

Relational reasoning using concurrent separation logic

Robbert Krebbers¹

Delft University of Technology, The Netherlands

January 20, 2020 @ ADSL, New Orleans, USA

¹This is joint work with Dan Frumin (Radboud University) and Lars Birkedal (Aarhus University)

Why prove relational properties of programs?

- ▶ Specifying programs

implementation \approx_{ctx} specification

Why prove relational properties of programs?

- ▶ Specifying programs

`implementation` \simeq_{ctx} `specification`

- ▶ Optimized versions of data structures

`hash_table` \simeq_{ctx} `assoc_list`

Why prove relational properties of programs?

- ▶ Specifying programs

`implementation` \simeq_{ctx} `specification`

- ▶ Optimized versions of data structures

`hash_table` \simeq_{ctx} `assoc_list`

- ▶ Proving program transformations

$\forall e_{source}. \text{compile}(e_{source}) \simeq_{ctx} e_{source}$

Language features that complicate refinements

- ▶ Mutable state

$$\left(\text{let } x = f() \text{ in } (x, x) \right) \not\sim_{\text{ctx}} \left(f(), f() \right)$$

Language features that complicate refinements

- ▶ Mutable state

$$\left(\text{let } x = f() \text{ in } (x, x) \right) \not\sim_{ctx} \left(f(), f() \right)$$

- ▶ Higher-order functions

$$\left(\lambda().1 \right) \not\sim_{ctx} \left(\text{let } x = \text{ref}(0) \text{ in } (\lambda().x \leftarrow (1 + !x); !x) \right)$$

Language features that complicate refinements

- ▶ Mutable state

$$\left(\text{let } x = f() \text{ in } (x, x) \right) \not\sim_{ctx} \left(f(), f() \right)$$

- ▶ Higher-order functions

$$\left(\lambda().1 \right) \not\sim_{ctx} \left(\text{let } x = \text{ref}(0) \text{ in } (\lambda().x \leftarrow (1 + !x); !x) \right)$$

- ▶ Concurrency

$$\left(x \leftarrow 10; x \leftarrow 11 \right) \not\sim_{ctx} \left(x \leftarrow 11 \right)$$

What do such relational properties mean mathematically?

Contextual refinement

Contextual refinement: the “gold standard” of program refinement:

$$e_1 \lesssim_{\text{ctx}} e_2 : \tau \triangleq \forall (C : \tau \rightarrow \text{int}). \forall v. C[e_1] \downarrow v \implies C[e_2] \downarrow v$$

“Any behavior of a (well-typed) client C using e_1 can be matched by a behavior of the same client using e_2 ”

Contextual refinement

Contextual refinement: the “gold standard” of program refinement:

$$e_1 \lesssim_{\text{ctx}} e_2 : \tau \triangleq \forall (\mathcal{C} : \tau \rightarrow \text{int}). \forall v. \mathcal{C}[e_1] \downarrow v \implies \mathcal{C}[e_2] \downarrow v$$

“Any behavior of a (well-typed) client \mathcal{C} using e_1 can be matched by a behavior of the same client using e_2 ”

Very hard to prove: Quantification over all clients

Logical relations to the rescue!

Do not prove contextual refinement directly, but use a **binary logical relation**:

$$e_1 \lesssim e_2 : \tau$$

- ▶ $e_1 \lesssim e_2 : \tau$ is defined structurally on the type τ
- ▶ Does not involve quantification over all clients \mathcal{C}
- ▶ Soundness $e_1 \lesssim e_2 : \tau \implies e_1 \lesssim_{\text{ctx}} e_2 : \tau$ proved once and for all

A bit of history

Logical relations $e_1 \simeq e_2 : \tau$ are notoriously hard to define when having recursive types, higher-order state (type-world circularity), ...

A bit of history

Logical relations $e_1 \simeq e_2 : \tau$ are notoriously hard to define when having recursive types, higher-order state (type-world circularity), ...

Solutions to define such logical relations:

A bit of history

Logical relations $e_1 \simeq e_2 : \tau$ are notoriously hard to define when having recursive types, higher-order state (type-world circularity), ...

Solutions to define such logical relations:

- ▶ **Step-indexing** (Appel-McAllester, Ahmed, ...)

Solve circularities by stratifying everything by a natural number corresponding to the number of computation steps

A bit of history

Logical relations $e_1 \simeq e_2 : \tau$ are notoriously hard to define when having recursive types, higher-order state (type-world circularity), ...

Solutions to define such logical relations:

- ▶ **Step-indexing** (Appel-McAllester, Ahmed, ...)

Solve circularities by stratifying everything by a natural number corresponding to the number of computation steps

- ▶ **Logical approach** (LSLR, LADR, CaReSL, Iris, ...)

Hide step-indexing using modalities to obtain clearer definitions and proofs

A bit of history

Logical relations $e_1 \simeq e_2 : \tau$ are notoriously hard to define when having recursive types, higher-order state (type-world circularity), ...

Solutions to define such logical relations:

- ▶ **Step-indexing** (Appel-McAllester, Ahmed, ...)

Solve circularities by stratifying everything by a natural number corresponding to the number of computation steps

- ▶ **Logical approach** (LSLR, LADR, CaReSL, Iris, ...)

Hide step-indexing using modalities to obtain clearer definitions and proofs


We tried to take this one step further

Prove program refinements using inference rules à la concurrent separation logic

Instead of Hoare triples $\{P\} e \{Q\}$ we have refinement judgments $e_1 \preceq e_2 : \tau$


- ▶ Refinement proofs by symbolic execution as we know from separation logic
- ▶ Modular and conditional specifications
- ▶ Modeled using the “logical approach”

ReLoC: mechanized separation logic for interactive refinement proofs of fine-grained concurrent programs




ReLoC: mechanized separation logic for interactive refinement proofs of **fine-grained concurrent programs**

- ▶ **Fine-grained concurrency**: programs use low-level synchronization primitives for more granular parallelism




ReLoC: **mechanized** separation logic for interactive refinement proofs of fine-grained concurrent programs

- ▶ **Fine-grained concurrency:** programs use low-level synchronization primitives for more granular parallelism
- ▶ **Mechanized:** soundness proven sound using the Iris framework in Coq



ReLoC: mechanized separation logic for **interactive refinement proofs** of fine-grained concurrent programs

- ▶ **Fine-grained concurrency**: programs use low-level synchronization primitives for more granular parallelism
- ▶ **Mechanized**: soundness proven sound using the Iris framework in Coq
- ▶ **Interactive refinement proofs**: using high-level tactics in Coq



ReLoC: (simplified) grammar

$$P, Q \in \text{Prop} ::= \forall x. P \mid \exists x. P \mid P \vee Q \mid \dots$$

ReLoC: (simplified) grammar

$$P, Q \in \text{Prop} ::= \forall x. P \mid \exists x. P \mid P \vee Q \mid \dots$$
$$\mid P * Q \mid P \multimap Q \mid \ell \mapsto_i v \mid \ell \mapsto_s v$$

Separation logic for handling mutable state

- ▶ $\ell \mapsto_i v$ for the left-hand side (implementation)
- ▶ $\ell \mapsto_s v$ for the right-hand side (specification)

