

Operational Semantics of a Weak Memory Model with Channel Communication

Daniel Fava, Martin Steffen, Volker Stolz

1st March 2018



In this talk

- *operational* semantics for a WMM
 - inspired by Go
 - channel communication
 - based on happens-before
- proof of basic *correctness* property
- executable within the \mathbb{K} rewriting framework



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References


Memory model

MCM

A specification what to expect from from a shared mamory, what may be observed (by reads) and what not.

Rest

- bottom-line: *sequential consistency* Lamport (interleaving of reads and writes),
- *weak* or *relaxed*: basically weaker than that.



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus


Correctness

Conclusion

References

How to specify a MM

- in prose (as in <https://golang.org/ref/mem>)
- litmus tests
- **axiomatic** (candidate executions)
- **operational** (SOS)



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References

Go sales pitch



- “language for the 21st century”
- relatively new language (with some not so new features?)
- a lot of *fanfare* & backed by Google no less
- existing show-case applications
 - docker
 - dropbox ...



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References

Go's stated design principles

- appealing to C programmers
- KISS: “keep it simple, stupid”
- built-in **concurrency**
- “strongly typed”
- efficient
- fast *compilation*, appealing for scripting



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References

Go's non-revolutionary feature mix

- imperative
- object-oriented (?)
- compiled
- concurrent (goroutines)
- “strongishly” typed
- garbage collected
- portable
- higher-order functions and closures



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References



- simple concurrent calculus with “goroutines”
- A-normal form
- channels:
 - dynamically created
 - “higher-order” channels (à la π ...)
 - *bounded* channels
 - mixed choice
 - with “channel guards”
 - default-clause



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References

Syntax

$v ::= r \mid \underline{n}$	values
$e ::= t \mid v \mid \text{load } z \mid z := v \mid \text{if } v \text{ then } t \text{ else } t \mid \text{go } t$ $\mid \text{make } (\text{chan } T, v) \mid \leftarrow v \mid v \leftarrow v \mid \text{close } v$	expression
$g ::= v \leftarrow v \mid \leftarrow v \mid \text{default}$	guards
$t ::= \text{let } r = e \text{ in } t \mid \sum_i \text{let } r_i = g_i \text{ in } t_i$	threads

- \sum : choice (select, case, default)

Go's concurrency model

- **only** sync-primitive: *channel* communication (but read the fine-print)
- shared variable communication possible \Rightarrow
- simple **happens-before** memory model

Mantra

Don't communicate by sharing memory; share memory by communicating. (R. Pike)

Rest

- straightforward and simple model, still they advise:

"If you must read the rest of this document [= the Go MM] to understand the behavior of your program, you are being too clever. Don't be clever."



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References

Happens-before

- dates back to Lamport
- “unrelated” to actually “happening before”

Observational + “liberal”

a read **can observe** a write W **unless**

1. read **definitely** “too late”
2. a different write **definitely** “overwrites” W (shadows)



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References

Nature of synchronization

- generally:¹ restricting otherwise possible *interleavings*
- in connection with shared memory: intuition often (cf. *write buffers*)

*“Data Memory Barrier (DMB). This forces all earlier-in-program-order memory accesses to **become** globally **visible** before any subsequent accesses.” (random quote, some ARM programmer’s guide)*



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References

¹independent from shared memory

Nature of synchronization

- generally:¹ restricting otherwise possible *interleavings*
- in connection with shared memory: intuition often (cf. *write buffers*)

*“Data Memory Barrier (DMB). This forces all earlier-in-program-order memory accesses to **become** globally **visible** before any subsequent accesses.” (random quote, some ARM programmer’s guide)*

Happens-before

synchronization = making things **INVISIBLE**

¹independent from shared memory



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References

Two(*) ingredients for HB only

1. *program* order
2. channel communication
 - 2.1 sending \rightarrow_{hb} receiving
 - 2.2 full buffer

Sends and receives

- A *send* on a channel happens-before the corresponding *receive* from that channel completes.
- The i th *receive* on a channel with capacity k happens-before the $i + k$ th *send* from that channel completes.

