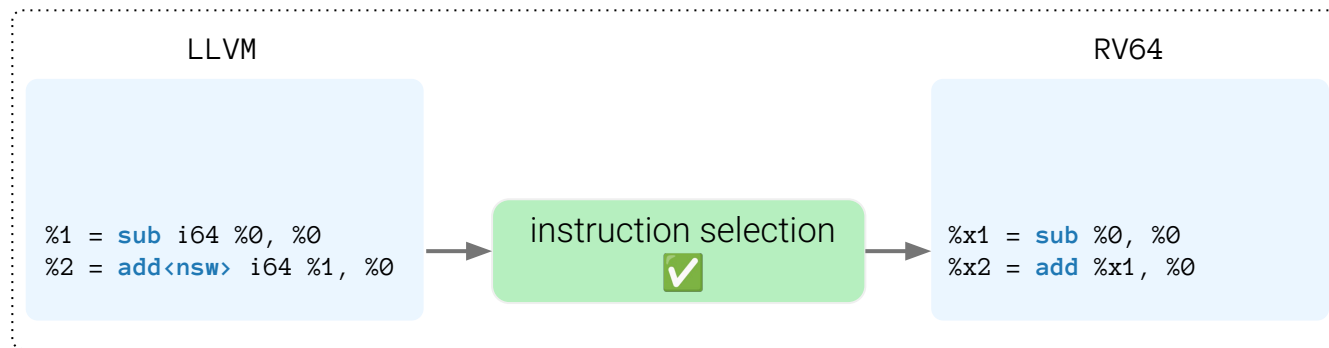
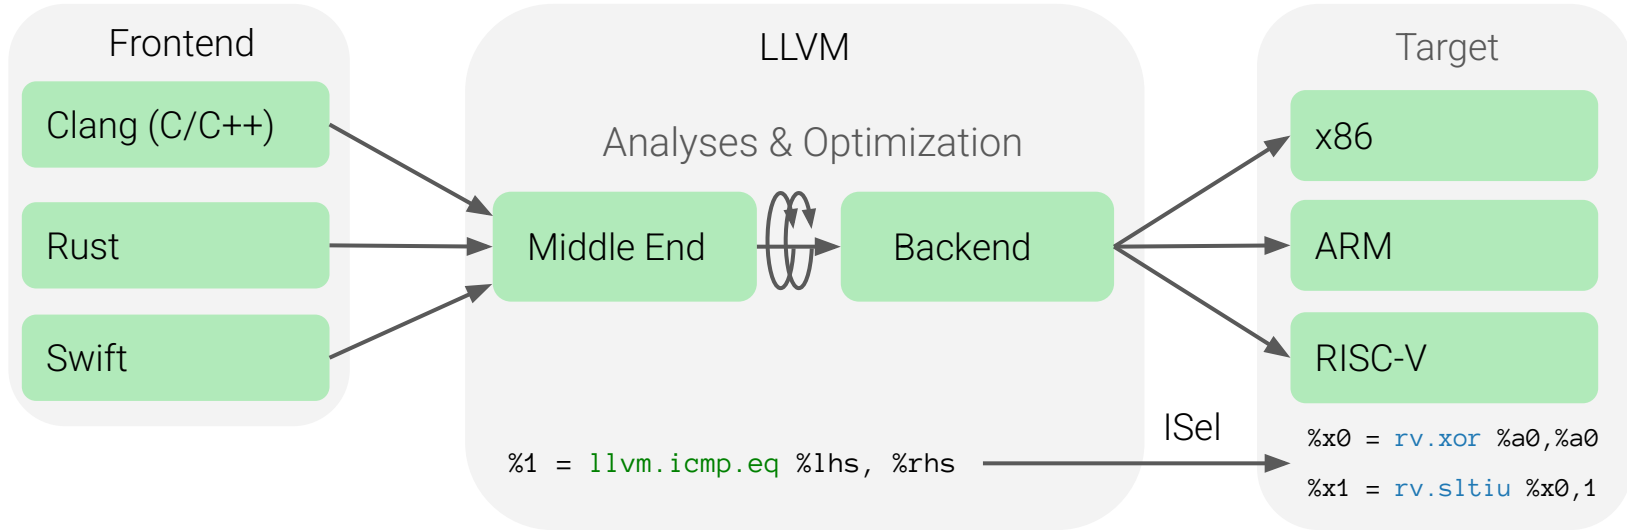


# Scaling Certified Instruction Selection For LLVM IR Through Bitblasting

**Sarah Linh Kuhn, Luisa Cicolini, Alex Keizer, Osman Yasar, Tobias Grosser**



# Many Programming Languages Rely On LLVM




# A look at LLVM's *backend*



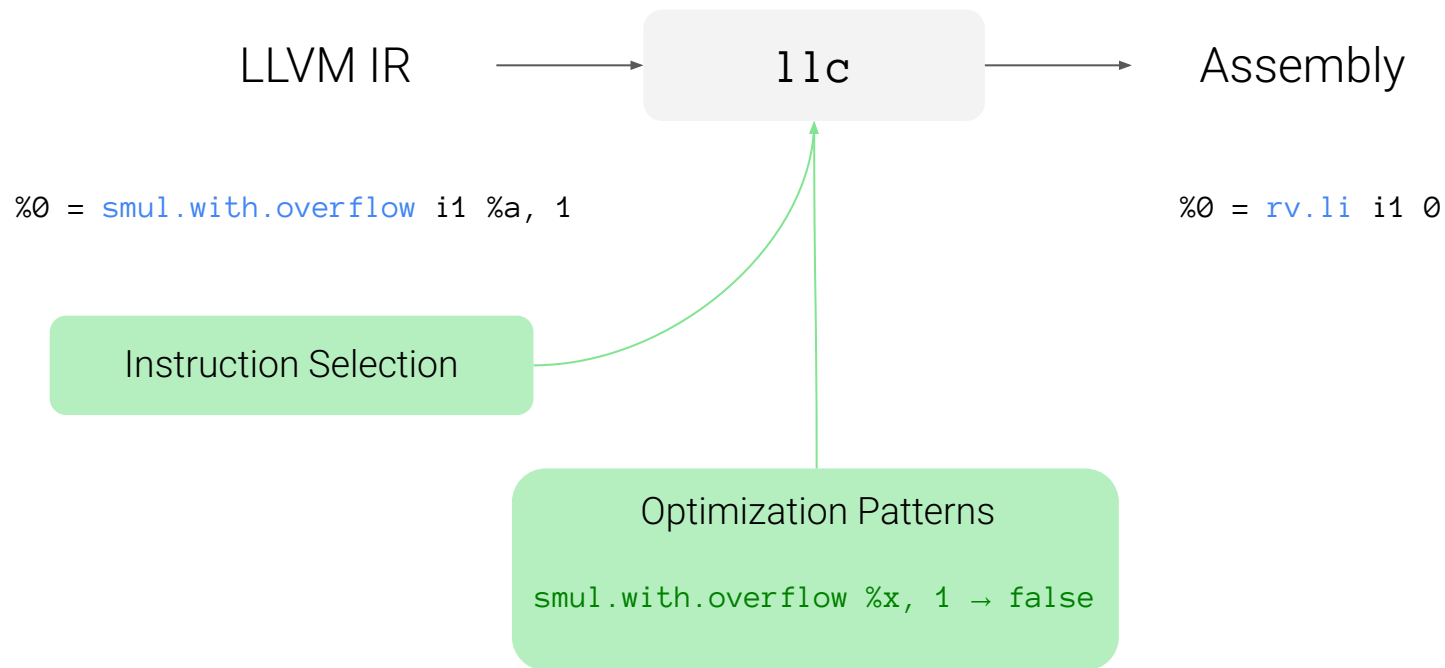
```
%0 = smul.with.overflow i1 %a, 1
```

```
%0 = rv.li i1 0
```