ReLoC: (simplified) grammar

$$\begin{aligned} P, Q \in \text{Prop} ::= & \forall x. P \mid \exists x. P \mid P \vee Q \mid \dots \\ & \mid P * Q \mid P \multimap Q \mid \ell \mapsto_i v \mid \ell \mapsto_s v \\ & \mid (\Delta \models e_1 \lesssim e_2 : \tau) \mid \llbracket \tau \rrbracket_{\Delta}(v_1, v_2) \mid \dots \end{aligned}$$

Separation logic for handling mutable state

- ▶ $\ell \mapsto_i v$ for the left-hand side (implementation)
- ▶ $\ell \mapsto_s v$ for the right-hand side (specification)

Logic with first-class refinement propositions to allow conditional refinements

- ▶ $(\ell_1 \mapsto_i v) \multimap (e_1 \lesssim e_2 : \tau)$
- ▶ $(e_1 \lesssim e_2 : \text{unit} \rightarrow \tau) \multimap (f(e_1) \lesssim e_2(); e_2() : \tau)$

Proving refinements of pure programs

Some rules for pure programs

Symbolic execution rules

$$\frac{\Delta \models K[e'_1] \lesssim e_2 : \tau \quad e_1 \rightarrow_{\text{pure}} e'_1}{\Delta \models K[e_1] \lesssim e_2 : \tau}^*$$

Some rules for pure programs

Symbolic execution rules

$$\frac{\Delta \models K[e'_1] \lesssim e_2 : \tau \quad e_1 \rightarrow_{\text{pure}} e'_1}{\Delta \models K[e_1] \lesssim e_2 : \tau}^*$$

$$\frac{\Delta \models e_1 \lesssim K[e'_2] : \tau \quad e_2 \rightarrow_{\text{pure}} e'_2}{\Delta \models e_1 \lesssim K[e_2] : \tau}^*$$

Some rules for pure programs

Structural rules

$$\frac{\Delta \Vdash e_1 \lesssim e_2 : \tau \quad * \quad \Delta \Vdash e'_1 \lesssim e'_2 : \tau'}{\Delta \Vdash (e_1, e'_1) \lesssim (e_2, e'_2) : \tau \times \tau'}^*$$

Some rules for pure programs

Structural rules

$$\frac{\Delta \Vdash e_1 \lesssim e_2 : \tau \quad * \quad \Delta \Vdash e'_1 \lesssim e'_2 : \tau'}{\Delta \Vdash (e_1, e'_1) \lesssim (e_2, e'_2) : \tau \times \tau'}^*$$

$$\frac{\exists(R : Val \times Val \rightarrow Prop). \quad [\alpha := R], \Delta \Vdash e_1 \lesssim e_2 : \tau}{\Delta \Vdash \text{pack}(e_1) \lesssim \text{pack}(e_2) : \exists \alpha. \tau}^*$$

Some rules for pure programs

Structural rules

$$\frac{\Delta \Vdash e_1 \simeq e_2 : \tau \quad * \quad \Delta \Vdash e'_1 \simeq e'_2 : \tau'}{\Delta \Vdash (e_1, e'_1) \simeq (e_2, e'_2) : \tau \times \tau'}^*$$

$$\frac{\exists(R : Val \times Val \rightarrow Prop). \quad [\alpha := R], \Delta \Vdash e_1 \simeq e_2 : \tau}{\Delta \Vdash \text{pack}(e_1) \simeq \text{pack}(e_2) : \exists \alpha. \tau}^*$$

$$\frac{[\tau]_{\Delta}(v_1, v_2)}{\Delta \Vdash v_1 \simeq v_2 : \tau}^* \quad \frac{\square \left(\forall v_1 v_2. [\tau]_{\Delta}(v_1, v_2) \rightarrow \Delta \Vdash e_1[v_1/x_1] \simeq e_2[v_2/x_2] : \sigma \right)}{\Delta \Vdash \lambda x_1. e_1 \simeq \lambda x_2. e_2 : \tau \rightarrow \sigma}^*$$

Example

A bit interface:

$$\text{bitT} \triangleq \exists \alpha. \alpha \times (\alpha \rightarrow \alpha) \times (\alpha \rightarrow \text{bool})$$

- ▶ constructor
- ▶ flip the bit
- ▶ view the bit as a Boolean

Example

A bit interface:

$$\text{bitT} \triangleq \exists \alpha. \alpha \times (\alpha \rightarrow \alpha) \times (\alpha \rightarrow \text{bool})$$

- ▶ constructor
- ▶ flip the bit
- ▶ view the bit as a Boolean

Two implementations:

$$\text{bit_bool} \triangleq \text{pack}(\text{true}, (\lambda b. \neg b), (\lambda b. b))$$

$$\text{bit_nat} \triangleq \text{pack}(1, (\lambda n. \text{if } n = 0 \text{ then } 1 \text{ else } 0), (\lambda n. n = 1))$$

Example

A bit interface:

$$\text{bitT} \triangleq \exists \alpha. \alpha \times (\alpha \rightarrow \alpha) \times (\alpha \rightarrow \text{bool})$$

- ▶ constructor
- ▶ flip the bit
- ▶ view the bit as a Boolean

Two implementations:

$$\text{bit_bool} \triangleq \text{pack}(\text{true}, (\lambda b. \neg b), (\lambda b. b))$$

$$\text{bit_nat} \triangleq \text{pack}(1, (\lambda n. \text{if } n = 0 \text{ then } 1 \text{ else } 0), (\lambda n. n = 1))$$

Refinement (and vice versa):

$$\text{bit_bool} \preceq \text{bit_nat} : \text{bitT}$$

Proof of the refinement

`pack(true, ($\lambda b. \neg b$), ($\lambda b. b$))`

\approx

`pack(1, ($\lambda n. \text{if } n = 0 \text{ then } 1 \text{ else } 0$), ($\lambda n. n = 1$))`

:

$\exists \alpha. \alpha \times (\alpha \rightarrow \alpha) \times (\alpha \rightarrow \text{bool})$

Proof of the refinement

`pack(true, ($\lambda b. \neg b$), ($\lambda b. b$))`

\approx

`pack(1, ($\lambda n. \text{if } n = 0 \text{ then } 1 \text{ else } 0$), ($\lambda n. n = 1$))`

:

$\exists \alpha. \alpha \times (\alpha \rightarrow \alpha) \times (\alpha \rightarrow \text{bool})$

Need to come with with an $R : \text{Val} \times \text{Val} \rightarrow \text{Prop}$

Proof of the refinement

where $R \triangleq \{(\text{true}, 1), (\text{false}, 0)\}$

$\text{pack}(\text{true}, (\lambda b. \neg b), (\lambda b. b))$

\approx

$\text{pack}(1, (\lambda n. \text{if } n = 0 \text{ then } 1 \text{ else } 0), (\lambda n. n = 1))$

:

$\exists \alpha. \alpha \times (\alpha \rightarrow \alpha) \times (\alpha \rightarrow \text{bool})$

Need to come with with an $R : \text{Val} \times \text{Val} \rightarrow \text{Prop}$

Proof of the refinement

$[\alpha := R] \models$ where $R \triangleq \{(\text{true}, 1), (\text{false}, 0)\}$

$(\text{true}, (\lambda b. \neg b), (\lambda b. b))$

\approx

$(1, (\lambda n. \text{if } n = 0 \text{ then } 1 \text{ else } 0), (\lambda n. n = 1))$

