-order reads

**T1**  
→ store D := 42  
store F := 1

→ **T2**  
r1 = load F //reads 1  
if (r1 == 1) then  
  r2 = load D //reads 0  
else  
  ..

**T1.regs:** ...  
**T1.promises:** ...  
...

**T2.regs:** r1 ↦ 0@0, r2 ↦ 0@0 ...  
**T2.promises:** ...  
...

**memory:** (init)@0, (D = 42)@1

**Idea 1:** out-of-order read execution  
by reading from message history

```

T1
store D := 42
store F := 1
  
```

→

```

T2
r1 = load F      //reads 1
if (r1 == 1) then
  r2 = load D    //reads 0
else
  ..
  
```

**T1.regs:** ...

T1.promises: ...

...

**T2.regs:** r1 ↦ 0@0, r2 ↦ 0@0 ...

T2.promises: ...

...

**memory:** (init)@0, (D = 42)@1, (F = 1)@2

**Idea 1:** out-of-order read execution  
by reading from message history

**T1**

store D := 42

store F := 1

**T2**

r1 = load F //reads 1

if (r1 == 1) then

    r2 = load D //reads 0

else

..

**T1.regs:** ...

T1.promises: ...

...

**T2.regs:** r1 ↦ 1<sub>00</sub>, r2 ↦ 0<sub>00</sub> ...

T2.promises: ...

...

**memory:** (init)<sub>00</sub>, (D = 42)<sub>01</sub>, (F = 1)<sub>02</sub>

**Idea 1: out-of-order read execution**  
by reading from message history

**T1**

```
store D := 42
store F := 1
```

**T2**

```
r1 = load F      //reads 1
if (r1 == 1) then
  r2 = load D    //reads 0
else
  ..
```

**T1.regs:** ...

T1.promises: ...

...

**T2.regs:** r1 ↦ 1@0, r2 ↦ 0@0 ...

T2.promises: ...

...

**memory:** (init)@0, (D = 42)@1, (F = 1)@2

## Idea 1: out-of-order read execution by reading from message history

## Out-of-order reads

**T1**

store D := 42  
store F := 1

**T2**

r1 = load F //reads 1  
if (r1 == 1) then  
  r2 = load D //reads 0  
else  
  ..

**T1.regs:** ...

T1.promises: ...

...

**T2.regs:** r1 ↦ 1@0, r2 ↦ 0@0 ...

T2.promises: ...

...

**memory:** (init)@0, (D = 42)@1, (F = 1)@2

**Idea 1: out-of-order read execution**  
by reading from message history

**T1**

```
store D := 42
store F := 1
```

**T2**

```
r1 = load F      //reads 1
if (r1 == 1) then
  r2 = load D    //reads 0
else
  ..
```

**T1.regs:** ...

T1.promises: ...

...

**T2.regs:** r1 ↦ 1<sub>@0</sub>, r2 ↦ 0<sub>@0</sub> ...

T2.promises: ...

...

**memory:** (init)<sub>@0</sub>, (D = 42)<sub>@1</sub>, (F = 1)<sub>@2</sub>



## Idea 2: ordering reads with views.

A view is a timestamp of a “seen” write.

### T1

```
store D := 42
store F := 1
```

regstate: Reg  $\mapsto$  Val  $\times$  View

### T2

```
r1 = load F           //reads 1
if (r1 == 1) then
  r2 = load (D+r1-r1) //reads 0
else
  ..
```

address dependency



## Idea 2: ordering reads with views.

A view is a timestamp of a “seen” write.

regstate: Reg  $\mapsto$  Val  $\times$  View

**T1**

store D := 42

store F := 1

**T2**

r1 = load F //reads 1

if (r1 == 1) then

    r2 = load (D+r1-r1) //reads 0

else

..

**T1.reg:** ...

T1.promises: ...

...

**T2.reg:** r1  $\mapsto$  0@0, r2  $\mapsto$  0@0 ...

T2.promises: ...

...

**memory:** (init)@0, (D = 42)@1, (F = 1)@2

## Idea 2: ordering reads with views.

A view is a timestamp of a “seen” write.

regstate: Reg  $\mapsto$  Val  $\times$  View

**T1**

store D := 42

store F := 1

**T2**

r1 = load F //reads 1

if (r1 == 1) then

    r2 = load (D+r1-r1) //reads 0

else

..

