

# An introduction to Iris

Jean-Marie Madiot & François Pottier

JFLA 2026

## 1 What is Iris (About)?

## 2 Basic Connectives

## 3 Mutable State

## 4 Locks (Primitive)

## 5 Invariants

## 6 Locks (User-Defined)

# A (Very Partial) History

To prove the *safety* and *correctness* of programs,

# A (Very Partial) History

To prove the *safety* and *correctness* of programs,

- in the beginning there was Floyd-Hoare logic (1967–1969)
  - *propositions* about the machine's state

# A (Very Partial) History

To prove the *safety* and *correctness* of programs,

- in the beginning there was Floyd-Hoare logic (1967–1969)
  - *propositions* about the machine's state
- then there was Separation Logic (1999–2002)
  - *assertions* about *fragments* of the machine's state
  - *separation* and *ownership*
  - *[reasoning should be] confined to the cells that the program actually accesses* — O'Hearn, Reynolds, Yang (2001)

# A (Very Partial) History

Then it became apparent that SL could be pushed much further.

# A (Very Partial) History

Then it became apparent that SL could be pushed much further.

- Concurrent Separation Logic (2004–2007)
  - shared *locks* mediating access to exclusive assertions
  - guaranteed *data race freedom*

# A (Very Partial) History

Then it became apparent that SL could be pushed much further.

- Concurrent Separation Logic (2004–2007)
  - shared *locks* mediating access to exclusive assertions
  - guaranteed *data race freedom*
- Iris (2015–2017)
  - separation never truly exists; a *fiction* of separation suffices
  - *stability* of assertions is key
  - monolithic machine state, separable *ghost state*, and *invariants*

# This Introduction to Iris

Iris is a large and complex system ([paper](#); [lecture notes](#); [tutorial](#)).

- As of today, [145 Iris-related papers](#) listed

We wish to

- introduce just the key ideas
- give demonstrations of Iris at work

Two lectures:

- #1 (FP): basic concepts; locks; invariants
- #2 (JMM): user-defined separable ghost state

# What is Logic About?

Logic involves *propositions* about an unchanging mathematical world.

A proposition has a *truth value*: it is either *true* or *false*, and forever so.

$\text{even}(1)$  — *false*

$\text{even}(2)$  — *true*

$\forall n : \mathbb{N}. \exists p : \mathbb{N}. n \leq p \wedge \text{prime}(p)$  — *true*

$\forall x : \mathbb{N}. \text{even}(x) \rightarrow \text{odd}(x + 1)$  — *true*

The rules of logic ensure that only true propositions have proofs.

# What Logic for a Changing World?

Can one make *assertions* about a changing world?

There is nobody in the street.

# What Logic for a Changing World?

Can one make *assertions* about a changing world?

There is nobody in the street.

— *may be true now*

# What Logic for a Changing World?

Can one make *assertions* about a changing world?

There is nobody in the street.

- *may be true now*
- *could become false at any time*
- *somebody could turn the corner*
- *an unstable assertion about a changing world*

# What Logic for a Changing World?

Can one make *stable, local* assertions about a changing world?

My room is painted white.

— *true now*

# What Logic for a Changing World?

Can one make *stable, local* assertions about a changing world?

My room is painted white.

- *true now*
- *perhaps not true forever*
- *I might decide to paint it a different color*

# What Logic for a Changing World?

Can one make *stable, local* assertions about a changing world?

My room is painted white.

- *true now*
- *perhaps not true forever*
- *I might decide to paint it a different color*
- *but no one else may do so (I own this room)*

# What Logic for a Changing World?

Can one make *stable, local* assertions about a changing world?

My room is painted white.

- *true now*
- *perhaps not true forever*
- *I might decide to paint it a different color*
- *but no one else may do so (I own this room)*
- *a **stable** assertion*
- *expressing **knowledge** about the world,*
- ***permission** to change the world,*
- *and **absence of permission** for others to change it*

## More Examples of Stable Assertions

Can one make *stable, local* assertions about a changing world?

I was born on a Monday.

## More Examples of Stable Assertions

Can one make *stable, local* assertions about a changing world?

I was born on a Monday.

— *true*

## More Examples of Stable Assertions

Can one make *stable, local* assertions about a changing world?

I was born on a Monday.

- *true*
- *was not true 60 years ago*

## More Examples of Stable Assertions

Can one make *stable, local* assertions about a changing world?

