

Towards a separation logic for Multicore OCaml

Glen Mével, Jacques-Henri Jourdan, François Pottier

May 25, 2020

PPS seminar, Paris

CNRS & Inria, Paris, France

The weak memory model

Multicore OCaml

Extension of the OCaml language with **multicore programming**.

Research project at OCaml Labs (Cambridge), will be merged eventually.

Strengths:

- brings multicore abilities to a functional, statically typed, memory-safe programming language;
- (gives the programmer a simpler memory model than that of C11, hopefully;)
- limited performance drop for sequential code.

Goals of this PhD:

- Build a proof system for Multicore OCaml programs.
- Prove interesting concurrent data structures.

Sequential consistency

Consider this concurrent program:

```
x := 0
y := 0
x := 1 || y := 1
A := y || B := x
```

Possible outcomes (A, B): (0, 1), (1, 0), (1, 1).

Observed (M.OCaml on 2-core x86-64): 91%, 9%, 0.001%.

Sequential consistency

Consider this concurrent program:

```
x := 0
y := 0
x := 1 || y := 1
A := y || B := x
```

Possible outcomes (A, B): (0, 1), (1, 0), (1, 1).

Observed (M.OCaml on 2-core x86-64): 91%, 9%, 0.001%.

Sequential consistency

Consider this concurrent program:

```
x := 0
y := 0
x := 1 || y := 1
A := y || B := x
```

Possible outcomes (A, B): (0, 1), (1, 0), (1, 1).

Observed (M.OCaml on 2-core x86-64): 91%, 9%, 0.001%.

Sequential consistency

Consider this concurrent program:

```
x := 0
y := 0
x := 1  || y := 1
A := y  || B := x
```

Possible outcomes (A, B): (0, 1), (1, 0), (1, 1).

Observed (M.OCaml on 2-core x86-64): 91%, 9%, 0.001%.

A weaker memory model

Consider this concurrent program:

```
x := 0
y := 0
x := 1 ← y := 1
A := y — B := x
```

Possible outcomes (A, B): (0, 1), (1, 0), (1, 1), (0, 0).

Observed (M.OCaml on 2-core x86-64): 91%, 9%, 0.001%, 0.1%.

A weaker memory model

Consider this concurrent program:

```
x := 0
y := 0
x := 1 ← y := 1
A := y — B := x
```

Possible outcomes (A, B): (0, 1), (1, 0), (1, 1), (0, 0).

Observed (M.OCaml on 2-core x86-64): 91%, 9%, 0.001%, 0.1%.

The compiler may reorder a write after a read.
(The processor too.)

Weak memory models

Sequential consistency is unrealistic.

We need a **weaker memory model**, where different threads have different **views** of the shared state.

The model should be specific to our language.

Existing works: Java (2000s), C11 (2010s; also Rust).

Candidate model for Multicore OCaml:

Dolan, Sivaramakrishnan, Madhavapeddy.
Bounding Data Races in Space and Time.
PLDI 2018.

Two access modes: non-atomic, atomic.

An operational model for Multicore OCaml: non-atomics

$$\begin{array}{l} x := x_1 \\ A := !y \end{array} \parallel \begin{array}{l} y := y_1 \\ B := !x \end{array}$$

An operational model for Multicore OCaml: non-atomics


$$\begin{array}{l} x := x_1 \quad \parallel \quad y := y_1 \\ A := !y \quad \parallel \quad B := !x \end{array}$$

Each **non-atomic** location has a **history**, *i.e.* a map from timestamps to values (timestamps are per location).

x : x_0

y : y_0


An operational model for Multicore OCaml: non-atomics

$$\begin{array}{l} x := x_1 \quad \parallel \quad y := y_1 \\ A := !y \quad \parallel \quad B := !x \end{array}$$


Each **non-atomic** location has a **history**, *i.e.* a map from timestamps to values (timestamps are per location).

Each thread has its own **view** of the non-atomic store, *i.e.* a map from non-atomic locations to timestamps.


