



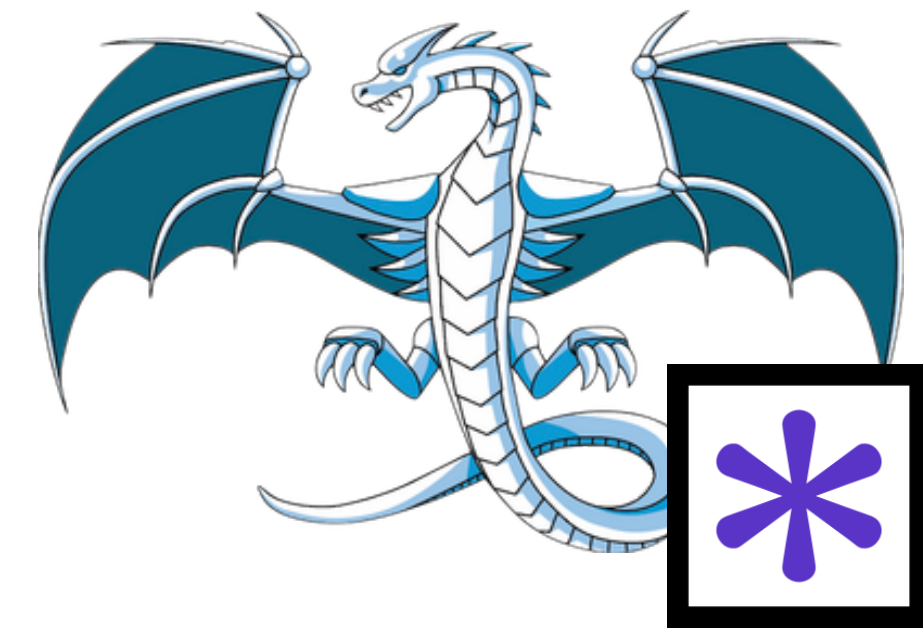
Modular Semantics ▷ (and Meta-theory) for LLVM IR

Cambium Seminar

Inria

Irene Yoon
12 / 04 / 2024

Joint work with collaborators at...



Velliris

A Relational Separation Logic for LLVM IR

Irene Yoon, Simon Spies, Youngju Song, Lennard Gäher, Derek Dreyer, Steve Zdancewic

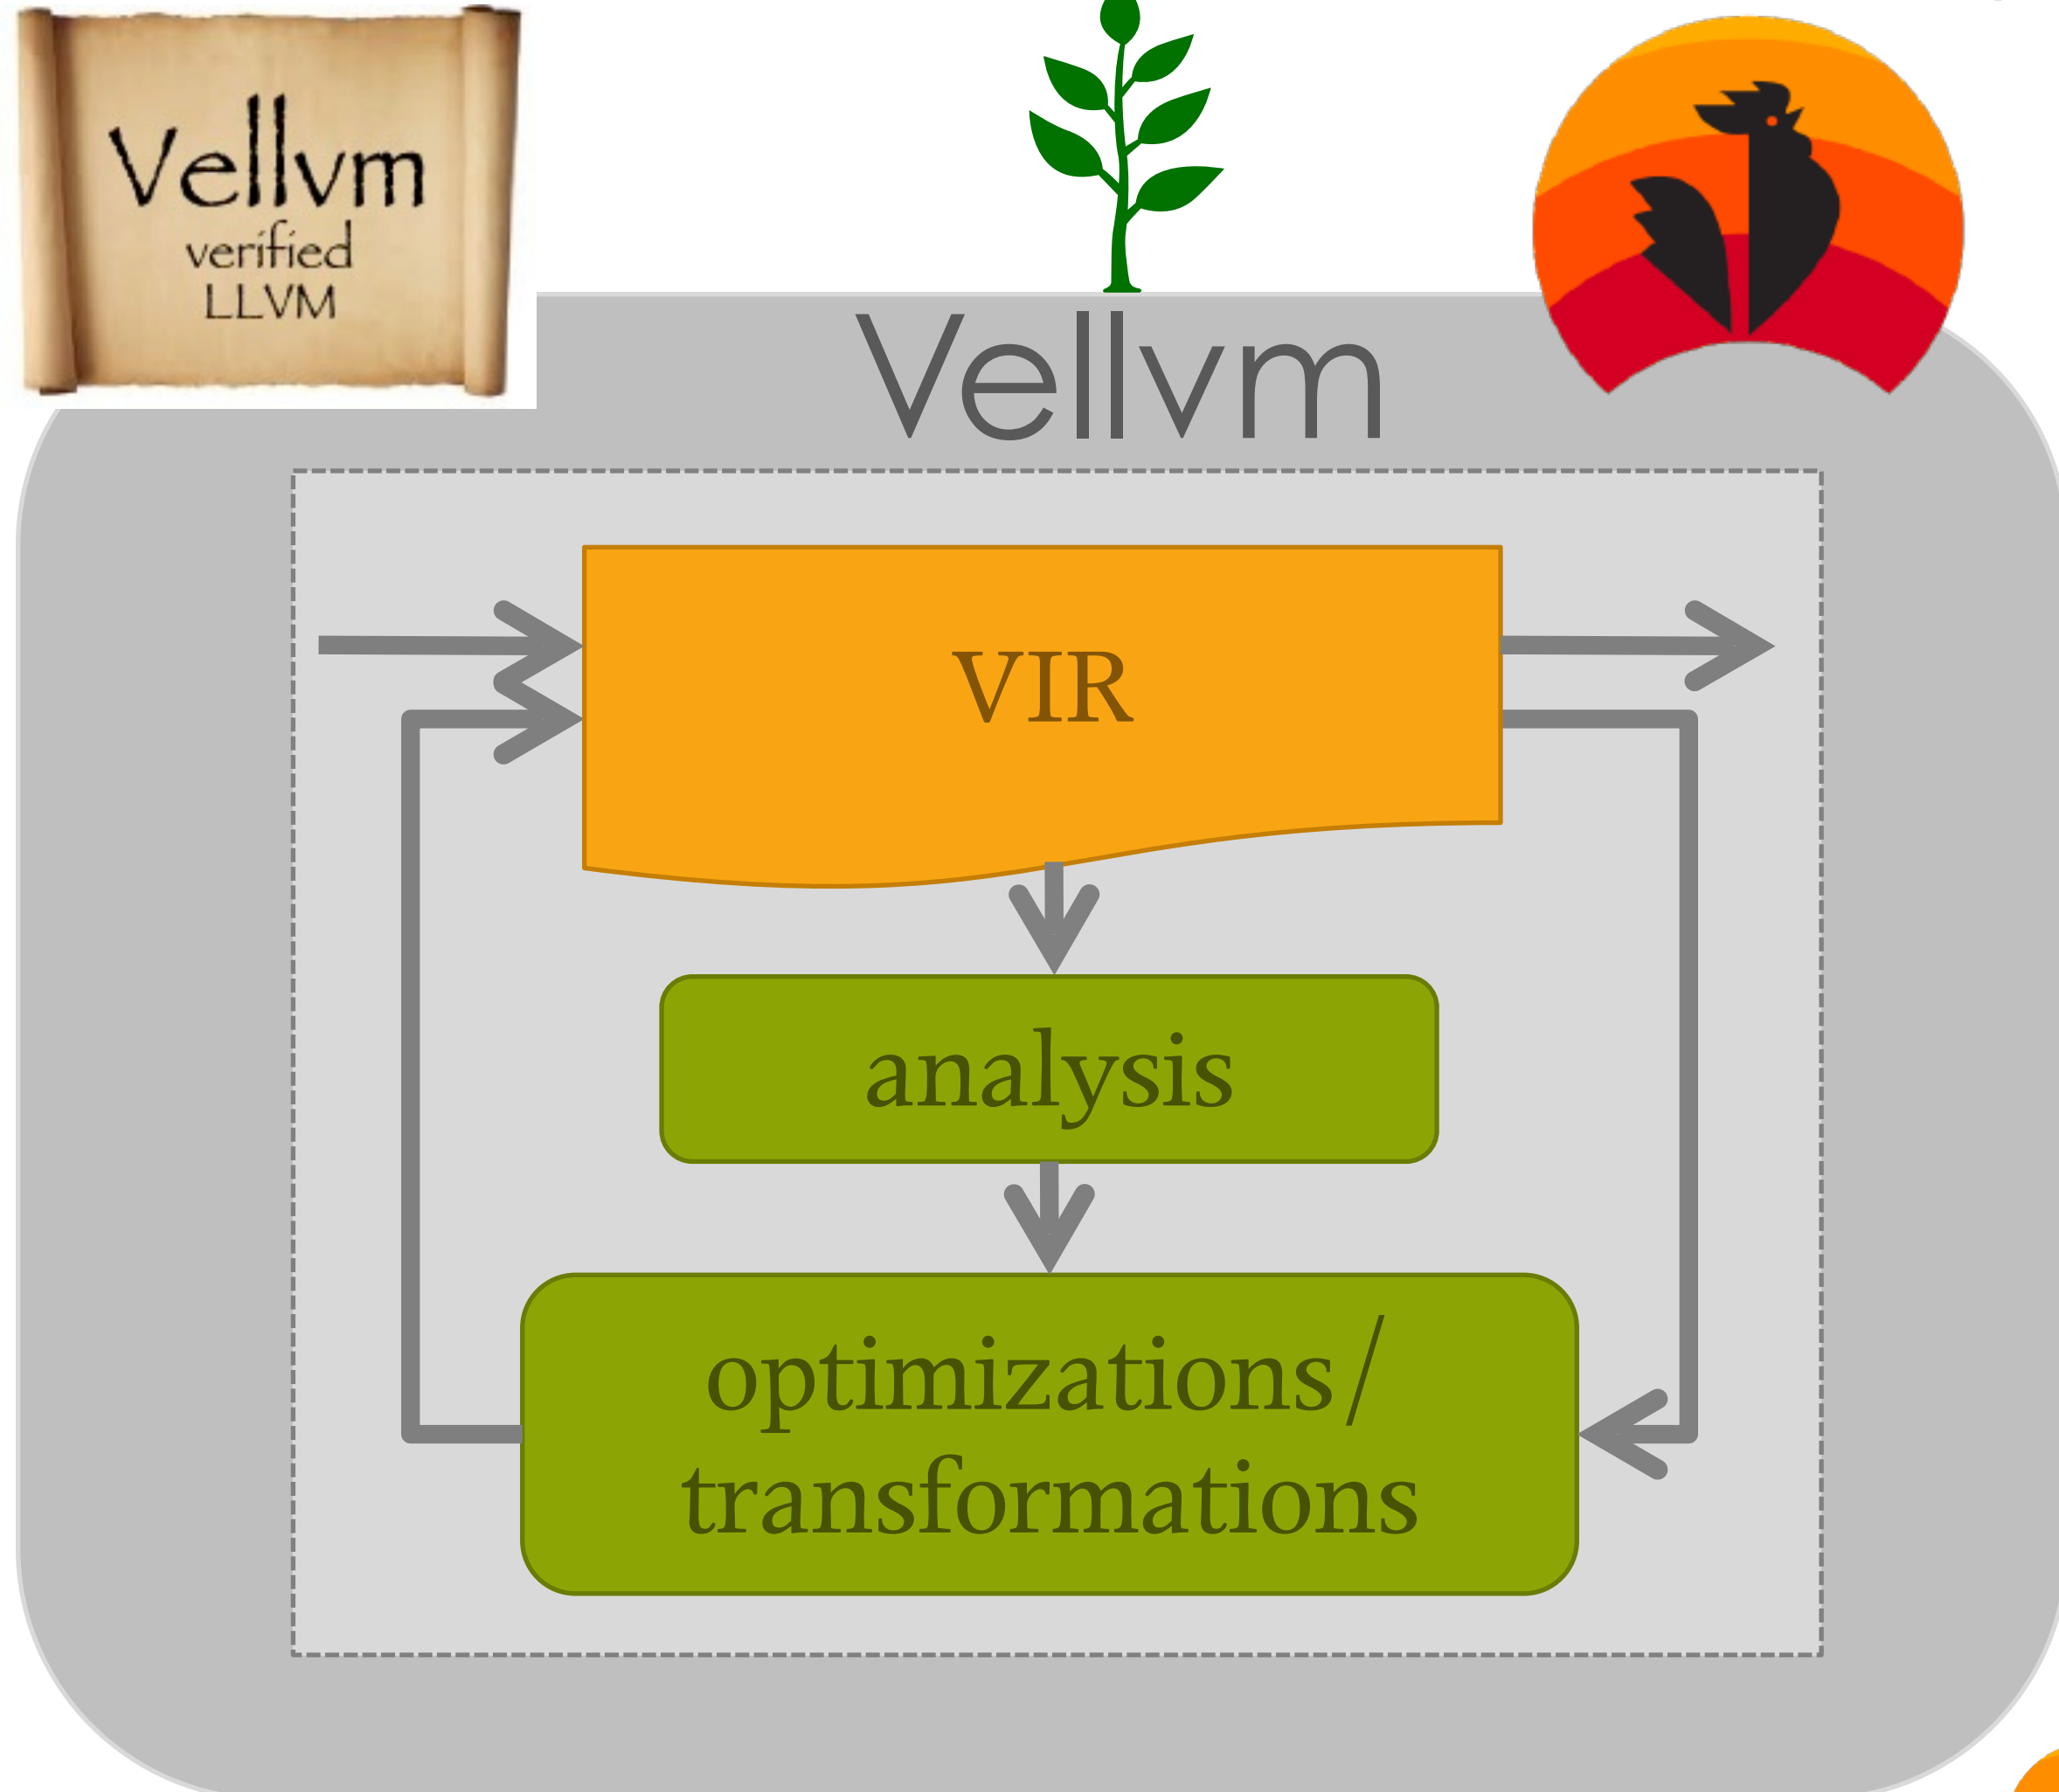


Irene Yoon | Cambium Seminar | 24/05/24

Vellvm 2.0 Overview



<https://github.com/vellvm/vellvm>



Selected publications and drafts*

[Zakowski et al. - ICFP 2021]

Modular and executable semantics for LLVM IR

[Yoon et al. - ICFP 2022]

Meta-theory for layered monadic interpreters

[Zaliva et al.]

