

# Cosmo: a concurrent separation logic for Multicore OCaml

---

**Glen Mével**, Jacques-Henri Jourdan, François Pottier

August, 2020

ICFP, “New York”

LRI & Inria, Paris, France

## Our aim:

- Verifying
- fine-grained concurrent programs
- in the setting of Multicore OCaml's memory model.

**Our contribution:** A concurrent separation logic with views.

**Multicore OCaml:** OCaml language with multicore programming.

**Weak memory model for Multicore OCaml:**

- Formalized in PLDI 2018.
- Two flavours of locations: “atomic”, “non-atomic”.

(Also at ICFP 2020: *Retrofitting Parallelism onto OCaml*)

# In traditional fine-grained concurrent separation logics...

We can assert ownership of a location and specify its value:

$$\{x \mapsto 42\}$$
$$x := 44$$
$$\{x \mapsto 44\}$$

Ownership can be shared between all threads via an invariant:

$$\boxed{\exists n \in \mathbb{N}, x \mapsto n * n \text{ is even}} \vdash$$
$$\{\text{True}\}$$
$$x := 44$$
$$\{\text{True}\}$$

# The challenge: subjectivity

With weak memory, each thread has its own **view** of memory.

Some assertions are **subjective**:

- Their validity relies on the thread's view.

Invariants are **objective**:

- They cannot share subjective assertions.

**How to keep a simple and powerful enough logic?**

## Key idea: modeling subjectivity with views

A thread knows a subset  $\mathcal{U}$  of all writes to the memory.

- Affects how the thread interacts with memory.
- $\mathcal{U}$  is the thread's **view**.

New assertions:

- $\uparrow\mathcal{U}$  : we **have seen**  $\mathcal{U}$ , i.e. we know all writes in  $\mathcal{U}$ .
- $P @ \mathcal{U}$  : having seen  $\mathcal{U}$  is **objectively** enough for  $P$  to hold.

# Key idea: decomposing subjective assertions

Decompose subjective assertions:

$$P \iff \exists \mathcal{U}. \underbrace{P @ \mathcal{U}}_{\text{objective}} * \underbrace{\uparrow \mathcal{U}}_{\text{subjective}}$$

Share parts via distinct mechanisms:

- $P @ \mathcal{U}$  : via **objective invariants**, as usual.
- $\uparrow \mathcal{U}$  : via **synchronization** offered by the memory model.

## Our program logic

---



# Rules of atomic locations, simplified

$x \mapsto_{\text{at}} v$  : the **atomic** location  $x$  **stores** the value  $v$ .

- **Sequentially consistent.**
- **Objective.**
- Standard rules:

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

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

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

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

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

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

# Rules of non-atomic locations

$x \mapsto v$  : **we know** the latest value  $v$  of the **non-atomic** location  $x$ .

- **Relaxed**.
- **Subjective** — cannot appear in an invariant.
- Standard rules too!

$$\{x \mapsto v\}$$
$$x := v'$$
$$\{\lambda(). x \mapsto v'\}$$
$$\{x \mapsto v\}$$
$$!x$$
$$\{\lambda v'. v' = v * x \mapsto v\}$$

## Example: transferring an assertion through a spin lock

```
// release lock:  
lock :=at false
```

```
// acquire lock:  
while CAS lock false true = false  
do () done
```

## Example: transferring an assertion through a spin lock

```
{P}
```

```
// release lock:
```

```
lock :=at false
```

```
// acquire lock:
```

```
while CAS lock false true = false
```

```
do () done
```

```
{P}
```

## Example: transferring an assertion through a spin lock

$\{P\}$

*// release lock:*

lock :=<sub>at</sub> false

*// acquire lock:*

CAS lock false true

*// CAS succeeds*

$\{P\}$

## Example: transferring an assertion through a spin lock

```
{P}  
{ $\exists U. P @ U * \uparrow U$ }  
// release lock:  
lock :=at false
```

```
// acquire lock:  
    CAS lock false true  
// CAS succeeds  
{ $\exists U. P @ U * \uparrow U$ }  
{P}
```

## Example: transferring an assertion through a spin lock

```
{P}
{ $\exists U. P @ U * \uparrow U$ }
// release lock:
lock :=at false

// acquire lock:
      CAS lock false true
// CAS succeeds
{ $\exists U. P @ U * \uparrow U$ }
{P}
```

- $P @ U$  : transferred via **objective invariants**, as usual.