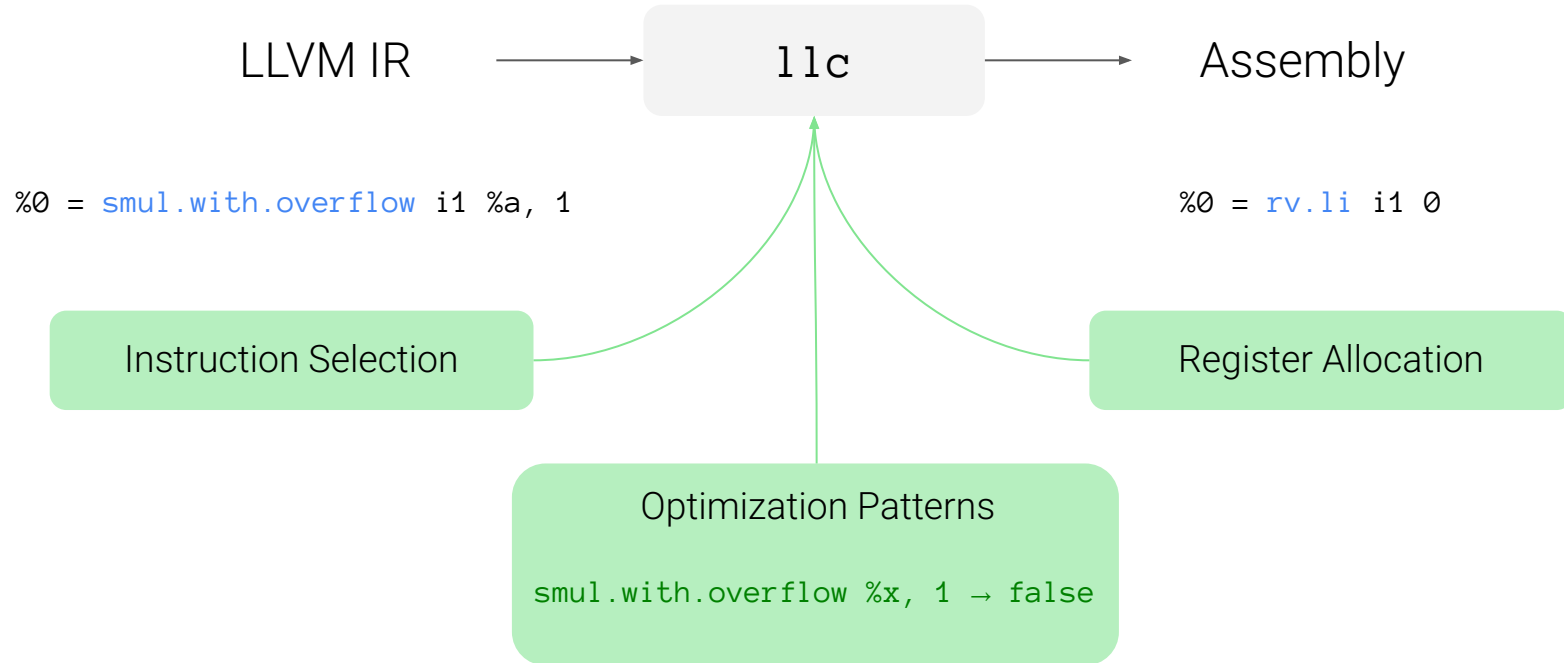
# A look at LLVM's *backend*



# A look at LLVM's *backend*



# A look at LLVM's *backend*



## Optimization Patterns

```
smul.with.overflow %x, 1 → false
```

# LLVM has *optimizations*

`x : i2 → signed integer`

Optimization Patterns

`smul.with.overflow %x, 1 → false`

# LLVM has *optimizations*

$x : i2 \rightarrow \text{signed integer}$

*	-2	-1	0	1
-2	4	2	0	-2
-1	2	1	0	-1
0	0	0	0	0
1	-2	-1	0	1

Optimization Patterns

```
smul.with.overflow %x, 1 → false
```

# LLVM has *optimizations...*

`x : i1 → signed integer`

*	-2	-1	0	1
-2	4	2	0	-2
-1	2	1	0	-1
0	0	0	0	0
1	-2	-1	0	1

Optimization Patterns

`smul.with.overflow %x, 1 → false`



# LLVM has *bugs*

`x : i1 → signed integer`

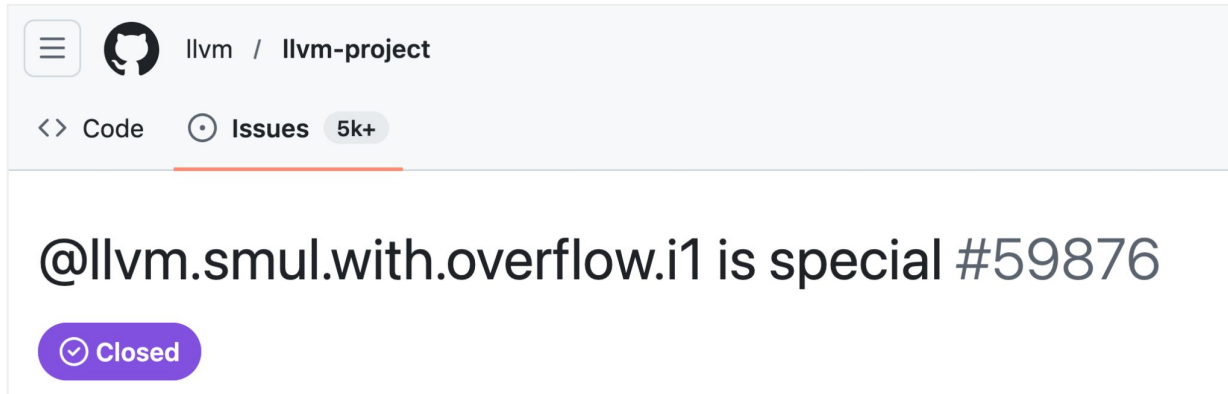
*	-2	-1	0	1
-2	4	2	0	-2
-1	2	1	0	-1
0	0	0	0	0
1	-2	-1	0	1

1 is interpreted as -1!


