

# A Concurrency Model based on Monadic Interpreters: Executable semantics for a concurrent subset of LLVM IR


PhD defense

Nicolas Chappe (Inria Lyon, LIP)


Supervised by Ludovic Henrio (CNRS, LIP) and Yannick Zakowski (Inria Lyon, LIP)

November 22, 2024

Computer programs are everywhere, including in critical systems that have to be bug-free



Computer programs are everywhere, including in critical systems that have to be bug-free




The code the machine understands is hard to understand for humans

---

```
    cmpl    $0, a(%rip)
    jl      .LBB0_2
# %bb.1:
    leaq    .L.str(%rip), %rdi
    movb    $0, %al
    callq   printf@PLT
.LBB0_2:
```

---

Computer programs are everywhere, including in critical systems that have to be bug-free



The code the machine understands is hard to understand for humans

---

```
    cmpl    $0, a(%rip)
    jl      .LBB0_2
# %bb.1:
    leaq    .L.str(%rip), %rdi
    movb    $0, %al
    callq   printf@PLT
.LBB0_2:
```

---

*Verified compilation* can reconcile these two facts

---

```
if (a >= 0)  
    say_hello()
```

---

Human-readable code  
written in a *programming language*

---

```
if (a >= 0)
    say_hello()
```

---

Human-readable code  
written in a *programming language*

Compiler  
→


---

```
        cmpl    $0, a(%rip)
        jl      .LBB0_2
# %bb.1:
        leaq    .L.str(%rip), %rdi
        movb    $0, %al
        callq   printf@PLT
.LBB0_2:
```

---


Machine-readable code  
in *assembly language*

- Is the compiler output correct?



- 1 Formalize the behaviors of programs in the source language


- Is the compiler output correct?



- 1 Formalize the behaviors of programs in the source language
- 2 Formalize the behaviors of programs in the target language

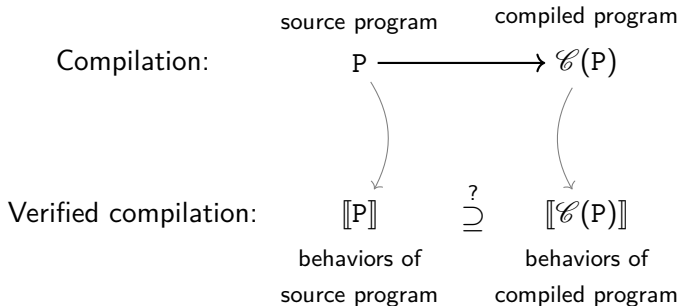


- Is the compiler output correct?



- 1 Formalize the behaviors of programs in the source language
- 2 Formalize the behaviors of programs in the target language
- 3 Check that the behaviors of the compiled program are also behaviors of the source program

- Is the compiler output correct?



- 1 Formalize the behaviors of programs in the source language
  - 2 Formalize the behaviors of programs in the target language
  - 3 Check that the behaviors of the compiled program are also behaviors of the source program
- Notable example: CompCert in Coq/Rocq, compiles C to optimized machine code

---

```
if (a >= 0)
  say_hello()
```

---

Human-readable code  
that satisfies some safety  
properties

Verified compiler



---

```
    cmpl    $0, a(%rip)
    jl      .LBB0_2
# %bb.1:
    leaq    .L.str(%rip), %rdi
    movb    $0, %al
    callq   printf@PLT
.LBB0_2:
```

---

Machine-readable code  
that satisfies the same safety  
properties

Many programming languages

C

Rust

Swift

Many programming languages

C

Rust

Swift


Many kinds of machine  
(*architectures*)