I was born on a Monday.

- *true*
- *was not true 60 years ago*
- *nobody can change this fact*

## More Examples of Stable Assertions

Can one make *stable, local* assertions about a changing world?

I was born on a Monday.

- *true*
- *was not true 60 years ago*
- *nobody can change this fact*
- a *stable assertion about a changing world*

An example of an assertion that *becomes true* at some point in time and thereafter *persists* forever.

## More Examples of Stable Assertions

Can one make *stable, local* assertions about a changing world?

Over 129,864,880 books have been published.

- *true*
- *was not true 60 years ago*
- *nobody can invalidate this fact*
- a *stable assertion about a changing world*

## More Examples of Stable Assertions

Can one make *stable, local* assertions about a changing world?

Over 129,864,880 books have been published.

- *true*
- *was not true 60 years ago*
- *nobody can invalidate this fact*
- a *stable assertion about a changing world*
- *though anyone has permission to publish new books*

Stable because this aspect of the world evolves in a *monotonic* way.

# What is a Stable Assertion?

An assertion should

- express *knowledge* about (a fragment of) the world
- represent *permission* to change (this fragment of) the world
- represent *interdiction* for others to make incompatible changes

An assertion is *stable* if it contains *enough interdiction* to justify the knowledge and permission that it offers.

# What is Separation Logic?

Separation Logic (SL) is a logic where *every assertion is stable*.

- SL = Stability Logic?

# What is Separation Logic?

Separation Logic enables *local reasoning* about a composite system.

- each participant has *partial knowledge* of the world and *partial permission* to change the world
- one participant's knowledge is never invalidated by another participant's actions
- the share (knowledge and permissions) of one participant is compatible with the share of every other participant
- at all times, *the conjunction of all shares* is consistent

1 What is Iris (About)?

2 Basic Connectives

Conjunction

Implication

Persistence

Update

Execution

3 Mutable State

4 Locks (Primitive)

5 Invariants

6 Locks (User-Defined)

The world is partly *physical*, partly *ghost*.

Typical examples of basic assertions:

- a *physical memory cell*,  $x \mapsto v$ 
  - the points-to assertion (Reynolds, 2002)
- an *immutable* physical memory cell,  $x \mapsto_{\square} v$ 
  - the persistent points-to assertion (Friis Vindum and Birkedal, 2021)
- a *ghost memory cell*,  $\boxed{a}^\gamma$ 
  - new in Iris 1 (Jung et al., 2015)

I want to describe five fundamental connectives:

- *conjunction*,  $A * B$ 
  - decomposes a view of the world into several parts
- *implication*,  $A \rightarrow B$ 
  - change one's view of the world – not the world itself
- *persistence*,  $\square A$ 
  - means “forever  $A$ ”
- *update*,  $\Rightarrow B$ 
  - changes the ghost world
  - the binary form  $A \Rightarrow B$  is sugar for  $\square(A \rightarrow \Rightarrow B)$
- *execution*,  $\text{ex } s \{B\}$ 
  - changes the ghost and physical world
  - the Hoare triple  $\{A\} s \{B\}$  is sugar for  $\square(A \rightarrow \text{ex } s \{B\})$

I will not discuss today:

- *pure* assertions  $\top P \top$  where  $P$  is a proposition
- *quantifiers*  $\forall x.A, \exists x.A$
- the *later* modality  $\triangleright A$
- user-defined assertions, which can
  - *inductive*: linked list (segment), tree, iterated conjunction
  - *co-inductive*
  - *guarded recursive*: ex

I will discuss later today:

- *locks*, first considered primitive, then user-defined
- *invariants*

## 2 Basic Connectives

Conjunction

Implication

Persistence

Update

Execution

Conjunction  $A * B$  means

- $A$  holds and  $B$  holds
- and one can act on one side *without disturbing* the other—*stability*.

This is visible in the way  $\rightarrow$  and  $\Rightarrow$  and  $\text{ex}$  interact with  $*$ .

It is sometimes called “separating” conjunction

- because  $x \mapsto v * y \mapsto v'$  implies  $\lceil x \neq y \rceil$

but the key point is stability.

Conjunction is associative and commutative. *True* is its unit.

It is *not idempotent*:

- Some assertions are not duplicable: in general,  $A \not\vdash A * A$
- Every persistent assertion is duplicable:  $\Box A \vdash \Box A * \Box A$

The logic is *affine*, as opposed to linear:  $A \vdash \text{True}$ .