Optimization Patterns

```
smul.with.overflow %x, 1 → false
```

# LLVM has *bugs*

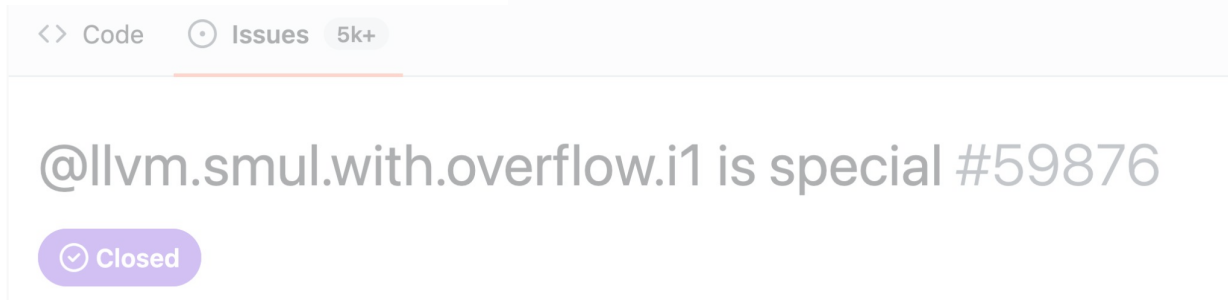


# LLVM has *bugs*

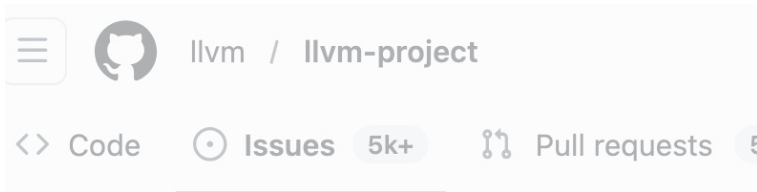


is:issue state:open type:Bug

Open 1.700 Closed 3.419



# LLVM has *bugs*



## High-Throughput, Formal-Methods-Assisted Fuzzing for LLVM

is:issue state:open type:Bug

Open 1.700 Closed

Yuyou Fan  
University of Utah  
USA  
yuyou.fan@utah.edu

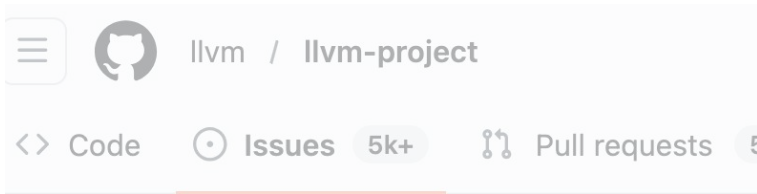
John Regehr  
University of Utah  
USA  
regehr@cs.utah.edu

zing throughput is 12x higher, on average, than a flow that runs mutation, optimization, and formal n separate processes. So far we have used alive- and report 33 previously unknown bugs in LLVM.

@llvm.s

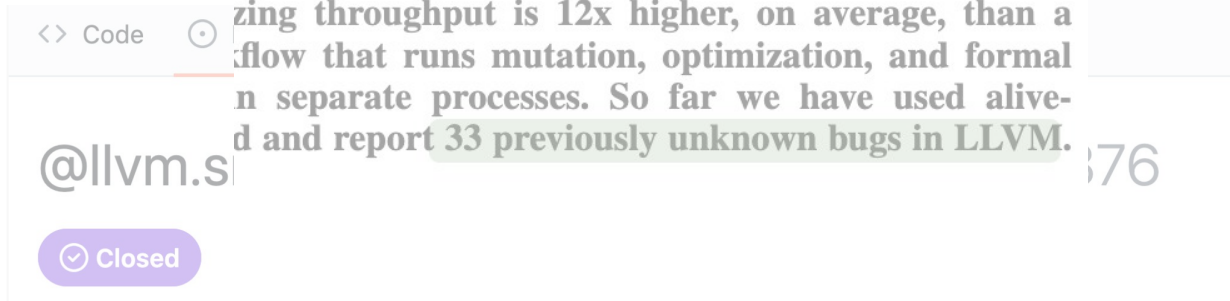
Closed

# LLVM has *bugs*



is:issue state:open type:Bug

Open 1.700 Closed



## Alive2: Bounded Translation Validation for LLVM

Nuno P. Lopes  
nlopes@microsoft.com  
Microsoft Research  
UK

Juneyoung Lee  
juneyoung.lee@sf.snu.ac.kr  
Seoul National University  
South Korea

Chung-Kil Hur  
gil.hur@sf.snu.ac.kr  
Seoul National University  
South Korea

Zhengyang Liu  
liuz@cs.utah.edu  
University of Utah  
USA

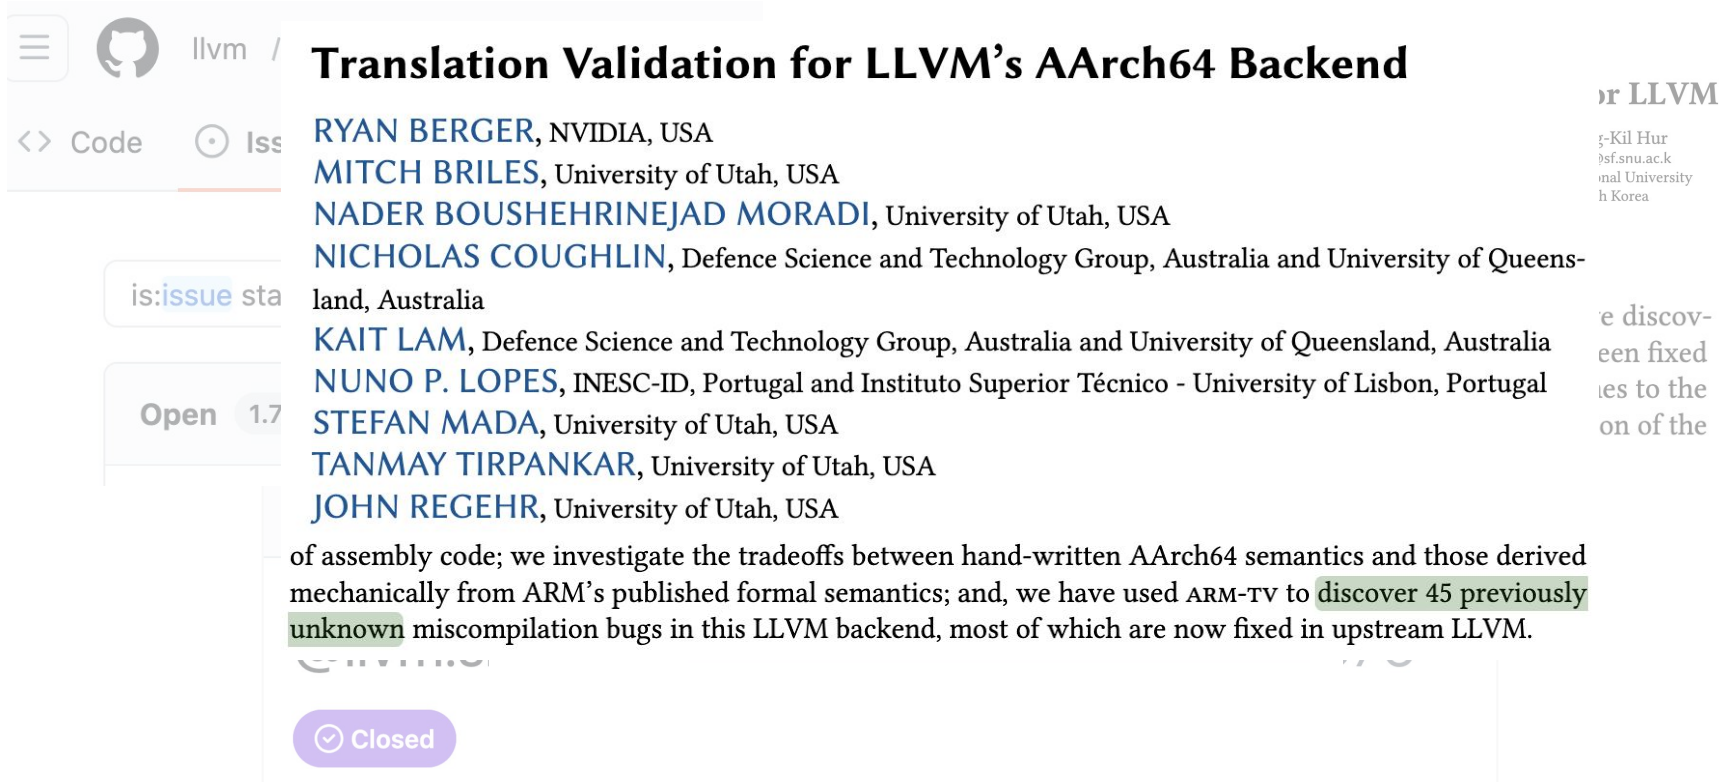
John Regehr  
regehr@cs.utah.edu  
University of Utah  
USA