**T1.reg:** ...

T1.promises: ...

...

**T2.reg:** r1  $\mapsto$  0@0, r2  $\mapsto$  0@0...

T2.promises: ...

...

**memory:** (init)@0, (D = 42)@1, (F = 1)@2

## Idea 2: ordering reads with views.

A view is a timestamp of a “seen” write.

regstate:  $\text{Reg} \mapsto \text{Val} \times \text{View}$

**T1**

store D := 42

store F := 1

**T2**

r1 = load F //reads 1

if (r1 == 1) then

    r2 = load (D+r1-r1) //reads 0

else

..

**T1.reg:** ...

T1.promises: ...

...

**T2.reg:** r1  $\mapsto$  1@2, r2  $\mapsto$  0@0...

T2.promises: ...

...

**memory:** (init)@0, (D = 42)@1, (F = 1)@2

## Idea 2: ordering reads with views.

A view is a timestamp of a “seen” write.

regstate: Reg  $\mapsto$  Val  $\times$  View

**T1**

store D := 42

store F := 1



**T2**

r1 = load F //reads 1

if (r1 == 1) then

    r2 = load (D+r1-r1) //reads 0

else

..



**T1.reg:** ...

T1.promises: ...

...

**T2.reg:** r1  $\mapsto$  1@2, r2  $\mapsto$  0@0...

T2.promises: ...

...

**memory:** (init)@0, (D = 42)@1, (F = 1)@2

## Idea 2: ordering reads with views.

A view is a timestamp of a “seen” write.

regstate: Reg  $\mapsto$  Val  $\times$  View

**T1**

store D := 42

store F := 1



**T2**

r1 = load F //reads 1

if (r1 == 1) then

    r2 = load (D+r1-r1) //reads 0

else

..



**T1.reg:** ...

T1.promises: ...

...

**T2.reg:** r1  $\mapsto$  1@2, r2  $\mapsto$  0@0...

T2.promises: ...

...

**memory:** (init)@0, (D = 42)@1, (F = 1)@2

### Idea 3: out-of-order writes with promises.

A thread can promise any write at any time, if it can later fulfil the promise.

**T1**

```
store D := 42  
store F := 1
```

**T2**

```
r1 = load F           //reads 1  
if (r1 == 1) then  
    r2 = load (D+r1-r1) //reads 0  
else  
    ..
```

## Idea 3: out-of-order writes with promises.

A thread can promise any write at any time, if it can later fulfil the promise.

→ **T1**  
 store D := 42  
 store F := 1

→ **T2**  
 r1 = load F //reads 1  
 if (r1 == 1) then  
 r2 = load (D+r1-r1) //reads 0  
 else  
 ..

**T1.regs:** ...

**T1.promises:** ∅

...

**T2.regs:** r1 ↦ 0@0, r2 ↦ 0@0 ...

**T2.promises:** ∅

...

**memory:** (init)@0



## Idea 3: out-of-order writes with promises.

A thread can promise any write at any time, if it can later fulfil the promise.

→ **T1**  
 store D := 42  
 store F := 1

→ **T2**  
 r1 = load F //reads 1  
 if (r1 == 1) then  
   r2 = load (D+r1-r1) //reads 0  
 else  
 ..

**T1.regs:** ...

**T1.promises:** (F = 1)<sub>@1</sub>

...

**T2.regs:** r1 ↦ 0<sub>@0</sub>, r2 ↦ 0<sub>@0</sub> ...

**T2.promises:** ∅

...

**memory:** (init)<sub>@0</sub>, (F = 1)<sub>@1</sub>

## Idea 3: out-of-order writes with promises.

A thread can promise any write at any time, if it can later fulfil the promise.

**T1**  
 → store D := 42  
 store F := 1

**T2**  
 → r1 = load F //reads 1  
 if (r1 == 1) then  
     r2 = load (D+r1-r1) //reads 0  
 else  
 ..

**T1.regs:** ...

**T1.promises:** (F = 1)<sub>@1</sub>

...

**T2.regs:** r1 ↦ 0<sub>@0</sub>, r2 ↦ 0<sub>@0</sub> ...


**T2.promises:** ∅

...