x86

ARM

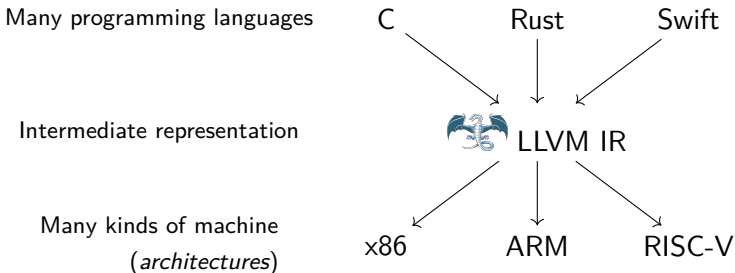
RISC-V

Many programming languages



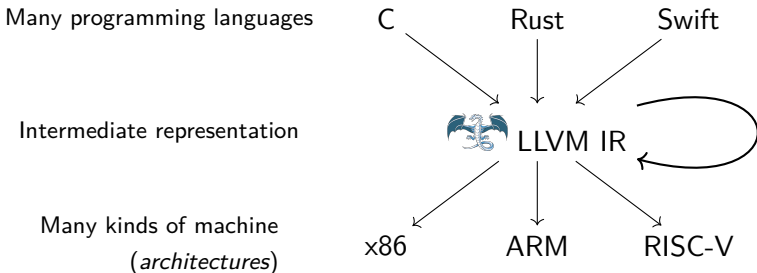
Many kinds of machine  
(*architectures*)

This is getting complex!



- LLVM IR: intermediate language that factors out many compilers<sup>1</sup>

<sup>1</sup>Chris Lattner and Vikram Adve. "LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation". In: CGO '04. USA: IEEE Computer Society, 2004.




- LLVM IR: intermediate language that factors out many compilers<sup>1</sup>
- Many program transformations and optimizations defined on LLVM IR
- Interesting starting point for a verified compilation infrastructure


<sup>1</sup>Chris Lattner and Vikram Adve. “LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation”. In: CGO '04. USA: IEEE Computer Society, 2004.



Vellvm's goal: *efficiently model the semantics of LLVM IR in Rocq*




Vellvm's goal: *efficiently model the semantics of LLVM IR in Rocq*<sup>2</sup>



---

<sup>2</sup>Yannick Zakowski, Calvin Beck, Irene Yoon, Ilia Zaichuk, Vadim Zaliva, and Steve Zdancewic. "Modular, Compositional, and Executable Formal Semantics for LLVM IR". In: *Proc. ACM Program. Lang.* 5.ICFP (Aug. 2021). DOI: 10.1145/3473572.

Vellvm's goal: *efficiently model the semantics of LLVM IR in Rocq*<sup>2</sup>




Limitation: *no support for concurrency*


---

<sup>2</sup>Yannick Zakowski, Calvin Beck, Irene Yoon, Ilia Zaichuk, Vadim Zaliva, and Steve Zdancewic. "Modular, Compositional, and Executable Formal Semantics for LLVM IR". In: *Proc. ACM Program. Lang.* 5.ICFP (Aug. 2021). DOI: 10.1145/3473572.


Standard program: series of instructions for the processor




Standard program: series of instructions for the processor



**Concurrent** program: contains several tasks (or threads)




What happens when two tasks access the same memory cell?




- Many other subtleties, cf. *weak memory models*

Vellvm's goal: *efficiently model the semantics of LLVM IR in Rocq*



Our goal: *efficiently model the semantics of **concurrent** LLVM IR in Rocq*




# Context: Interaction Trees



- Tree model for representing programs
- Modular, reusable semantics
- Executable semantics
- Mechanized as a Rocq (formerly Coq) library



Return a value



Interact with the  
environment



Perform an internal  
computation

<sup>3</sup>Li-yao Xia, Yannick Zakowski, Paul He, Chung-Kil Hur, Gregory Malecha, Benjamin C. Pierce, and Steve Zdancewic. “Interaction trees: representing recursive and impure programs in Coq”. In: *Proc. ACM Program. Lang.* 4.POPL (Dec. 2019). DOI: 10.1145/3371119.


- Each **Vis** node represents an effect of the program
- Branches of a **Vis** node represent the possible answers from the environment
- **Ret** nodes are leaves returning a value

LLVM IR example with memory access effects:

---

```
%x = load @a  
store %x, @b  
ret 0
```

---



ITrees are *coinductive*, they can be infinitely deep.

```
while true do print
```