An operational model for Multicore OCaml: non-atomics

$$\begin{array}{l} x := x_1 \\ A := !y \end{array} \parallel \begin{array}{l} y := y_1 \\ B := !x \end{array}$$


non-atomic write

- Timestamp must be fresh.
- Timestamp must be newer than current thread's view.
- Current thread's view is updated.

An operational model for Multicore OCaml: non-atomics


$$\begin{array}{l} x := x_1 \longrightarrow y := y_1 \\ A := !y \parallel B := !x \end{array}$$


non-atomic write

- Timestamp must be fresh.
- Timestamp must be newer than current thread's view.
- Current thread's view is updated.

An operational model for Multicore OCaml: non-atomics

$x := x_1$ \parallel $y := y_1$
 $A := !y$ \parallel $B := !x$




non-atomic read

- Returns **any** value at least as recent as current thread's view.
- Current thread's view is unchanged.

An operational model for Multicore OCaml: non-atomics

$x := x_1 \parallel y := y_1$
 $A := !y \rightarrow B := !x$



non-atomic read

- Returns **any** value at least as recent as current thread's view.
- Current thread's view is unchanged.

Non-atomic locations are useful for updating the state locally, but they don't provide synchronization.

Atomic locations allow the message-passing idiom.

An operational model for Multicore OCaml: atomics

```
x := x1 | REPEAT  
a :=at true |   C := !at a  
                | UNTIL C = true  
                | B := !x
```

thread B thread A

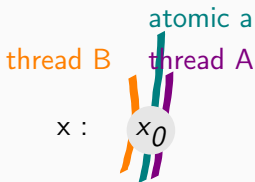
x :



a: false

An operational model for Multicore OCaml: atomics

```
x := x1
a :=at true
|| REPEAT
||   C := !at a
|| UNTIL C = true
|| B := !x
```



x :

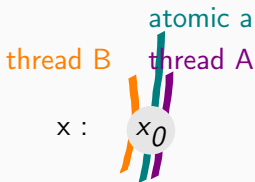
x₀

a : false

Each **atomic** location stores **one** value,
and **one view** of the non-atomic store.

An operational model for Multicore OCaml: atomics

```
x := x1
a :=at true
REPEAT
  C := !at a
UNTIL C = true
B := !x
```



x :

x₀

a : false

An operational model for Multicore OCaml: atomics

```
x := x1
```


```
a :=at true
```

```
REPEAT
```

```
  C := !at a
```

```
UNTIL C = true
```

```
  B := !x
```



```
x :
```

```
x0
```

```
x1
```

```
a : false
```

An operational model for Multicore OCaml: atomics

```
x := x1
```

```
a :=at true
```

```
REPEAT
```

```
  C := !at a
```

```
UNTIL C = true
```

```
  B := !x
```

thread B

x :

x₀

x₁

atomic a
thread A


a : true

atomic write


Merges the writer's view into the atomic location's view.

An operational model for Multicore OCaml: atomics

```
x := x1  
a :=at true
```



```
REPEAT  
  C := !_at a  
UNTIL C = true  
B := !x
```




a : true

atomic read

Merges the atomic location's view into the reader's view.

An operational model for Multicore OCaml: atomics

```
x := x1
a :=at true
REPEAT
  C := !at a
UNTIL C = true
B := !x
```



Our program logic

Rules of non-atomic locations

The predicate $x \mapsto v$ means that we own the non-atomic location x and that we have **seen** its latest value, which is v .

Non-atomic write:

$$\{x \mapsto v\}$$

$$x := v'$$

$$\{\lambda(). x \mapsto v'\}$$

Non-atomic read:

$$\{x \mapsto v\}$$

$$!x$$

$$\{\lambda v'. v' = v * x \mapsto v\}$$

Impact of the weak memory model on our CSL

Invariants are the mechanism by which threads can share propositions in a Concurrent Separation Logic such as Iris:

$$\frac{\{P * I\} e \{Q * I\} \quad e \text{ atomic}}{\boxed{I} \vdash \{P\} e \{Q\}}$$

The proposition $x \mapsto v$ is **subjective**: its truth depends on the thread's view of memory.

