Iris: Monoids and Invariants as an Orthogonal basis for Concurrent Reasoning

Kasper Svendsen

joint work with Ralf Young, David Swasey, Filip Sieczkowski, Aaron Turon, Lars Birkedal and Derek Dreyer



. . . .

# A uniform framework for describing interference





- Supports encoding of existing reasoning principles
  - Monoids for **expressing** protocols on shared state
  - Invariants for **enforcing** protocols on shared state



- Invariants and monoids are orthogonal
  - Treating them as such, leads to a simpler logic, and a **model** simple enough to **formalize in Coq**



- Supports a notion of logical atomicity
  - extends reasoning principles usually reserved for atomic code to code that **appears** to be atomic
  - we can **define** logical atomicity in Iris

### Iris



- extends reasoning principles usually reserved for atomic code to code that **appears** to be atomic
- we can **define** logical atomicity in Iris

### Iris



- extends reasoning principles usually reserved for atomic code to code that **appears** to be atomic
- we can **define** logical atomicity in Iris

Part 1 Iris

$$\frac{\{\triangleright R * P\} e \{\triangleright R * Q\}_{\mathcal{E}} e \text{ atomic}}{\left\{ R^{\iota} * P \right\} e \left\{ Q \right\}_{\mathcal{E} \uplus \{\iota\}}}$$
The set of invariants that we may open







### Monoids

- Iris is parameterised by a notion of ghost resources
- Ghost resources consists of
  - Information about the current ghost state
  - **Rights** to update ghost state
- We use monoids to model ghost resources

### Monoids

- Ghost resource [m] asserts ownership of m fragment
- Ghost resources can be split arbitrarily

$$[\underline{m_1} \cdot \underline{m_2}] \Leftrightarrow [\underline{m_1}] * [\underline{m_2}]$$

and support frame-preserving updates

$$\frac{\forall a_f. \ (a \cdot a_f) \downarrow \Rightarrow (b \cdot a_f) \downarrow}{\left[\bar{a}\right] \Rightarrow \left[\bar{b}\right]}$$

Part 2 Recovering existing reasoning principles

# Deriving small-footprint specifications

- **Example**: recovering small-footprint specifications from large-footprint specifications
- Same idea as in Superficially Substructural Types (ICFP12) and Fictional Separation Logic (ESOP12)

### A $\lambda$ -calculus with channels

• We instantiate Iris with a  $\lambda$ -calculus with channels

 $e ::= \dots \mid \mathsf{newch} \mid \mathsf{send}(e, e) \mid \mathsf{tryrecv}(e) \mid \mathsf{fork}(e)$ 

• with the following per-thread reduction semantics

 $C[c \mapsto M]; \operatorname{send}(c, v) \to C[c \mapsto M \uplus \{v\}]; ()$  $C[c \mapsto \emptyset]; \operatorname{tryrecv}(c) \to C[c \mapsto \emptyset]; \operatorname{none}$  $C[c \mapsto M \uplus \{v\}]; \operatorname{tryrecv}(c) \to C[c \mapsto M]; \operatorname{some}(v)$ 

## Large-footprint specs

- Reduction relation lifts directly to large-footprint specs
- The reduction

$$C[c \mapsto M]; \mathbf{send}(c, v) \to C[c \mapsto M \uplus \{v\}]; ()$$

yields the following axiom

 $\{\lfloor C[c\mapsto M] \rfloor\} \operatorname{send}(c,v) \{r. \ r = () \land \lfloor C[c\mapsto M \uplus \{v\}] \rfloor\}$ 

Asserts exclusive ownership of entire physical state

### Small-footprint specs

• Large-footprint spec requires global reasoning

 $\{\lfloor C[c\mapsto M] \rfloor\} \operatorname{send}(c,v) \{r. \ r = () \land \lfloor C[c\mapsto M \uplus \{v\}] \rfloor\}$ 

• **Goal:** Derive small-footprint specification that only mentions channels affected by each operation

### Small-footprint specs

#### • Idea

- Introduce appropriate channel ghost resources
- Introduce an invariant that owns the physical state (so that it can be shared) and ties ghost resources to physical state
- Extends to a general construction