Vis Print



Step



Vis Print



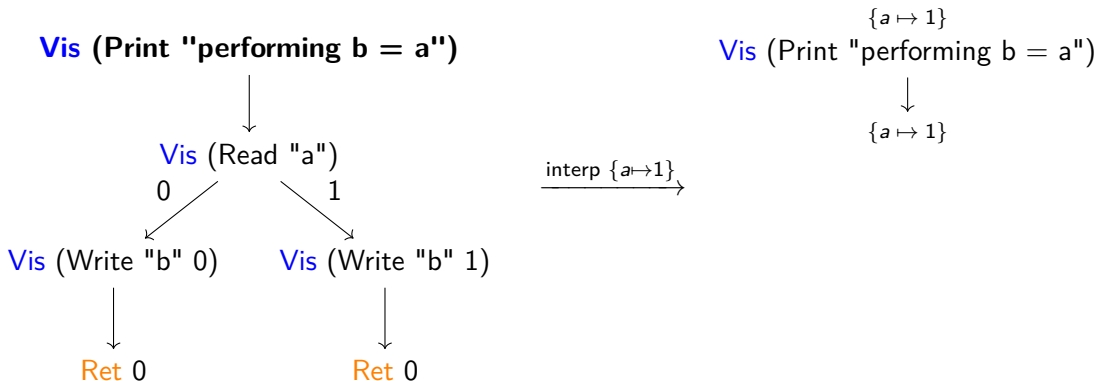
Step



- Unary **Step** nodes mark internal events such as loop iterations

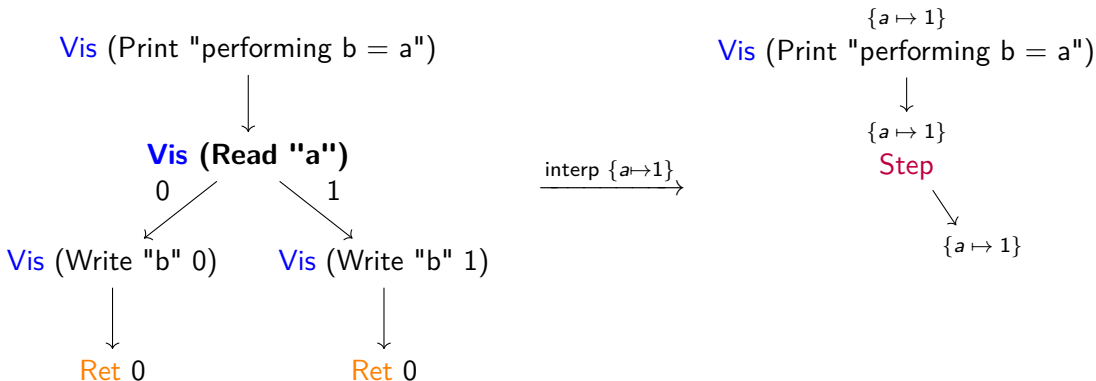
We can make the semantics appear by providing a stateful interpreter for **Vis** nodes.

- Write events update a memory state
- Read events return the corresponding value from memory



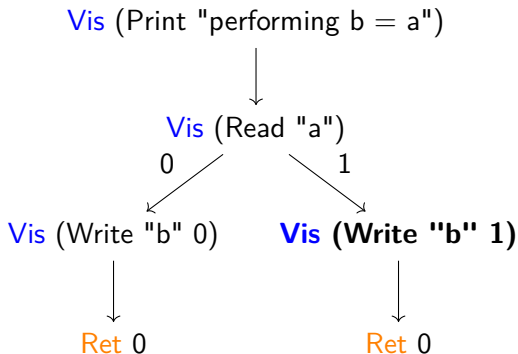
We can make the semantics appear by providing a stateful interpreter for **Vis** nodes.

- Write events update a memory state
- Read events return the corresponding value from memory

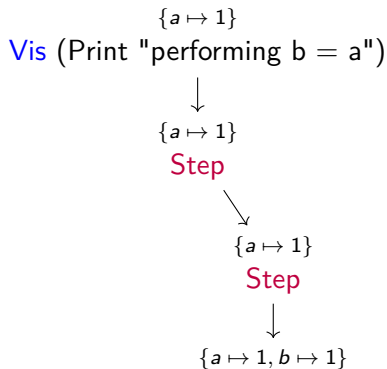


We can make the semantics appear by providing a stateful interpreter for **Vis** nodes.