It is unsound to share it *via* an invariant.

Propositions which are true in all threads are called **objective**:

- “pure” facts, such as $v = 5$
- ghost state, such as $\gamma \hookrightarrow \circ 5$
- atomic state, such as $a \mapsto_{\text{at}} (v, \mathcal{V})$

Only objective propositions can be put in an invariant.

Rules of atomic locations (simplified)

The predicate $a \mapsto_{\text{at}} v$ means that we own the atomic location a , which stores the value v .

It is **objective**.

Atomic write:

$$\{a \mapsto_{\text{at}} v\}$$

$$a :=_{\text{at}} v'$$

$$\{\lambda(). a \mapsto_{\text{at}} v'\}$$

Atomic read:

$$\{a \mapsto_{\text{at}} v\}$$

$$!_{\text{at}} a$$

$$\{\lambda v'. v' = v * a \mapsto_{\text{at}} v\}$$

Rules of atomic locations (simplified)

The predicate $a \mapsto_{\text{at}} v$ means that we own the atomic location a , which stores the value v .

It is **objective**.

Atomic write:

$$\{a \mapsto_{\text{at}} v\}$$

$$a :=_{\text{at}} v'$$

$$\{\lambda(). a \mapsto_{\text{at}} v'\}$$

Atomic read:

$$\{a \mapsto_{\text{at}} v\}$$

$$!_{\text{at}} a$$

$$\{\lambda v'. v' = v * a \mapsto_{\text{at}} v\}$$

views

Views are ordered by inclusion.

The predicate $\uparrow \mathcal{V}$ means “the current thread’s view includes \mathcal{V} ”.

Rules of atomic locations

The predicate $a \mapsto_{\text{at}} (v, \mathcal{V})$ means that we own the atomic location a , which stores the value v and a view (at least) \mathcal{V} .

It is **objective**.

Atomic write:

$$\{a \mapsto_{\text{at}} (v, \mathcal{V}) * \uparrow \mathcal{V}'\}$$

$$a :=_{\text{at}} v'$$

$$\{\lambda(). a \mapsto_{\text{at}} (v', \mathcal{V}')\}$$

Atomic read:

$$\{a \mapsto_{\text{at}} (v, \mathcal{V})\}$$

$$!_{\text{at}} a$$

$$\{\lambda v'. v' = v * a \mapsto_{\text{at}} (v, \mathcal{V}) * \uparrow \mathcal{V}\}$$

views

Views are ordered by inclusion.

The predicate $\uparrow \mathcal{V}$ means “the current thread’s view includes \mathcal{V} ”.

Rules of atomic locations

The predicate $a \mapsto_{\text{at}} (v, \mathcal{V})$ means that we own the atomic location a , which stores the value v and a view (at least) \mathcal{V} .

It is **objective**.

Atomic write:

$$\{a \mapsto_{\text{at}} (v, \mathcal{V}) * \uparrow \mathcal{V}'\}$$

$$a :=_{\text{at}} v'$$

$$\{\lambda(). a \mapsto_{\text{at}} (v', \mathcal{V}' \sqcup \mathcal{V}) * \uparrow \mathcal{V}\}$$

Atomic read:

$$\{a \mapsto_{\text{at}} (v, \mathcal{V})\}$$

$$!_{\text{at}} a$$

$$\{\lambda v'. v' = v * a \mapsto_{\text{at}} (v, \mathcal{V}) * \uparrow \mathcal{V}\}$$

views

Views are ordered by inclusion.

The predicate $\uparrow \mathcal{V}$ means “the current thread’s view includes \mathcal{V} ”.

Propositions are monotonic

Subjective propositions are **monotonic** w.r.t. the thread's view.

One reason: the frame rule:

$$\{a \mapsto_{\text{at}} v * P\}$$

$$a :=_{\text{at}} v'$$

$$\{\lambda(). a \mapsto_{\text{at}} v' * P\}$$

Propositions are monotonic

Subjective propositions are **monotonic** w.r.t. the thread's view.