:

$\alpha \times (\alpha \rightarrow \alpha) \times (\alpha \rightarrow \text{bool})$

Proof of the refinement

$[\alpha := R] \models$ where $R \triangleq \{(\text{true}, 1), (\text{false}, 0)\}$

$(\text{true}, (\lambda b. \neg b), (\lambda b. b))$

\approx

$(1, (\lambda n. \text{if } n = 0 \text{ then } 1 \text{ else } 0), (\lambda n. n = 1))$

:

$\alpha \times (\alpha \rightarrow \alpha) \times (\alpha \rightarrow \text{bool})$

Use structural rule for products

Proof of the refinement

$[\alpha := R] \models$ where $R \triangleq \{(\text{true}, 1), (\text{false}, 0)\}$

$\text{true} \lesssim 1$: α

$(\lambda b. \neg b) \lesssim (\lambda n. \text{if } n = 0 \text{ then } 1 \text{ else } 0)$: $\alpha \rightarrow \alpha$

$(\lambda b. b) \lesssim (\lambda n. n = 1)$: $\alpha \rightarrow \text{bool}$

Proof of the refinement

$[\alpha := R] \models$ where $R \triangleq \{(\text{true}, 1), (\text{false}, 0)\}$

$\text{true} \approx 1$: α

$(\lambda b. \neg b) \approx (\lambda n. \text{if } n = 0 \text{ then } 1 \text{ else } 0)$: $\alpha \rightarrow \alpha$

$(\lambda b. b) \approx (\lambda n. n = 1)$: $\alpha \rightarrow \text{bool}$

Proof of the refinement

$[\alpha := R] \models$ where $R \triangleq \{(\text{true}, 1), (\text{false}, 0)\}$

$\text{true} \approx 1$: α (by def of R)

$(\lambda b. \neg b) \approx (\lambda n. \text{if } n = 0 \text{ then } 1 \text{ else } 0)$: $\alpha \rightarrow \alpha$

$(\lambda b. b) \approx (\lambda n. n = 1)$: $\alpha \rightarrow \text{bool}$

Proof of the refinement

$[\alpha := R] \models$ where $R \triangleq \{(\text{true}, 1), (\text{false}, 0)\}$

$\text{true} \rightsquigarrow 1$: α (by def of R)

$(\lambda b. \neg b) \rightsquigarrow (\lambda n. \text{if } n = 0 \text{ then } 1 \text{ else } 0)$: $\alpha \rightarrow \alpha$ (λ -rule + symb. exec.)

$(\lambda b. b) \rightsquigarrow (\lambda n. n = 1)$: $\alpha \rightarrow \text{bool}$

After using the λ rule and case analysis on R :

$\neg \text{true} \rightsquigarrow \text{if } 1 = 0 \text{ then } 1 \text{ else } 0$: α

$\neg \text{false} \rightsquigarrow \text{if } 0 = 0 \text{ then } 1 \text{ else } 0$: α

Proof of the refinement

$[\alpha := R] \models$ where $R \triangleq \{(\text{true}, 1), (\text{false}, 0)\}$

$\text{true} \rightsquigarrow 1$: α (by def of R)

$(\lambda b. \neg b) \rightsquigarrow (\lambda n. \text{if } n = 0 \text{ then } 1 \text{ else } 0)$: $\alpha \rightarrow \alpha$ (λ -rule + symb. exec.)

$(\lambda b. b) \rightsquigarrow (\lambda n. n = 1)$: $\alpha \rightarrow \text{bool}$ (λ -rule + symb. exec.)

Proof of the refinement

$[\alpha := R] \models$ where $R \triangleq \{(\text{true}, 1), (\text{false}, 0)\}$

$\text{true} \rightsquigarrow 1$: α (by def of R)

$(\lambda b. \neg b) \rightsquigarrow (\lambda n. \text{if } n = 0 \text{ then } 1 \text{ else } 0)$: $\alpha \rightarrow \alpha$ (λ -rule + symb. exec.)

$(\lambda b. b) \rightsquigarrow (\lambda n. n = 1)$: $\alpha \rightarrow \text{bool}$ (λ -rule + symb. exec.)

□

Reasoning about mutable state

Separation logic to the rescue!

“Vanilla” separation logic [O’Hearn, Reynolds, Yang; CSL’01]

Propositions P, Q denote ownership of resources

Points-to connective $\ell \mapsto v$:

Exclusive ownership of location ℓ with value v

Separating conjunction $P * Q$:

The resources consists of separate parts satisfying P and Q

Basic example:

$$\{\ell_1 \mapsto v_1 * \ell_2 \mapsto v_2\} \text{swap}(\ell_1, \ell_2) \{\ell_1 \mapsto v_2 * \ell_2 \mapsto v_1\}$$

the $*$ ensures that ℓ_1 and ℓ_2 are different memory locations

Mutable state and separation logic for refinements

There are two versions of the **points-to connective**:

- ▶ $\ell \mapsto_i v$ for the left-hand side/implementation
- ▶ $\ell \mapsto_s v$ for the right-hand side/specification

Mutable state and separation logic for refinements

There are two versions of the **points-to connective**:

- ▶ $l \mapsto_i v$ for the left-hand side/implementation
- ▶ $l \mapsto_s v$ for the right-hand side/specification

Example:

$$l_1 \mapsto_i 4 \quad * \quad l_2 \mapsto_s 0 \quad * \quad (!l_1) \lesssim (l_2 \leftarrow 4; !l_2) : \text{int}$$

Some rules for mutable state

Symbolic execution

$$\frac{l_1 \mapsto_i - \quad * \quad (l_1 \mapsto_i v_1 \quad * \quad \Delta \models K[()] \lesssim e_2 : \tau)}{\Delta \models K[l_1 \leftarrow v_1] \lesssim e_2 : \tau}^*$$

Some rules for mutable state

Symbolic execution

$$\frac{\ell_1 \mapsto_i - \quad * \quad (\ell_1 \mapsto_i v_1 \quad * \quad \Delta \Vdash K[()] \lesssim e_2 : \tau)}{\Delta \Vdash K[\ell_1 \leftarrow v_1] \lesssim e_2 : \tau}^*$$

$$\frac{\ell_2 \mapsto_s - \quad * \quad (\ell_2 \mapsto_s v_2 \quad * \quad \Delta \Vdash e_1 \lesssim K[()] : \tau)}{\Delta \Vdash e_1 \lesssim K[\ell_2 \leftarrow v_2] : \tau}^*$$

Some rules for mutable state

Symbolic execution

$$\frac{l_1 \mapsto_i - \quad * \quad (l_1 \mapsto_i v_1 \quad * \quad \Delta \Vdash K[()] \lesssim e_2 : \tau)}{\Delta \Vdash K[l_1 \leftarrow v_1] \lesssim e_2 : \tau}^*$$

$$\frac{l_2 \mapsto_s - \quad * \quad (l_2 \mapsto_s v_2 \quad * \quad \Delta \Vdash e_1 \lesssim K[()] : \tau)}{\Delta \Vdash e_1 \lesssim K[l_2 \leftarrow v_2] : \tau}^*$$

$$\frac{\forall l_1. l_1 \mapsto_i v_1 \quad * \quad \Delta \Vdash K[l_1] \lesssim e_2 : \tau}{\Delta \Vdash K[\text{ref}(v_1)] \lesssim e_2 : \tau}^* \quad \frac{\forall l_2. l_2 \mapsto_s v_2 \quad * \quad \Delta \Vdash e \lesssim K[l_2] : \tau}{\Delta \Vdash e_1 \lesssim K[\text{ref}(v_2)] : \tau}^*$$