## 2 Basic Connectives

Conjunction

Implication

Persistence

Update

Execution

# Implication (Magic Wand)

Implication  $A \rightarrow B$  means:

- by *consuming*  $A$
- and by *consuming*  $A \rightarrow B$  as well
- you can get  $B$ .

Think of *two puzzle pieces* that fit together.

Implication changes *your view* of the world, not the world itself.

- $x \mapsto 0 \rightarrow \exists y. x \mapsto y * \lceil 0 \leq y \rceil$
- $x \mapsto y * y \mapsto z \rightarrow \text{listseg}(x, z)$
- $x \mapsto y \rightarrow (y \mapsto z \rightarrow \text{listseg}(x, z))$

## Basic Laws of Implication

Here is one formulation (Ishtiaq and O'Hearn, 2001):

$$\frac{R * A \vdash B}{R \vdash A -* B}$$

$$\frac{R \vdash A -* B \quad R' \vdash A}{R * R' \vdash B}$$

$$\frac{A \vdash A' \quad B \vdash B'}{A * B \vdash A' * B'}$$

# Basic Laws of Implication

Here is one formulation (Ishtiaq and O'Hearn, 2001):

$$\frac{R * A \vdash B}{R \vdash A \multimap B}$$

$$\frac{R \vdash A \multimap B \quad R' \vdash A}{R * R' \vdash B}$$

$$\frac{A \vdash A' \quad B \vdash B'}{A * B \vdash A' * B'}$$

This should be easier to read:

$$\begin{aligned} & R \multimap A \multimap B \\ \equiv & (R * A) \multimap B \quad \text{currying/uncurrying} \end{aligned}$$

$$\begin{aligned} & (A \multimap B) * A \multimap B \\ & \text{application} \end{aligned}$$

$$\begin{aligned} & (A \multimap B) \\ \multimap & (A * R \multimap B * R) \\ & \text{stability (frame)} \end{aligned}$$

$$\begin{aligned} & (R * A \multimap B) * R \\ \multimap & A \multimap B \\ & \text{partial application} \end{aligned}$$

## Non-Separating Conjunction

$A \wedge B$  is an external choice:

- you can have  $A$  *and* you can have  $B$
- but you can have *only one* of them.

$A \wedge B$  is equivalent to

$$\exists S. \quad S * (S \multimap A) * (S \multimap B)$$

Here is a proof.

## 2 Basic Connectives

Conjunction

Implication

**Persistence**

Update

Execution

$\square A$  means that  $A$  is forever true.

An assertion is *persistent* if it can be written in the form  $\square A$ .

- by definition,  $\text{Persistent}(P)$  means  $P \vdash \square P$

Intuitively, a proof of  $\square A$  is a proof of  $A$  that uses persistent facts only.

$$\frac{\square A \vdash B}{\square A \vdash \square B}$$

introduction

$$\frac{\square A}{\neg A}$$

elimination

## 2 Basic Connectives

Conjunction

Implication

Persistence

Update

Execution

A binary update  $A \Rightarrow B$  means:

- by consuming  $A$
- and by *changing the world*
- you can get  $B$ .

It is sugar for  $\square(A \multimap B)$ .

A unary update  $\Rightarrow B$  means

- permission to *change the world* to get  $B$ .

Caveat: there are several notions of update; I am blurring the distinction.

An update is used to allocate a new ghost cell:

- $\text{True} \Rightarrow \exists \gamma. \boxed{a}^\gamma$

and to update a ghost cell (simplified rule—see JMM's lecture):

- $\boxed{a}^\gamma \Rightarrow \boxed{b}^\gamma$

An update is used to open and close an invariant (later today).