Rest

- channel close, init, thread creation, packages, locks, once



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References

Operational semantics (weak)



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Configuration

$$P ::= n\langle\sigma, t\rangle \mid n(z:=v) \mid \bullet \mid P \parallel P \mid n[q] \mid \nu n P. \quad (1)$$

Title

Introduction

Calculus

Correctness

Conclusion

References

Operational semantics (weak)



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Configuration

$$P ::= n\langle\sigma, t\rangle \mid n(z:=v) \mid \bullet \mid P \parallel P \mid n[q] \mid \nu n P. \quad (1)$$

thread

$$n\langle\sigma, t\rangle$$

Title

Introduction

Calculus

Correctness

Conclusion

References

Operational semantics (weak)



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Configuration

$$P ::= n\langle\sigma, t\rangle \mid n(z:=v) \mid \bullet \mid P \parallel P \mid n[q] \mid \nu n P. \quad (1)$$

thread

$$n\langle\sigma, t\rangle$$

channel

$$n[q]$$

Title

Introduction

Calculus

Correctness

Conclusion

References

Operational semantics (weak)



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Configuration

$$P ::= n\langle\sigma, t\rangle \mid n(z:=v) \mid \bullet \mid P \parallel P \mid n[q] \mid \nu n P. \quad (1)$$

thread

 $n\langle\sigma, t\rangle$

write event

 $n(z:=v)$

channel

 $n[q]$

Title

Introduction

Calculus

Correctness

Conclusion

References

Write & read steps

$$\frac{\dots}{p\langle\sigma, z := v; t\rangle \rightarrow p\langle\sigma', t\rangle \parallel n(z:=v)}$$

$$\frac{\dots}{p\langle\sigma, \text{let } r = \text{load } z \text{ in } t\rangle \parallel n(z:=v) \rightarrow p\langle\sigma, \text{let } r = v \text{ in } t\rangle \parallel n(z:=v)}$$

Synchronization = making things unobservable

- reads and writes: no **synchronization**
- **program order**
 - the only component of happens-before
 - for $x:=1; x:=2$: value 1 unobservable

but only locally

- channel communication; only (interesting) means of synchronization



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References

Synchronization = making things unobservable

- reads and writes: no **synchronization**
- **program order**
 - the only component of happens-before
 - for $x:=1; x:=2$: value 1 unobservable

but only locally

- channel communication; only (interesting) means of synchronization

Channel communication

- send the communicated value from sender to receiver



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References

Synchronization = making things unobservable

- reads and writes: no **synchronization**
- **program order**
 - the only component of happens-before
 - for $x:=1; x:=2$: value 1 unobservable

but only locally

- channel communication; only (interesting) means of synchronization

Channel communication

- send the communicated value from sender to receiver
- inform receiver of local knowledge of **UNobservable write events**



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References

Shadow sets and local information

- every write even: unique identifier

thread local information

- which events are locally **known** to be unobservable (**shadowed**)
- which events are locally **known** to have **happened-before** (at the current point)

Rest

$$n\langle\sigma, t\rangle$$

- local “state” tuple (E_{hb}, E_s) , $\sigma : 2^{(N \times X)} \times 2^N$.



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References

Read & write once more

- of course: shadow sets used to make writes invisible
- local update of “program order”

$$\sigma = (E_{hb}, E_s) \quad \sigma' = (E_{hb} + (\mathbf{n}, \mathbf{z}), E_s + \mathbf{E}_{hb}(\mathbf{z}))$$

$$p\langle \sigma, z := v; t \rangle \rightarrow p\langle \sigma', t \rangle \parallel n(z:=v)$$

$$\sigma = (_, E_s) \quad \mathbf{n} \notin \mathbf{E}_s$$

$$p\langle \sigma, \text{let } r = \text{load } z \text{ in } t \rangle \parallel n(z:=v) \rightarrow p\langle \sigma, \text{let } r = v \text{ in } t \rangle \parallel n(z:=v)$$

Channel communication

- sending values + knowledge about σ

$$\frac{\neg \text{closed}(c_f[q_2]) \quad \sigma' = \sigma}{p\langle \sigma, c \leftarrow v; t \rangle \parallel c[q_2] \rightarrow p\langle \sigma', t \rangle \parallel c[(v, \sigma) :: q_2]} \text{R-SEND}$$

$$\frac{v \neq \perp \quad \sigma' = \sigma + \sigma''}{c_b[q_1] \parallel p\langle \sigma, \text{let } r = \leftarrow c \text{ in } t \rangle \parallel c_f[q_2 :: (v, \sigma'')] \rightarrow c_b[\sigma :: q_1] \parallel p\langle \sigma', \text{let } r = v \text{ in } t \rangle \parallel c_f[q_2]} \text{R-REC}$$

Bounded channels

- A *send* on a channel happens-before the corresponding *receive* from that channel completes.
- The i th *receive* on a channel with capacity k happens-before the $i + k$ th *send* from that channel completes.