High-Throughput, Formal Fuzzing for LLVM

By running Alive2 over LLVM's unit test suite, we discovered and reported 47 new bugs, 28 of which have been fixed already. Moreover, our work has led to eight patches to the LLVM Language Reference—the definitive description of the

Fuzzing throughput is 12x higher, on average, than a workflow that runs mutation, optimization, and formal verification in separate processes. So far we have used alive2 and report 33 previously unknown bugs in LLVM.

# LLVM has *bugs*



The screenshot shows a GitHub issue page for the LLVM project. The issue title is "Translation Validation for LLVM's AArch64 Backend". The issue is marked as "Closed" and has 17 open discussions. The issue description, which is highlighted in the image, reads: "of assembly code; we investigate the tradeoffs between hand-written AArch64 semantics and those derived mechanically from ARM's published formal semantics; and, we have used ARM-TV to discover 45 previously unknown miscompilation bugs in this LLVM backend, most of which are now fixed in upstream LLVM." The issue is attributed to RYAN BERGER, NVIDIA, USA; MITCH BRILES, University of Utah, USA; NADER BOUSHEHRINEJAD MORADI, University of Utah, USA; NICHOLAS COUGHLIN, Defence Science and Technology Group, Australia and University of Queensland, Australia; KAIT LAM, Defence Science and Technology Group, Australia and University of Queensland, Australia; NUNO P. LOPES, INESC-ID, Portugal and Instituto Superior Técnico - University of Lisbon, Portugal; STEFAN MADA, University of Utah, USA; TANMAY TIRPANKAR, University of Utah, USA; and JOHN REGEHR, University of Utah, USA.


**Translation Validation for LLVM's AArch64 Backend**

RYAN BERGER, NVIDIA, USA  
MITCH BRILES, University of Utah, USA  
NADER BOUSHEHRINEJAD MORADI, University of Utah, USA  
NICHOLAS COUGHLIN, Defence Science and Technology Group, Australia and University of Queensland, Australia  
KAIT LAM, Defence Science and Technology Group, Australia and University of Queensland, Australia  
NUNO P. LOPES, INESC-ID, Portugal and Instituto Superior Técnico - University of Lisbon, Portugal  
STEFAN MADA, University of Utah, USA  
TANMAY TIRPANKAR, University of Utah, USA  
JOHN REGEHR, University of Utah, USA

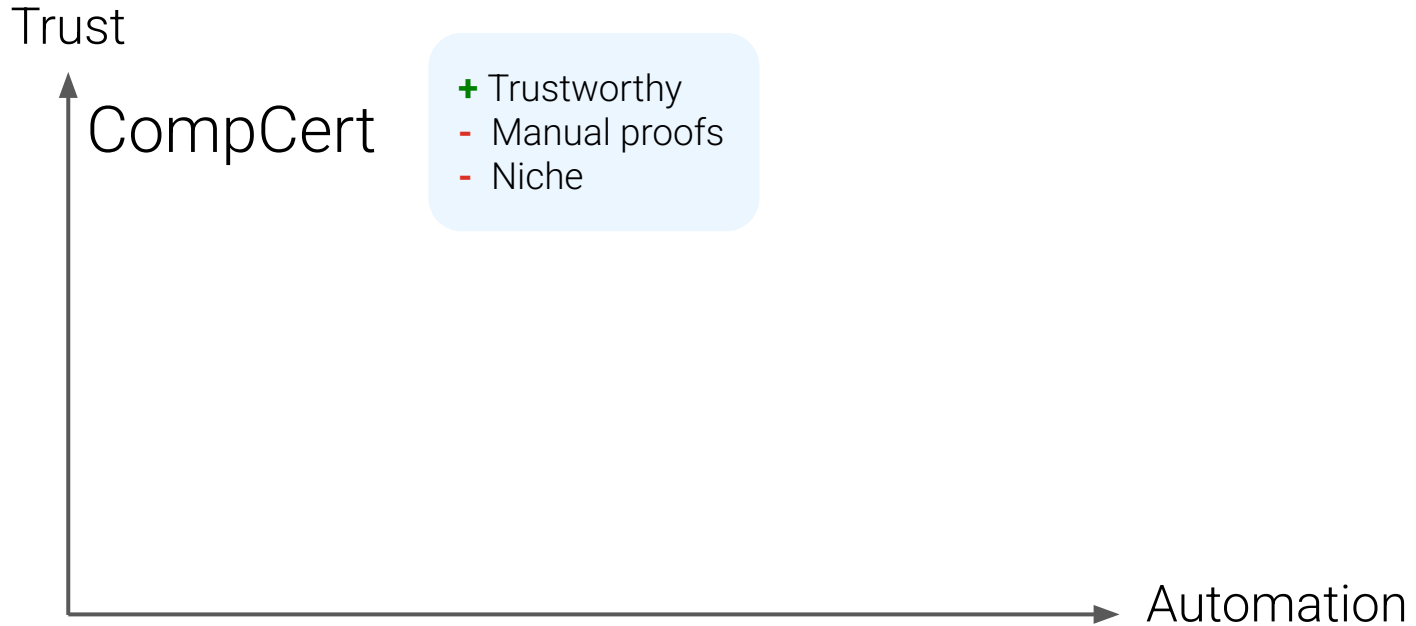
of assembly code; we investigate the tradeoffs between hand-written AArch64 semantics and those derived mechanically from ARM's published formal semantics; and, we have used ARM-TV to discover 45 previously unknown miscompilation bugs in this LLVM backend, most of which are now fixed in upstream LLVM.