# Basic Laws of Binary Update

Binary update behaves very much like implication:

$$\begin{array}{ccc} R \Rightarrow A \Rightarrow B & & (A \Rightarrow B) \\ \equiv (R * A) \Rightarrow B & (A \Rightarrow B) * A \Rightarrow B & -* (A * R \Rightarrow B * R) \\ \text{currying/uncurrying} & \text{application} & \text{stability (frame)} \end{array}$$

$$\begin{array}{c} (R * A \Rightarrow B) * R \\ -* A \Rightarrow B \\ \text{partial application} \end{array}$$

# Basic Laws of Unary Update

It is easier to remember just the laws of unary update:

$$A \dashv\ast \Rightarrow A$$

return (reflexivity)

$$A \dashv\ast \Rightarrow A \dashv\ast \Rightarrow A$$

join (transitivity)

$$A \dashv\ast B$$
$$A \dashv\ast (\Rightarrow A) \dashv\ast (\Rightarrow B)$$

covariance (map)

$$A \dashv\ast (\Rightarrow B)$$
$$A \dashv\ast \Rightarrow (A \dashv\ast B)$$

*stability* (strength)

One sums up these laws by saying: unary update is a strong monad.

# An Interesting Non-Law

Update *does not commute* with universal quantification:

$$\forall x. \Rightarrow A \quad \not\vdash \quad \Rightarrow \forall x. A$$

An intuitive explanation is: a ghost cell can be updated *in any way* you wish but not *in all ways* at once:

$$\begin{array}{ccc} \boxed{a}^\gamma & \text{entails} & \forall b. \Rightarrow \boxed{b}^\gamma \\ \boxed{a}^\gamma & \text{does not entail} & \Rightarrow \forall b. \boxed{b}^\gamma \end{array}$$

This is the same reason why the *value restriction* exists in ML.

This also explains the lack of an *intersection rule* in Iris:

$$\forall x. \boxed{\text{ex}} \in \{A\} \quad \not\vdash \quad \boxed{\text{ex}} \in \{\forall x. A\}$$

## 2 Basic Connectives

Conjunction

Implication

Persistence

Update

Execution

# A Programming Language

So far everything has been about logic, not about programming.

Now assume that *a programming language* (syntax, semantics) is given.

For example, it might be

- a WHILE language, whose statements return no result;
- a  $\lambda$ -calculus, whose expressions return a value.

The assertion `ex s {B}` means

- there is *permission* to *execute* the statement *s* *once*
- this is safe (execution won't crash)
- this may change the (physical and ghost) world
- and if/once execution terminates, *B* will hold.

In the Iris literature, `ex` is named *wp* for “weakest precondition”.

In *dynamic logic* (Pratt, 1974) it is written  $[s]B$ .

I like `ex`  $s \{B\}$  because it can seem to mean

- “out of  $s$  one gets  $B$ ”
- “executing  $s$  establishes  $B$ ”

The Hoare triple, or Separation Logic triple,

$$\{A\} \; s \; \{B\}$$

is sugar for

$$\square(A \rightarrow * \text{ ex } s \; \{B\})$$

# Basic Laws of Execution

An update is a special case of an execution assertion.

$$\begin{aligned} & \Rightarrow B \\ \equiv & \text{ex } \text{skip } \{B\} \\ & \text{skip (return)} \end{aligned}$$

$$\begin{aligned} & \text{ex } s_1 \{ \text{ex } s_2 \{B\} \} \\ \equiv & \text{ex } (s_1; s_2) \{B\} \\ & \text{sequencing (join)} \end{aligned}$$

$$\begin{aligned} & A \multimap B \\ \multimap & \text{ex } s \{A\} \multimap \text{ex } s \{B\} \\ & \text{weakening (map)} \end{aligned}$$

$$\begin{aligned} & A * \text{ex } s \{B\} \\ \multimap & \text{ex } s \{A * B\} \\ & \text{stability / frame (strength)} \end{aligned}$$

# Execution Absorbs Updates

Execution absorbs updates:

$$\Rightarrow \text{ex } s \{B\}$$

-\*  $\text{ex } s \{B\}$

update before execution

$$\text{ex } s \{\Rightarrow B\}$$

-\*  $\text{ex } s \{B\}$

update after execution

These laws because the *definition* of  $\text{ex}$  involves  $\Rightarrow$ .

*Structured parallel composition* and *thread creation* are easy to describe:

$\text{ex } s_1 \{B_1\} * \text{ex } s_2 \{B_2\}$	$\text{ex } s \{B\}$
$\text{ex } (s_1 \parallel s_2) \{B_1 * B_2\}$	$\text{ex } (\text{fork } s) \{ \text{True} \}$

The second rule offers no way of waiting for the child thread to finish so as to obtain  $B$ . It is up to the user to implement such a mechanism using channels, references, etc.

If the programming language has *expressions* (which return values) then one uses  $\text{ex } e \{ \psi \}$  where  $\psi : \text{Val} \rightarrow i\text{Prop}$ .