One reason: the frame rule:

$\{a \mapsto_{\text{at}} v * P \text{ this holds at the thread's current view}\}$

$a :=_{\text{at}} v'$

$\{\lambda(). a \mapsto_{\text{at}} v' * P \text{ this holds at the thread's now extended view}\}$

The message passing idiom

The objective proposition “ P at \mathcal{V} ” is the subjective proposition P seen at a fixed view \mathcal{V} .


$$P \iff \exists \mathcal{V}. (\uparrow \mathcal{V}) * (P \text{ at } \mathcal{V})$$

(\Rightarrow) If P holds now, then it holds at the current view.

(\Leftarrow) If P holds at some earlier view, then it holds now.

The message passing idiom

$$P \iff \exists \mathcal{V}. (\uparrow \mathcal{V}) * (P \text{ at } \mathcal{V})$$



The message passing idiom

$$P \iff \exists \mathcal{V}. (\uparrow \mathcal{V}) * (P \text{ at } \mathcal{V})$$

$\{a \mapsto_{\text{at}} (\text{false}, \emptyset) * P\}$

$a :=_{\text{at}} \text{true}$

$C := !_{\text{at}} a$

$\{a \mapsto_{\text{at}} (\text{true}, \mathcal{V}) * P\}$

The message passing idiom

$$P \iff \exists \mathcal{V}. (\uparrow \mathcal{V}) * (P \text{ at } \mathcal{V})$$

$\{a \mapsto_{\text{at}} (\text{false}, \emptyset) * P\}$

$\{a \mapsto_{\text{at}} (\text{false}, \emptyset) * P \text{ at } \mathcal{V} * \uparrow \mathcal{V}\}$

$a :=_{\text{at}} \text{true}$

$C := !_{\text{at}} a$

$\{a \mapsto_{\text{at}} (\text{true}, \mathcal{V}) * P\}$

The message passing idiom

$$P \iff \exists \mathcal{V}. (\uparrow \mathcal{V}) * (P \text{ at } \mathcal{V})$$

$\{a \mapsto_{\text{at}} (\text{false}, \emptyset) * P\}$

$\{a \mapsto_{\text{at}} (\text{false}, \emptyset) * P \text{ at } \mathcal{V} * \uparrow \mathcal{V}\}$

$a :=_{\text{at}} \text{true}$

$\{a \mapsto_{\text{at}} (\text{true}, \mathcal{V}) * P \text{ at } \mathcal{V}\}$

$C := !_{\text{at}} a$

$\{a \mapsto_{\text{at}} (\text{true}, \mathcal{V}) * P\}$

The message passing idiom

$$P \iff \exists \mathcal{V}. (\uparrow \mathcal{V}) * (P \text{ at } \mathcal{V})$$

$\{a \mapsto_{\text{at}} (\text{false}, \emptyset) * P\}$

$\{a \mapsto_{\text{at}} (\text{false}, \emptyset) * P \text{ at } \mathcal{V} * \uparrow \mathcal{V}\}$

$a :=_{\text{at}} \text{true}$

$\{a \mapsto_{\text{at}} (\text{true}, \mathcal{V}) * P \text{ at } \mathcal{V}\}$

This proposition is objective:
it can be put in an invariant.

$C := !_{\text{at}} a$

$\{a \mapsto_{\text{at}} (\text{true}, \mathcal{V}) * P\}$

The message passing idiom

$$P \iff \exists \mathcal{V}. (\uparrow \mathcal{V}) * (P \text{ at } \mathcal{V})$$

$\{a \mapsto_{\text{at}} (\text{false}, \emptyset) * P\}$

$\{a \mapsto_{\text{at}} (\text{false}, \emptyset) * P \text{ at } \mathcal{V} * \uparrow \mathcal{V}\}$

$a :=_{\text{at}} \text{true}$

$\{a \mapsto_{\text{at}} (\text{true}, \mathcal{V}) * P \text{ at } \mathcal{V}\}$

inv

$\{a \mapsto_{\text{at}} (\text{true}, \mathcal{V}) * P \text{ at } \mathcal{V}\}$

$C := !_{\text{at}} a$

$\{a \mapsto_{\text{at}} (\text{true}, \mathcal{V}) * P\}$

This proposition is objective:
it can be put in an invariant.