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References

Bounded channels

- A *send* on a channel happens-before the corresponding *receive* from that channel completes.
- The i th *receive* on a channel with capacity k happens-before the $i + k$ th *send* from that channel completes.

Bounded channels

- There is also a “backward synchronization”
 - from an “earlier” receive to a sender

Rest <2>

- *forward channel* (as shown)
- *backward channel*, propagating local σ knowledge



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References

Send and receive

$$\frac{\neg \text{closed}(c_f[q_2]) \quad \sigma' = \sigma + \sigma''}{c_b[q_1 :: \sigma''] \parallel p\langle \sigma, c \leftarrow v; t \rangle \parallel c_f[q_2] \rightarrow c_b[q_1] \parallel p\langle \sigma', t \rangle \parallel c_f[(v, \sigma) :: q_2]} \text{R-SEND}$$

$$\frac{v \neq \perp \quad \sigma' = \sigma + \sigma''}{c_b[q_1] \parallel p\langle \sigma, \text{let } r = \leftarrow c \text{ in } t \rangle \parallel c_f[q_2 :: (v, \sigma'')] \rightarrow c_b[\sigma :: q_1] \parallel p\langle \sigma', \text{let } r = v \text{ in } t \rangle \parallel c_f[q_2]} \text{R-REC}$$

Delayed reads

- so far: delayed or buffered writes
- also *delayed reads* (load buffers) possible \Rightarrow “read events”

More (and more complex) “events”

$$m(\sigma, z := n_1)_p \quad \text{and} \quad m[\sigma, ?n_4]_p$$

Rest

- chain of “future references”
- symbolic execution
- nota bene:
 - the write itself is not “delayed”
 - it’s the negative information (invisibility of other writes via the shadow sets, that travels slow)



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

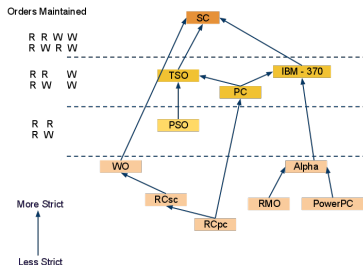
Correctness

Conclusion

References

Memory models

- WMMs, it's a jungle
- out-of-thin air
 - should be avoided (or should it?)
 - not even crystal clear what it is.



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References

... but there's a bottom line

No matter how “relaxed” you want your memory model one thing is non-negotiable:

DRF-SF

Data-race free programs have to be sequentially consistent
Manson et al. [8]



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction


Calculus

Correctness

Conclusion

References

Simulation



“weak simulates strong”

sure thing



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion


References

Simulation



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz



Title

Introduction

Calculus

Correctness

Conclusion

References

“weak simulates strong”

sure thing

“strong simulates weak”


definitely not

Simulation



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz



Title

Introduction

Calculus

Correctness

Conclusion

References

“weak simulates strong”

sure thing

“strong simulates weak”

conditionally, for RF programs

- “simultaneous” access to a shared location, where at least one is a write access

Manifest race (case W/W)

config C with

$$C \xrightarrow{p_1(z!)} \rightarrow_s \xrightarrow{p_2(z!)} \rightarrow_s$$

Rest

- race: reachable configuration with manifest race



Core of the proof

abstraction function/relation

“weak config” \rightarrow “strong config”

Rest

- problem: configs contains “alternatives”

$$n_1(\lfloor z:=v_1 \rfloor) \parallel n_2(\lfloor z:=v_2 \rfloor)$$

- strong semantics: exactly **one** value of z



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References

From local to global view \Rightarrow consensus



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Lemma (Consensus possible)

Weak configurations obey the following invariant

$$\bigcap_{p \in P} W_P^o(z @ p) \neq \emptyset . \quad (2)$$

- adding also *read events* to configurations

Title

Introduction

Calculus

Correctness

Conclusion

References

RF programs \Rightarrow stronger consensus



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Lemma (Race-free consensus when it counts)

Assume $P_0 \xrightarrow{*_w} P$ with P_0 race-free. If $P \xrightarrow{p(z?)}_w$ or $P \xrightarrow{p(z!)}_w$, then

$$\bigcap_{p_i} W_P^o(z@p_i) = \{n\}, \quad (3)$$

where the intersection ranges over an arbitrary set of processes which includes p .