$\text{ex } e \{ y.B \}$  means

- there is *permission* to *execute* the expression  $e$  *once*
- and (if it terminates then) it returns a value  $y$  such that  $B$  holds.

The skip and sequencing rules become

$$\begin{array}{ll} \Rightarrow \psi v & \equiv \text{ex } e_1 \{v. \text{ex } [v/x]e_2 \{\psi\}\} \\ \equiv \text{ex } v \{\psi\} & \equiv \text{ex } (\text{let } x = e_1 \text{ in } e_2) \{\psi\} \\ \text{return} & \text{bind} \end{array}$$

These rules are used to reason *step by step* about a program.

They allow *symbolic execution* inside the proof assistant.

1 What is Iris (About)?

2 Basic Connectives

**3 Mutable State**

4 Locks (Primitive)

5 Invariants

6 Locks (User-Defined)

The assertion  $x \mapsto v$  describes a *reference* (a mutable memory block).

This assertion means:

- the reference at address  $x$  *currently contains* the value  $v$
- *permission* to read and write this reference
- interdiction for anyone else to read or write this reference

This assertion is *exclusive*:  $x \mapsto v * x \mapsto v \vdash \text{False}$ .

More generally, there is *separation*:  $x \mapsto v * y \mapsto v' \vdash \neg x \neq y$ .

This assertion is *not duplicable*:  $x \mapsto v \not\vdash x \mapsto v * x \mapsto v$ .

# Operations on References

References can be *allocated*, *read*, and *written*.

$\text{ex } (\text{ref } v) \{v'. \exists \ell. \lceil v' = \ell \rceil * \ell \mapsto v\}$   
create

$\ell \mapsto v$   
 $\text{---*--- } \text{ex } (!\ell) \{v'. \lceil v' = v \rceil * \ell \mapsto v\}$  read

$\ell \mapsto v$   
 $\text{---*--- } \text{ex } (\ell := v') \{_. \ell \mapsto v'\}$  write

In a language without GC, there would be a *deallocation* operation.

# Operations on References, Texan Style

This postcondition-passing style makes the rules easier to apply:

$$\begin{array}{ll} \mathit{True} * (\forall \ell. \ell \mapsto v \rightarrow \psi \ell) & \ell \mapsto v * (\ell \mapsto v \rightarrow \psi v) \\ -* \text{ ex } (\text{ref } v) \{ \psi \} & -* \text{ ex } (!\ell) \{ \psi \} \\ & \text{create} & \text{read} \end{array}$$

$$\begin{array}{l} \ell \mapsto v * (\ell \mapsto v' \rightarrow \psi ()) \\ -* \text{ ex } (\ell := v') \{ \psi \} \\ & & \text{write} \end{array}$$

Instead of saying: the postcondition of `write` is  $\ell \mapsto v'$ ,  
say: it is anything you want, provided it is implied by  $\ell \mapsto v'$ .

# Making a Reference Immutable

A mutable reference can be forever turned into an immutable one.

$$\begin{array}{c} \ell \mapsto v \\ \Rightarrow \ell \mapsto_{\square} v \\ \text{freeze} \end{array} \quad \begin{array}{c} \ell \mapsto_{\square} v \\ \rightarrow * \text{ ex } (!\ell) \{ v'. \lceil v' = v \rceil \} \\ \text{read frozen} \end{array}$$

The two views cannot co-exist:  $\ell \mapsto v * \ell \mapsto_{\square} v'$  implies *False*.

A write access to a reference requires a mutable points-to assertion.

A read access requires a (mutable or immutable) points-to assertion.

Therefore a write and a read *can never* be simultaneously enabled!

- *Data-race freedom* is guaranteed. (Good!)
- Communication between threads is *impossible*. (Bad!)

These points hold even if *read-modify-write* operations (FAA, CAS, etc.) are allowed, as they also require an exclusive points-to assertion.

To allow threads to interact, one must introduce

- synchronisation primitives: for example, *locks*; or
- shared *invariants*.

1 What is Iris (About)?

2 Basic Connectives

3 Mutable State

4 Locks (Primitive)

5 Invariants

6 Locks (User-Defined)

In OCaml, an abstract type of locks could look like this:

```
type lock
(* a lock can be shared between several threads *)
val newlock : unit -> lock
val acquire : lock -> unit (* acquire access permission *)
val release : lock -> unit (* release access permission *)
```

The type-checker does not know *what data structure* a lock protects, so cannot check that acquire and release are correctly used.