- Write events update a memory state
- Read events return the corresponding value from memory

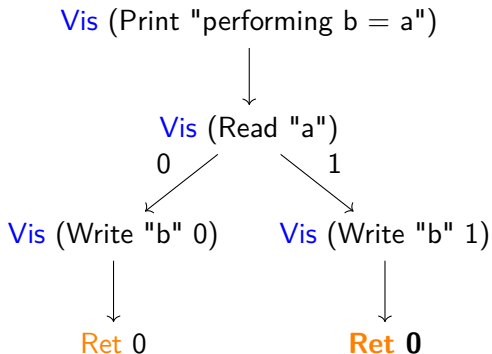


$\xrightarrow{\text{interp } \{a \mapsto 1\}}$

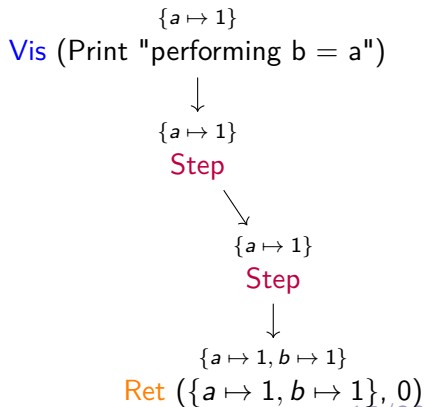


We can make the semantics appear by providing a stateful interpreter for **Vis** nodes.


- Write events update a memory state
- Read events return the corresponding value from memory




interp  $\{a \mapsto 1\}$



- General methodology: successively refine each kind of event node.



- In the end, we can execute the tree




Vellvm applies this approach to LLVM, with 6 layers of interpretation




# Choice Trees

Consider a Rand event that generates a random number. How to interpret it?




- No proper way to interpret it with ITrees

- Not interpreting it is semantically wrong




Consider a Rand event that generates a random number. How to interpret it?





- No proper way to interpret it with ITrees

- Not interpreting it is semantically wrong



- Special Br nodes represent nondeterministic choices, we will give them appropriate semantics later

- Infinite (coinductive) tree model for representing programs
- Rocq library that takes inspiration from ITrees, with nondeterminism support
- Collaboration between LIP and UPenn



- Infinite (coinductive) tree model for representing programs
- Rocq library that takes inspiration from ITrees, with nondeterminism support
- Collaboration between LIP and UPenn




We can give nondeterministic semantics to **Br** nodes by defining a fitting notion of equivalence of CTrees

# Building a labeled transition system (LTS) from a CTree

An LTS consists of a set of states and a set of labeled transitions.


↓  
Ret v

→



# Building a labeled transition system (LTS) from a CTree

An LTS consists of a set of states and a set of labeled transitions.





# Building a labeled transition system (LTS) from a CTree


An LTS consists of a set of states and a set of labeled transitions.

↓  
**Ret**  $v$

→




→



↓  
**Step**  
↓  
 $t$

→







# Building a labeled transition system (LTS) from a CTree


An LTS consists of a set of states and a set of labeled transitions.

↓  
**Ret** v

→





→



↓  
**Step**  
↓  
t


→





→

- Collapsed in the LTS
- Does not generate a transition



## CTrees and the underlying LTS: example




# CTrees and the underlying LTS: example





# CTrees and the underlying LTS: example



→





# CTree equivalence: Strong bisimulation on the LTS






Definition of a strong bisimulation  $t \mathcal{R} u$ , with two half-games

# CTree equivalence: Strong bisimulation on the LTS



Definition of a strong bisimulation  $t \mathcal{R} u$ , with two half-games

# CTree equivalence: Strong bisimulation on the LTS





Definition of a strong bisimulation  $t \mathcal{R} u$ , with two half-games

- Br nodes verify the algebraic laws of non-determinism
- Equational proof principles for coinductive proofs
- Many up-to principles, using the coinduction library from Damien Pous


- We cannot reason about collapsed Br nodes, which complicated some of my proofs.
- We can build an LTS in which they appear.



- We cannot reason about collapsed Br nodes, which complicated some of my proofs.
- We can build an LTS in which they appear.




Br as  $\epsilon$  transitions: *explicit LTS*




Br nodes collapsed in the original LTS

- We cannot reason about collapsed Br nodes, which complicated some of my proofs.
- We can build an LTS in which they appear.





How do we define bisimulation on this alternative explicit LTS?

It should relate the same CTrees as the original definition.