Verified HELIX front-end

[Beck et al. - ICFP 2024 (conditionally accepted)]

Infinite/finite memory model for LLVM IR

[Yoon et al. - **Today's talk**]

Relational separation logic for LLVM IR



* : all results mechanized in the Coq Proof Assistant

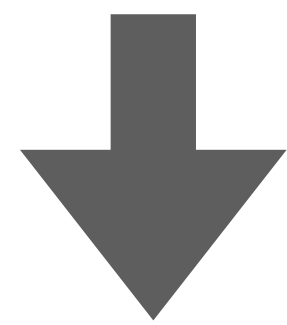
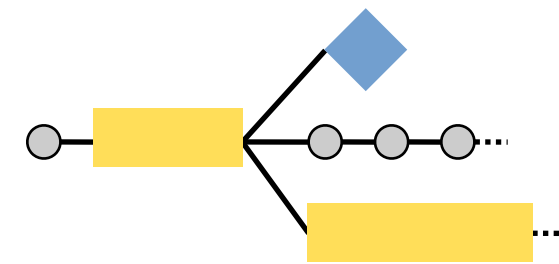
Vellvm 2.0

Interaction Tree based semantics for LLVM IR

Interaction Trees (itrees)

[Xia et al. 2020]

github.com/DeepSpec/InteractionTrees



Used to build

(Re)Vellvm

github.com/vellvm/vellvm



A generic toolkit to **define and reason** about the semantics of interactive systems

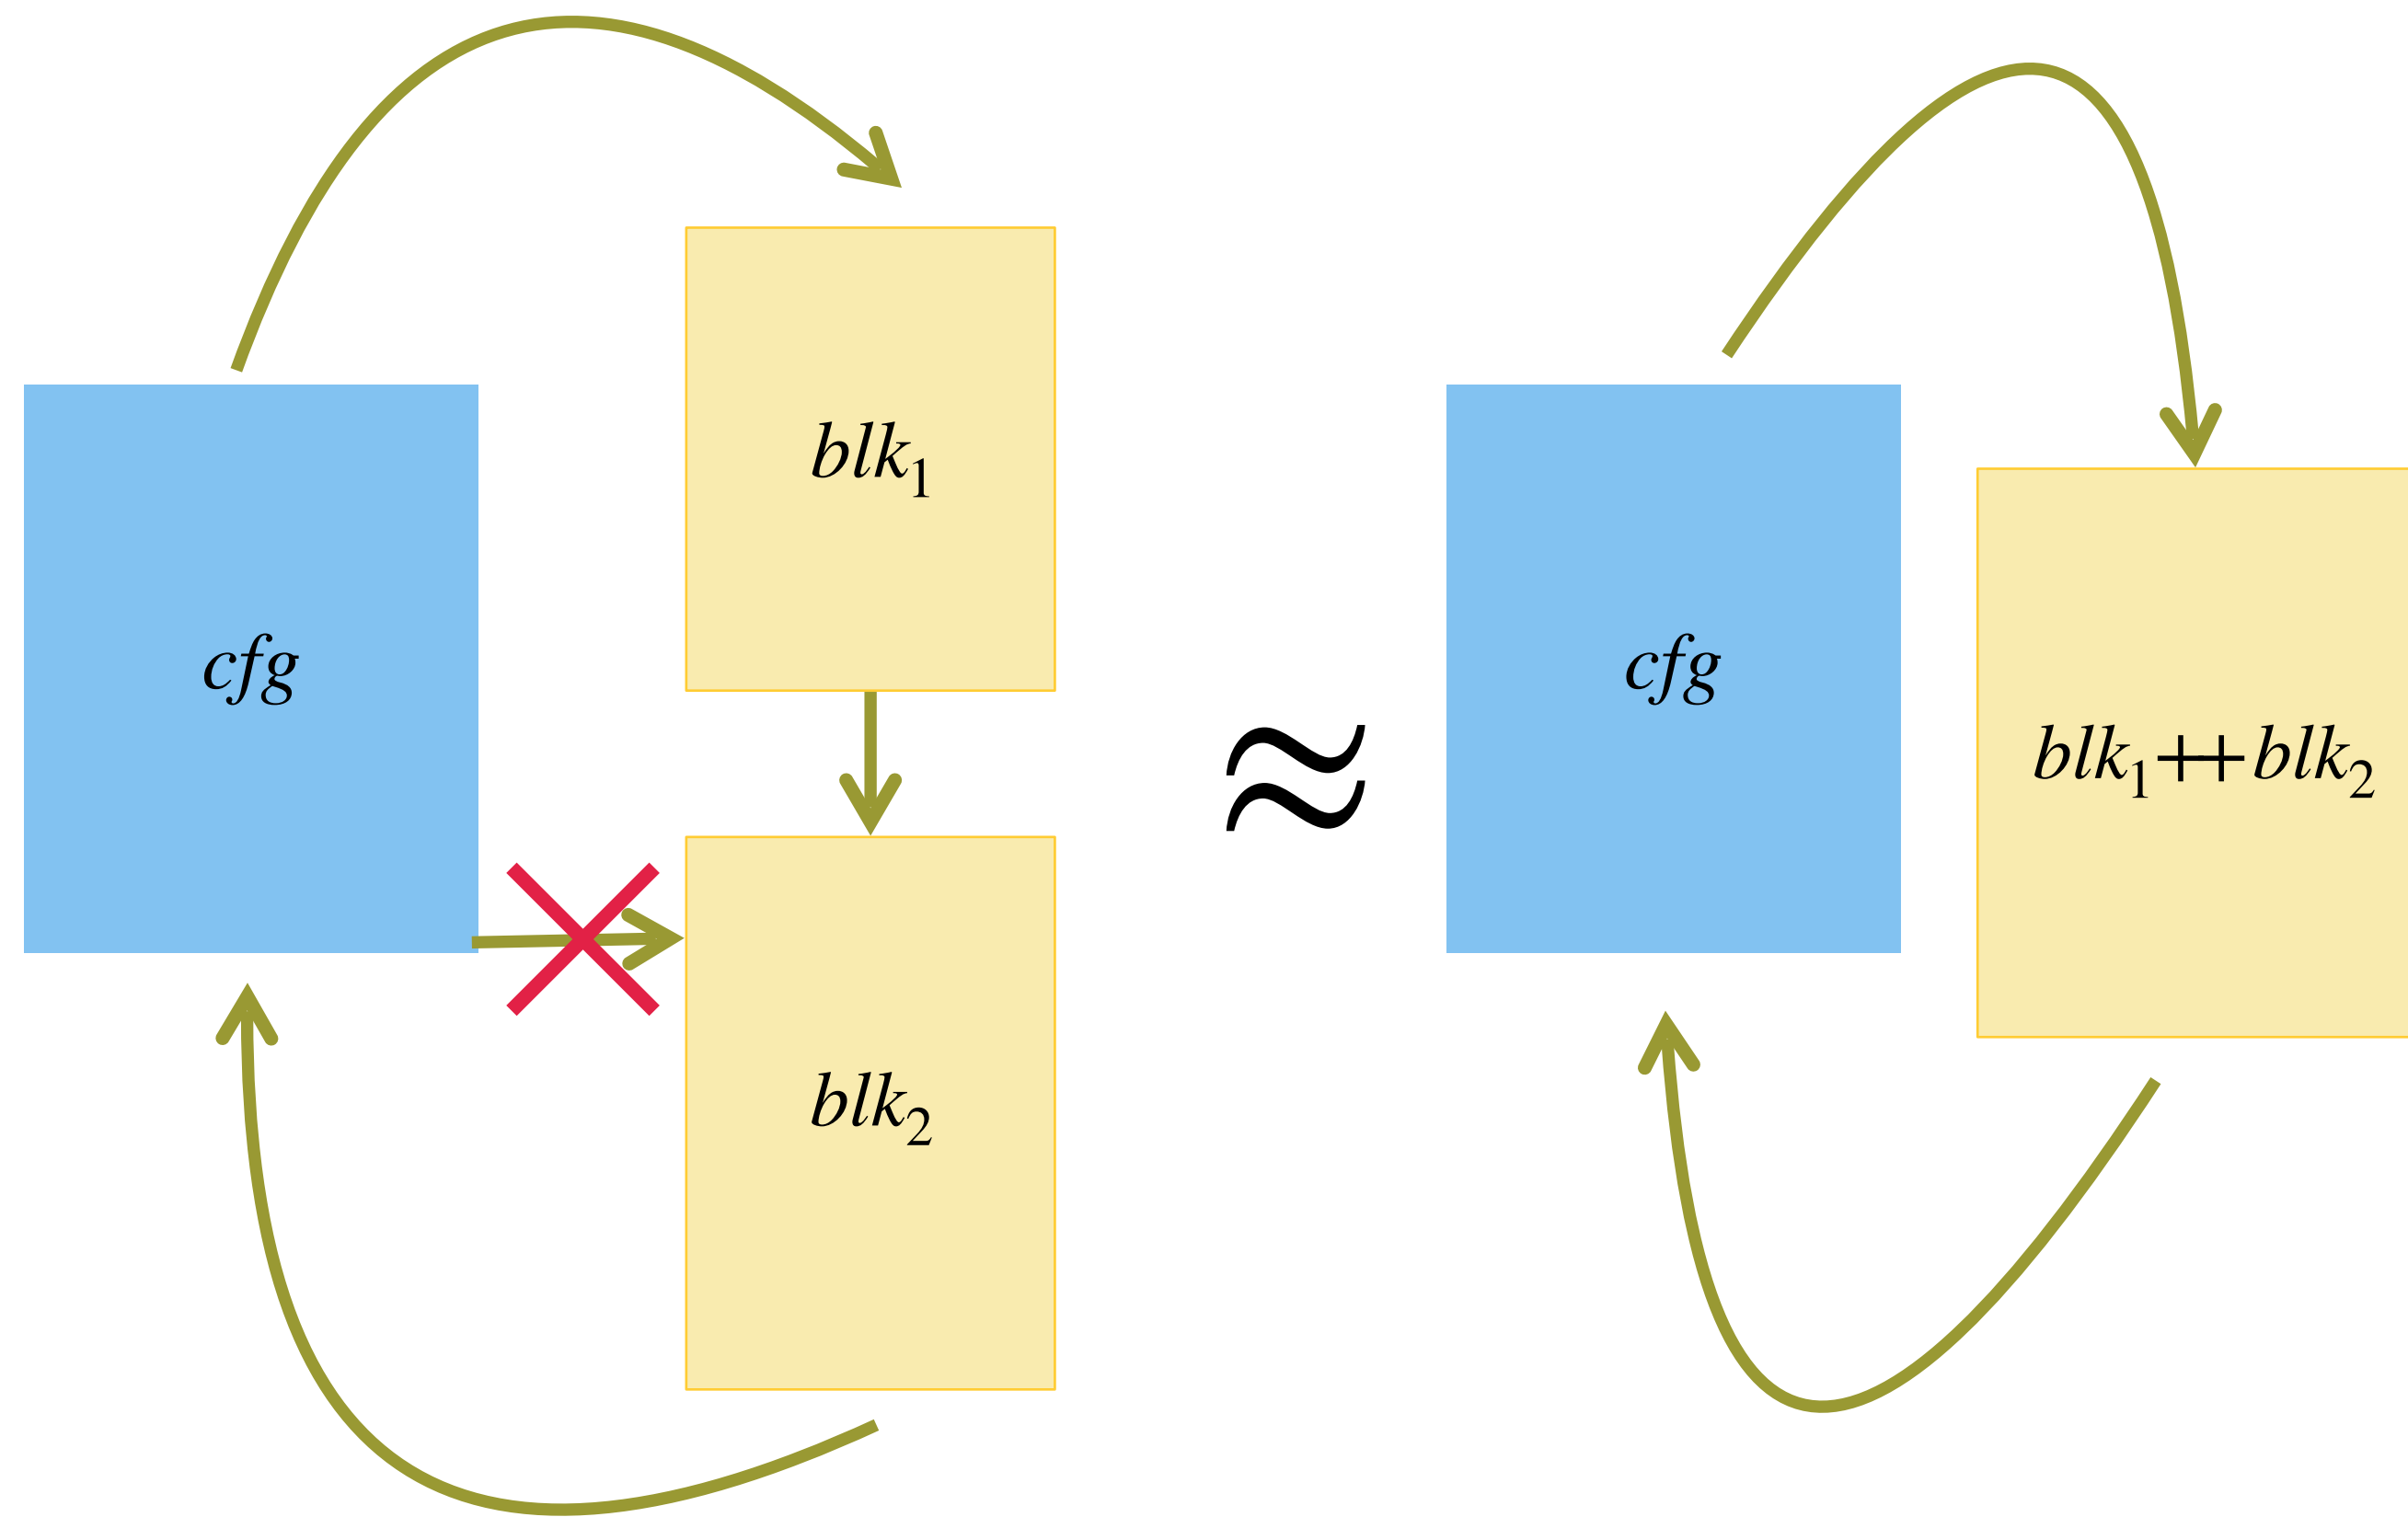
Semantics: **Compositional**, **Modular**, **Executable**

Reasoning: Equational, termination sensitive

VIR: a **compositional**, **modular** and **executable** formal semantics for (sequential) LLVM IR

Benefits of ITree-based reasoning

Reasoning about control flow

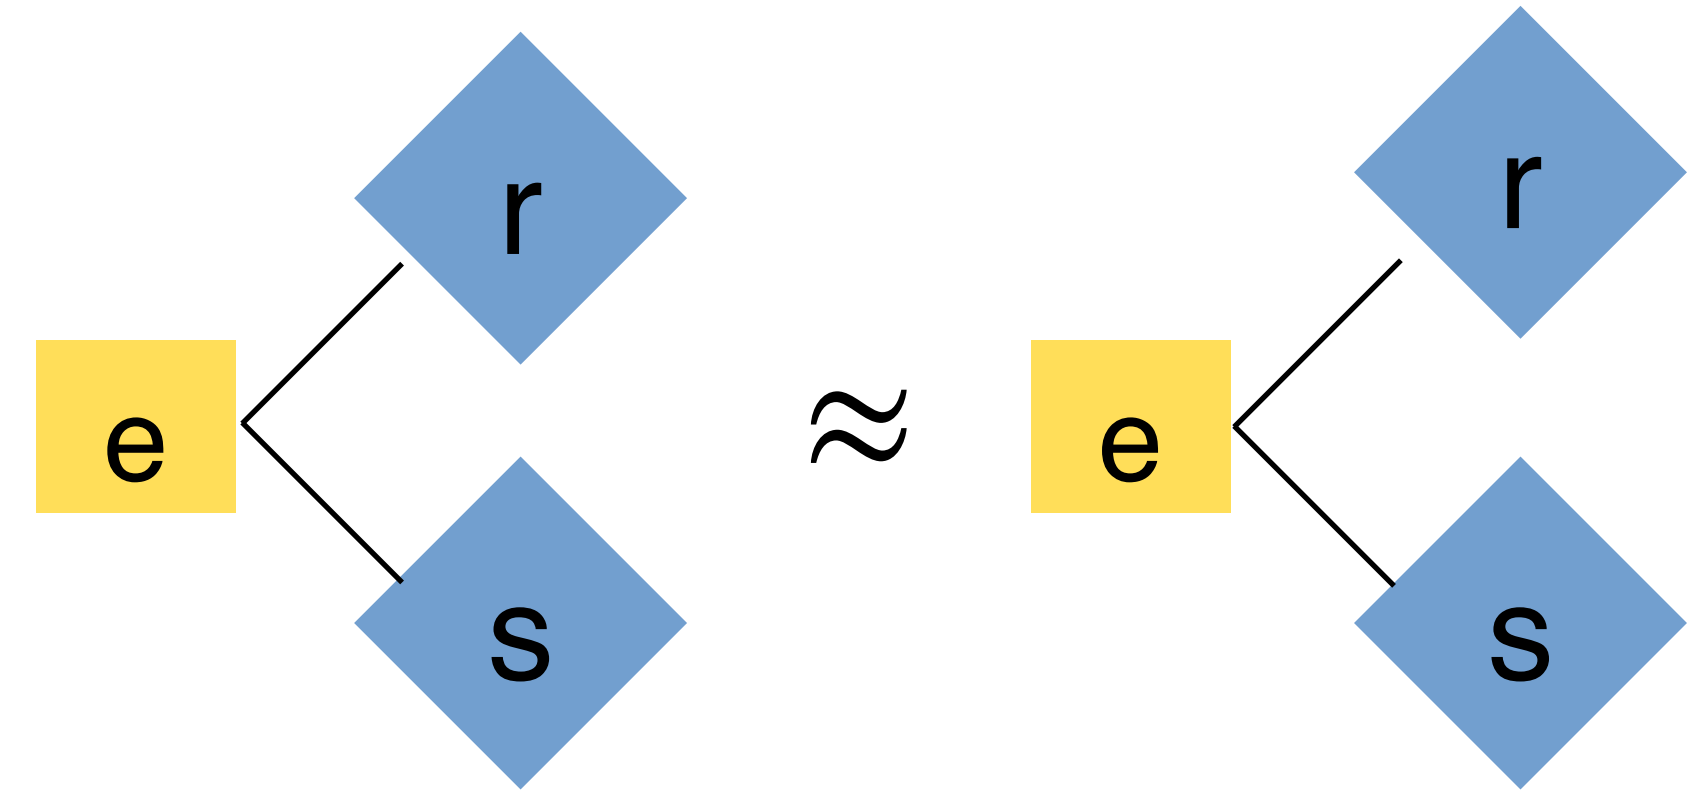


- Proof of block-merging optimization
- Reasoning about control flow is simple (if it does not change the trace of events)

Weak bisimulation on Interaction Trees

Termination-sensitive relational reasoning

$$e_t \approx_R e_s \quad \text{"eutt"}$$



- Program e_t and e_s are related to each other (i.e. bisimilar) w.r.t. a value relation R
 - If one program diverges, the other must diverge in a similar way (& vice versa.)
 - If one program terminates, the other must terminate in a similar way (& vice versa.) and the returned values are related by the value relation R

Benton-style relational Hoare reasoning



Let's reason about programs via Hoare triples (quadruple)!