### Channel-local monoid

- **Goal:** ghost channels resources that support exclusive ownership of individual channels
- Use partial channel "heaps"

$$|\operatorname{NET}| = Chan \stackrel{\text{fin}}{\rightharpoonup} MsgBag$$
$$f \cdot g = f \cup g, \quad \text{if } \operatorname{dom}(f) \cap \operatorname{dom}(g) = \emptyset$$

•  $[c \mapsto M]$  asserts exclusive ownership of ghost channel c and that contains messages M

### Authoritative monoid

- Goal: a monoid with
  - An authoritative element  $m \bullet$  that asserts that the current ghost state is exactly m
  - A partial element  $m \circ$  that asserts ownership of an *m* fragment of the authoritative state
- s.t. all fragments combine to the authoritative state

 $\{c \prec M\}$ 

 $\operatorname{send}(c,m)$ 

Channel resource asserts ownership of corresponding fragment:

$$c \prec M \triangleq [ [c \mapsto M] _{\circ}]$$

 $\{c \prec M\}$ 

#### $\operatorname{send}(c,m)$

#### **Invariant**: the physical state is authoritative ghost state

 $\exists C. \left[ \stackrel{-}{\underline{C}} \bullet \right] * \left\lfloor C \right\rfloor$ 

Channel resource asserts ownership of corresponding fragment:

$$c \prec M \triangleq [ [c \mapsto M] \circ ]$$

 $\{c \prec M\}$ 

#### $\operatorname{send}(c,m)$

#### **Invariant**: the physical state is authoritative ghost state

 $\exists C. \left[ \stackrel{-}{\underline{C}} \bullet \right] * \left\lfloor C \right\rfloor$ 

Channel resource asserts ownership of corresponding fragment:

$$c \prec M \triangleq [ [c \mapsto M] \circ ]$$

 $\{c \prec M\}$   $\{ [c \mapsto M] \circ ] *$  send(c, m)

$$\{c \prec M \uplus \{m\}\}$$

**Invariant**: the physical state is authoritative ghost state

 $\exists C. \left[ \stackrel{-}{\underline{C}} \bullet \right] * \left\lfloor C \right\rfloor$ 

Channel resource asserts ownership of corresponding fragment:

$$c \prec M \triangleq [ [c \mapsto M] \circ ]$$

 $\{c \prec M\} \\ \{ [c \mapsto M] \circ ] * [C \bullet] * [C \downarrow\} \\ \mathbf{send}(c, m)$ 

$$\{c \prec M \uplus \{m\}\}$$

**Invariant**: the physical state is authoritative ghost state

 $\exists C. \left[ \stackrel{-}{\underline{C}} \bullet \right] * \left\lfloor C \right\rfloor$ 

Channel resource asserts ownership of corresponding fragment:

$$c \prec M \triangleq [ [c \mapsto M] \circ ]$$

 $\begin{aligned} \{c \prec M\} \\ \{[c \mapsto M] \circ] * [C \bullet] * [C] \} \\ send(c, m) \\ \{ & * |C[c \mapsto C(c) \uplus \{m\}| \} \end{aligned}$ 

**Invariant**: the physical state is authoritative ghost state

 $\exists C. \left[ \stackrel{-}{\underline{C}} \bullet \right] * \left\lfloor C \right\rfloor$ 

Channel resource asserts ownership of corresponding fragment:

$$c \prec M \triangleq [ [c \mapsto M] \circ ]$$

 $\begin{aligned} \{c \prec M\} \\ \{\left[c \mapsto M\right] \circ\right] * \left[c \bullet\right] * \left[C \bullet\right] \\ & \text{send}(c, m) \\ \{\left[c \mapsto M\right] \circ\right] * \left[c \bullet\right] * \left[C \bullet\right] * \left[C[c \mapsto C(c) \uplus \{m\}\right] \end{aligned}$ 

**Invariant**: the physical state is authoritative ghost state

 $\exists C. \left[ \stackrel{-}{\underline{C}} \bullet \right] * \left\lfloor C \right\rfloor$ 