The message passing idiom

$$P \iff \exists \mathcal{V}. (\uparrow \mathcal{V}) * (P \text{ at } \mathcal{V})$$

$\{a \mapsto_{\text{at}} (\text{false}, \emptyset) * P\}$

$\{a \mapsto_{\text{at}} (\text{false}, \emptyset) * P \text{ at } \mathcal{V} * \uparrow \mathcal{V}\}$

$a :=_{\text{at}} \text{true}$

$\{a \mapsto_{\text{at}} (\text{true}, \mathcal{V}) * P \text{ at } \mathcal{V}\}$

inv

$\{a \mapsto_{\text{at}} (\text{true}, \mathcal{V}) * P \text{ at } \mathcal{V}\}$

$C := !_{\text{at}} a$

$\{a \mapsto_{\text{at}} (\text{true}, \mathcal{V}) * P \text{ at } \mathcal{V} * \uparrow \mathcal{V}\}$

$\{a \mapsto_{\text{at}} (\text{true}, \mathcal{V}) * P\}$

This proposition is objective:
it can be put in an invariant.

Example: spin lock

A spin lock implements a lock using an atomic boolean variable.

```
let rec acquire lk =
  if CAS lk false true
  then ()
  else acquire lk

let release lk =
  lk :={at} false
```

The invariant in the sequentially consistent model is:

$$\text{lockInv } lk \ P \triangleq \\ lk \mapsto_{\text{at}} \text{true} \ \vee \ (\quad lk \mapsto_{\text{at}} \text{false} \quad * \ P)$$

Example: spin lock

A spin lock implements a lock using an atomic boolean variable.

```
let rec acquire lk =
  if CAS lk false true
  then ()
  else acquire lk

let release lk =
  lk :={at} false
```

The invariant in the weak model is:

$$\text{lockInv } lk \ P \triangleq \\ lk \mapsto_{\text{at}} \text{true} \ \vee \ (\exists \mathcal{V}. lk \mapsto_{\text{at}} (\text{false}, \mathcal{V}) * P \text{ at } \mathcal{V})$$

Example: ticket lock

A ticket lock implements a lock using two atomic integer variables.

The invariant in the sequentially consistent model is:

$$\text{lockInv turn next } \gamma P \triangleq$$

$$\exists t, n.$$

$$\text{turn} \mapsto_{\text{at}} t$$

$$* \text{ next} \mapsto_{\text{at}} n$$

$$* (\text{ticket } \gamma t \vee (\text{locked } \gamma * P))$$

$$* \gamma \hookrightarrow (\bullet \dots)$$

Example: ticket lock

A ticket lock implements a lock using two atomic integer variables.

The invariant in the weak model is:

$$\text{lockInv turn next } \gamma P \triangleq$$

$$\exists t, n, \mathcal{V}.$$

$$\text{turn} \mapsto_{\text{at}} (t, \mathcal{V})$$

$$* \text{ next} \mapsto_{\text{at}} n$$

$$* (\text{ticket } \gamma t \vee (\text{locked } \gamma * P \text{ at } \mathcal{V}))$$

$$* \gamma \hookrightarrow (\bullet \dots)$$

Example: Dekker's mutual exclusion

Dekker's algorithm solves the mutual exclusion problem using three atomic variables.

The invariant and representation predicate in the SC model are:

$$\begin{aligned} \text{DekkerInv } \text{turn } \text{flag}_0 \text{ flag}_1 \gamma P &\triangleq \\ &\exists t, f_0, f_1, c_0, c_1. \\ &(\forall i \in \{0, 1\}. \text{flag}_i \mapsto_{\text{at}} f_i \quad * \gamma_i \hookrightarrow \bullet c_i) \\ &* ((\neg c_0 \wedge \neg c_1) \text{ -* } P) \\ &* \dots \\ \text{isDekker } i \gamma &\triangleq \\ &\gamma_i \hookrightarrow \circ \text{false} \end{aligned}$$

Example: Dekker's mutual exclusion

Dekker's algorithm solves the mutual exclusion problem using three atomic variables.