Reasoning about higher-order functions and concurrency

State encapsulation

Modules with **encapsulated state**:

```
let x = ref(e) in (λy. ...)
```

The reference can only be used in the closure

State encapsulation

Modules with **encapsulated state**:

```
let x = ref(e) in (λy. ...)
```

The reference can only be used in the closure

Simple example:

```
counter  $\triangleq$  (λ(). let x = ref(1) in (λ(). FAA(x, 1))) : unit → (unit → int)
```

- ▶ `counter()` constructs an instance `c : unit → int` of the counter module
- ▶ Calling `c()` in subsequently gives `1, 2, ...`
- ▶ The reference `x` is private to the module

The problem

Modules with **encapsulated state**:

$$\text{let } x = \text{ref}(e) \text{ in } \underbrace{(\lambda y. \dots)}_f$$

The problem

Modules with **encapsulated state**:

$$\text{let } x = \text{ref}(e) \text{ in } \underbrace{(\lambda y. \dots)}_f$$

Reasoning about such modules is challenging:

- ▶ f can be called multiple times by clients

So, the value of x can change in each call

The problem

Modules with **encapsulated state**:

$$\text{let } x = \text{ref}(e) \text{ in } \underbrace{(\lambda y. \dots)}_f$$

Reasoning about such modules is challenging:

- ▶ f can be called multiple times by clients
So, the value of x can change in each call
- ▶ f can even be called even in parallel!
So, f cannot get exclusive access to $x \mapsto v$

The problem

Modules with **encapsulated state**:

$$\text{let } x = \text{ref}(e) \text{ in } \underbrace{(\lambda y. \dots)}_f$$

Reasoning about such modules is challenging:

- ▶ f can be called multiple times by clients
So, the value of x can change in each call
- ▶ f can even be called even in parallel!
So, f cannot get exclusive access to $x \mapsto v$

We need to guarantee that closures do not get access to exclusive resources

Persistent resources

The “persistent” modality \Box in Iris/ReLoC:

$\Box P \triangleq$ “ P holds **without** assuming exclusive resources”

Examples:

- ▶ Equality is persistent: $(x = y) \vdash \Box(x = y)$
- ▶ Points-to connectives are not: $((\ell \mapsto v) \not\vdash \Box(\ell \mapsto v))$
- ▶ More examples later. . .

ReLoC's λ -rule again

The \square modality makes sure no exclusive resources can escape into closures:

$$\frac{\square \left(\begin{array}{c} \forall v_1 v_2. \llbracket \tau \rrbracket_{\Delta}(v_1, v_2) \multimap * \\ \Delta \Vdash e_1[v_1/x_1] \lesssim e_2[v_2/x_2] : \sigma \end{array} \right)}{\Delta \Vdash \lambda x_1. e_1 \lesssim \lambda x_2. e_2 : \tau \rightarrow \sigma}^*$$

ReLoC's λ -rule again

The \square modality makes sure no exclusive resources can escape into closures:

$$\frac{\square \left(\frac{\forall v_1 v_2. \llbracket \tau \rrbracket_{\Delta}(v_1, v_2) \multimap \Delta \Vdash e_1[v_1/x_1] \lesssim e_2[v_2/x_2] : \sigma}{\Delta \Vdash \lambda x_1. e_1 \lesssim \lambda x_2. e_2 : \tau \rightarrow \sigma} \right)^*}{\Delta \Vdash \lambda x_1. e_1 \lesssim \lambda x_2. e_2 : \tau \rightarrow \sigma}^*$$

Prohibits “wrong” refinements, for example:

$$\left(\lambda(). 1 \right) \not\lesssim_{ctx} \left(\text{let } x = \text{ref}(0) \text{ in } (\lambda(). x \leftarrow (1 + !x); !x) \right)$$

Due to \square , the resource $x \mapsto_s 0$ cannot be used to prove the closure

But it should be possible to use resources in closures

For example:

$$\left(\lambda(). \text{let } x = \text{ref}(1) \text{ in } (\lambda(). \text{FAA}(x, 1)) \right)$$
$$\approx$$
$$\left(\begin{array}{l} \lambda(). \text{let } x = \text{ref}(1), l = \text{newlock } () \text{ in} \\ \quad \lambda(). \text{acquire}(l); \\ \quad \quad \text{let } v = !x \text{ in} \\ \quad \quad x \leftarrow v + 1; \\ \quad \quad \text{release}(l); v \end{array} \right)$$

Iris-style Invariants

The invariant connective \boxed{R}

expresses that R is maintained as an invariant on the state

Iris-style Invariants

The invariant connective \boxed{R}

expresses that R is maintained as an invariant on the state

Invariants allow to share resources:

- ▶ A resource R can be turned into \boxed{R} at any time
- ▶ Invariants are persistent: $\boxed{R} \vdash \Box \boxed{R}$
- ▶ ... thus can be used to prove closures

Iris-style Invariants

The invariant connective \boxed{R}

expresses that R is maintained as an invariant on the state

Invariants allow to share resources:

- ▶ A resource R can be turned into \boxed{R} at any time
- ▶ Invariants are persistent: $\boxed{R} \vdash \Box \boxed{R}$
- ▶ ... thus can be used to prove closures

But that comes with a cost:

- ▶ Invariants \boxed{R} can only be accessed during atomic steps on the left-hand side
- ▶ ... while multiple steps on the right-hand side can be performed

Example

```
let x = ref(1) in (λ(). FAA(x, 1))
```

\rightsquigarrow

```
let x = ref(1) , l = newlock () in  
  (λ(). acquire(l);  
    let v = !x in  
    x ← v + 1;  
    release(l); v)
```

Example

```
let x = ref(1) in (λ(). FAA(x, 1))
```

\rightsquigarrow

```
let x = ref(1), l = newlock () in  
  (λ(). acquire(l);  
    let v = !x in  
    x ← v + 1;  
    release(l); v)
```

Example

$x_1 \mapsto_i 1$

$(\lambda(). \text{FAA}(x_1, 1))$

\rightsquigarrow

```
let x = ref(1) , l = newlock () in
  (\lambda(). acquire(l);
    let v = !x in
      x ← v + 1;
      release(l); v)
```

Example

$x_1 \mapsto_i 1$

$(\lambda(). \text{FAA}(x_1, 1))$

\approx

```
let x = ref(1), l = newlock () in
  (\lambda(). acquire(l);
    let v = !x in
      x ← v + 1;
      release(l); v)
```

Example

$x_1 \mapsto_i 1$

$x_2 \mapsto_s 1$

$(\lambda(). \text{FAA}(x_1, 1))$

\rightsquigarrow

```
let l = newlock () in
  (\lambda(). acquire(l);
    let v = !x2 in
      x2 ← v + 1;
      release(l); v)
```