**memory:** (init)<sub>@0</sub>, (F = 1)<sub>@1</sub>, (D = 42)<sub>@2</sub>

## Idea 3: out-of-order writes with promises.

A thread can promise any write at any time, if it can later fulfil the promise.



**T1.regs:** ...  
**T1.promises:**  $\emptyset$   
 ...

**T2.regs:**  $r1 \mapsto 0@0, r2 \mapsto 0@0 \dots$   
**T2.promises:**  $\emptyset$   
 ...

**memory:** (init) $@0$ , (F = 1) $@1$ , (D = 42) $@2$

## Idea 3: out-of-order writes with promises.

A thread can promise any write at any time, if it can later fulfil the promise.

**T1**  
 store D := 42  
 → store F := 1

**T2**  
 → r1 = load F //reads 1  
 if (r1 == 1) then  
   r2 = load (D+r1-r1) //reads 0  
 else  
 ..

**T1.regs:** ...  
**T1.promises:** ∅  
 ...

**T2.regs:** r1 ↦ 1@1, r2 ↦ 0@0 ...  
**T2.promises:** ∅  
 ...

**memory:** (init)@0, (F = 1)@1, (D = 42)@2

## Idea 3: out-of-order writes with promises.

A thread can promise any write at any time, if it can later fulfil the promise.

**T1**  
 store D := 42  
 → store F := 1

**T2**  
 r1 = load F //reads 1  
 → if (r1 == 1) then  
   r2 = load (D+r1-r1) //reads 0  
 else  
 ..

**T1.regs:** ...  
**T1.promises:** ∅  
 ...

**T2.regs:** r1 ↦ 1@1, r2 ↦ 0@0 ...  
**T2.promises:** ∅  
 ...

**memory:** (init)@0, (F = 1)@1, (D = 42)@2

## Full model details include

- **write ordering:** with views  $T.V_{wm}, T.V_{wp}$
- **certification:** thread-locally prevent executions with unfulfilled promises
- **coherence**  $T.coh$
- **write forwarding**  $T.fwdb$
- **barriers, release/acquire:** uniformly with views  $T.V_{rm}, T.V_{rp}, T.V_{wm}, T.V_{wp}, T.V_{rel}$
- **load/store exclusive instructions**  $T.exclb$

## Executable tool

---

### Executable tool for interactive and exhaustive execution of ARMv8 or RISC-V concurrent assembly programs.

Building on Sail, Sail ISA models, rmem

[Sarkar et al 2011/12, Gray et al 2015, Flur et al 2016/17, Pulte et al 2018, Armstrong et al 2019]

- **exhaustively:** enumerate all possible final states allowed by model, each with witnessing trace
- **interactively:** interactively replay witnessing trace for debugging



## Optimisation 1: promises first (optimisation proved in Coq)

For every trace  $tr$ , there is an equivalent trace  $tr'$  s.th:

$$tr' = \underbrace{- \xrightarrow{t_1} - \xrightarrow{t_2} - \xrightarrow{t_3} - \dots}_{promise} \underbrace{- \xrightarrow{t_n} - \xrightarrow{t_{n+1}} - \xrightarrow{t_{n+2}} - \dots}_{non-promise}$$

## Optimisation 2: explore final thread states in parallel

## Run text book concurrent data structure/lock implementations:

1. **write data structure in C++ or Rust**  
write test code, testing operations and logging data
2. **compile with gcc/llvm -O3** to ARMv8 assembly,  
map assembly into litmus format with script
3. **run tool to enumerate final states**