Closed

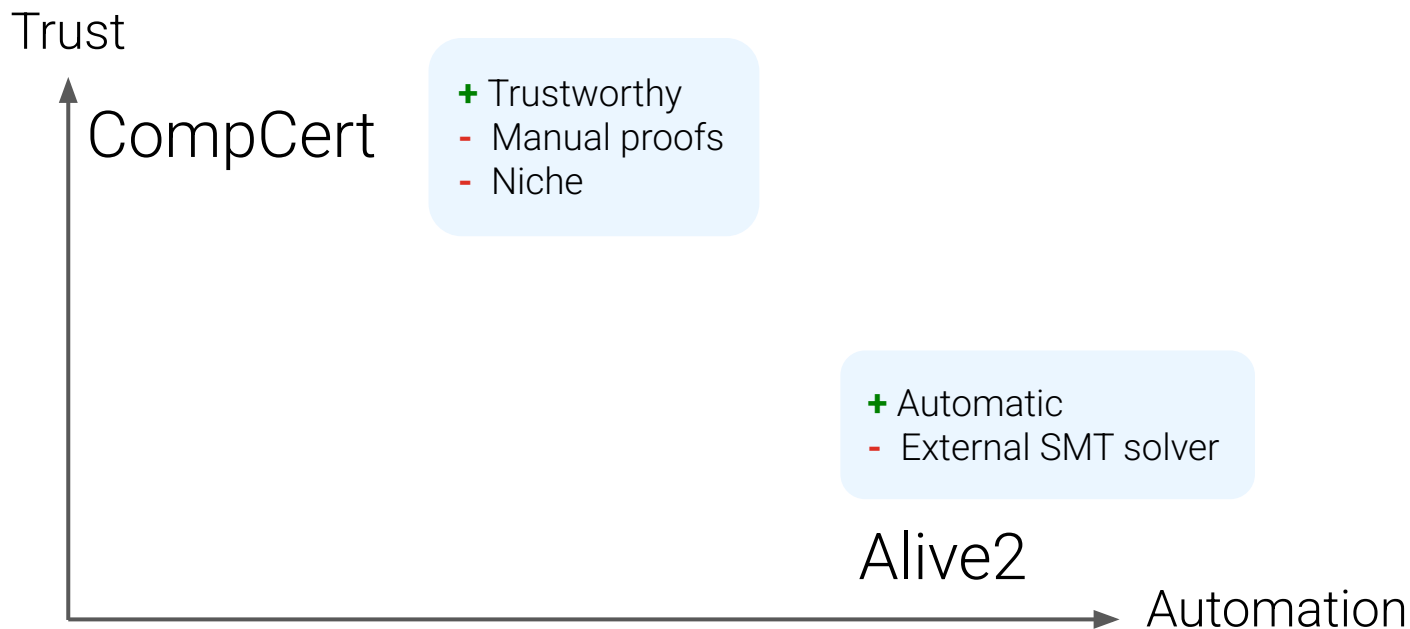
## A *Trade-off*: Trust vs. Automation



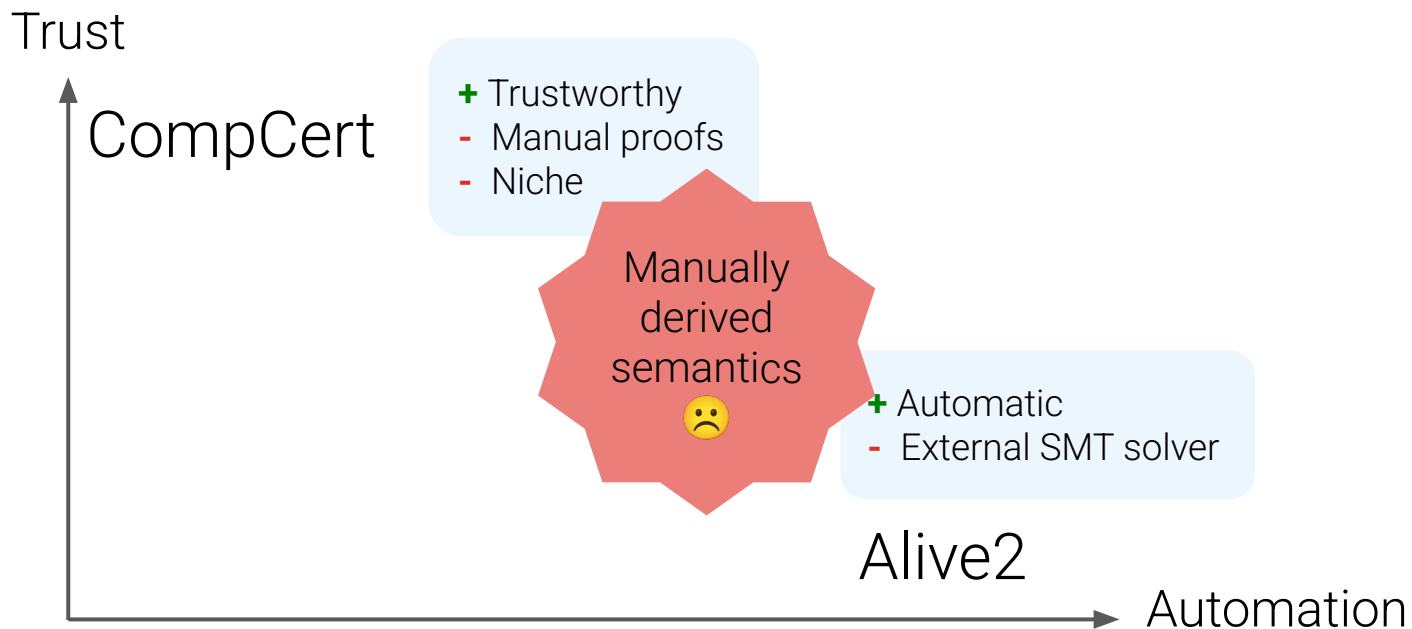
# *Fully-verified* compilation



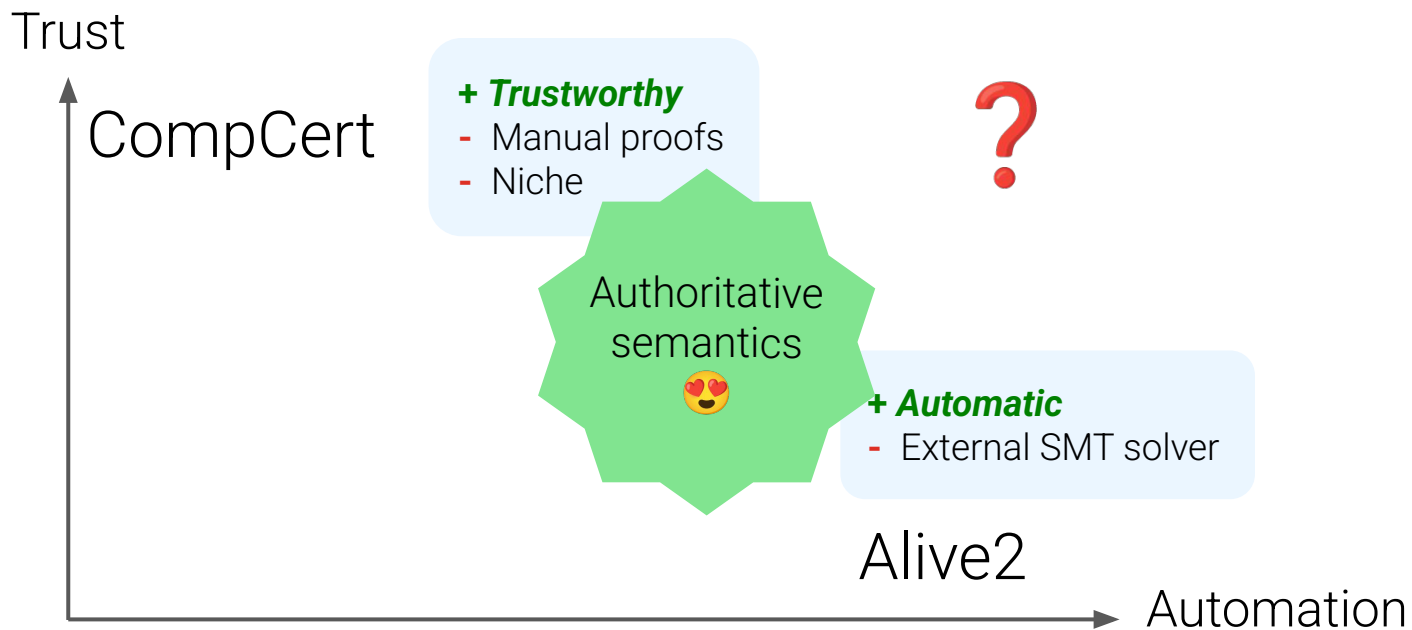
# *Fully-verified* compilation vs. *automated* translation-validation