Example

$x_1 \mapsto_i 1$

$x_2 \mapsto_s 1$

$(\lambda(). \text{FAA}(x_1, 1))$

\rightsquigarrow

`let l = newlock () in`

`(\lambda(). acquire(l);`

`let v = !x2 in`

`x2 ← v + 1;`

`release(l); v)`

Example

$x_1 \mapsto_i 1$

$x_2 \mapsto_s 1$

`isLock(l, unlocked)`

$(\lambda(). \text{FAA}(x_1, 1))$

\rightsquigarrow

$(\lambda(). \text{acquire}(l);$

`let $v = !x_2$ in`

$x_2 \leftarrow v + 1;$

`release(l); v)`

Example

$\exists n.$

$x_1 \mapsto_i n$

$x_2 \mapsto_s n$

`isLock(l, unlocked)`

`($\lambda().$ FAA(x_1 , 1))`

\rightsquigarrow

`($\lambda().$ acquire(l);`

`let $v = !x_2$ in`

`$x_2 \leftarrow v + 1$;`

`release(l); v)`

Example

| |
|---|
| $\exists n. x_1 \mapsto_i n *$ |
| $x_2 \mapsto_s n *$ |
| <code>isLock(<i>l</i>, unlocked)</code> |

$(\lambda(). \text{FAA}(x_1, 1))$

\rightsquigarrow

```
( $\lambda().$  acquire(l);  
  let  $v = !x_2$  in  
   $x_2 \leftarrow v + 1$ ;  
  release(l);  $v$ )
```

Example

| |
|---|
| $\exists n. x_1 \mapsto_i n *$ |
| $x_2 \mapsto_s n *$ |
| <code>isLock(<i>l</i>, unlocked)</code> |

`FAA(x1, 1)`

\sim

`acquire(l);`

`let v = !x2 in`

`x2 ← v + 1;`

`release(l); v`

Example

| |
|----------------------------------|
| $\exists n. x_1 \mapsto_i n *$ |
| $x_2 \mapsto_s n *$ |
| <code>isLock(l, unlocked)</code> |

`FAA(x1, 1)`

\approx

```
acquire(l);  
let v = !x2 in  
x2 ← v + 1;  
release(l); v
```

Example

```
∃n. x1 ↦i n *  
    x2 ↦s n *  
    isLock(l, unlocked)
```

```
x1 ↦i n  
x2 ↦s n  
isLock(l, unlocked)
```

```
FAA(x1, 1)
```

≈

```
acquire(l);  
let v = !x2 in  
x2 ← v + 1;  
release(l); v
```

Example

```
∃n. x1 ↦i n *  
    x2 ↦s n *  
    isLock(l, unlocked)
```

x₁ ↦_i n + 1

x₂ ↦_s n

isLock(l, unlocked)

n

↗

acquire(l);

let v = !x₂ in

x₂ ← v + 1;

release(l); v

Example

```
∃n. x1 ↦i n *  
    x2 ↦s n *  
    isLock(l, unlocked)
```

x₁ ↦_i n + 1

x₂ ↦_s n

isLock(l, unlocked)

n

↗

acquire(l);

let v = !x₂ in

x₂ ← v + 1;

release(l); v

Example

```
∃n. x1 ↦i n *  
    x2 ↦s n *  
    isLock(l, unlocked)
```

x₁ ↦_i n + 1

x₂ ↦_s n

isLock(l, locked)

n

↗

let v = !x₂ in

x₂ ← v + 1;

release(l); v

Example

```
 $\exists n. x_1 \mapsto_i n *$   
 $x_2 \mapsto_s n *$   
isLock(l, unlocked)
```

$x_1 \mapsto_i n + 1$

$x_2 \mapsto_s n$

`isLock(l, locked)`

`n`

\approx

`let v = !x2 in`

`x2 ← v + 1;`

`release(l); v`

Example

$\exists n. x_1 \mapsto_i n *$
 $x_2 \mapsto_s n *$
 $\text{isLock}(l, \text{unlocked})$

$x_1 \mapsto_i n + 1$

$x_2 \mapsto_s n$

$\text{isLock}(l, \text{locked})$

n

\simeq

$x_2 \leftarrow n + 1;$

$\text{release}(l); n$

Example

$\exists n. x_1 \mapsto_i n *$
 $x_2 \mapsto_s n *$
 $\text{isLock}(l, \text{unlocked})$

$x_1 \mapsto_i n + 1$

$x_2 \mapsto_s n$

$\text{isLock}(l, \text{locked})$

n

\sim

$x_2 \leftarrow n + 1;$

$\text{release}(l); n$

Example

$\exists n. x_1 \mapsto_i n *$
 $x_2 \mapsto_s n *$
 $\text{isLock}(l, \text{unlocked})$

$x_1 \mapsto_i n + 1$

$x_2 \mapsto_s n + 1$

$\text{isLock}(l, \text{locked})$

n

\surd

$\text{release}(l); n$

Example

$\exists n. x_1 \mapsto_i n *$
 $x_2 \mapsto_s n *$
 $\text{isLock}(l, \text{unlocked})$

$x_1 \mapsto_i n + 1$

$x_2 \mapsto_s n + 1$

$\text{isLock}(l, \text{locked})$

n

\sim

$\text{release}(l); n$

Example

$\exists n. x_1 \mapsto_i n *$
 $x_2 \mapsto_s n *$
 $\text{isLock}(l, \text{unlocked})$

$x_1 \mapsto_i n + 1$

$x_2 \mapsto_s n + 1$

$\text{isLock}(l, \text{unlocked})$

n

\surd

n

Example

$\exists n. x_1 \mapsto_i n *$

$x_2 \mapsto_s n *$

$\text{isLock}(l, \text{unlocked})$


n

γ_2

n

Wrapping up...

- ▶ ReLoC provides rules allowing this kind of simulation reasoning, formally
- ▶ The example can be done in Coq in almost the same fashion
- ▶ The approach scales to: lock-free concurrent data structures, generative ADTs, examples from the logical relations literature



Implementation in Coq

ReLoC is build on top of the **Iris framework**, so we can inherit:

- ▶ Iris's invariants
- ▶ Iris's ghost state
- ▶ Iris's Coq infrastructure
- ▶ ...



The proofs we have done in Coq

ReLoC judgments $e_1 \lesssim e_2 : \tau$ are modeled as a shallow embedding using the “logical approach” to logical relations

Proved in Coq:

- ▶ Proof rules: All the ReLoC rules hold in the shallow embedding
- ▶ Soundness: $e_1 \lesssim e_2 : \tau \implies e_1 \lesssim_{ctx} e_2 : \tau$
- ▶ Actual program refinements: concurrent data structures, and examples from the logical relations literature

Need to reason in separation logic!

Iris Proof Mode (IPM) [Krebbers *et al.*; POPL'17]

Lemma test {A} (P Q : iProp) ($\Psi : A \rightarrow \text{iProp}$) :
P * ($\exists a, \Psi a$) * Q \multimap Q * $\exists a, P * \Psi a$.

Proof.

```
iIntros "[H1 [H2 H3]]".  
iDestruct "H2" as (x) "H2".  
iSplitL "H3".  
- iAssumption.  
- iExists x.  
  iFrame.
```

Qed.

Iris Proof Mode (IPM) [Krebbers et al.; POPL'17]

Lemma test {A} (P Q : iProp) ($\Psi : A \rightarrow \text{iProp}$) :
P * ($\exists a, \Psi a$) * Q \rightarrow Q * $\exists a, P * \Psi a$.