A stack, protected by a lock, could look like this:

```
type 'a stack =
  { data: 'a list ref; lock: lock } (* lock protects data *)

let make () =
  let data = ref [] in
  let lock = newlock() in
  { data; lock }

let push x stack =
  acquire stack.lock;           (* acquire permission *)
  stack.data := x :: !stack.data; (* access the data *)
  release stack.lock            (* release permission *)
```

To verify the safety of this code, *reasoning rules for locks* are needed.

There exists  $isLock : Val \rightarrow iProp \rightarrow iProp$  such that:

$$\begin{array}{c}
 R \\
 \begin{array}{ccc}
 \text{Persistent} (isLock \vee R) & \xrightarrow{*} & \text{ex } (\text{newlock}()) \{v. isLock \vee R\} \\
 \text{share} & & \text{create}
 \end{array} \\
 \\[10pt]
 \begin{array}{ccc}
 isLock \vee R & & isLock \vee R * R \\
 \xrightarrow{*} \text{ex } (\text{acquire } v) \{R\} & \xrightarrow{*} & \text{ex } (\text{release } v) \{True\} \\
 \text{acquire} & & \text{release}
 \end{array}
 \end{array}$$

From the user's point of view, acquire *produces*  $R$ ; release *consumes*  $R$ .

A stack, protected by a lock, could look like this:

```
type 'a stack =
  { data: 'a list ref; lock: lock } (* lock protects data *)

let make () =
  let data = ref [] in
  let lock = newlock() in
  { data; lock }

let push x stack =
  acquire stack.lock;           (* acquire permission *)
  stack.data := x :: stack.data; (* access the data     *)
  release stack.lock            (* release permission *)
```

See [a proof of safety](#) of this code.

The permission to access the data appears only within critical sections

- between release and acquire

so *data-race freedom* is still guaranteed

- even though interactions between threads are now possible

One could improve this Lock API in several ways:

- separating the *creation* of the lock and the *initialization* of the assertion  $R$
- use *fractions* to keep track of sharing and allow *canceling* a lock whose fraction is 1
- introduce an assertion *isLocked*  $v$  to *prevent releasing a lock that one does not hold*
  - under our API, such a mistake is possible if  $R * R \not\vdash \text{False}$
- view *isLocked*  $v$  as an *obligation* to eventually release the lock
  - current Iris does not allow this, as it is affine
  - current Iris does not guarantee absence of deadlocks

In this formulation, acquire yields a *unique permission* to release:

$$\text{Persistent}(\text{isLock} \vee R) \quad \begin{matrix} & R \\ \text{share} & \rightarrow * \text{ ex } (\text{newlock}()) \{v. \text{isLock} \vee R\} \\ & \text{create} \end{matrix}$$

$$\begin{matrix} & \text{isLock} \vee R \\ \rightarrow * \text{ ex } (\text{acquire } v) \{R * (R \rightarrow * \text{ ex } (\text{release } v) \{ \text{True} \})\} \\ & \text{acquire then release} \end{matrix}$$

This prevents releasing a lock that one does not hold.

The proof is left as an **exercise**.

In this formulation, the assertion  $isLock \vee R$  is not needed.

$$R \\ \text{ex } (\text{newlock}()) \\ -* \quad \left\{ v. \square \left( \begin{array}{l} \text{ex } (\text{acquire } v) \\ \{ R * (R -* \text{ex } (\text{release } v) \{ \text{True} \}) \} \end{array} \right) \right\} \\ \text{create then forever (acquire then release)}$$

The entire Lock API is described by just one rule!

The proof is left as an **exercise**.

1 What is Iris (About)?

2 Basic Connectives

3 Mutable State

4 Locks (Primitive)

5 Invariants

6 Locks (User-Defined)

# Many Synchronization Operations

We have described *locks* as primitive objects that allow synchronization.

But there are many more:

- semaphores,
- barriers,
- condition variables,
- channels (concurrent FIFO queues),
- concurrent data structures of all kinds.

We want to *construct* and *verify* them, not view them all as primitive.