Channel resource asserts ownership of corresponding fragment:

$$c \prec M \triangleq [ [ c \mapsto M ] \circ ]$$

$$\begin{aligned} \{c \prec M\} \\ \{\left[\overline{c} \mapsto \overline{M}\right] \circ\right] * \left[\overline{C} \bullet\right] * \left[C\right] \} \\ & \underbrace{\operatorname{send}(c, m)}_{\left\{\left[\overline{c} \mapsto \overline{M}\right] \circ\right]} * \left[\overline{C} \bullet\right] * \left[\overline{C} \circ\right] * \left[\overline{C[c \mapsto C(c) \uplus \{m\}]} \right] \end{aligned}$$

**Invariant**: the physical state is authoritative ghost state

 $\exists C. \left[ \stackrel{-}{\underline{C}} \bullet \right] * \left\lfloor C \right\rfloor$ 

Channel resource asserts ownership of corresponding fragment:

$$c \prec M \triangleq [ [ c \mapsto M ] \circ ]$$

 $\begin{aligned} \{c \prec M\} \\ \{\left[\overline{c} \mapsto \overline{M}\right] \circ\right] * \left[\overline{C} \bullet\right] * \lfloor C \rfloor \} \\ & \underbrace{\operatorname{send}(c, m)}_{\left\{\left[\overline{c} \mapsto \overline{M}\right] \circ\right]} * \left[\overline{C} \bullet\right] * \left[\overline{C} [c \mapsto C(c) \uplus \{m\}\right] \} \\ \{\left[\overline{c} \mapsto \overline{M} \uplus \{m\}\right] \circ\right] * \left[\overline{C'} \bullet\right] * \lfloor C' \rfloor \} \\ \{c \prec M \uplus \{m\} \} \end{aligned}$ 

# Deriving small-footprint specifications

- Channel monoid encodes small-footprint channel resources
- Invariant relates ghost and physical state using authoritative monoid to allow ownership of channel fragments

Recovering existing reasoning techniques

- We saw how to recover reasoning principles from Superficially Substructural Types and Fictional Separation
- One can also recover reasoning principles from CaReSL and iCAP through a encoding of STSs as monoids

Part 3 Logical atomicity

• In part 2 we used the invariant rule to access the shared physical resource

$$\frac{\{\triangleright R * P\} e \{\triangleright R * Q\}_{\mathcal{E}} \quad e \text{ atomic}}{\left\{ \left[ R \right]^{\iota} * P \right\} e \left\{ Q \right\}_{\mathcal{E} \uplus \{\iota\}}}$$

- This rule only applies to **atomic** expressions
- Iris allows us to extend this reasoning principle to logically atomic code

• In part 2 we used the invariant rule to access the shared physical resource

$$\{ \triangleright R * P \} e \{ \triangleright R * Q \}_{\mathcal{E}} e \text{ atomic}$$

$$\{ \boxed{R}^{t} * P \} e \text{ We can define logically atomic triples}$$

$$\langle P \rangle e \langle Q \rangle$$

- This rule only applies to atom
- Iris allows us to extend this reasoning principle to logically atomic code

• **Example:** a blocking receive operation

recv 
$$\triangleq$$
 rec  $recv(c)$ . let  $v = tryrecv(c)$  in  
case  $v$  of none  $= recv(c) \mid some(m) = m$ 

- Spins (without side effects) until a msg is received
- The linearisation point is the first successful tryrecv

#### · Ideas

- Let clients reason about the state immediately before and after the linearisation point
- Let clients open invariants around the linearisation point

#### Ideas

Parameterise our specifications with view shifts

- Let clients reason about the state immediately before and after the linearisation point
- Let clients open invariants around the linearisation point

Let view shifts open and close invariants

### Mask-changing view shifts

• Index view shifts with the set of invariants enabled before and after the view shift

$$P \stackrel{\mathcal{E}_1}{\Rightarrow} \stackrel{\mathcal{E}_2}{\Rightarrow} Q$$