Br as  $\epsilon$  transitions: *explicit LTS*





# Intertwined bisimulation on the explicit LTS



# Intertwined bisimulation on the explicit LTS







# Intertwined bisimulation on the explicit LTS



① Left:  $\xrightarrow{a}$

Intertwined bisim. game  $\mathcal{R}_l \cap \mathcal{R}_r$

# Intertwined bisimulation on the explicit LTS




① Left:  $\xrightarrow{a}$

Intertwined bisim. game  $\mathcal{R}_l \cap \mathcal{R}_r$

② Left:  $\xrightarrow{\epsilon}$

Left half-game  $\mathcal{R}_l$  only

# Intertwined bisimulation on the explicit LTS



① Left:  $\xrightarrow{a}$

① Left:  $\xrightarrow{a}$



Intertwined bisim. game  $\mathcal{R}_l \cap \mathcal{R}_r$

② Left:  $\xrightarrow{\epsilon}$

Left half-game  $\mathcal{R}_l$  only

② Left:  $\xrightarrow{\epsilon}$

# Intertwined bisimulation on the explicit LTS





We define intertwined bisimulation using two *mutually coinductive* relations  $\mathcal{R}_l$  and  $\mathcal{R}_r$ .

- Proved equivalent to strong bisimulation on the original LTS
- Makes some proofs significantly easier, especially around the interpretation combinator
- Equational theory similar to the original definition

- CTrees capture a wide class of nondeterministic LTSs
- Effects of a program can be handled in a modular way with interpretation
- CTrees are executable
- Paper at POPL'23<sup>4</sup>

## Other personal contributions in the thesis

- Notions of refinement: Strong similarity, complete similarity
- Meta-theoretical results on interpretation and other combinators
- ...

---

<sup>4</sup>Nicolas Chappe, Paul He, Ludovic Henrio, Yannick Zakowski, and Steve Zdancewic. “Choice Trees: Representing Nondeterministic, Recursive, and Impure Programs in Coq”. In: *Proc. ACM Program. Lang.* 7.POPL (Jan. 2023). DOI: 10.1145/3571254.

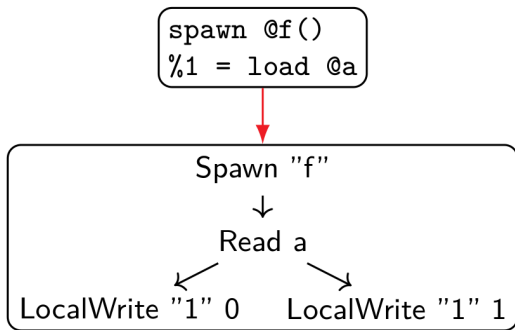
# Concurrent $\mu$ Vellvm

```
1  define @main() {  
2      store atomic 0, @x monotonic  
3      store atomic 0, @y monotonic  
4      call @thrd_create(@f)  
5      %1 = atomicrmw xchg @x, 1 monotonic  
6      %2 = atomicrmw xchg @y, 1 release  
7      ret 0  
8  }  
9  
10 define @f() {  
11     %y1 = atomicrmw xchg @y, 2 acquire  
12     %x1 = load atomic @x monotonic  
13     ret 0  
14 }
```


Syntax:

- Restricted subset of LLVM IR
- Thread management  
(thrd\_create and thrd\_join)
- Concurrent memory accesses

- We represent LLVM IR functions as CTrees with (VarE + ThreadE + MemE) events
- Reusable interpretation passes give semantics to these events




- We represent LLVM IR functions as CTrees with (VarE + ThreadE + MemE) events
- Reusable interpretation passes give semantics to these events




```
Variant VarE : Type → Type :=  
| LocalWrite (id: ident) (v: value) : VarE unit  
| LocalRead (id: ident) : VarE value  
| GlobalRead (id: ident) : VarE value
```

- We represent LLVM IR functions as CTrees with (VarE + ThreadE + MemE) events
- Reusable interpretation passes give semantics to these events



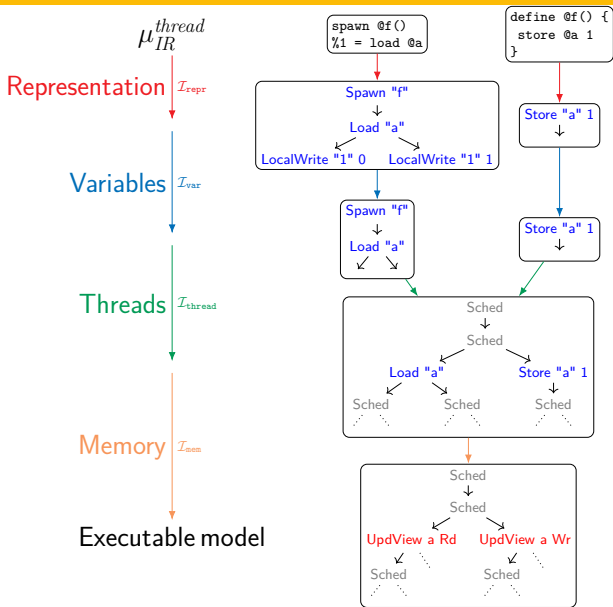
```
Variant ThreadE : Type → Type :=  
| Spawn (f: fid) (arg: value) : ThreadE thread_id
```