# A General Principle, or Recipe

To describe a runtime mechanism  
that involves multiple participant threads and transfers of permissions,

- ① define custom *ghost state*  
to represent each participant's view
- ② prove *ghost update* lemmas  
describing how the participants' views can evolve
- ③ install an *invariant*  
to relate the physical state and the ghost state

Points 1 and 2 will be covered by JMM. Now what is an invariant?

I am *not* talking about

- a data structure invariant
  - a user-defined assertion such as *isLinkedList*  $\ell$  vs
- a loop invariant
  - the precondition of a recursive function

An Iris *invariant* is an assertion that *everyone agrees to maintain, forever*.

An invariant can be

- *created* (established) at a certain point in time
  - an invariant is part of the ghost state
- *opened* (temporarily violated), then *closed* (established) again
  - everyone can *depend* on the invariant
  - everyone must *preserve* the invariant
  - violations must be short-lived: at most *one atomic instruction*
- *shared* between participants
  - an invariant is never destroyed
  - its existence can be advertised to all participants

## First Attempt (Unsound)

One can dynamically *create*, *share*, *open*, and *close* invariants.

$$\begin{array}{lll} P \Rightarrow \boxed{P} & \text{Persistent}(\boxed{P}) & \begin{array}{c} \boxed{P} \\ \dashv \Rightarrow (P * (P \dashv \Rightarrow \text{True})) \end{array} \\ \text{create} & \text{share} & \text{open / close} \end{array}$$

This is analogous to creating, sharing, acquiring, releasing a lock, but the whole thing is *ghost*—there is no runtime machinery.

These rules are *unsound*.

With these rules,  
an invariant can be *opened twice* simultaneously,  
by the same thread or by two distinct threads,  
duplicating  $P$ .

This simplified presentation differs from Iris and is not machine-checked.

Introduce a (ghost) assertion  $W$ , for *world satisfaction*.

- $W$  is a witness that all invariants in the world are satisfied (closed)
- $W$  can also be viewed as *permission to open* and exploit invariants

Restrict the rule open / close :

$P \Rightarrow \boxed{P}$   
creation

$\text{Persistent}(\boxed{P})$   
share

$\rightarrow * \boxed{P} \Rightarrow (P * (P \rightarrow \Rightarrow \boxed{W}))$   
open / close

Now, the question is,

- how can the token  $W$  be *obtained*?
- when and how must it be *surrendered*?

Now, the question is,

- how can the token  $W$  be *obtained*?
- when and how must it be *surrendered*?

We want  $W$  to appear/disappear before/after every *atomic expression*.

One can think of  $W$  as a token that is

- given by the scheduler to the active thread
- taken from the active thread by the scheduler

Parameterize the `execution` assertion with a *mask*  $m \in \{0, 1\}$ .

- $\text{ex}_0 \ e \ \{\psi\}$  means  $e$  is safe even if some invariants are violated
  - interleaving with other threads forbidden
  - $e$  must be atomic
- $\text{ex}_1 \ e \ \{\psi\}$  means  $e$  is safe provided all invariants hold
  - the proof can exploit (open and close) invariants
  - interleaving with other threads permitted

All of the rules for  $\text{ex}_m$  are polymorphic in  $m$  *except sequencing*, which requires  $m = 1$ :

$$\begin{aligned} & \Gamma \vdash m = 1 \\ \rightarrow & \text{ex}_m \ e_1 \ \{v. \text{ex}_m \ [v/x]e_2 \ \{\psi\}\} \\ \rightarrow & \text{ex}_m \ (\text{let } x = e_1 \text{ in } e_2) \ \{\psi\} \end{aligned}$$

bind

In other words,  $\text{ex}_0$  cannot reason about composite expressions; it is restricted to *atomic expressions*.

$ex_0$  and  $ex_1$  are related as follows:

$$\begin{array}{ll} \text{---} \ast \begin{array}{c} ex_0 \ e \ \{\psi\} \\ ex_1 \ e \ \{\psi\} \end{array} & \begin{array}{c} (W \rightarrow ex_0 \ e \ \{W * \psi\}) \\ \ast \begin{array}{c} ex_1 \ e \ \{\psi\} \\ \text{weaken} \end{array} \end{array} \\ & \qquad \qquad \qquad \text{atomic} \end{array}$$

The second rule states that *during the execution of an atomic expression* the token  $W$  appears out of thin air.