- "eutt" can be seen as a partial (e.g. without a proof of termination) relational Hoare triple with a trivial precondition

$$e_t \approx_R e_s \quad \{ \top \} e_t \approx e_s \{ R \}$$

- A partial Hoare triple can be instantiated by taking the diagonal:

$$\{ \top \} e \{ Q \} := e \approx_{\lambda x, y. x=y \wedge Q x} e$$

Benton-style relational Hoare reasoning



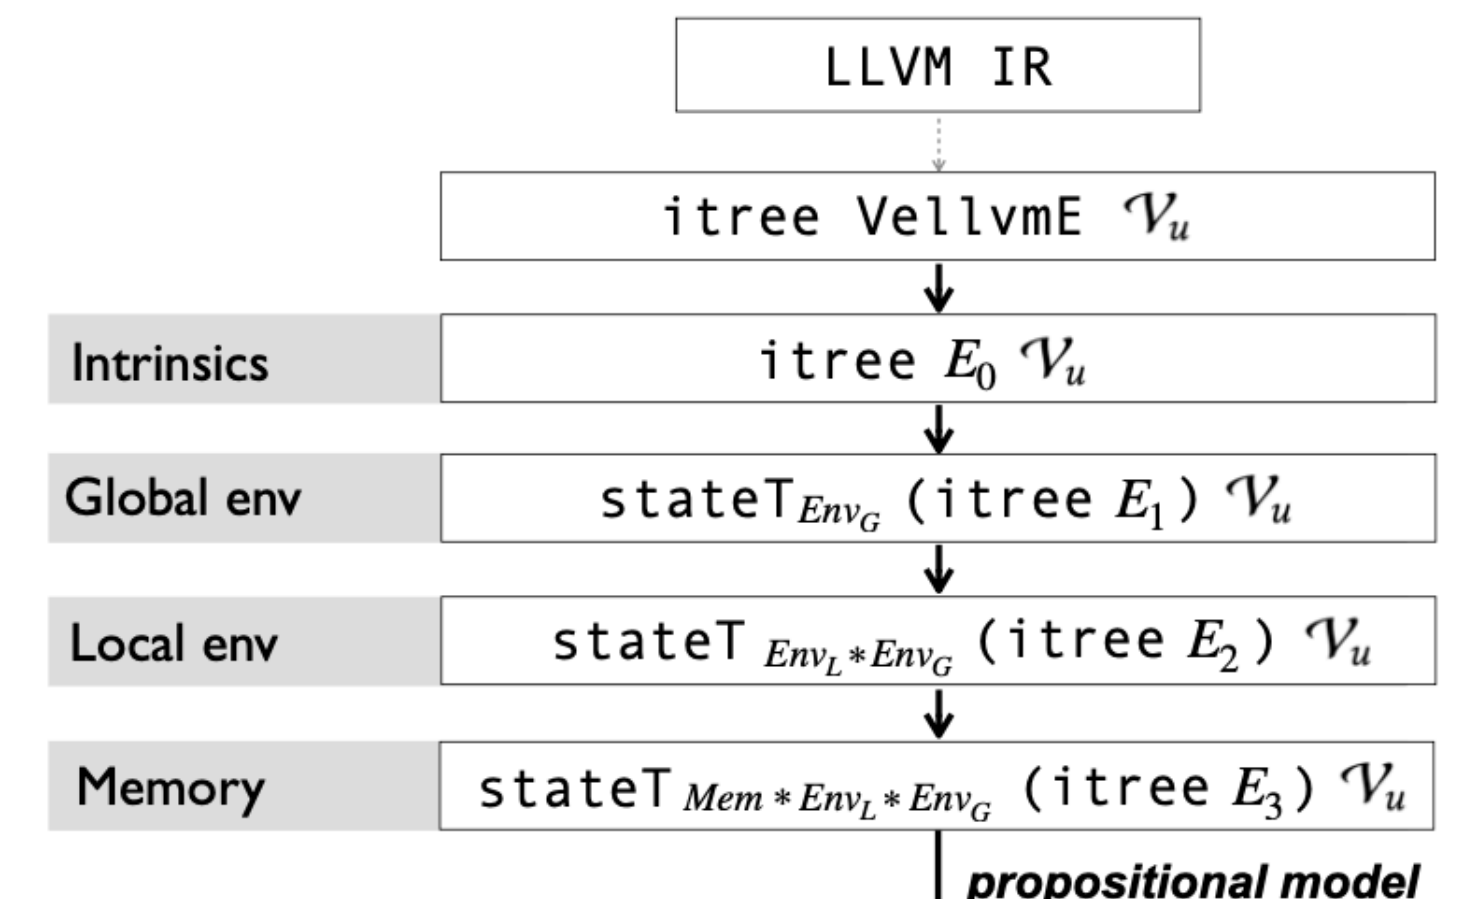
Let's reason about programs via Hoare triples (quadruple)!

- There are not many transformations that preserve interactions with state
- After one step of interpretation (of a state with carrier type S) $\llbracket - \rrbracket$, we see a state-passing tree (S \rightarrow itree F (S * A)). The post-condition on eutt relates both values and the final state.

$$\llbracket e_t \rrbracket \sigma_t \approx_Q \llbracket e_s \rrbracket \sigma_s$$

- We can have relational pre- and post-conditions about the initial and final states of programs

$$\{P\}e_t \approx e_s\{Q\} := \forall \sigma_t, \sigma_s. P(\sigma_t, \sigma_s) \Rightarrow \llbracket e_t \rrbracket \sigma_t \approx_Q \llbracket e_s \rrbracket \sigma_s$$



Global state invariants and assumptions

A need for modularity

- The pre- and post- condition must carry global invariants about the state

Given a stateful interpretation function $\llbracket - \rrbracket$,

$$\{P\}e_t \approx e_s\{Q\} := \forall \sigma_t, \sigma_s. P(\sigma_t, \sigma_s) \Rightarrow \llbracket e_t \rrbracket \sigma_t \approx_Q \llbracket e_s \rrbracket \sigma_s$$

... and global invariants are difficult to work with!

In particular, in the setting of LLVM IR..

1. LLVM IR transformations make assumptions over memory regions that call for localized reasoning
2. Stack-allocated regions allocated using "alloca" are automatically collected upon function return, and local variables live in the scope of a function: it would be nice for the logic to be aware of this stack discipline

Example: Loop invariant code motion

What invariants does the compiler assume for its optimizations?

```
1 void increment(int *n);
2 int get_int (int *x) {
3     int *n; int i = 0; n = &i;
4     while (*n < *x) { increment(n); }
5     return *n;
6 }
```

```
1 void increment(int *n);
2 int get_int_opt (int *x) {
3     int *n; int i = 0;
4     n = &i; int y = *x;
5     while (*n < y) { increment(n); }
6     return *n;
7 }
```

- LLVM optimizations (1) reorder (or modify/remove) memory-related instructions, and (2) often make certain assumptions about external calls while doing so
- By adding an annotation at the generated LLVM IR (function attribute) for the C code above, one can specify that the function only accesses memory through its arguments

```
declare i32 @increment(i32*) argmemonly
```

” function can only affect memory accessible by the arguments passed on to the function ”

Another example: Load-after-store on "promotable" locations

Compilers want to use assumptions from static analysis passes

- LLVM IR transformation often uses assumptions derived from analysis passes (e.g. alias analysis), and from the perspective of verifying optimizations, we need a way to state these assumptions

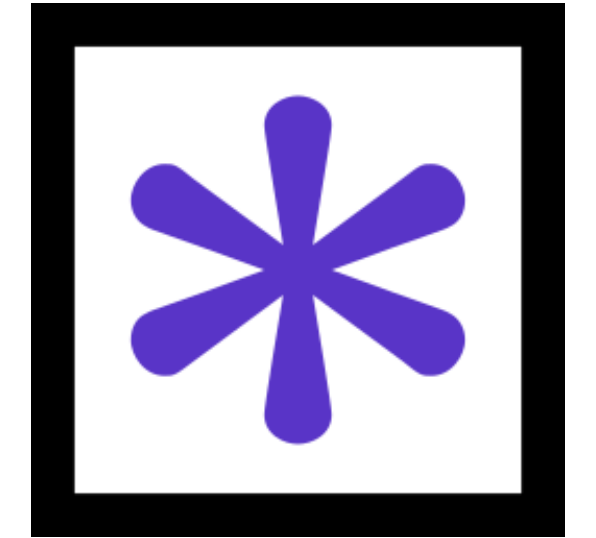
```
%a = alloca i32
...
store i32 6, i32 %a
...
%b = load %a
ret %b
...
```

```
%a = alloca i32
...
store i32 6, i32 %a
...
ret 6
...
```

- "promotable" register: no aliasing, no storing to memory

Separation Logic

Localized reasoning about resources for all!



- Separation logic [O. Hearn et al.]