Lemma (Race-free consensus)

Weak configurations for race-free programs obey the following invariant

$$\bigcap_{p_i \in P} W_P^o(z@p_i) = \{n\}. \quad (4)$$

Title

Introduction


Calculus

Correctness

Conclusion

References

Conditional simulation



- augment the configuration with additional read-events
- \Rightarrow consensus lemmas
- \Rightarrow DRF-SC



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References

- rewrite-based engine
- used variously for executable semantics
 - C++ memory models
 - “Ethereum” smart contracts platform
 - ...
- see <https://github.com/dfava/mmgo>



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References

Receive in \mathbb{K}

R-Rec

$$v \neq \perp \quad \sigma' = \sigma + \sigma''$$

$$\begin{array}{l} c_b[q_1] \parallel p\langle\sigma, \text{let } r = \leftarrow c \text{ in } t\rangle \parallel c_f[q_2 :: (v, \sigma'')] \rightarrow \\ c_b[\sigma :: q_1] \parallel p\langle\sigma', \text{let } r = v \text{ in } t\rangle \parallel c_f[q_2] \end{array}$$

Rest

```
rule <goroutine>
  <k> <- channel(Ref:Int) => V ... </k>
  <sigma>
    <HB> HMap:Map => mergeHB(HMap, HMapDP) </HB>
    <S> SSet:Set => SSet SSetDP </S>
  </sigma>
  <id> _ </id>
</goroutine>
<chan>
  <ref> Ref </ref>
  <type> _ </type>
  <forward> ListItem( ListItem(V)
    ListItem(HMapDP)
    ListItem(SSetDP) ) => .List </forward>
  <backward> BQ:List => ListItem( ListItem(HMap)
    ListItem(SSet)) BQ </backward>
</chan>
requires notBool( V ==K $eot )
```

Related work

- loads of material on *axiomatic* semantics
- operational:
 - Boudol and Petri based on *rewriting theory*
 - Kang et al.: “promising” semantics with “*clocks*”
 - Flanagan and Freund: *adversarial* memory
 - Demange et al. Plan B (Java buffered write semantics BMM)
 - Pichon-Pharabod and Sewell: operational semantics avoiding OOTA
 - Alrahman et al.
 - Matthias Perner et al: parametrized semantics for NI (earlier today)
 - ...



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References

Conclusion

- formalizing WMM for some calculus with channel
- DRF-SC simulation proof
- read delays under work



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Title

Introduction

Calculus

Correctness

Conclusion

References

References I



Operational
Semantics of a
Weak Memory
Model with
Channel
Communication

Daniel Fava,
Martin Steffen,
Volker Stolz

Bibliography

- [1] Alrahman, Y. A., Andric, M., Beggiato, A., and Lluch-Lafuente, A. (2014). Can we efficiently check concurrent programs under relaxed memory models in Maude? In Escobar, S., editor, *Rewriting Logic and Its Applications – 10th International Workshop, WRLA 2014, Held as a Satellite Event of ETAPS, Grenoble, France, April 5-6, 2014, Revised Selected Papers*, volume 8663 of *Lecture Notes in Computer Science*, pages 21–41. Springer Verlag.
- [2] Boudol, G. and Petri, G. (2009). Relaxed memory models: An operational approach. In *Proceedings of POPL '09*, pages 392–403. ACM.
- [3] Demange, D., Laporte, V., Zhao, L., Jagannathan, S., Pichardie, D., and Vitek, J. (2013). Plan B: A buffered memory model for Java. In *Proceedings of POPL '13*, pages 329–342. ACM.
- [4] Flanagan, C. and Freund, S. N. (2010). Adversarial memory for detecting destructive races. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*. ACM.
- [5] Kang, J., Hur, C., Lahav, O., Vafeiadis, V., and Dreyer, D. (2017). A promising semantics for relaxed-memory concurrency. In Castagna, G. and Gordon, A. D., editors, *Proceedings of POPL '17*, pages 175–189. ACM.
- [6] Lamport, L. (1978). Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558–565.
- [7] Lamport, L. (1979). How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Transactions on Computers*, C-28(9):690–691.
- [8] Manson, J., Pugh, W., and Adve, S. V. (2005). The Java memory memory. In *Proceedings of POPL '05*. ACM.
- [9] Pichon-Pharabod, J. and Sewell, P. (2016). A concurrency-semantics for relaxed atomics that permits optimisation and avoids out-of-thin-air executions. In *Proceedings of POPL '16*. ACM.

Title

Introduction

Calculus

Correctness

Conclusion

References