Proof.

iInt **Lemma in the Iris logic**
iDesolve H2 as (x) H2'.
iSplitL "H3".
- iAssumption.
- iExists x.
iFrame.

Qed.

Iris Proof Mode (IPM) [Krebbers et al.; POPL'17]

```
Lemma test {A} (P Q : iProp) ( $\Psi : A \rightarrow iProp$ ) :  
  P * ( $\exists a, \Psi a$ ) * Q -* Q *  $\exists a, P * \Psi a$ .
```

Proof.

```
iIntros "[H1 [H2 H3]]".  
iDestruct "H2" as (x) "H2".  
iSplitL "H3".  
- iAssumption.  
- iExists x.  
  iFrame.
```

Qed.

1 subgoal

A : Type

P, Q : iProp

$\Psi : A \rightarrow iProp$

x : A

----- (1/1)

"H1" : P

"H2" : Ψx

"H3" : Q

-----*

Q * ($\exists a : A, P * \Psi a$)

Iris Proof Mode (IPM) [Krebbers et al.; POPL'17]

```
Lemma test {A} (P Q : iProp) ( $\Psi$  : A  $\rightarrow$  iProp) :  
  P * ( $\exists$  a,  $\Psi$  a) * Q  $\multimap$  Q *  $\exists$  a, P *  $\Psi$  a.
```

Proof.

```
iIntros "[H1 [H2 H3]]".  
iDestruct "H2" as (x) "H2".  
iSplitL "H3".  
- iAssumption.  
- iExists x.  
  iFrame.
```

Qed.

1 subgoal

A : Type

P, Q : iProp

Ψ : A \rightarrow iProp

x : A

----- (1/1)

"H1" : P

"H2" : Ψ x

"H3" : Q

-----*

Q * (\exists a : A, P * Ψ a)

* means: resources should be split

Iris Proof Mode (IPM) [Krebbers et al.; POPL'17]

```
Lemma test {A} (P Q : iProp) ( $\Psi : A \rightarrow iProp$ ) :  
  P * ( $\exists a, \Psi a$ ) * Q  $\multimap$  Q *  $\exists a, P * \Psi a$ .
```

Proof.

```
iIntros "[H1 [H2 H3]]".  
iDestruct "H2" as (x) "H2".  
iSplitL "H3".  
- iAssumption
```

The hypotheses for the left conjunct

Qed.

1 subgoal

A : Type

P, Q : iProp

$\Psi : A \rightarrow iProp$

x : A

----- (1/1)

"H1" : P

"H2" : Ψx

"H3" : Q

-----*

Q * ($\exists a : A, P * \Psi a$)

* means: resources should be split

Iris Proof Mode (IPM) [Krebbers et al.; POPL'17]

```
Lemma test {A} (P Q : iProp) ( $\Psi : A \rightarrow iProp$ ) :  
  P * ( $\exists a, \Psi a$ ) * Q  $\multimap$  Q *  $\exists a, P * \Psi a$ .
```

Proof.

```
iIntros "[H1 [H2 H3]]".  
iDestruct "H2" as (x) "H2".  
iSplitL "H3".  
- iAssumption
```

The hypotheses for the left conjunct

Qed.

2 subgoals

A : Type

P, Q : iProp

$\Psi : A \rightarrow iProp$

x : A

----- (1/2)

"H3" : Q

----- *

Q

----- (2/2)

"H1" : P

"H2" : Ψx

----- *

$\exists a : A, P * \Psi a$

Iris Proof Mode (IPM) [Krebbers *et al.*; POPL'17]

Lemma test {A} (P Q : iProp) ($\Psi : A \rightarrow \text{iProp}$) :
P * ($\exists a, \Psi a$) * Q \multimap Q * $\exists a, P * \Psi a$.

Proof.

```
iIntros "[H1 [H2 H3]]".  
iDestruct "H2" as (x) "H2".  
iSplitL "H3".  
- iAssumption.  
- iExists x.  
  iFrame.
```

Qed.

Iris Proof Mode (IPM) [Krebbers et al.; POPL'17]

```
Lemma test {A} (P Q : iProp) ( $\Psi : A \rightarrow iProp$ ) :  
  P * ( $\exists a, \Psi a$ ) * Q -* Q *  $\exists a, P * \Psi a$ .
```

Proof.

```
iIntros "[H1 [H2 H3]]".  
by iFrame.
```

Qed.

No more subgoals.

We can also solve this lemma automatically

ReLoC in Iris Proof Mode

- ▶ The ReLoC rules are just lemmas that can be `iApplied`
- ▶ We have more automated support for symbolic execution
- ▶ Iris Proof Mode features a special context for persistent hypotheses, which is crucial for dealing with invariants

Ongoing work:

Proving security properties using
relational reasoning in separation logic

Language-based security

Program variables are divided into two groups:

- ▶ **low-sensitivity** variables l_1, l_2, \dots
- ▶ **high-sensitivity** variables h_1, h_2, \dots

Confidentiality: the data stored in high-sensitivity variables should not leak to low-sensitivity variables, e.g., $l_1 \leftarrow !h_1 + 1$ does not happen

Proved via **non-interference**: changing the values of h_1, h_2, \dots and running the program does not affect the resulting values of l_1, l_2, \dots

Type systems for non-interference

Type system where types are annotated with labels from a lattice $\mathbf{L} \sqsubseteq \mathbf{H}$

$$\vdash l_i : \text{ref int}^{\mathbf{L}}$$

$$\vdash h_i : \text{ref int}^{\mathbf{H}}$$

$$\frac{\vdash e : \text{ref int}^{\chi} \quad \vdash e' : \text{int}^{\xi} \quad \xi \sqsubseteq \chi}{\vdash e \leftarrow e' : \text{unit}}$$

$$\frac{\vdash e : \text{int}^{\chi} \quad \vdash e' : \text{int}^{\xi}}{\vdash e + e' : \text{int}^{\chi \sqcup \xi}}$$

Type systems for non-interference

Type system where types are annotated with labels from a lattice $\mathbf{L} \sqsubseteq \mathbf{H}$

$$\frac{\vdash l_i : \text{ref int}^{\mathbf{L}} \quad \vdash e' : \text{int}^{\xi} \quad \xi \sqsubseteq \chi}{\vdash e \leftarrow e' : \text{unit}} \quad \frac{\vdash h_i : \text{ref int}^{\mathbf{H}} \quad \vdash e : \text{int}^{\chi} \quad \vdash e' : \text{int}^{\xi}}{\vdash e + e' : \text{int}^{\chi \sqcup \xi}}$$

Example:

$$\frac{\vdash l_1 : \text{ref int}^{\mathbf{L}} \quad \frac{\vdash !h_1 : \text{int}^{\mathbf{H}} \quad \vdash 1 : \text{int}^{\mathbf{L}}}{\vdash !h_1 + 1 : \text{int}^{\mathbf{H}}}}{\vdash l_1 \leftarrow !h_1 + 1 : \text{unit}} \quad \mathbf{H} \sqsubseteq \mathbf{L}$$

Shortcomings of type systems

- ▶ Type systems can be extended to cover more PL features (dynamic references, higher-order functions, exceptions), although it is not straightforward
- ▶ Type systems are too weak: in many situations non-interference depends on functional correctness