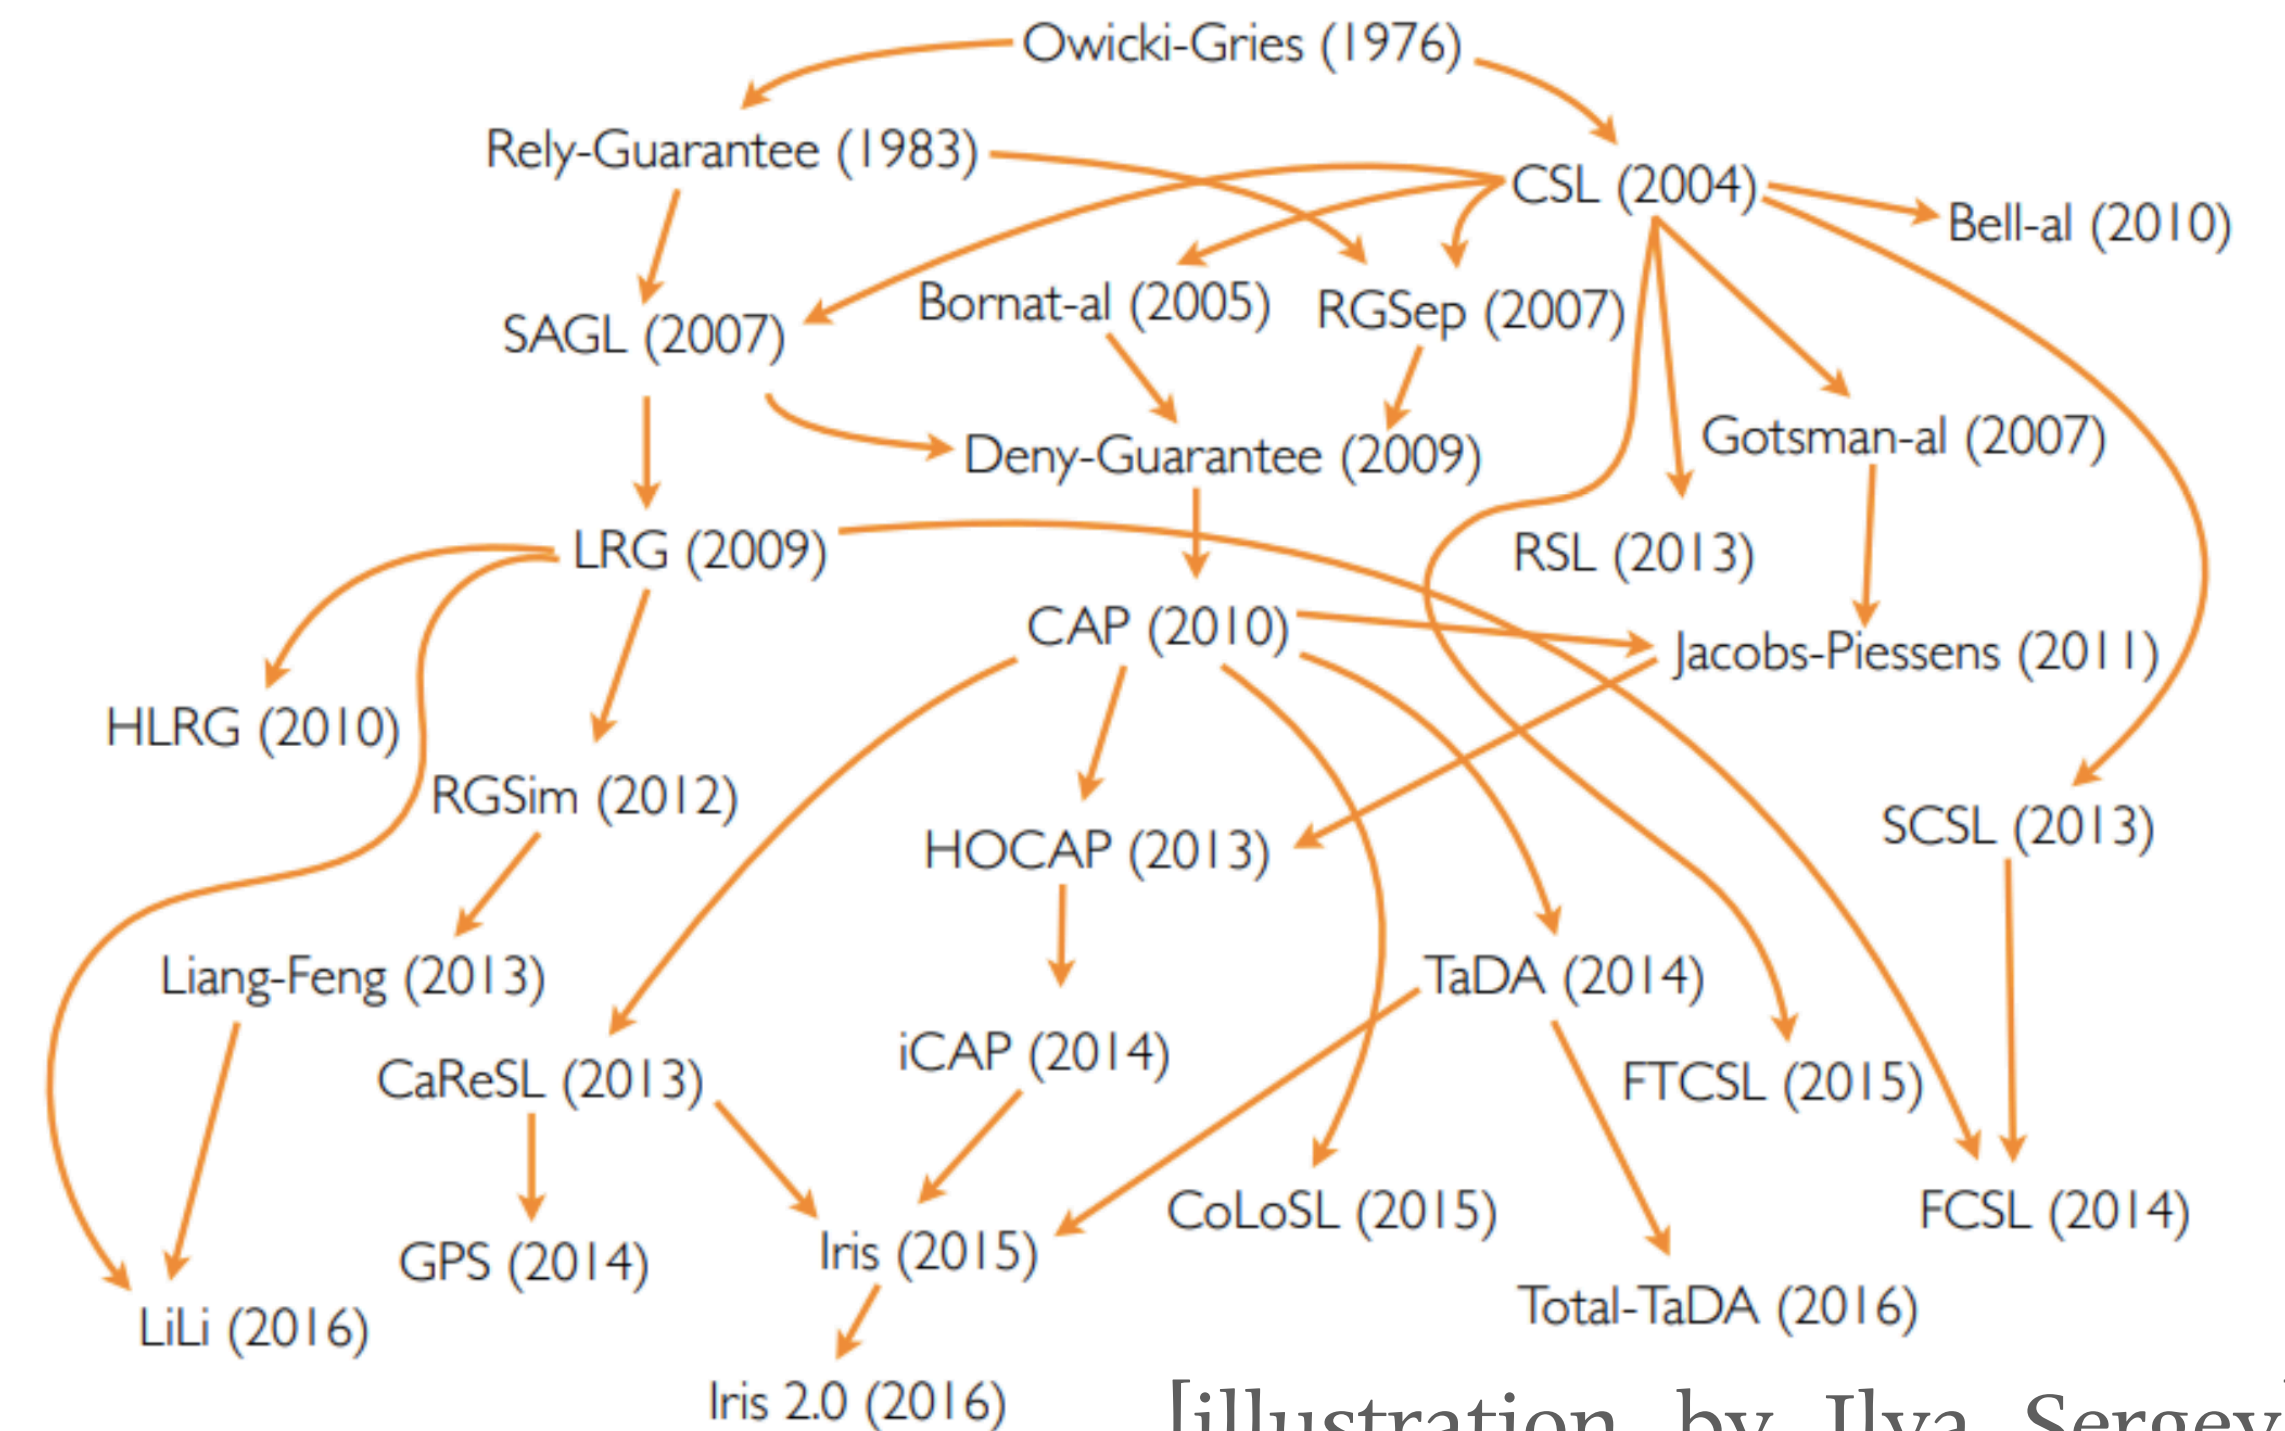
$$P * Q \quad \frac{\{P\} C \{Q\}}{\{P * R\} C \{Q * R\}}$$

- Iris [Jung et al.]: a higher-order concurrent separation logic framework

- Highly reusable and influential in consolidating variants of separation logics

- Used for various other realistic semantics (RustBelt, RefinedC, Iris-WASM, etc).

The genealogy of separation logics



[illustration by Ilya Sergey]

Can we bring the niceties of separation logic to LLVMIR?

Semantics*



VIR

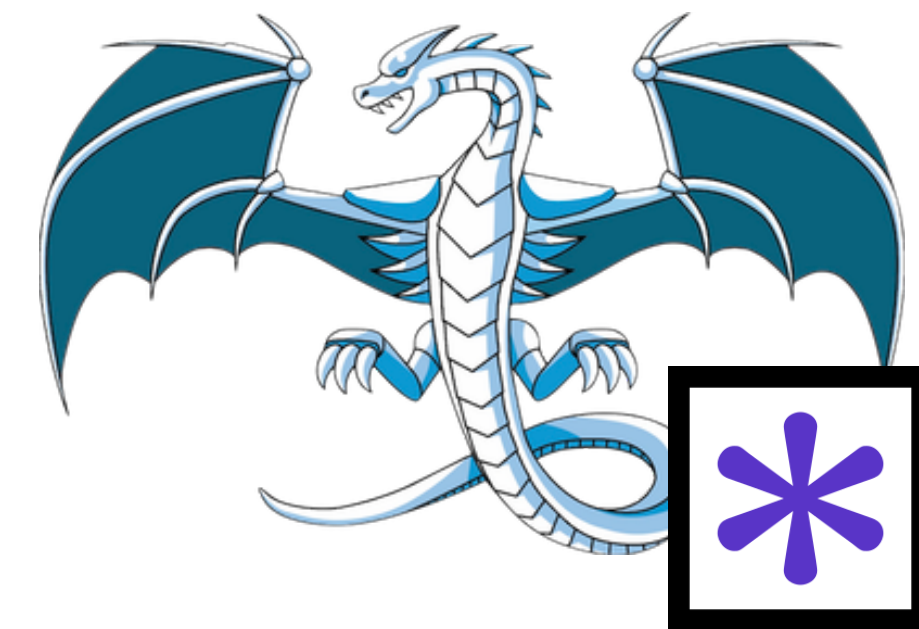
Program logic



Iris

+

=

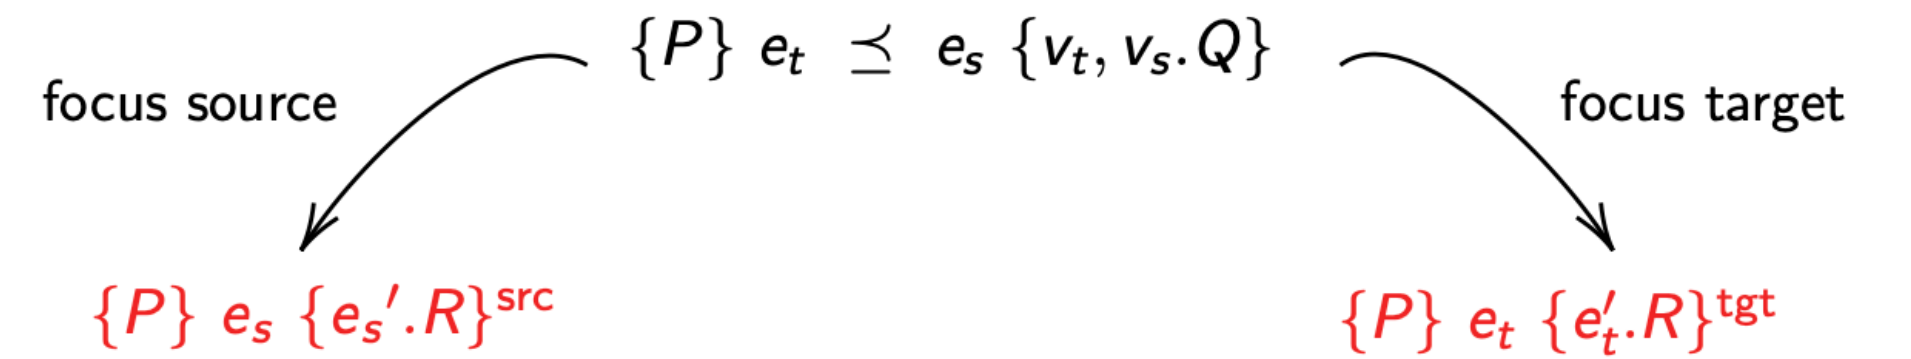


Velliris

Relational separation logic for LLVM IR!

Simuliris : relational Hoare logic in Iris

- Focusing rules on source and target programs
- Termination-sensitive simulations in Iris



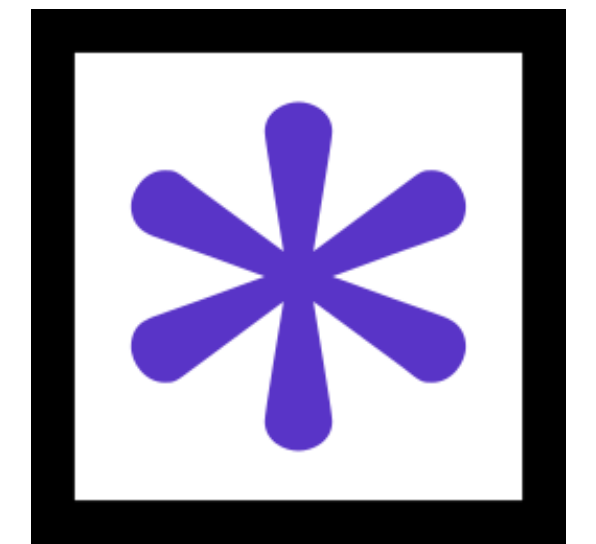
$$\frac{\{P\} e_s \{v_s. \Psi v_s\}^{\text{src}} \quad \forall v_s. \Psi v_s \multimap e_t \leq k_s v_s \{\Phi\}}{\{P\} e_t \leq x \leftarrow e_s ;; k_s x \{\Phi\}} \text{SOURCEFOCUS}$$

$$\frac{\{P\} e_t \{v_t. \Psi v_t\}^{\text{tgt}} \quad \forall v_t. \Psi v_t \multimap k_t v_t \leq e_s \{\Phi\}}{\{P\} x \leftarrow e_t ;; k_t x \leq e_s \{\Phi\}} \text{TARGETFOCUS}$$

$$\frac{\{P\} e_t \leq e_s \{v_t v_s. \Psi v_t v_s\} \quad \forall v_t v_s. \Psi v_t v_s \multimap k_t v_t \leq k_s v_s \{\Phi\}}{\{P\} (x \leftarrow e_t ;; k_t x) \leq (x \leftarrow e_s ;; k_s x) \{\Phi\}} \text{SIMBIND}$$

$$\{ \Phi v_t v_s \} \text{Ret } v_t \leq \text{Ret } v_s \{\Phi\} \text{ SIMVALUE} \quad \frac{\{P\} e_t \leq e_s \{\Phi\}}{\{P * R\} e_t \leq e_s \{v_t v_s. \Phi v_t v_s * R\}} \text{SIMFRAME}$$

(Typical) Recipe to use Iris



(1) Ingredient: an abstract view on state (ghost theory) using separation logic resources

Iris has a notion of resource algebras and generic constructions of resource algebras suitable for read-only map, permission-based ownership, etc.

For expository purposes, an example of a much simpler resource algebra.

Given a partial commutative monoid $(R, (\odot), \epsilon)$ we can define a ghost theory :

Given a heap which is a partial map from addresses to integers, we can define an ownership predicate $\ell \mapsto v$ where

$$(\ell \mapsto v) \odot (\ell' \mapsto w) := (\ell \mapsto v; \ell' \mapsto w) \quad \text{where } \ell \neq \ell'$$

$$L \odot L' := (L ++ L') \quad \text{where } \text{dom}(L) \cap \text{dom}(L') = \emptyset$$

$$(\ell \mapsto v) \odot (\ell \mapsto w) := \perp$$

$$L \odot L' := \perp \quad \text{where } \text{dom}(L) \cap \text{dom}(L') \neq \emptyset$$



(Typical) Recipe to use Iris

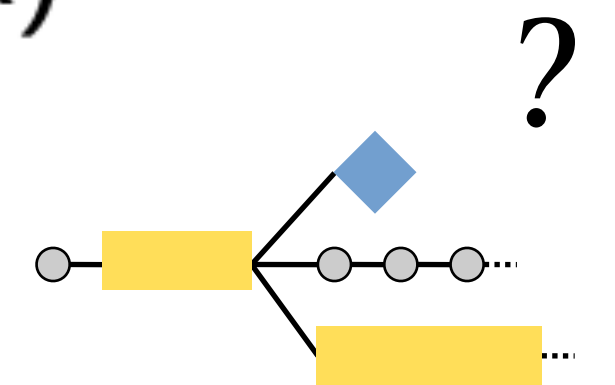


(2) Ingredient: a small-step semantics

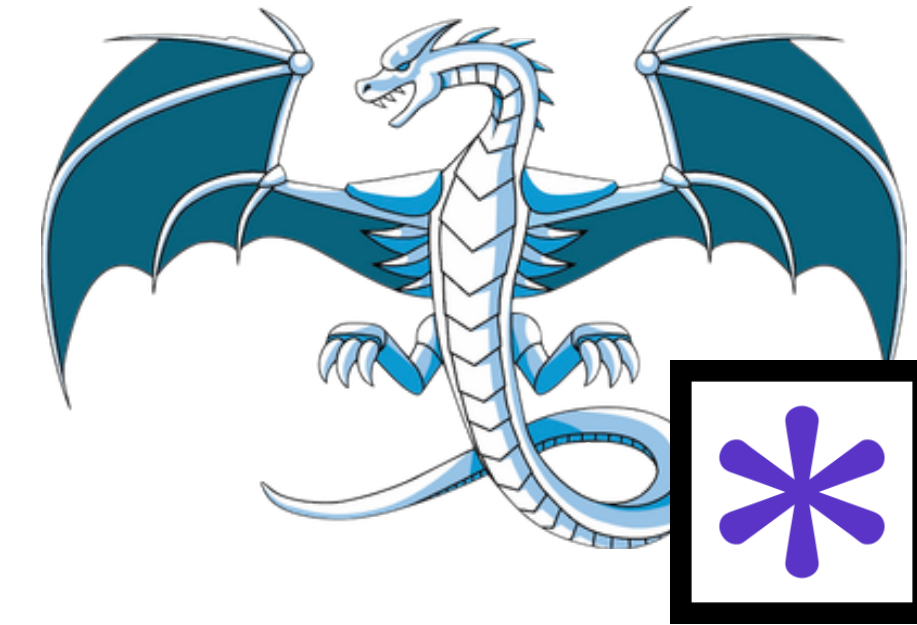
Given the small-step semantics and ghost theory, a Hoare triple can be derived via the typical weakest precondition model of Iris

$$\{P\}e\{\Phi\} \triangleq \Box(P \rightarrow_{*wp} e\{\Phi\})$$

Given a postcondition Φ , $wp\ e\{\Phi\}$ gives the *weakest precondition* under which all executions of e are *safe* (i.e. does not get stuck) and all return values v satisfy $\Phi(v)$



Roadmap



1. A taste of the stack-based ghost theory for LLVM IR
2. Memory-relevant attributes in LLVM IR
3. Model : A relational weakest-precondition model for Interaction Trees in Iris
4. Adequacy

VIR State

- The state of VIR: $Global * (Local * LocalStack) * (Mem * FrameStack)$

$Global ::= id \hookrightarrow \mathcal{V}$

$Local ::= id \hookrightarrow \mathcal{V}$

$Frame ::= list Addr$

$LocalStack ::= list Local$

$id ::= string \quad Addr ::= \mathcal{Z} * \mathcal{Z}$

$FrameStack ::= list Frame$

$Allocated ::= list Addr$

$Mem ::= Allocated * (Addr \hookrightarrow list byte)$

$byte ::= SUnDef \mid Byte \mathcal{Z} \mid Ptr Addr \mid PtrFrag$

- The stack resources need to be managed upon function entry and exit
- The ghost theory in Velliris deals with this deallocation by keeping track of the stack frame and the associated set of stack-allocated locations.