#### • Asserts

- that we can update the instrumented state from *P* to *Q* without changing the physical state
- where the invariants in  $\mathcal{E}_1$  are enabled before the view shift
- and the invariants in  $\mathcal{E}_2$  are enabled after the view shift

### Mask-changing view shifts

• We can change the invariant mask around atomic expressions, provided we restore it again

$$\begin{array}{ccc}
e \text{ atomic} \\
P^{\{\iota\}} \Rightarrow^{\emptyset} P' & \{P'\} e \{v. Q'\}_{\emptyset} & \forall v. Q' \stackrel{\emptyset}{\Rightarrow}^{\{\iota\}} Q \\
& \{P\} e \{v. Q\}_{\{\iota\}}
\end{array}$$

We can open and close invariants using view shifts

$$\boxed{P}^{\iota} {}^{\{\iota\}} \Longrightarrow^{\emptyset} \triangleright P \qquad \qquad \boxed{P}^{\iota} \ast \triangleright P {}^{\emptyset} \Longrightarrow^{\{\iota\}} \top$$

• Idea: Let clients open and close invariants around linearisation point and update instrumented state

$$\langle P \rangle \ e \ \langle Q \rangle_{\mathcal{E}} \approx \forall R_p, R_q, \mathcal{E}_R. \ \mathcal{E} \cap \mathcal{E}_R = \emptyset \land (R_p \iff^{-\mathcal{E}_R} P) \land (Q \Longrightarrow_{-\mathcal{E}_R} R_q) \Rightarrow \ \{R_p\} \ e \ \{R_q\}$$

• This allows us to open invariants around logically atomic code

$$\frac{\langle \triangleright R * P \rangle \ e \ \langle \triangleright R * Q \rangle_{\mathcal{E}}}{\langle \left[ R \right]^{\iota} * P \rangle \ e \ \langle Q \rangle_{\mathcal{E} \uplus \{\iota\}}}$$

• Idea: Let clients open and close invariants around linearisation point and update instrumented state

$$\langle P \rangle \ e \ \langle Q \rangle_{\mathcal{E}} \approx \forall R_p, R_q, \mathcal{E}_R. \ \mathcal{E} \cap \mathcal{E}_R = \emptyset \land (R_p \iff^{-\mathcal{E}_R} P) \land (Q \Longrightarrow_{-\mathcal{E}_R} R_q) \Rightarrow \ \{R_p\} \ e \ \{R_q\}$$

This allows us to open invariants around logically atomic code

$$\frac{\langle \triangleright R * P \rangle \ e \ \langle \triangleright R * Q \rangle_{\mathcal{E}}}{\langle \left[ R \right]^{\iota} * P \rangle \ e \ \langle Q \rangle_{\mathcal{E} \uplus \{\iota\}}}$$

From the client's point of view it looks like we have access to the invariant R for the duration of e.

 Idea: Let clients open and c linearisation point and update instru From the module's point of view we only access the invariant in the linearisation point.

$$\langle P \rangle \ e \ \langle Q \rangle_{\mathcal{E}} \approx \forall R_p, R_q, \mathcal{E}_R. \ \mathcal{E} \cap \mathcal{E}_R = \emptyset \land$$

$$(R_p \iff^{-\mathcal{E}_R} P) \land (Q \Longrightarrow_{-\mathcal{E}_R} R_q)$$

$$\Rightarrow \ \{R_p\} \ e \ \{R_q\}$$

This allows us to open invariants around logically atomic code

$$\frac{\langle \triangleright R * P \rangle \ e \ \langle \triangleright R * Q \rangle_{\mathcal{E}}}{\langle \left[ R \right]^{\iota} * P \rangle \ e \ \langle Q \rangle_{\mathcal{E} \uplus \{\iota\}}}$$

From the client's point of view it looks like we have access to the invariant R for the duration of e.

### Case study



 Logical atomicity is not built into Iris, but Iris is sufficiently expressive that we can **define** it in Iris.

### Conclusions

- Iris is
  - simpler than previous logics
  - can encode reasoning principles from previous logics
  - and can do some fancy new stuff (logical atomicity)
- Monoids and invariants are all you need