## Example: transferring an assertion through a spin lock

```
{P}
{ $\exists U. P @ U * \uparrow U$ }
// release lock:
lock :=at false

// acquire lock:
      CAS lock false true
// CAS succeeds
{ $\exists U. P @ U * \uparrow U$ }
{P}
```

- $\uparrow U$  : transferred via synchronization.



## Example: transferring an assertion through a spin lock

```
{P}
{ $\exists U. P @ U * \uparrow U$ }
// release lock:
lock :=at false
```

happens before

```
// acquire lock:
CAS lock false true
// CAS succeeds
{ $\exists U. P @ U * \uparrow U$ }
{P}
```

- $\uparrow U$  : transferred via “atomic” accesses.

# Rules of atomic locations, simplified

$x \mapsto_{\text{at}} v$  : the atomic location  $x$  stores the value  $v$ .

- Sequentially consistent behavior for  $v$ .
- Objective.
- Rules:

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

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

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

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

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

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

# Rules of atomic locations

$x \mapsto_{\text{at}} (v, \mathcal{U})$  : the atomic location  $x$  stores the value  $v$  **and a view** (at least)  $\mathcal{U}$ .

- Sequentially consistent behavior for  $v$ .
- **Release/acquire** behavior for  $\mathcal{U}$ .
- Objective (still).
- Rules:

$$\{x \mapsto_{\text{at}} (v, \mathcal{U}) * \uparrow \mathcal{U}'\}$$

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

$$\{\lambda(). x \mapsto_{\text{at}} (v', \mathcal{U}')\}$$

$$\{x \mapsto_{\text{at}} (v, \mathcal{U})\}$$

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

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

# Rules of atomic locations

$x \mapsto_{\text{at}} (v, \mathcal{U})$  : the atomic location  $x$  stores the value  $v$  **and a view** (at least)  $\mathcal{U}$ .

- Sequentially consistent behavior for  $v$ .
- **Release/acquire** behavior for  $\mathcal{U}$ .
- Objective (still).
- Rules:

$$\frac{\{x \mapsto_{\text{at}} (v, \mathcal{U}) * \uparrow \mathcal{U}'\} \quad \{x \mapsto_{\text{at}} (v, \mathcal{U})\}}{x :=_{\text{at}} v' \quad \text{!}_{\text{at}} x \quad \text{release}} \frac{\{x \mapsto_{\text{at}} (v, \mathcal{U}) * \uparrow \mathcal{U}'\}}{\{\lambda(). x \mapsto_{\text{at}} (v', \mathcal{U}')\}} \quad \{\lambda v'. v' = v * x \mapsto_{\text{at}} (v, \mathcal{U}) * \uparrow \mathcal{U}\}$$

# Rules of atomic locations

$x \mapsto_{\text{at}} (v, \mathcal{U})$  : the atomic location  $x$  stores the value  $v$  **and a view** (at least)  $\mathcal{U}$ .

- Sequentially consistent behavior for  $v$ .
- **Release/acquire** behavior for  $\mathcal{U}$ .
- Objective (still).
- Rules:

$$\{x \mapsto_{\text{at}} (v, \mathcal{U}) * \uparrow \mathcal{U}'\}$$

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

$$\{\lambda(). x \mapsto_{\text{at}} (v', \mathcal{U}')\}$$

$$\{x \mapsto_{\text{at}} (v, \mathcal{U})\}$$

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

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

acquire



## Application: the spin lock

---

# The spin lock

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

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

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

Interface:

$$\boxed{\text{isLock } lk \ P} \vdash$$
$$\left\{ \begin{array}{l} \{P\} \text{ release } lk \ \{\text{True}\} \\ \{\text{True}\} \text{ acquire } lk \ \{P\} \end{array} \right.$$

# The spin lock

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

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

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

Invariant in traditional CSL:

$$\boxed{lk \mapsto_{\text{at}} \text{true} \quad \vee \quad (lk \mapsto_{\text{at}} \text{false} \quad * \quad P)} \vdash$$
$$\left\{ \begin{array}{l} \{P\} \text{ release } lk \{ \text{True} \} \\ \{ \text{True} \} \text{ acquire } lk \{ P \} \end{array} \right.$$



# The spin lock

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

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

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

Invariant in our logic (where  $P$  is subjective!):

$$\boxed{lk \mapsto_{\text{at}} \text{true} \quad \vee \quad (\exists \mathcal{U}. lk \mapsto_{\text{at}} (\text{false}, \mathcal{U}) * P @ \mathcal{U})} \vdash$$
$$\left\{ \begin{array}{l} \{P\} \text{ release } lk \{ \text{True} \} \\ \{ \text{True} \} \text{ acquire } lk \{ P \} \end{array} \right.$$

More case studies:

- Ticket lock
- Dekker mutual exclusion algorithm
- Peterson mutual exclusion algorithm

Method for proving correctness under weak memory:

1. Start with the invariant under sequential consistency;
2. Identify how information flows between threads;
  - i.e. where are the synchronization points;
3. Refine the invariant with corresponding views.