Reasoning about local environments

Stack frame rules (not to be confused with the frame rule)

Ghost resources

$\text{Frame}^{\text{src}} i$ We are at index "i" on the stack frame.

$\langle \text{id} := v \rangle_i^{\text{src}}$ At frame index "i", we have access to local id "id" with value "v" stored on it.

$\text{Local}_i^{\text{src}} L$ At frame index "i", "L" is the domain of the local environment.

(Source side) Hoare triples

$$\frac{\{P\} e_s \{v_s. \Psi v_s\}^{\text{src}} \quad \forall v_s. \Psi v_s \multimap e_t \leq k_s v_s \{\Phi\}}{\{P\} e_t \leq x \leftarrow e_s ;; k_s x \{\Phi\}} \text{SOURCEFOCUS}$$

LOCALREAD

$$\begin{aligned} & \{ \langle \text{id} := v \rangle_i^{\text{src}} * \text{Frame}^{\text{src}} i \} \\ & \quad \text{trigger (LRd}^{\mathcal{V}u} (\uparrow \text{id})) \\ & \{ v'_s. v'_s = v * \langle \text{id} := v \rangle_i^{\text{src}} \}^{\text{src}} \end{aligned}$$

We are currently at stack frame "i", and we know that the local environment stores "v" for "id".

LOCALWRITE

$$\begin{aligned} & \{ (\text{id} \notin L) * \text{Frame}^{\text{src}} i * \text{Local}_i^{\text{src}} L \} \\ & \quad \text{trigger (LWr}^{()} (\uparrow \text{id}, v)) \\ & \{ v'_s. \langle \text{id} := v \rangle_i^{\text{src}} * \text{Frame}^{\text{src}} i * \text{Local}_i^{\text{src}} (\{ [\%id] \} \cup L) \}^{\text{src}} \end{aligned}$$

We are currently at stack frame "i", and we can extend the local domain and get a new local environment predicate.

Function calls in Vellvm

Function calls, stack-allocated resources

```
(* The denotation of an itree function is a coq function that takes  
a list of uvalues and returns the appropriate semantics. *)
```

```
Definition function_denotation : Type :=  
list uvalue → itree L0' uvalue.
```

```
Definition denote_function (df:definition dtyp (cfg dtyp)) : function_denotation :=  
λ (args : list uvalue) ⇒  
  (* We match the arguments variables to the inputs *)  
  bs ←  
    (* Allow only full application of functions *)  
    (if Nat.eqb (List.length (df_args df)) (List.length args) then  
      ret (List.combine (df_args df) args)  
      else raise ("Incorrect argument length for function")) ;;  
  (* generate the corresponding writes to the local stack frame *)  
  trigger MemPush ;;  
  trigger (StackPush bs) ;;  
  rv ← translate_instr_to_L0' (denote_cfg (df_instrs df)) ;;  
  trigger StackPop ;;  
  trigger MemPop ;;  
  ret rv.
```

Stack-allocated resources

Function calls, stack-allocated resources

- Each function call allocates a new stack frame

SOURCEPUSHFRAME

$\{ \text{Frame}^{\text{src}} i \}$

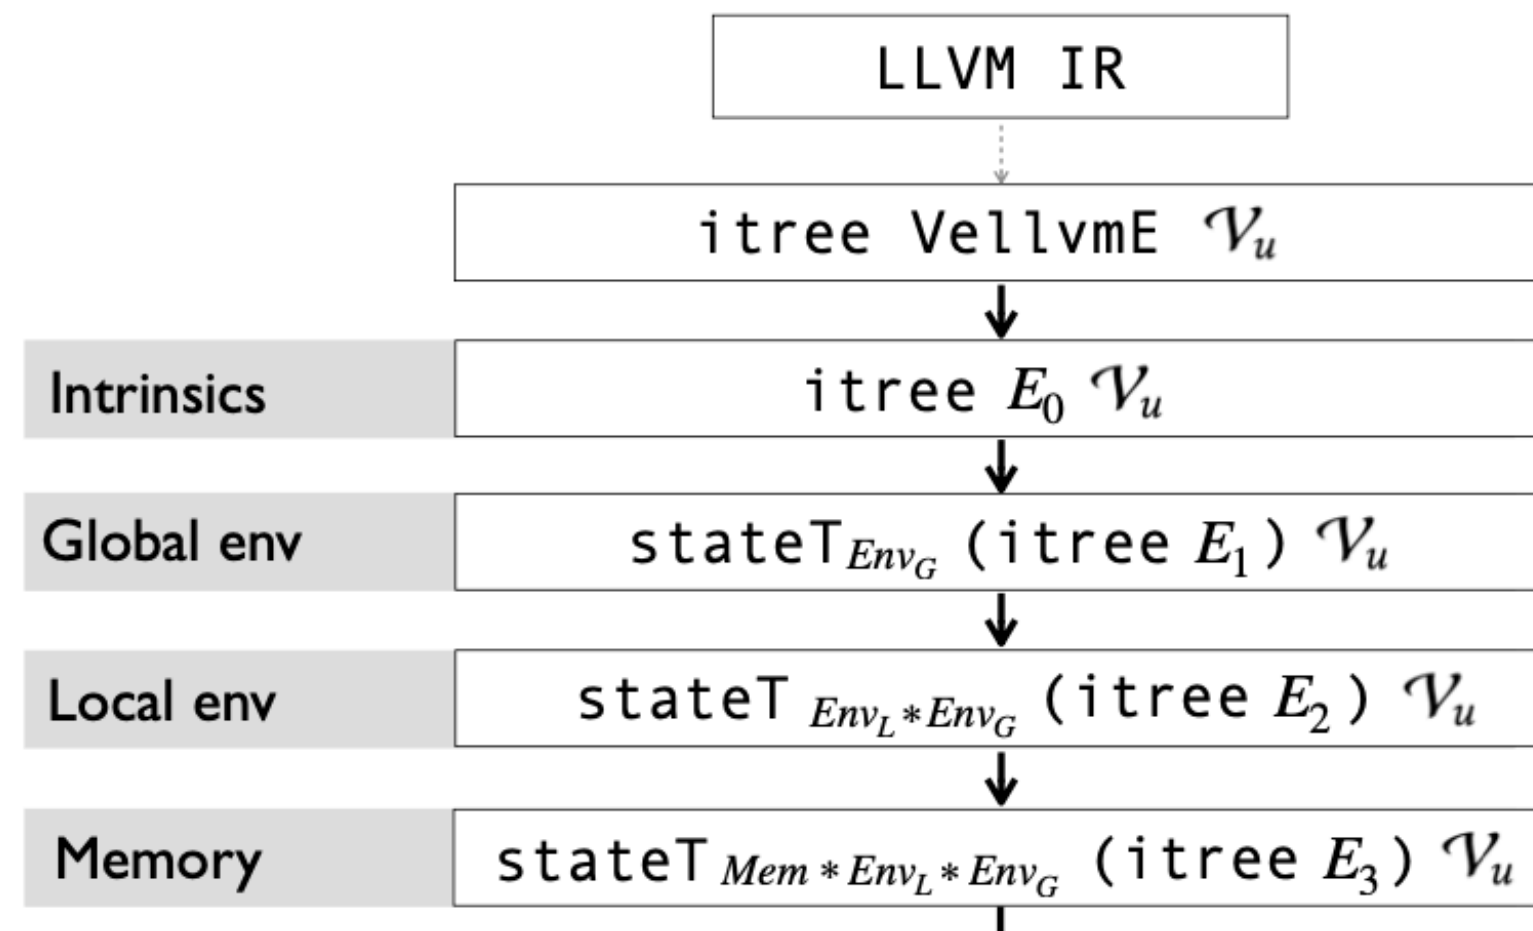
$\text{trigger} (\text{MPush}^{()}) ; ; \text{trigger} (\text{LPush}^{()} (args))$

$\{ v'_s. \exists f. \text{Frame}^{\text{src}} (f :: i) * \text{Alloca}_f^{\text{src}} \emptyset * \text{Local}_f^{\text{src}} (\text{dom}(args)) * (*_{(id,v) \in args} \langle \text{id} := v \rangle_f^{\text{src}}) \}^{\text{src}}$

We are currently at stack frame "i", and if we push a new memory frame and local frame with arguments "args", we update the current frame index, and get a empty memory frame and local domain and ownership over arguments "args" pushed onto the stack.

Event rules

Atomic proof rules over events



Need to build atomic proof rules over the sum of events on VIR

Memory-relevant event rules

SOURCEALLOCA

$$\{ \text{Alloca}_i^{\text{src}} S * \text{Frame}^{\text{src}} i \}$$

$$\text{trigger} (\text{Alloca}^{\mathcal{V}} (\tau))$$

$$\{ v'_s. \exists \ell_s. v_s = \ell_s * \ell_s \mapsto^{\text{src}} \text{new_block}^\tau * \text{Alloca}_i^{\text{src}} \{z\} \cup S * \text{Frame}^{\text{src}} i \}^{\text{src}}$$

SOURCELOAD

$$\{ (v \in \tau) * \ell_s \mapsto^{\text{src}} v \}$$

$$\text{trigger} (\text{Load}^{\mathcal{V}_u} (\tau, \text{addr}(\ell)))$$

$$\{ v'_s. v'_s = v * \ell_s \mapsto^{\text{src}} v \}^{\text{src}}$$

SOURCESTORE

$$\{ (v \in \tau) * \ell_s \mapsto^{\text{src}} v \}$$

$$\text{trigger} (\text{Store}^{()} (\text{addr}(\ell), v'))$$

$$\{ v'_s. \ell_s \mapsto^{\text{src}} v' \}^{\text{src}}$$

UB and Exception event rules

SIMUB $e \leq \text{trigger} (\text{UB}^\emptyset) \{ \Phi \}$

SIMEXC $e \leq \text{trigger} (\text{Throw}^\emptyset) \{ \Phi \}$

"Undefined behavior subsumes all behavior"

and

"The simulation holds only if the source program does not go wrong"

Instruction rules

Example: Alloca instruction

- Given atomic proof rules, it is straightforward to build rules over denotations on syntax

$$\begin{aligned} \llbracket (id, \text{alloca } (\tau)) \rrbracket_i &= dv \leftarrow \\ &\text{trigger } (\text{Alloca}^{\mathcal{V}} (\tau)) ;; \\ &\text{trigger } (\text{LWr}^{()} (\uparrow id, dv)) \end{aligned}$$

Denotation of an alloca instruction

SOURCEALLOCA

$$\{\text{Alloca}_i^{\text{src}} S * \text{Frame}^{\text{src}} i\}$$

$$\text{trigger } (\text{Alloca}^{\mathcal{V}} (\tau))$$

$$\{v'_s. \exists \ell_s. v_s = \ell_s * \ell_s \mapsto^{\text{src}} \text{new_block}^{\tau} * \text{Alloca}_i^{\text{src}} \{z\} \cup S * \text{Frame}^{\text{src}} i\}^{\text{src}}$$

LOCALWRITE

$$\{(id \notin L) * \text{Frame}^{\text{src}} i * \text{Local}_i^{\text{src}} L\}$$

$$\text{trigger } (\text{LWr}^{()} (\uparrow id, v))$$

$$\{v'_s. \langle id := v \rangle_i^{\text{src}} * \text{Frame}^{\text{src}} i * \text{Local}_i^{\text{src}} (\{[\%id]\} \cup L)\}^{\text{src}}$$

$$\frac{\{P\} e_t \leq e_s \{v_t v_s. \Psi v_t v_s\} \quad \forall v_t v_s. \Psi v_t v_s \multimap k_t v_t \leq k_s v_s \{\Phi\}}{\{P\} (x \leftarrow e_t ;; k_t x) \leq (x \leftarrow e_s ;; k_s x) \{\Phi\}} \text{SIMBIND}$$

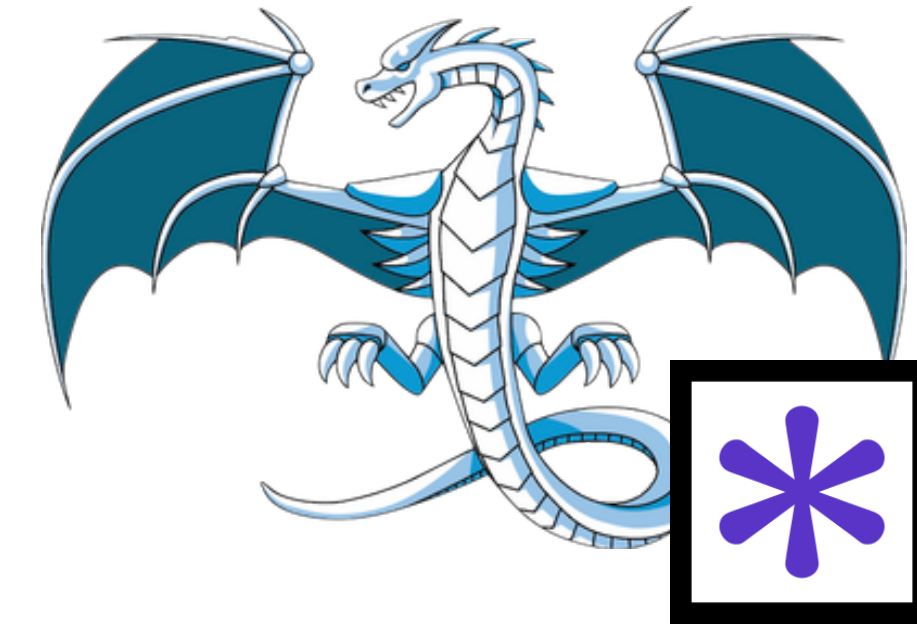
SOURCEINSTRALLOCA


$$\{(x \notin L) * \text{Frame}^{\text{src}} i * \text{Alloca}_i^{\text{src}} S * \text{Local}_i^{\text{src}} L\}$$

$$\%x^{\text{ld}} = \text{alloca } \tau$$

$$\{\exists \ell_s. \ell \mapsto^{\text{src}} \text{new_block}^{\tau} * \langle x := \text{addr}(\ell) \rangle_i^{\text{src}} * \text{Frame}^{\text{src}} i * \text{Alloca}_i^{\text{src}} S * \text{Local}_i^{\text{src}} (\{[\%x]\} \cup L)\}^{\text{src}}$$