# What about the *semantics*?



# Can we have it all?



We use Lean

# We use Lean

Functional  
Programming  
Language

Highest Level of  
Trust

Interactive  
Theorem  
Prover

# We use Lean

Functional  
Programming  
Language

Highest Level of  
Trust

Interactive  
Theorem  
Prover

```
def main : IO Unit := IO.println "Hello, world!"
```

# We use Lean

Functional  
Programming  
Language

Highest Level of  
Trust

Interactive  
Theorem  
Prover

```
def main : IO Unit := IO.println "Hello, world!"
```

```
def mySum : Nat := Id.run do
  let mut acc := 0
  for i in [0:4] do
    acc := acc + i
  pure acc
```

# We use Lean


Functional  
Programming  
Language

Highest Level of  
Trust

Interactive  
Theorem  
Prover

```
def main : IO Unit := IO.println "Hello, world!"
```

```
def mySum : Nat := Id.run do
  let mut acc := 0
  for i in [0:4] do
    acc := acc + i
  pure acc
```

```
theorem thm : mySum = 6 := by
  simp [mySum, List.foldl, List.range'] 
```


# Lean and its *bitvector* library and verified bitblaster

Latest updates: <https://dl.acm.org/doi/10.1145/3763167> Total Downloads: 333

RESEARCH-ARTICLE

**Interactive Bitvector Reasoning using Verified Bit-Blasting**

Published: 09 October 2025  
Accepted: 12 August 2025  
Received: 25 March 2025



`bv_decide` ✓


# Lean and its *bitvector* library and verified bitblaster

Latest updates: <https://dl.acm.org/doi/10.1145/3763167> Total Downloads: 333

RESEARCH-ARTICLE

**Interactive Bitvector Reasoning using Verified Bit-Blasting**

Published: 09 October 2025  
Accepted: 12 August 2025  
Received: 25 March 2025



`bv_decide` ✓

**example** (x : BitVec 32) : (y | x) + (y & x) = y + x

# Lean and its *bitvector* library and verified bitblaster

Latest updates: <https://dl.acm.org/doi/10.1145/3763167> Total Downloads: 333

RESEARCH-ARTICLE

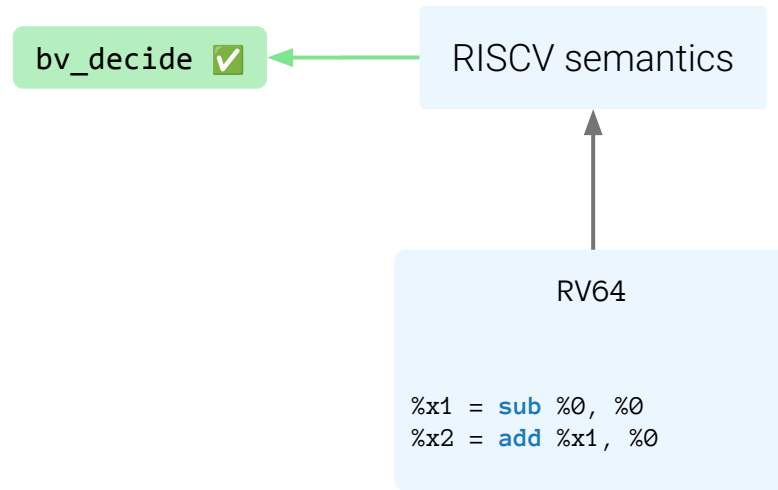
**Interactive Bitvector Reasoning using Verified Bit-Blasting**