Test	Promising	Flat
Linux Spinlock (asm)-7	0.61	9108.53
Linux Spinlock (cpp)-3	6.58	1472.74
Linux Spinlock (Rust)-3	4.88	52.52
Single-producer-single-consumer-3-3	1.36	249.26
Single-producer-multiple-consumers-3-3-3	71.12	ooT
Ticket lock/(opt)-3	18.08 / 20.13	ooT / ooT
Treiber stack (cpp)/(opt)-100-010-010	0.42 / 0.42	2144.52 / 5943.50
Treiber stack (cpp)/(opt)-100-100-010	8.70 / 8.70	ooT / ooT
Treiber stack (cpp)/(opt)-210-011-000	615.41 / 637.98	ooT / ooT
Treiber stack (Rust)-100-010-010	0.39	77.21
Treiber stack (Rust)-100-100-010	7.30	8940.03
Treiber stack (Rust)-210-011-000	522.19	ooT
Chase-Lev dequeue/(opt)-100-1-0	0.30 / 0.30	2.93 / 2.97
Chase-Lev dequeue/(opt)-110-1-0	0.44 / 0.44	1042.88 / 1114.39
Chase-Lev dequeue/(opt)-211-2-1	28.55 / 111.54	ooT / ooT
Chase-Lev dequeue/(opt)-100-000-000	1.34 / 2.95	2983.11 / ooT
Chase-Lev dequeue/(opt)-100-010-000	2.55 / 5.66	ooT / ooT
Chase-Lev dequeue/(opt)-110-100-010	2108.12 / ooT	ooT / ooT

ooT: out of time

### Promising-ARM/RISC-V:

- simpler operational concurrency model
- interactive checking and debugging tool
- promising performance results

#### **Promising and other rmem models:**

`https://github.com/rem-s-project/rmem`

#### **Sail and Sail models:**


`https://github.com/rem-s-project/sail`

#### **try model in online interface:**

`https://www.cl.cam.ac.uk/~pes20/rmem`

**Thank you**

## Why ARMv8/RISC-V?



### Language concurrency models:

no accepted thin-air free semantics

### Machine concurrency models:

better understood,  
compiler-independent tools  
that support hand-written assembly

**ARMv8:** widely used, subtle concurrency semantics

**RISC-V:** recently adopted similar model

## Write ordering with views

**T1**

```
store D := 42  
store[rel] F := 1
```

## Out-of-order writes

**T2**

```
r1 = load F //reads 1  
if (r1 == 1) then  
    r2 = load (D+r1-r1) //reads 0  
else  
    ..
```



## Write ordering with views

## Out-of-order writes

→ **T1**  
store D := 42  
store[rel] F := 1

→ **T2**  
r1 = load F //reads 1  
if (r1 == 1) then  
    r2 = load (D+r1-r1) //reads 0  
else  
    ..

**T1.regs:** ...

**T1.promises:** ∅

**T1.v<sub>wm</sub>:** 0

...

**T2.regs:** ...

**T2.promises:** ...

**T2.v<sub>wm</sub>:** ...

...

**memory:** (init)@0

## Write ordering with views

## Out-of-order writes

→ **T1**  
store D := 42  
store[rel] F := 1

→ **T2**  
r1 = load F //reads 1  
if (r1 == 1) then  
    r2 = load (D+r1-r1) //reads 0  
else  
    ..

**T1.regs:** ...

**T1.promises:** (F = 1)@1

**T1.v<sub>wm</sub>:** 0

...

**T2.regs:** ...

**T2.promises:** ...

**T2.v<sub>wm</sub>:** ...

...

**memory:** (init)@0, (F = 1)@1

## Write ordering with views

## Out-of-order writes

**T1**  
→ store D := 42  
store[rel] F := 1

→ **T2**  
r1 = load F //reads 1  
if (r1 == 1) then  
    r2 = load (D+r1-r1) //reads 0  
else  
    ..

**T1.regs:** ...

**T1.promises:** (F = 1)@1

**T1.v<sub>wm</sub>:** 2

...

**T2.regs:** ...

**T2.promises:** ...

**T2.v<sub>wm</sub>:** ...

...

**memory:** (init)@0, (F = 1)@1, (D = 42)@2

## Write ordering with views

## Out-of-order writes

**T1**  
→ store D := 42  
store[rel] F := 1

**T2**  
→ r1 = load F //reads 1  
if (r1 == 1) then  
    r2 = load (D+r1-r1) //reads 0  
else  
    ..

**T1.regs:** ...

**T1.promises:** (F = 1)@1

**T1.v<sub>wm</sub>:** 2

...

**T2.regs:** ...

**T2.promises:** ...

**T2.v<sub>wm</sub>:** ...

...

**memory:** (init)@0, (F = 1)@1, (D = 42)@2

Prevent executions with unfulfilled promises  
with certification.

For every step by thread T, do simple **thread-local** check: ensure  
**there exists trace by T executing in program-order,  
alone, under current memory fulfilling all its promises.**