Example: value-dependent classification

```
let r = { data = ref(secret);  
         is_classified = ref(true) } in  
  
while true do  
  if  $\neg$  ! r.is_classified  
  then out  $\leftarrow$  ! r.data  
  else ();  
   $\left\| \begin{array}{l} r.data \leftarrow 0; \\ r.is\_classified \leftarrow \text{false} \end{array} \right.$ 
```

The classification of $r.data$ depends on the run-time value $r.is_classified$

Example: value-dependent classification

```
let  $r = \left\{ \begin{array}{l} data = \text{ref}(\text{secret}); \\ is\_classified = \text{ref}(\text{true}) \end{array} \right\}$  in
```

```
while true do  
  if  $\neg !r.is\_classified$   
  then  $out \leftarrow !r.data$   
  else ();  
   $r.data \leftarrow 0;$   
   $r.is\_classified \leftarrow \text{false}$ 
```

The classification of $r.data$ depends on the run-time value $r.is_classified$

- ▶ Can we type this program with conventional type systems?
- ▶ Is this program secure?

Example: value-dependent classification

```
let  $r = \left\{ \begin{array}{l} data = \text{ref}(\text{secret}); \\ is\_classified = \text{ref}(\text{true}) \end{array} \right\}$  in
```

```
while true do  
  if  $\neg !r.is\_classified$   
  then  $out \leftarrow !r.data$   
  else ();  
   $r.data \leftarrow 0;$   
   $r.is\_classified \leftarrow \text{false}$ 
```

The classification of $r.data$ depends on the run-time value $r.is_classified$

- ▶ Can we type this program with conventional type systems? **No**
- ▶ Is this program secure? **Yes**

SeLoC: a relational extension of Iris for non-interference

- ▶ A relational variant of weakest preconditions
- ▶ A type system built on top of SeLoC using logical relations
- ▶ Soundness w.r.t. scheduler-independent notion of non-interference

SeLoC: a relational extension of Iris for non-interference

- ▶ A relational variant of weakest preconditions
 - ▶ to combine reasoning about non-interference with functional correctness
 - ▶ in the presence of fine-grained concurrency
- ▶ A type system built on top of SeLoC using logical relations

- ▶ Soundness w.r.t. scheduler-independent notion of non-interference

SeLoC: a relational extension of Iris for non-interference

- ▶ A relational variant of weakest preconditions
 - ▶ to combine reasoning about non-interference with functional correctness
 - ▶ in the presence of fine-grained concurrency
- ▶ A type system built on top of SeLoC using logical relations
 - ▶ Compatibility rules for composing typed programs
 - ▶ Can “drop down” to separation logic to prove more complicated programs
- ▶ Soundness w.r.t. scheduler-independent notion of non-interference

Double weakest precondition

The basic component of SeLoC:

$$\text{dwp } e_1 \ \& \ e_2 \ \{\Phi\}$$

Double weakest precondition

The basic component of SeLoC:

$$\text{dwp } e_1 \ \& \ e_2 \ \{\Phi\}$$

- ▶ Any two runs of the e_1 and e_2 are in a bisimulation and their results satisfy Φ

Double weakest precondition

The basic component of SeLoC:

$$\text{dwp } e_1 \ \& \ e_2 \ \{\Phi\}$$

- ▶ Any two runs of the e_1 and e_2 are in a bisimulation and their results satisfy Φ
- ▶ e_1 and e_2 have different secret data, but must produce the same public output

Double weakest precondition

The basic component of SeLoC:

$$\text{dwp } e_1 \ \& \ e_2 \ \{\Phi\}$$

- ▶ Any two runs of the e_1 and e_2 are in a bisimulation and their results satisfy Φ
- ▶ e_1 and e_2 have different secret data, but must produce the same public output
- ▶ Left-hand side and right-hand side resources: $\ell_1 \mapsto_L v_1$ and $\ell_2 \mapsto_R v_2$

Double weakest precondition

The basic component of SeLoC:

$$\text{dwp } e_1 \ \& \ e_2 \ \{\Phi\}$$

- ▶ Any two runs of the e_1 and e_2 are in a bisimulation and their results satisfy Φ
- ▶ e_1 and e_2 have different secret data, but must produce the same public output
- ▶ Left-hand side and right-hand side resources: $\ell_1 \mapsto_L v_1$ and $\ell_2 \mapsto_R v_2$
- ▶ Rules are similar to ReLoC, but require execution in lock-step

Double weakest precondition

The basic component of SeLoC:

$$\text{dwp } e_1 \ \& \ e_2 \ \{\Phi\}$$

- ▶ Any two runs of the e_1 and e_2 are in a bisimulation and their results satisfy Φ
- ▶ e_1 and e_2 have different secret data, but must produce the same public output
- ▶ Left-hand side and right-hand side resources: $\ell_1 \mapsto_L v_1$ and $\ell_2 \mapsto_R v_2$
- ▶ Rules are similar to ReLoC, but require execution in lock-step
- ▶ Soundness statement:

$$(\forall h_1, h_2 \in \mathbb{Z}. I_{out} \vdash \text{dwp } e[h_1/x] \ \& \ e[h_2/x] \ \{v_1 v_2. v_1 = v_2\}) \implies e \text{ is secure}$$

$$I_{out} \triangleq \boxed{\exists v \in \mathbb{Z}. out \mapsto_L v * out \mapsto_R v}$$

Proof of the example: value-dependent classification

```
let  $r = \left\{ \begin{array}{l} data = \text{ref}(\text{secret}); \\ is\_classified = \text{ref}(\text{true}) \end{array} \right\}$  in
```

```
while true do  
  if  $\neg !r.is\_classified$   
  then  $out \leftarrow !r.data$   
  else ();  
   $r.data \leftarrow 0;$   
   $r.is\_classified \leftarrow \text{false}$ 
```



Proof of the example: value-dependent classification

```
let r = { data = ref(secret);  
         is_classified = ref(true) } in
```

```
while true do  
  if  $\neg$  !r.is_classified  
  then out  $\leftarrow$  !r.data  
  else ();  
  r.data  $\leftarrow$  0;  
  r.is_classified  $\leftarrow$  false
```



Proof of the example: value-dependent classification

```
let  $r = \left\{ \begin{array}{l} data = \text{ref}(\text{secret}); \\ is\_classified = \text{ref}(\text{true}) \end{array} \right\}$  in
```


```
while true do  
  if  $\neg !r.is\_classified$   
  then  $out \leftarrow !r.data$   
  else ();  
   $r.data \leftarrow 0;$   
   $r.is\_classified \leftarrow \text{false}$ 
```



Proof of the example: value-dependent classification

```
let  $r = \left\{ \begin{array}{l} data = \text{ref}(\text{secret}); \\ is\_classified = \text{ref}(\text{true}) \end{array} \right\}$  in
```

```
while true do  
  if  $\neg !r.is\_classified$   
  then  $out \leftarrow !r.data$   
  else ();  
   $r.data \leftarrow 0;$   
   $r.is\_classified \leftarrow \text{false}$ 
```