By combining the previous rules, we obtain a simpler `open / close` rule, which does not mention  $W$ .

$$\begin{array}{l} \boxed{P} \\ -* (P -* \text{ex}_0 \ e \ \{P * \psi\}) \\ -* \text{ex}_1 \ e \ \{\psi\} \end{array}$$

`open / close`

By combining the previous rules, we obtain a simpler `open / close` rule, which does not mention  $W$ .

$$\begin{array}{l} \boxed{P} \\ -* (P -* \boxed{ex_0} e \{P * \psi\}) \\ -* \boxed{ex_1} e \{\psi\} \end{array}$$

`open / close`

Imagine  $e$  is a memory access (read, write, CAS, etc.). Then

- $P$  can be exploited to obtain  $\ell \mapsto v$  and *justify this access*
- the updated assertion  $\ell \mapsto v'$  must be used to reconstruct  $P$  thereby *proving that the invariant is preserved*

## Towards Sound Invariants

For example, specializing the rule for a write:

$\ell \mapsto v$	$\boxed{P}$
$\rightarrow \text{ex}_0 (\ell := v') \{ \_. \ell \mapsto v' \}$	$\rightarrow (P \rightarrow \text{ex}_0 \ e \{P * \psi\})$
write	open / close

yields the following rule:

- \*  $\boxed{P}$
- \*  $(P -* \exists v. \ell \mapsto v * (\ell \mapsto v' -* P * \psi ()))$
- \* **ex1**  $(\ell := v') \{\psi\}$

Invariants are a form of *higher-order ghost state*:

- assertions about the (physical and ghost) heap
- stored inside the (ghost) heap

In combination with ghost state,  
the rules that I have sketched are still *unsound*.

Two known paradoxes involve (roughly)

- storing at ghost address  $\gamma$  the proposition:  
“the proposition stored at address  $\gamma$  is false”
- creating an invariant whose content is the proposition:  
“it is impossible to initialize all invariants”

To avoid these paradoxes, the invariant opening rule must be weakened:

$P \Rightarrow \boxed{P}$	$Persistent(\boxed{P})$	$\rightarrow \boxed{P}$
create	share	$W \Rightarrow (\triangleright P * (\triangleright P \rightarrow \Rightarrow W))$ open / close

$P$  implies  $\triangleright P$ . The converse is false.

This prevents circular arguments where an invariant is exploited as part of its own initialization.

One defines  $\text{ex}$  so that every time one step of computation is taken  $\triangleright P$  can be transformed into  $P$ .

The rules that I have sketched can open *only one invariant at a time*.

- Opening an invariant consumes  $W$ .
- But, to open a second invariant,  $W$  is needed.

Iris has more complex rules, where a mask is not just one bit but a function of an infinite set of *names* to bits.

Then one can open two invariants simultaneously provided they have distinct names.

1 What is Iris (About)?

2 Basic Connectives

3 Mutable State

4 Locks (Primitive)

5 Invariants

6 Locks (User-Defined)

# Locks as a User-Defined Data Structure

A spin lock can be implemented as follows:

```
type lock = bool Atomic.t          (* true if lock is held *)
let newlock() = Atomic.make false
let try_acquire lock = Atomic.compare_and_set lock false true
let rec acquire lock = if not (try_acquire lock) then acquire lock
let release lock = Atomic.set lock false
```

# Locks as a User-Defined Data Structure

This data structure can be described by an invariant:

$$isLock \vee R \triangleq \exists \ell. \Gamma v = \ell \cap * \boxed{\ell \mapsto \text{true} \vee (\ell \mapsto \text{false} * R)}$$

Based on this definition of *isLock*,  
one can *prove* that the code satisfies the API shown earlier:

$Persistent(isLock \vee R)$	$\xrightarrow{*}$	$ex \text{ (newlock()) } \{v. isLock \vee R\}$
		<span style="border: 1px solid lightblue; padding: 2px;">share</span> <span style="border: 1px solid lightblue; padding: 2px;">create</span>
$isLock \vee R$		$isLock \vee R * R$
$\xrightarrow{*} ex \text{ (acquire } v \text{) } \{R\}$		$\xrightarrow{*} ex \text{ (release } v \text{) } \{True\}$
		<span style="border: 1px solid lightblue; padding: 2px;">acquire</span> <span style="border: 1px solid lightblue; padding: 2px;">release</span>

See *the proof* in all of its glory.

# That's all, folks!

Coming up next:

Everything you always wanted to know  
about ghost state but were afraid to ask