Roadmap



1. A taste of the stack-based ghost theory for LLVM IR 
2. Memory-relevant attributes in LLVM IR
3. Model : A relational weakest-precondition model for Interaction Trees in Iris
4. Adequacy

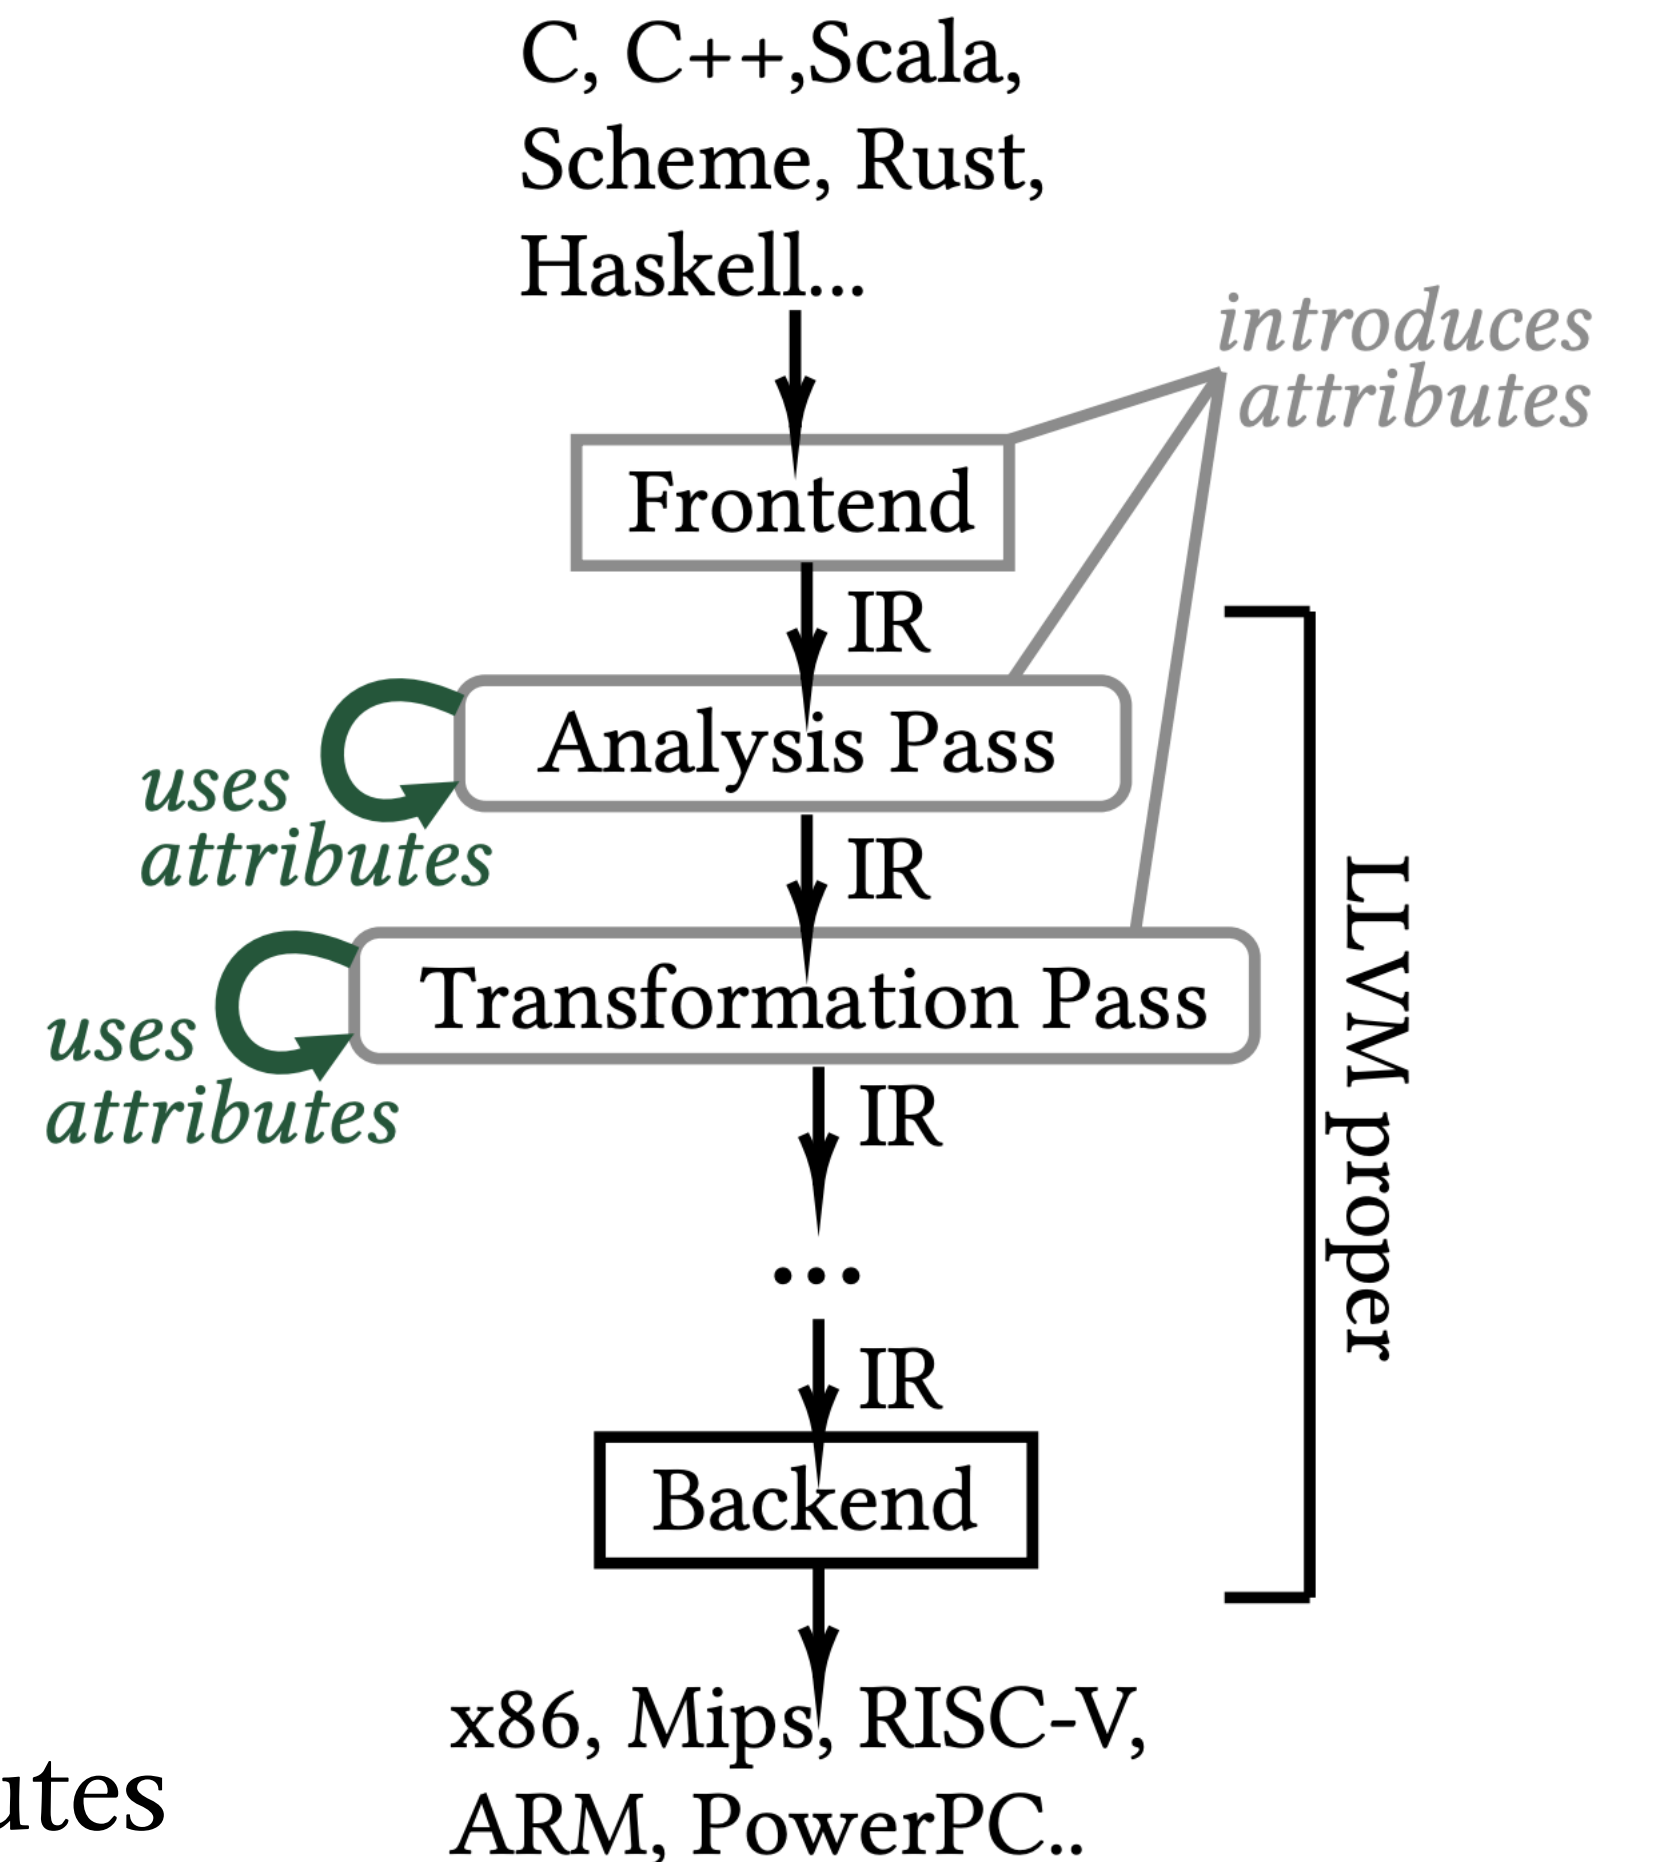
Memory attributes in LLVM IR

- LLVM optimization and analysis passes often use memory attributes, lightweight specifications about how a function may affect memory

```
define void @f(i32*) readonly argmemonly { ... }
```

” function f only reads from arguments passed on to the function ”

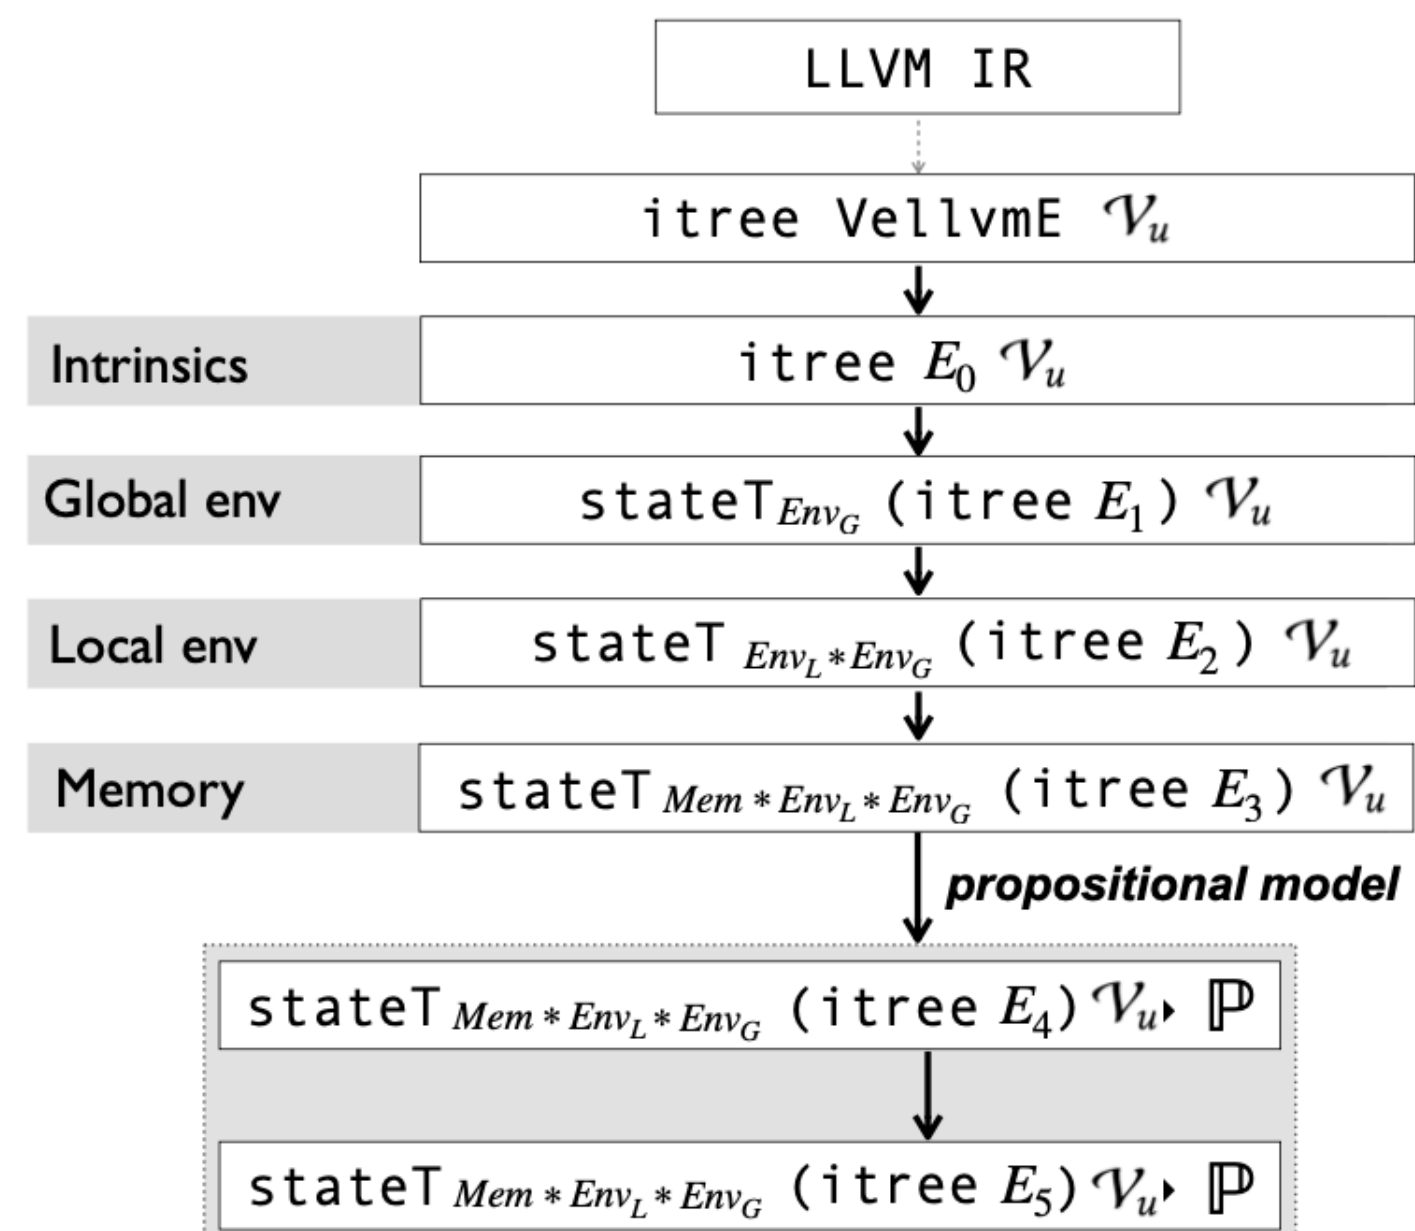
- Logical interpretation of memory attributes using permission-based ownership
- Can reason about reordering across calls and transformations that take advantage of memory attributes



External call semantics

Let's fix the naïve semantics!