## Conclusion

---

# Conclusion

**Key idea:** The logic of views enables concise and natural reasoning about how threads synchronize.

In the paper:

- Model of the logic.
- A lower-level logic.
- More case studies.

Fully mechanized in Coq with the Iris framework. 🧑🏻

Future work:

- Verify more shared data structures.
- Allow data races on non-atomics?

**Questions?**



## Verifying the spin lock

// release lk:

```
{ isLock lk P * P }
{ lk ↦at _ * P }
{ ∃U. lk ↦at _ *  $\overbrace{\uparrow U * P @ U}$  }
lk :=at false
{ ∃U. lk ↦at (false, U) * P @ U }
{ isLock lk P }
```

// acquire lk:

```
{ isLock lk P }
{ (∃U. lk ↦at (false, U) * P @ U) }
{ ∨ lk ↦at true }
if CAS lk false true
then
  { ∃U. lk ↦at true *  $\overbrace{\uparrow U * P @ U}$  }
  { lk ↦at _ * P }
  { isLock lk P * P }
else
  { lk ↦at true }
  { isLock lk P }
  acquire lk
  { isLock lk P * P }
```

# Model of the logic in Iris

Assertions are predicates on views:

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

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

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

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

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:

$$\text{WP } e \ \varphi \triangleq \lambda \mathcal{U} .$$

$$\text{valid } \mathcal{U} \multimap \text{WP } \langle e, \mathcal{U} \rangle \ (\lambda \langle v, \mathcal{U}' \rangle. \text{valid } \mathcal{U}' * \varphi \ v \ \mathcal{U}')$$

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



# Model of the logic in Iris

Assertions are **monotonic** predicates on views:

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

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

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

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

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:

$$\text{WP } e \varphi \triangleq \lambda \mathcal{U}_1. \forall \mathcal{U} \sqsupseteq \mathcal{U}_1.$$

$$\text{valid } \mathcal{U} \multimap \text{WP } \langle e, \mathcal{U} \rangle (\lambda \langle v, \mathcal{U}' \rangle. \text{valid } \mathcal{U}' * \varphi v \mathcal{U}')$$

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

# Model of the logic in Iris

Assertions are **monotonic** predicates on views:

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

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

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

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

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:

$$\text{WP } e \varphi \triangleq \lambda \mathcal{U}_1. \forall \mathcal{U} \sqsupseteq \mathcal{U}_1.$$

$$\text{valid } \mathcal{U} \multimap \text{WP } \langle e, \mathcal{U} \rangle (\lambda \langle v, \mathcal{U}' \rangle. \text{valid } \mathcal{U}' * \varphi v \mathcal{U}')$$

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

# Assertions are monotonic

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

One reason is the frame rule:

$$\begin{array}{c} \{x \mapsto v * P\} \\ x := v' \\ \{\lambda(). x \mapsto v' * P\} \end{array}$$

# Assertions are monotonic

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

One reason is the frame rule:

$$\{x \mapsto v * P \text{ — holds at the thread's current view}\}$$
$$x := v'$$
$$\{\lambda(). x \mapsto v' * P \text{ — holds at the thread's now extended view}\}$$

# Decomposition of subjective assertions

This theorem allows us to decompose a subjective assertion  $P$ :

$$P \iff \exists U. \underbrace{\uparrow U}_{\text{subjective}} * \underbrace{P @ U}_{\text{objective}}$$

We also have:

$$P @ U \implies \text{Objectively } \uparrow U * P$$

# Decomposition of subjective assertions

This theorem allows us to decompose a subjective assertion  $P$ :

$$P \iff \exists \mathcal{U}. \underbrace{\uparrow \mathcal{U}}_{\text{subjective}} * \underbrace{P @ \mathcal{U}}_{\text{objective}}$$

We also have:

$$P @ \mathcal{U} \iff \text{Objectively}(\uparrow \mathcal{U} * P)$$

where  $\text{Objectively } Q \iff (\forall \mathcal{U}. Q @ \mathcal{U}) \iff Q @ \emptyset$