The invariant and representation predicate in the weak model are:

$$\text{DekkerInv } \text{turn } \text{flag}_0 \text{ flag}_1 \gamma \ P \triangleq$$

$$\exists t, f_0, f_1, c_0, c_1, \mathcal{V}_0, \mathcal{V}_1.$$

$$\begin{aligned} & (\forall i \in \{0, 1\}. \text{flag}_i \mapsto_{\text{at}} (f_i, \mathcal{V}_i) * \gamma_i \hookrightarrow \bullet c_i * \gamma'_i \hookrightarrow \bullet \mathcal{V}_i) \\ & * ((\neg c_0 \wedge \neg c_1) \text{ at } (\mathcal{V}_0 \sqcup \mathcal{V}_1)) \\ & * \dots \end{aligned}$$

$$\text{isDekker } i \ \gamma \triangleq$$

$$\exists \mathcal{V}. \gamma_i \hookrightarrow \circ \text{false} * \gamma'_i \hookrightarrow \circ \mathcal{V} * \uparrow \mathcal{V}$$

Model of the logic in Iris

Propositions are predicates on views:

$$\text{vProp} \triangleq \text{view} \longrightarrow \text{iProp}$$

$$\uparrow \mathcal{V}_0 \triangleq \lambda \mathcal{V}. \mathcal{V}_0 \sqsubseteq \mathcal{V}$$

$$P * Q \triangleq \lambda \mathcal{V}. P \mathcal{V} * Q \mathcal{V}$$

$$P \multimap Q \triangleq \lambda \mathcal{V}. \quad P \mathcal{V} \multimap Q \mathcal{V}$$

Model of the logic in Iris

Propositions are **monotonic** predicates on views:

$$\text{vProp} \triangleq \text{view} \xrightarrow{\text{mon}} \text{iProp}$$

$$\uparrow \mathcal{V}_0 \triangleq \lambda \mathcal{V}. \mathcal{V}_0 \sqsubseteq \mathcal{V}$$

$$P * Q \triangleq \lambda \mathcal{V}. P \mathcal{V} * Q \mathcal{V}$$

$$P \multimap Q \triangleq \lambda \mathcal{V}_1. \forall \mathcal{V} \sqsupseteq \mathcal{V}_1. P \mathcal{V} \multimap Q \mathcal{V}$$

Model of the logic in Iris

Propositions are **monotonic** predicates on views:

$$\text{vProp} \triangleq \text{view} \xrightarrow{\text{mon}} \text{iProp}$$

$$\uparrow \mathcal{V}_0 \triangleq \lambda \mathcal{V}. \mathcal{V}_0 \sqsubseteq \mathcal{V}$$

$$P * Q \triangleq \lambda \mathcal{V}. P \ \mathcal{V} * Q \ \mathcal{V}$$

$$P \multimap Q \triangleq \lambda \mathcal{V}_1. \forall \mathcal{V} \sqsupseteq \mathcal{V}_1. P \ \mathcal{V} \multimap Q \ \mathcal{V}$$

We equip a language-with-view with an operational semantics:

$$\text{exprWithView} \triangleq \text{expr} \times \text{view}$$

Iris builds a WP calculus for `exprWithView` in `iProp`.

We derive a WP calculus for `expr` in `vProp` and prove adequacy:

$$\text{WP } e \ \varphi \triangleq \lambda \mathcal{V}_1. \forall \mathcal{V} \sqsupseteq \mathcal{V}_1. \text{WP } \langle e, \mathcal{V} \rangle \ (\lambda \langle v, \mathcal{V}' \rangle. \varphi \ v \ \mathcal{V}')$$

where $\varphi : \text{val} \rightarrow \text{vProp}$

Plans for the future:

- Prove more elaborate shared data structures
 - e.g. bounded queues with a circular buffer
- Data races on non-atomics:
 - How to allow them?
 - What are they useful for?