- Pre-existing VIR semantics: external calls could not affect memory



- Event transformer: transforms an event into a state-passing event

```
Variant stateEff (E : Type -> Type) : Type -> Type :=  
| StateEff {X} : S * E X -> stateEff (S * X).
```

With this simple change, external calls are aware of memory

How do we reason about function calls?

- `eutt` is not enough: we cannot rely on syntactic trace equivalence for function calls
 - e.g. `call foo (%p1)` can be related to `call foo (%p2)` if %p1 and %p2 store related pointers

Simuliris [Gäher et al.]-style public resources

- Locations in public bijection (related pointers) $l_t \leftrightarrow_h l_s$

LOCESCAPE

$$\frac{\{l_t \leftrightarrow_h l_s * P\} e_t \leq e_s \{\Phi\}}{\{l_t \mapsto^{\text{tgt}} v_t * l_s \mapsto^{\text{src}} v_s * \mathcal{V}(v_t, v_s) * P\} e_t \leq e_s \{\Phi\}}$$

How to reason about public resources

first approximation: Simuliris [Gäher et al.]-style public resources

- We can store and load from public resources if they haven't been checked out by others yet



Has this been checked out yet?

SIMSTORE

$$\{l_t \leftrightarrow_h l_s * \mathcal{V}(v_t, v_s) * \text{checkout } C * (l_t, l_s) \notin C\}$$
$$\text{trigger}(\text{Store}^{()}(\text{addr}(l_t), v_t)) \leq \text{trigger}(\text{Store}^{()}(\text{addr}(l_s), v_s))$$
$$\{v_t, v_s. \mathcal{V}(v_t, v_s) * \text{checkout } C\}$$

SIMLOAD

$$\{l_t \leftrightarrow_h l_s * \text{checkout } C * ((l_t, l_s) \notin C \vee C(l_t, l_s) < 1)\}$$
$$\text{trigger}(\text{Load}^{\mathcal{V}_u}(\tau, \text{addr}(l_t))) \leq \text{trigger}(\text{Load}^{\mathcal{V}_u}(\tau, \text{addr}(l_s)))$$
$$\{v_t, v_s. \mathcal{V}(v_t, v_s) * \text{checkout } C\}$$

Function attribute specifications

Attribute specifications

SIMCALL

$$\{\text{Frame}^{\text{tgt}} i_t * \text{Frame}^{\text{src}} i_s * \text{checkout } \emptyset * \vec{\mathcal{V}}(args_t, args_s)\}$$
$$\text{call } \tau f(args_t) \leq \text{call } \tau f(args_s)$$
$$\{v_t, v_s. \text{Frame}^{\text{tgt}} i_t * \text{Frame}^{\text{src}} i_s * \text{checkout } \emptyset * \mathcal{V}(v_t, v_s)\}$$


Nothing's been checked out, so anyone can have full access to public resources!

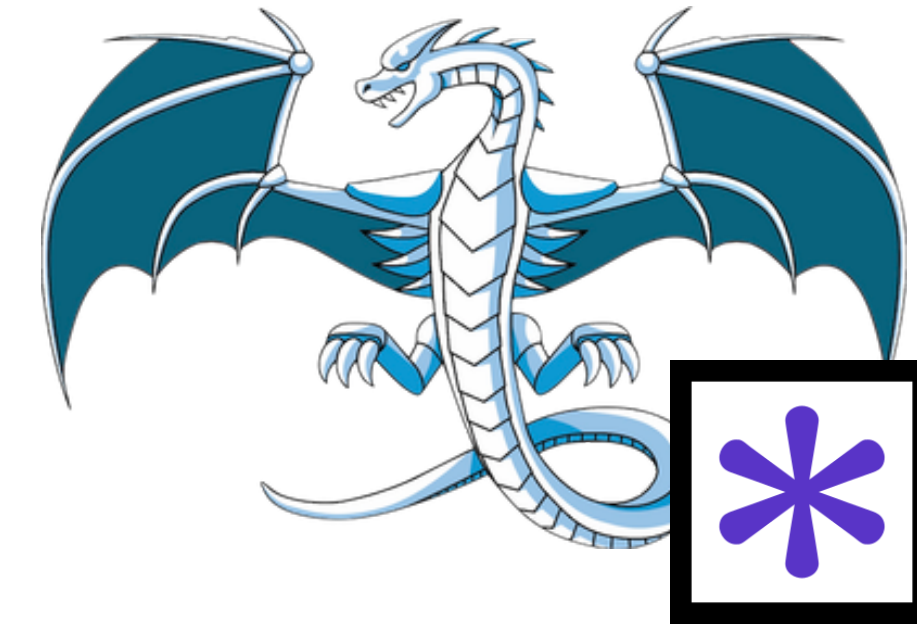
- Based on the function attribute, the simulation checks whether the patron should have full access to the material (or a partial scan) (i.e. permission-based ownership), and makes sure that all resources have been safely returned.



READONLY-CALL

$$\{\text{Frame}^{\text{tgt}} i_t * \text{Frame}^{\text{src}} i_s * \text{checkout } C *$$
$$* \vec{\mathcal{V}}(args_t, args_s) * (\forall (\ell_t, \ell_s) \in C. C(\ell_t, \ell_s) = q \wedge q < 1)\}$$
$$\text{call } \tau f(args_t) \text{ readonly } \leq \text{call } \tau f(args_s) \text{ readonly}$$
$$\{v_t, v_s. \text{Frame}^{\text{tgt}} i_t * \text{Frame}^{\text{src}} i_s * \mathcal{V}(v_t, v_s) * \text{checkout } C\}$$

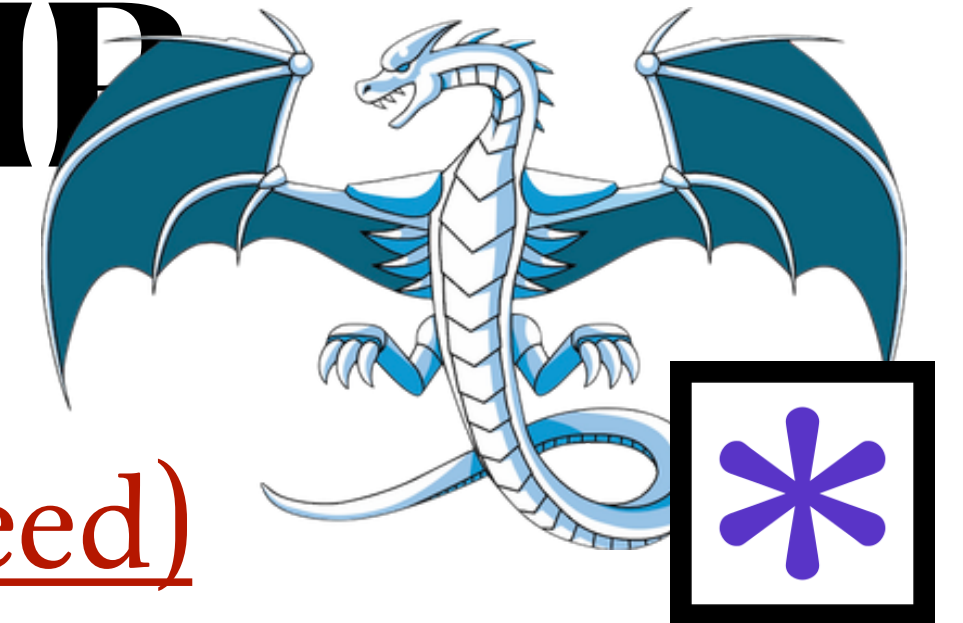

Let's check the access privileges ...

Roadmap



1. A taste of the stack-based ghost theory for LLVM IR 
2. Memory-relevant attributes in LLVM IR 
3. Model : A relational weakest-precondition model for Interaction Trees in Iris
4. Adequacy

Building an Iris framework for VIP



Typical recipe

(1) Ingredient: an abstract view on state
(ghost theory) using separation logic

resources

(2) Ingredient: a small-step semantics

Given a small-step semantics, a Hoare triple can be derived via the typical weakest precondition model* of Iris

(technically, a Banach guarded fixpoint)*

What we have (and need)

(1) Ingredient: A ghost theory for VIR
resources

(2) Ingredient: ITree-based semantics

A new weakest precondition model* of
Iris for stateful ITrees

(technically, a Knaster-Tarski mixed
fixpoint)*

Weakest precondition

Behind the scenes...



Definition `sim_expr_inner`

```

    (greatest_rec : st_expr_rel' -d> st_expr_rel')
    (least_rec : st_expr_rel' -d> st_expr_rel')
: st_expr_rel' -d> st_expr_rel' :=
λ ϕ st_t st_s ⇒
  (|==> ∃ (c : sim_case), ← (little trick with enums to avoid extra destruct on match cases,
                             since we don't have native variants or inductive types in Iris)
    match c with
    | BASE ⇒ ϕ st_t st_s
    | STUTTER_L ⇒ stutter_l least_rec ϕ st_t st_s
    | STUTTER_R ⇒ stutter_r least_rec ϕ st_t st_s
    | TAU_STEP ⇒ tau_step greatest_rec ϕ st_t st_s
    | VIS_STEP ⇒ vis_step greatest_rec ϕ st_t st_s
    | SOURCE_UB ⇒ source_ub st_t st_s
    | SOURCE_EXC ⇒ source_exc st_t st_s
    end)%I.

```

Definition `sim_expr_` : `expr_rel -d> expr_rel` :=

```

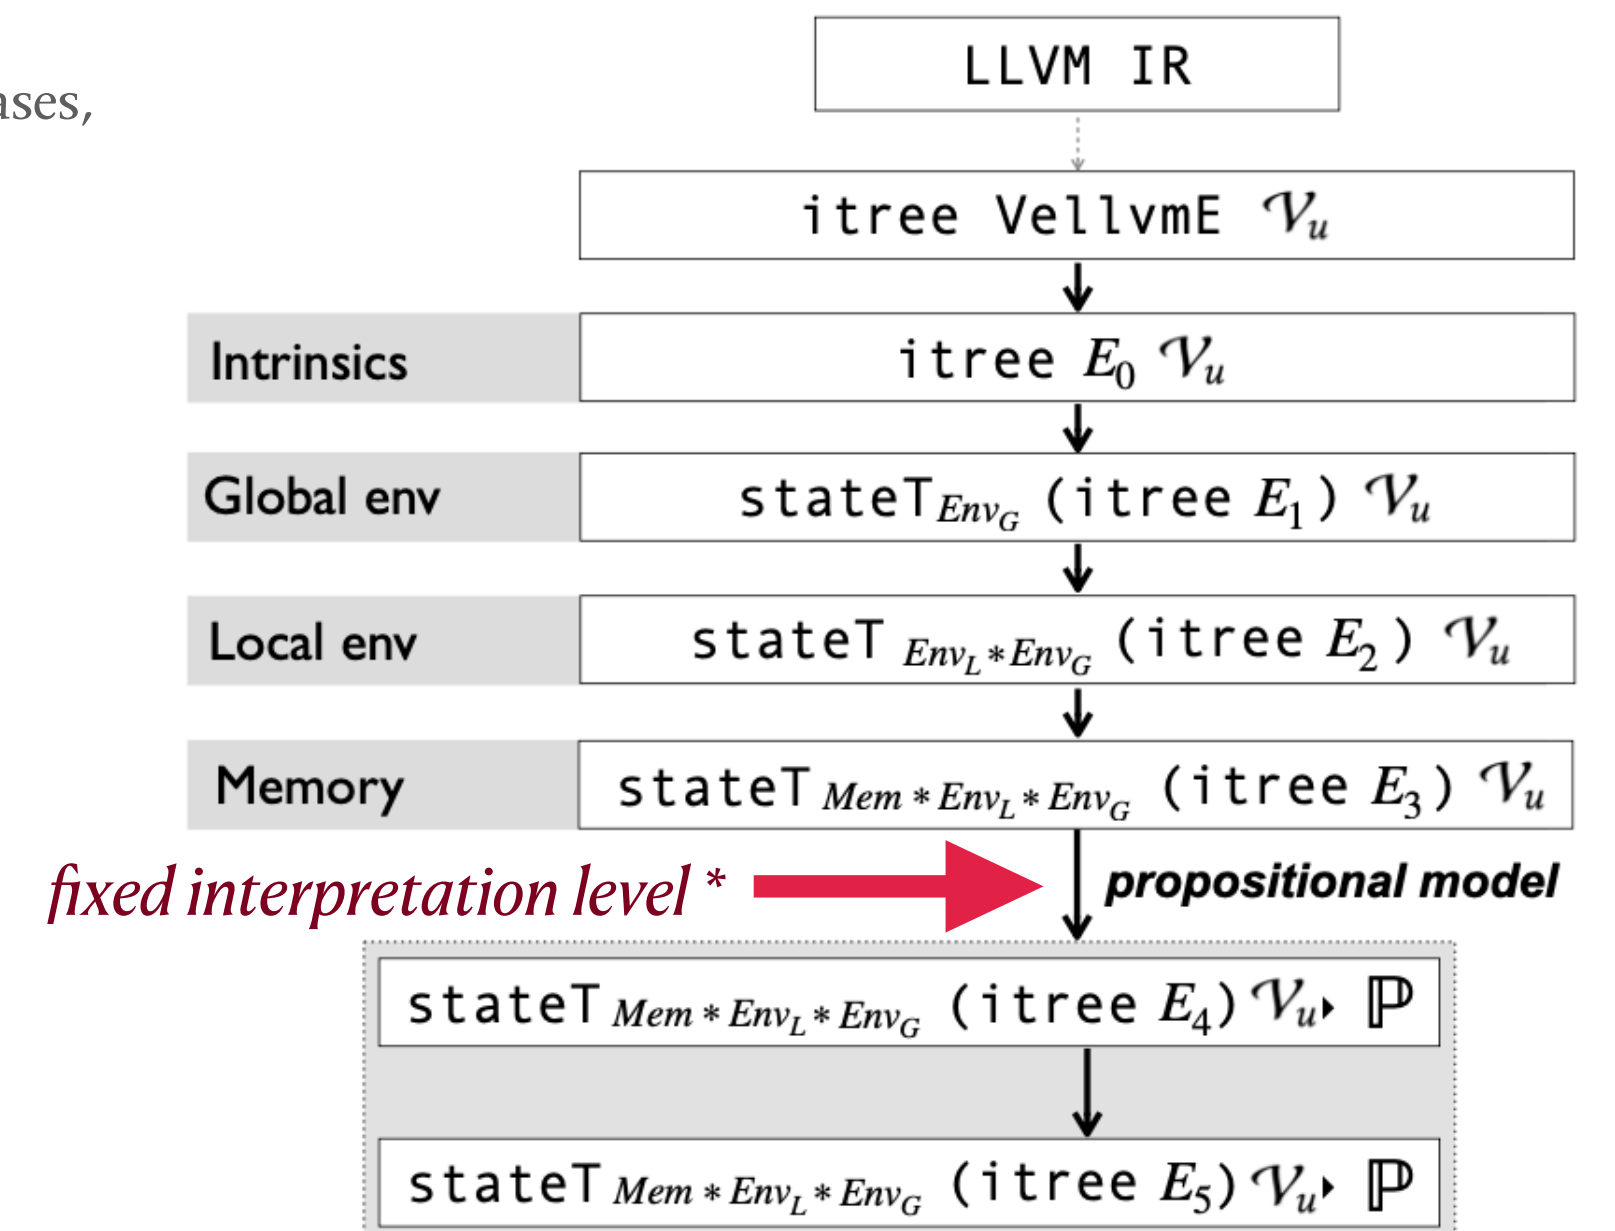
λ ϕ e_t e_s ⇒
  (∀ σ_t σ_s, state_interp σ_t σ_s ==*
    sim_coind ϕ ([[η⟨e_t⟩]] σ_t) ([[η⟨e_s⟩]] σ_s))%I.