Proof of the example: value-dependent classification

$$\begin{aligned} & (\text{in_state}(\text{Classified}) * \exists i_1, i_2. r_1.\text{is_classified} \mapsto_L \text{true} * \\ & \quad r_2.\text{is_classified} \mapsto_R \text{true} * r_1.\text{data} \mapsto_L i_1 * r_2.\text{data} \mapsto_R i_2) \\ \vee & (\text{in_state}(\text{Intermediate}) * \exists i. r_1.\text{is_classified} \mapsto_L \text{true} * \\ & \quad r_2.\text{is_classified} \mapsto_R \text{true} * r_1.\text{data} \mapsto_L i * r_2.\text{data} \mapsto_R i) \\ \vee & (\text{in_state}(\text{Declassified}) * \exists i. r_1.\text{is_classified} \mapsto_L \text{false} * \\ & \quad r_2.\text{is_classified} \mapsto_R \text{false} * r_1.\text{data} \mapsto_L i * r_2.\text{data} \mapsto_R i) \end{aligned}$$



Conclusions

ReLoC: contextual refinements

ReLoC: A Mechanised Relational Logic for Fine-Grained Concurrency

Dan Frumin
Radboud University
dfrumin@cs.ru.nl

Robbert Krebbers
Delft University of Technology
mail@robbertkrebbers.nl

Lars Birkedal
Aarhus University
birkedal@cs.au.dk

Abstract

We present ReLoC: a logic for proving refinements of programs in a language with higher-order state, fine-grained concurrency, polymorphism and recursive types. The core of our logic is a judgement $e \lesssim e' : \tau$, which expresses that a program e refines a program e' at type τ . In contrast to earlier work on refinements for languages with higher-order state and concurrency, ReLoC provides type- and structure-directed rules for manipulating this judgement, whereas previously, such proofs were carried out by unfolding the judgement into its definition in the model. These more abstract proof rules make it simpler to carry out refinement proofs.

Moreover, we introduce *logically atomic relational specifications*: a novel approach for relational specifications for compound expressions that take effect at a single instant in time. We demonstrate how to formalise and prove such relational specifications in ReLoC, allowing for more modular proofs.

ReLoC is built on top of the expressive concurrent separation logic Iris, allowing us to leverage features of Iris such as invariants and ghost state. We provide a mechanisation of our logic in Coq, which does not just contain a proof of soundness, but also facilitates for interactively carrying out refinements proofs. We have used these tactics to mechanise several examples, which demonstrates the practicality and modularity of our logic.

CCS Concepts • Theory of computation → Logic and verification; Separation logic; Concurrency; Program verification;

Keywords Separation logic, logical relations, fine-grained concurrency, Iris, atomicity

ACM Reference Format:

Dan Frumin, Robbert Krebbers, and Lars Birkedal. 2018. ReLoC: A Mechanised Relational Logic for Fine-Grained Concurrency. In *LICS '18: LICS '18: 33rd Annual ACM/IEEE Symposium on Logic in Computer Science*, July

```
read  $\triangleq \lambda x (). !x$ 
incx  $\triangleq \lambda x l. \text{acquire } l; \text{let } n = !x \text{ in } x \leftarrow 1 + n; \text{release } l; n$ 
counterx  $\triangleq \text{let } l = \text{newlock } () \text{ in let } x = \text{ref}(0) \text{ in}$ 
  (read  $x, \lambda (). \text{inc}_x x l$ )
incx  $\triangleq \text{rec inc } x = \text{let } e = !x \text{ in}$ 
  if CAS( $x, e, 1 + e$ ) then e else inc x
counterx  $\triangleq \text{let } x = \text{ref}(0) \text{ in (read } x, \lambda (). \text{inc}_x x)$ 
```

Figure 1. Two concurrent counter implementations.

are often referred to as the gold standards of equivalence and refinement of program expressions: contextual equivalence of e and e' means that it is safe for a compiler to replace any occurrence of e by e' , and contextual refinement is often used to specify the behaviour of programs, e.g., one can show the correctness of a fine-grained concurrent implementation of an abstract data type by proving that it contextually refines a coarse-grained implementation, which is understood as the specification.

A simple example is the specification of a fine-grained concurrent counter by a coarse-grained version, $\text{counter}_x \lesssim \text{counter}_x : (1 \rightarrow \mathbb{N}) \times (1 \rightarrow \mathbb{N})$; see Figure 1 for the code. The increment operation of the coarse-grained version, counter_x , is performed inside a critical section guarded by a lock, whereas the fine-grained version, counter_x , takes an “optimistic” lock-free approach to incrementing the value using a compare-and-set inside a loop. We will use the counter as a simple running example throughout the paper. Proving program refinements and equivalence directly is difficult because of the quantification over all contexts. As such, it is often the case that

SeLoC: non-interference

Compositional Non-Interference for Fine-Grained Concurrent Programs

Dan Frumin
Radboud University

Robbert Krebbers
Delft University of Technology

Lars Birkedal
Aarhus University

Abstract—We present SeLoC: a relational separation logic for verifying non-interference of fine-grained concurrent programs in a compositional way. SeLoC is more expressive than previous approaches, both in terms of the features of the target programming language, and in terms of the logic. The target programming language supports dynamically allocated references (pointers), higher-order functions, and fine-grained fork-based concurrency with low-level atomic operators like compare-and-set. The logic provides an invariant mechanism to establish protocols on data that is not protected by locks. This allows us to verify programs that were beyond the reach of previous approaches.

A key technical innovation in SeLoC is a relational version of weakest-preconditions to track information flow using separation logic resources. On top of these weakest-preconditions we build a type system-like abstraction, using invariants and logical relations. SeLoC has been mechanized on top of the Iris framework in the Coq proof assistant.

Index Terms—non-interference, fine-grained concurrency, invariants, logical relations, separation logic, Coq, Iris

I. INTRODUCTION

Non-interference is a form of *information flow control* (IFC) used to express security properties like confidentiality and secrecy, which guarantee that confidential information does not leak to attackers. In order to establish non-interference of programs used in practice, it is necessary to develop techniques that scale up to programming paradigms and programming constructs found in modern programming languages. Much effort has been put into that direction—e.g., to support dynamically allocated references and higher-order functions [11], [13] and

about each single run of a program, non-interference is stated in terms of multiple runs of the same program. One has to show that for different values of confidential inputs, the attacker cannot observe a different behavior.

Another reason for the discrepancy between the lack of expressiveness for techniques for non-interference compared to those for functional correctness is that a lot of prior work on non-interference has focused on type systems and type system-like logics, e.g., [1], [4], [6], [9], [10]. Such systems have the benefit of providing strong automation (by means of type checking), but lack capabilities to reason about functional correctness, and therefore to establish non-interference of more challenging programs.

In order to overcome aforementioned shortcomings, we take a different and more expressive approach that combines the power of type systems and concurrent separation logic. In our approach, one assigns flexible interfaces to individual program modules using types. The program modules can then be composed using typing rules, ensuring non-interference of the whole system. Individual programs can be verified against those interfaces using a relational concurrent separation logic, which allows one to carry out non-interference proofs intertwined with functional correctness proofs.

Although ideas from concurrent separation logic have been employed for establishing non-interference (for first-order programs) before, see [9], [10], we believe that the combination

At LICS'18

Under submission

Thank you!

Download **ReLoC** at <https://gitlab.mpi-sws.org/iris/reloc>

Download **Iris** at <https://iris-project.org/>