Published: 09 October 2025  
Accepted: 12 August 2025  
Received: 25 March 2025

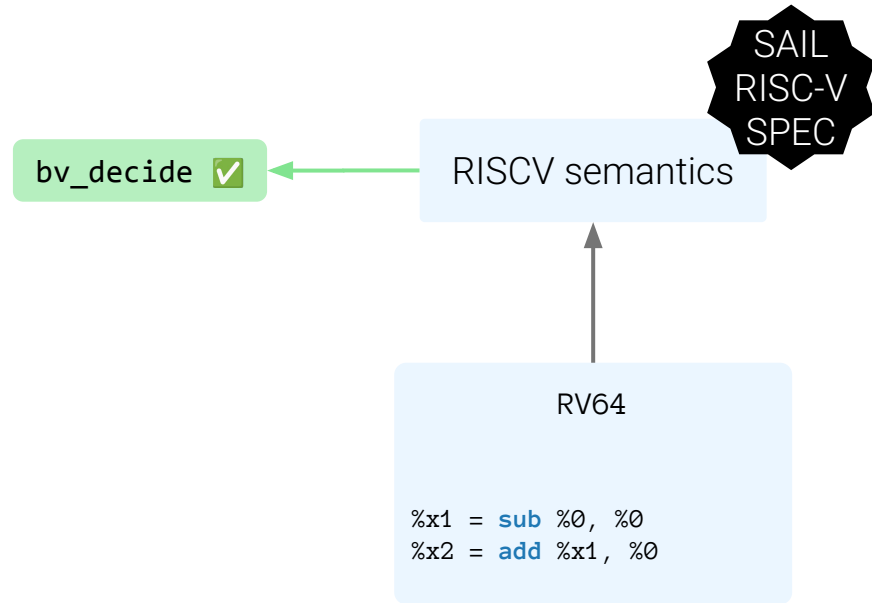
`bv_decide` ✓

**example** `(x : BitVec 32) : (y | x) + (y & x) = y + x` ✓

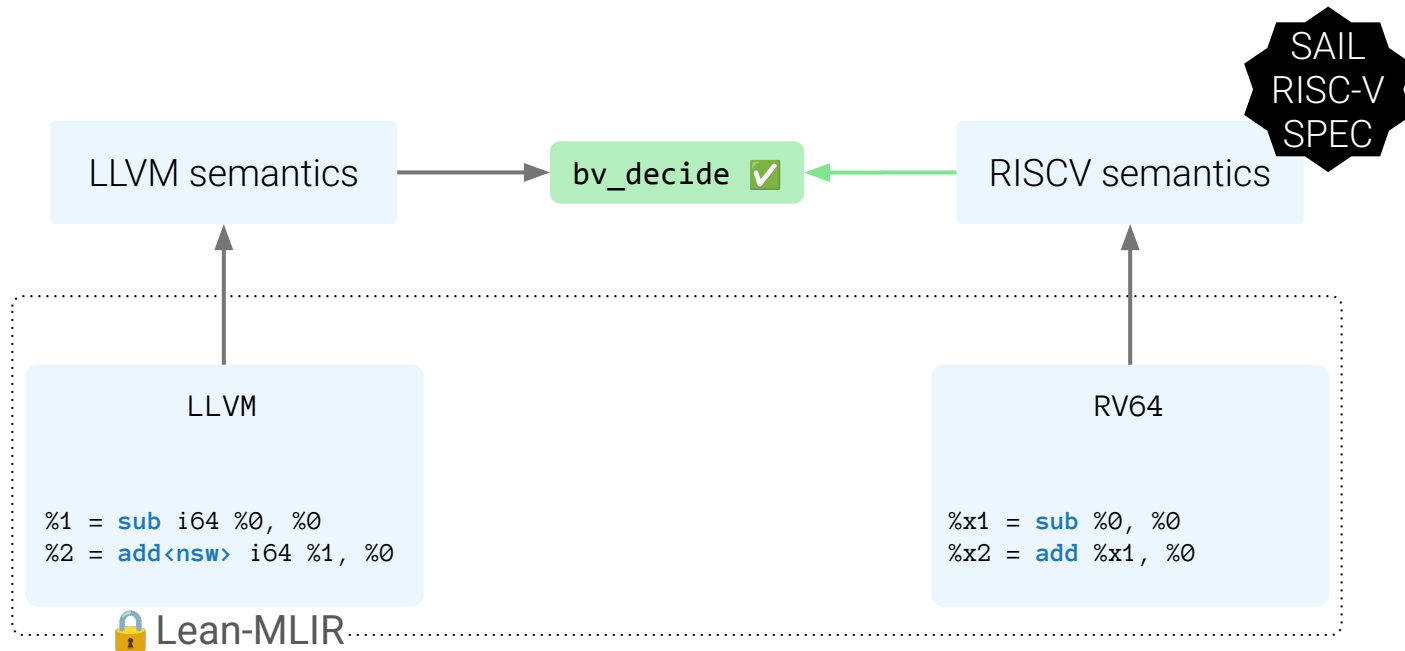
# A *bitblastable* mechanization of the RISC-V ISA



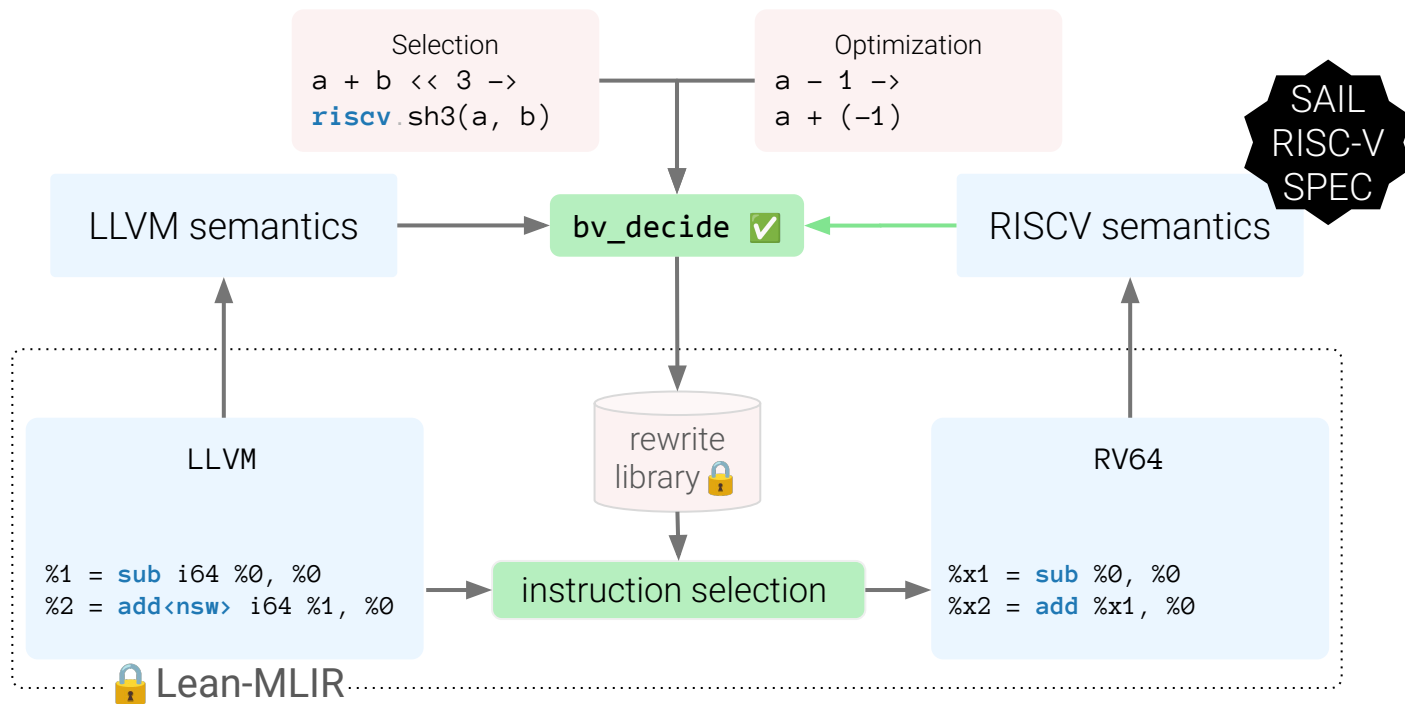
# Equivalent to the *Sail* processor model