```

(`sim_coind` takes the mixed greatest-least fixpoint of `sim_expr_inner`)

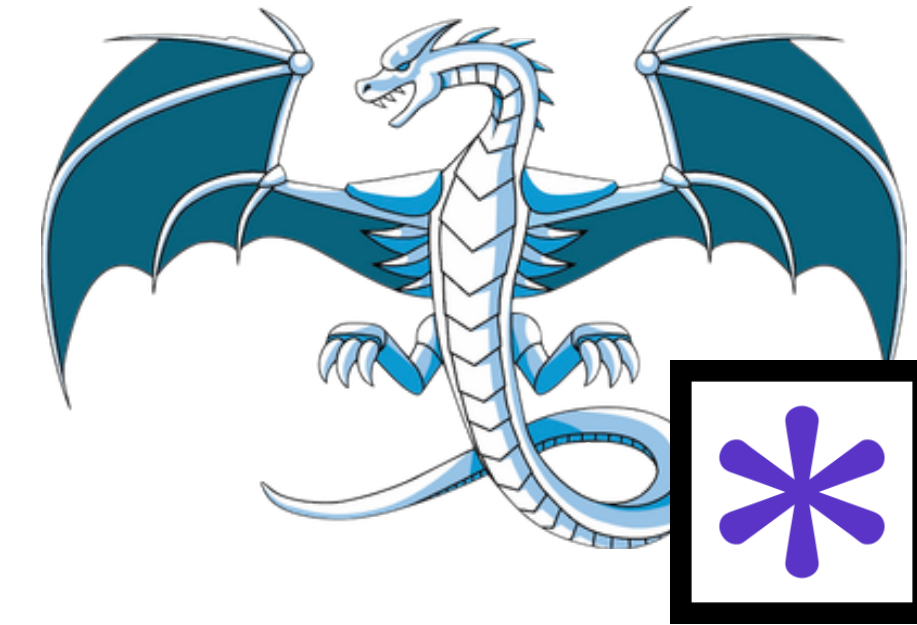
Q. What is the stateful interpretation function $\llbracket - \rrbracket$?




A.



*(with modified interpretation for calls)

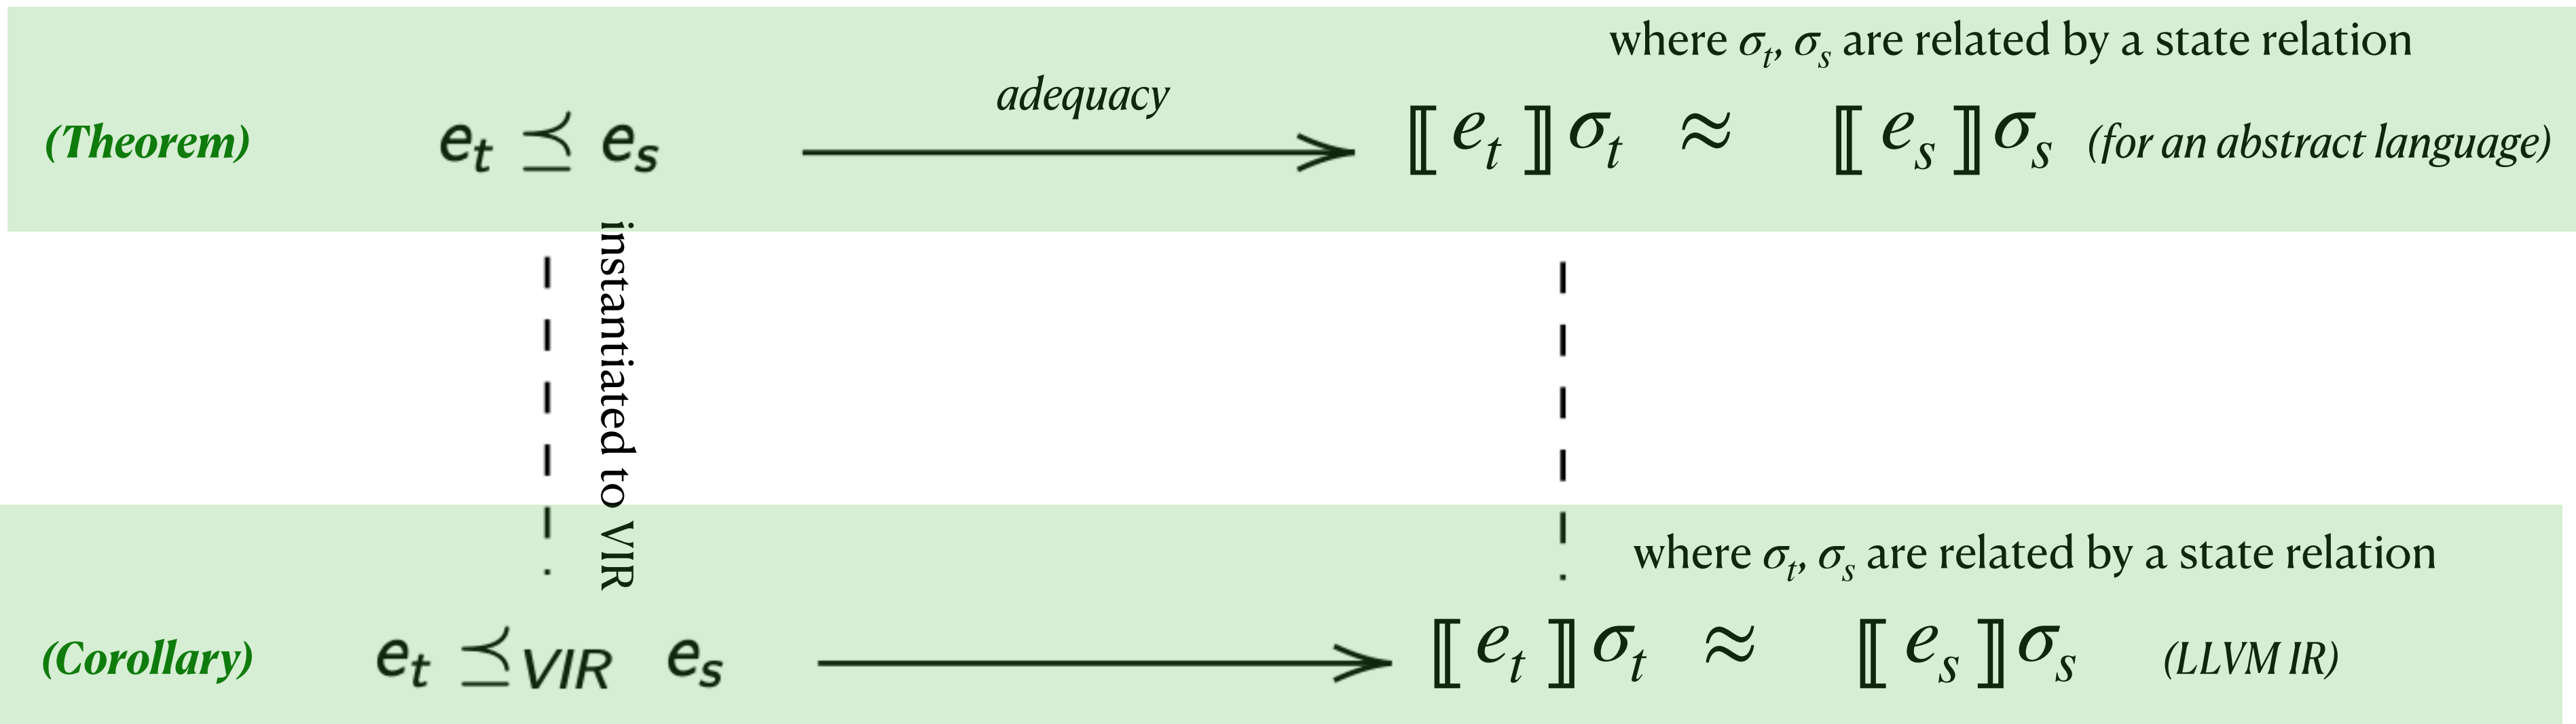
Roadmap



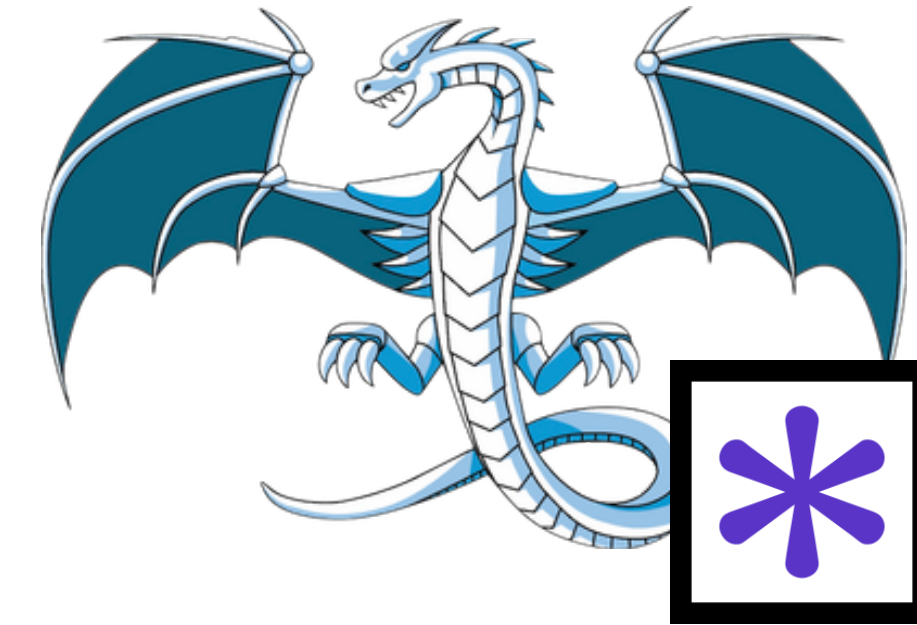
1. A taste of the stack-based ghost theory for LLVM IR 
2. Memory-relevant attributes in LLVM IR 
3. Model : A relational weakest-precondition model for Interaction Trees in Iris 
4. Adequacy





Adequacy

For ITrees e_t and e_s without external calls,



Roadmap



1. A taste of the stack-based ghost theory for LLVM IR 
2. Memory-relevant attributes in LLVM IR 
3. Model : A relational weakest-precondition model for Interaction Trees in Iris 
4. Adequacy 

Back to: Proving LICM

```
1 void increment(int *n);
2 int get_int (int *x) {
3     int *n; int i = 0; n = &i;
4     while (*n < *x) { increment(n); }
5     return *n;
6 }
```

```
1 void increment(int *n);
2 int get_int_opt (int *x) {
3     int *n; int i = 0;
4     n = &i; int y = *x;
5     while (*n < y) { increment(n); }
6     return *n;
7 }
```

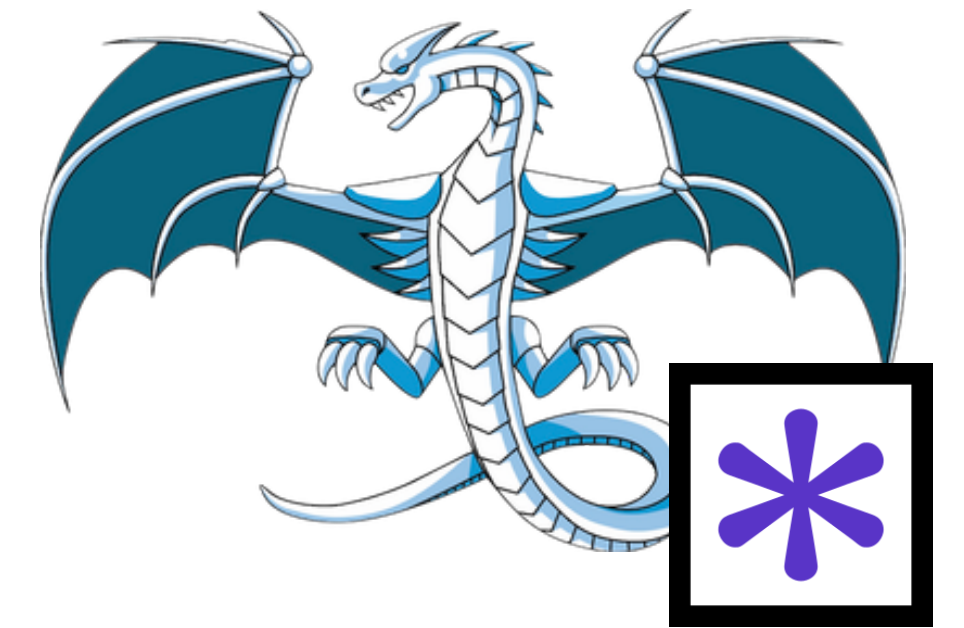
- Example: proof of simulation for a simple loop invariant code motion algorithm
- Benefits: can reorder memory-relevant instructions with function calls (could not be expressed before)
- Benefits: Hoare-style reasoning over loops; proof does not require explicit coinduction

Contributions



: results mechanized in the Coq Proof Assistant
Development: ~ 40k LOC (and ongoing...)

Velliris: A relational separation logic framework for LLVM IR



- A relational, coinductive weakest precondition model of Iris which supports a monadic semantics based on the Interaction Trees framework
- A relational separation logic and ghost theory for VIR resources
- Logical interpretation for memory-relevant attributes
- Examples: collection of simple examples and proof of simple loop invariant code motion algorithm
- Logical relation and contextual refinement (omitted)
- (Ongoing) case study: Verification of Mem2Reg algorithm

Thank you!