- We represent LLVM IR functions as CTrees with (VarE + ThreadE + MemE) events
- Reusable interpretation passes give semantics to these events

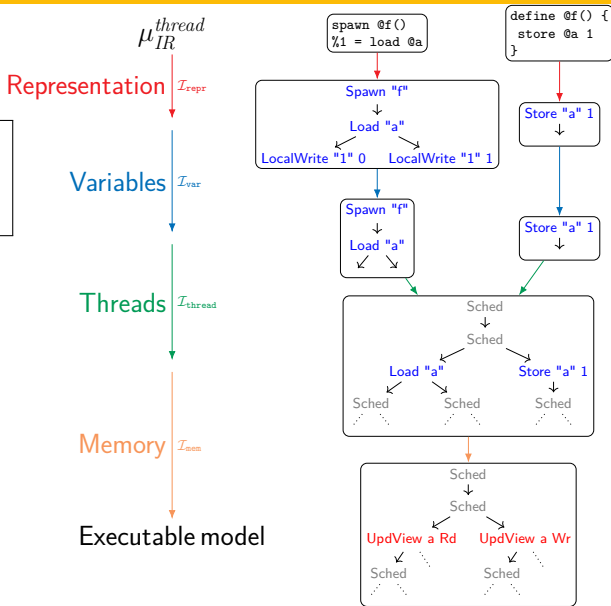


```
Variant MemE : Type → Type :=  
| Load (o: ordering) (k: addr) : MemE value  
| Store (o: ordering) (k: addr) (v: value) : MemE unit  
| RMW (o: ordering) (k: addr) (f: value → value) : MemE value  
| Fence (o: ordering) : MemE unit  
| Alloc (sz: nat) : MemE addr
```

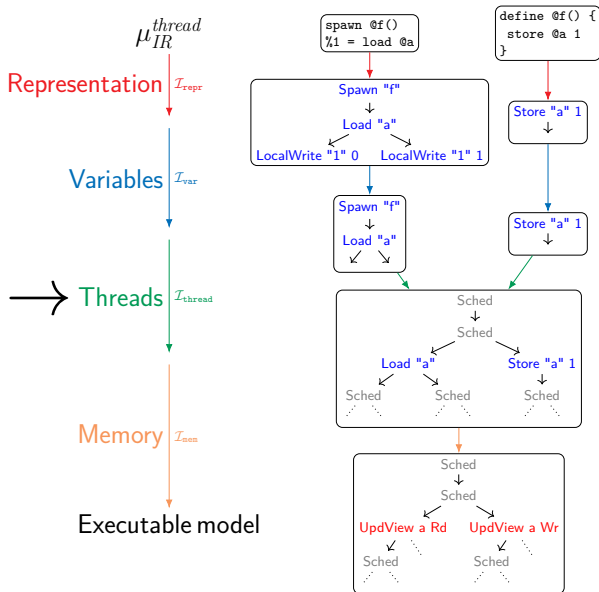




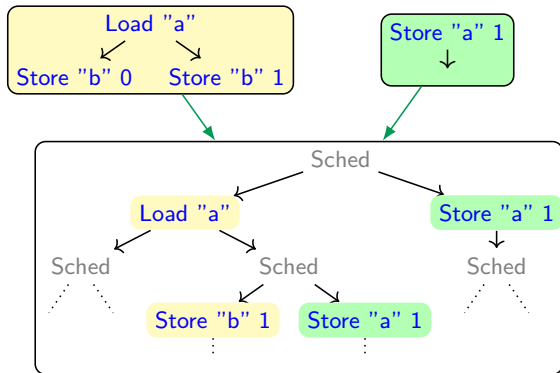
With a proof of transport of equivalences across the interpretation passes




# Interpreting multithreading



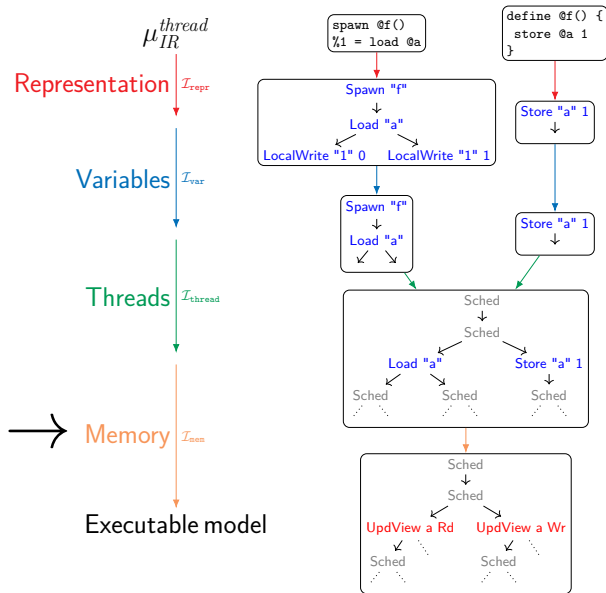
- To model the parallel composition of  $n$  threads, we interleave their CTrees, which gives a single CTree
- A `Br Sched` node chooses which thread to execute next



- To model the parallel composition of  $n$  threads, we interleave their CTrees, which gives a single CTree
- A `Br Sched` node chooses which thread to execute next
- Thread creation adds one CTree to the interleaving
- Based on a co-recursive combinator interleave fns tasks



# Interpreting memory events

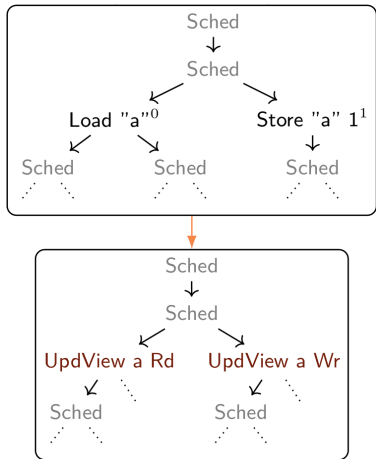


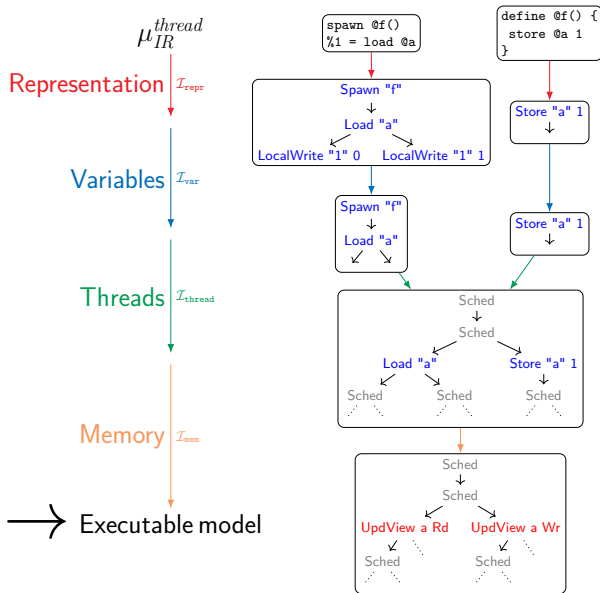
CTrees are modular! We support SC, TSO and a subset of Promising<sup>a</sup>, with simulation results.

## Promising-based model

- A very complete operational concurrent memory model
- We support read, write, read-modify-write and fence operations
- We support most levels of atomicity
- Monotonic accesses partly left to future work

<sup>a</sup>Jecheon Kang, Chung-Kil Hur, Ori Lahav, Viktor Vafeiadis, and Derek Dreyer. “A promising semantics for relaxed-memory concurrency”. In: *POPL'17*. ACM, 2017. DOI: 10.1145/3009837.3009850.





- We can still execute CTrees, but we have to schedule the remaining branching nodes
- Round-robin, random...
- Collecting interpreter written in OCaml



# Conclusion

## Contributions


- Choice Trees: modular and executable semantics for nondeterministic programs (POPL'23 paper)
- A novel notion of equivalence: intertwined bisimulation (paper in submission)
- Modular concurrency semantics, applied to LLVM (paper accepted at CPP'25)

## Future work


- Meta-theoretical results on the interpretation stack
- Integration into Vellvm
- A verified compilation chain



## Choice Trees



## Intertwined bisimulation



## LLVM IR concurrency

Thanks for your attention!