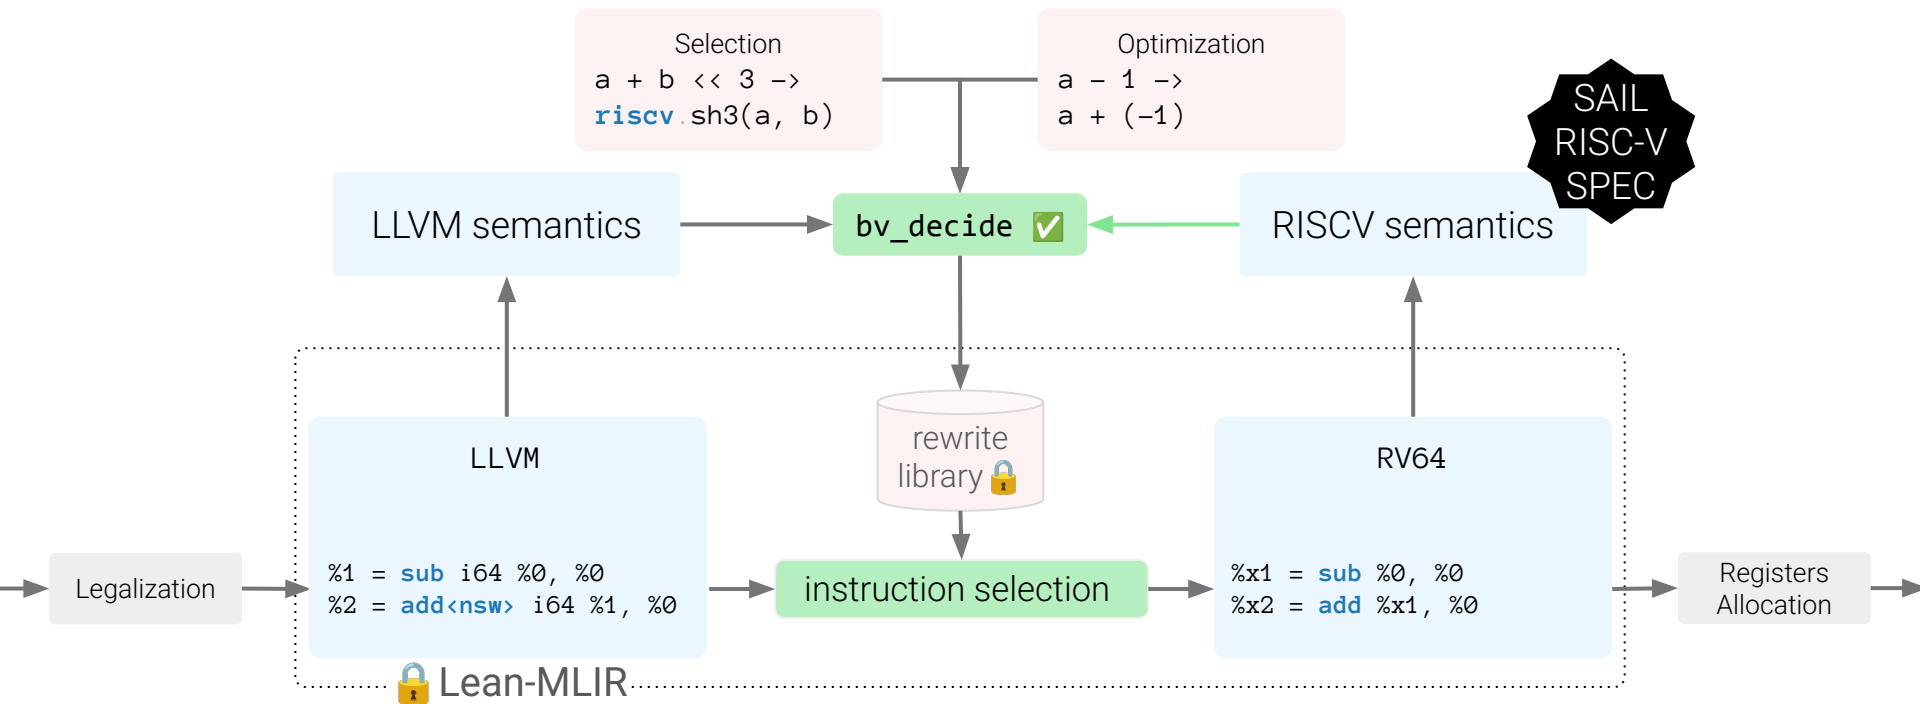
# In the *Lean-MLIR* framework



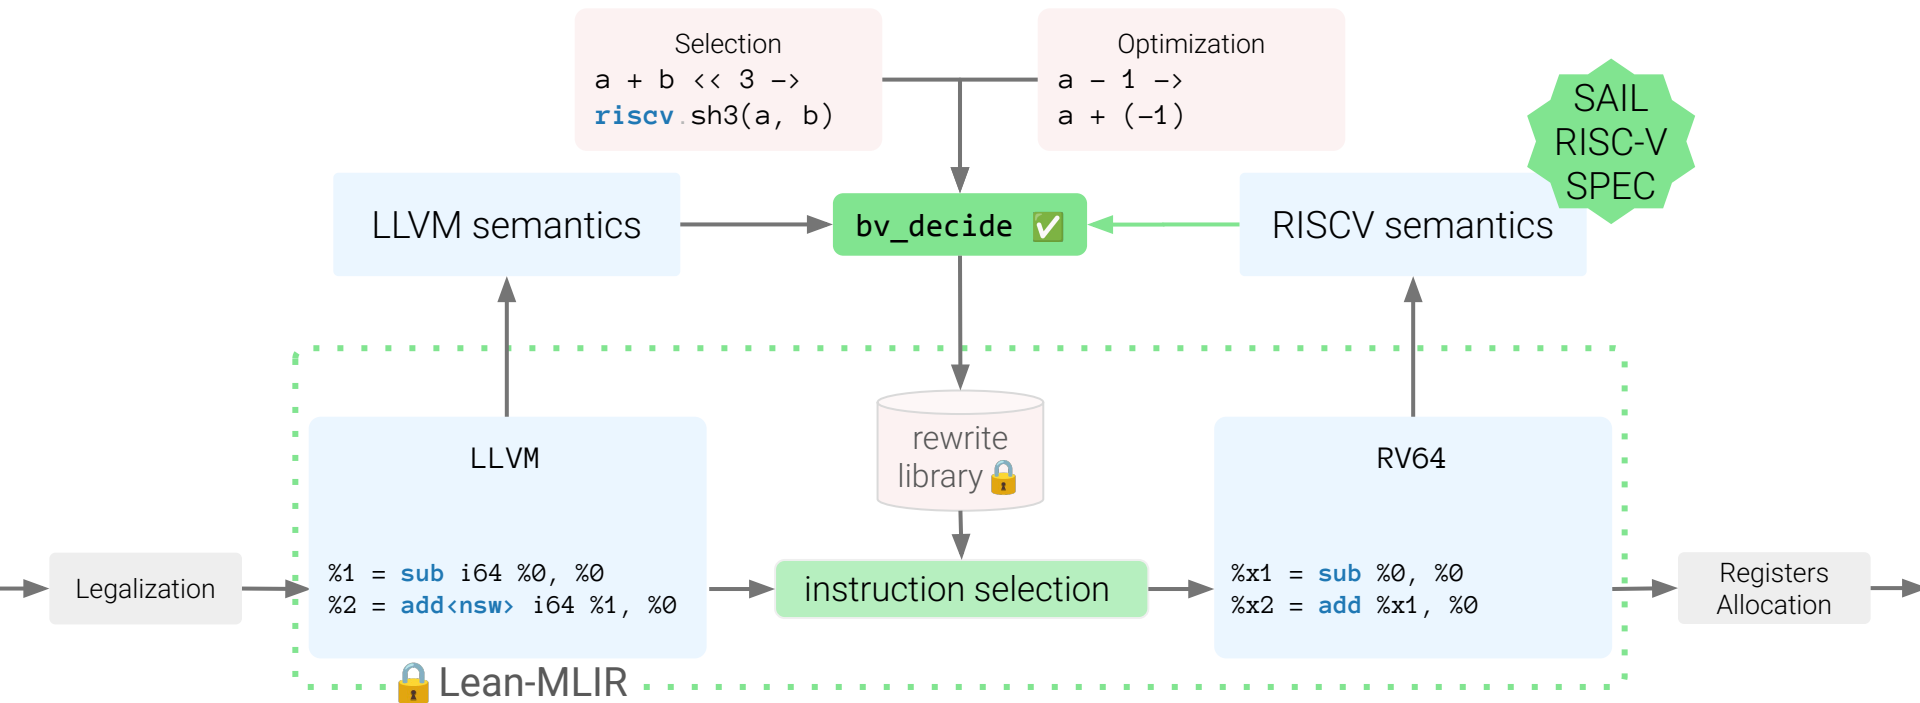
# Verifying LLVM's *instruction selection* and optimizations



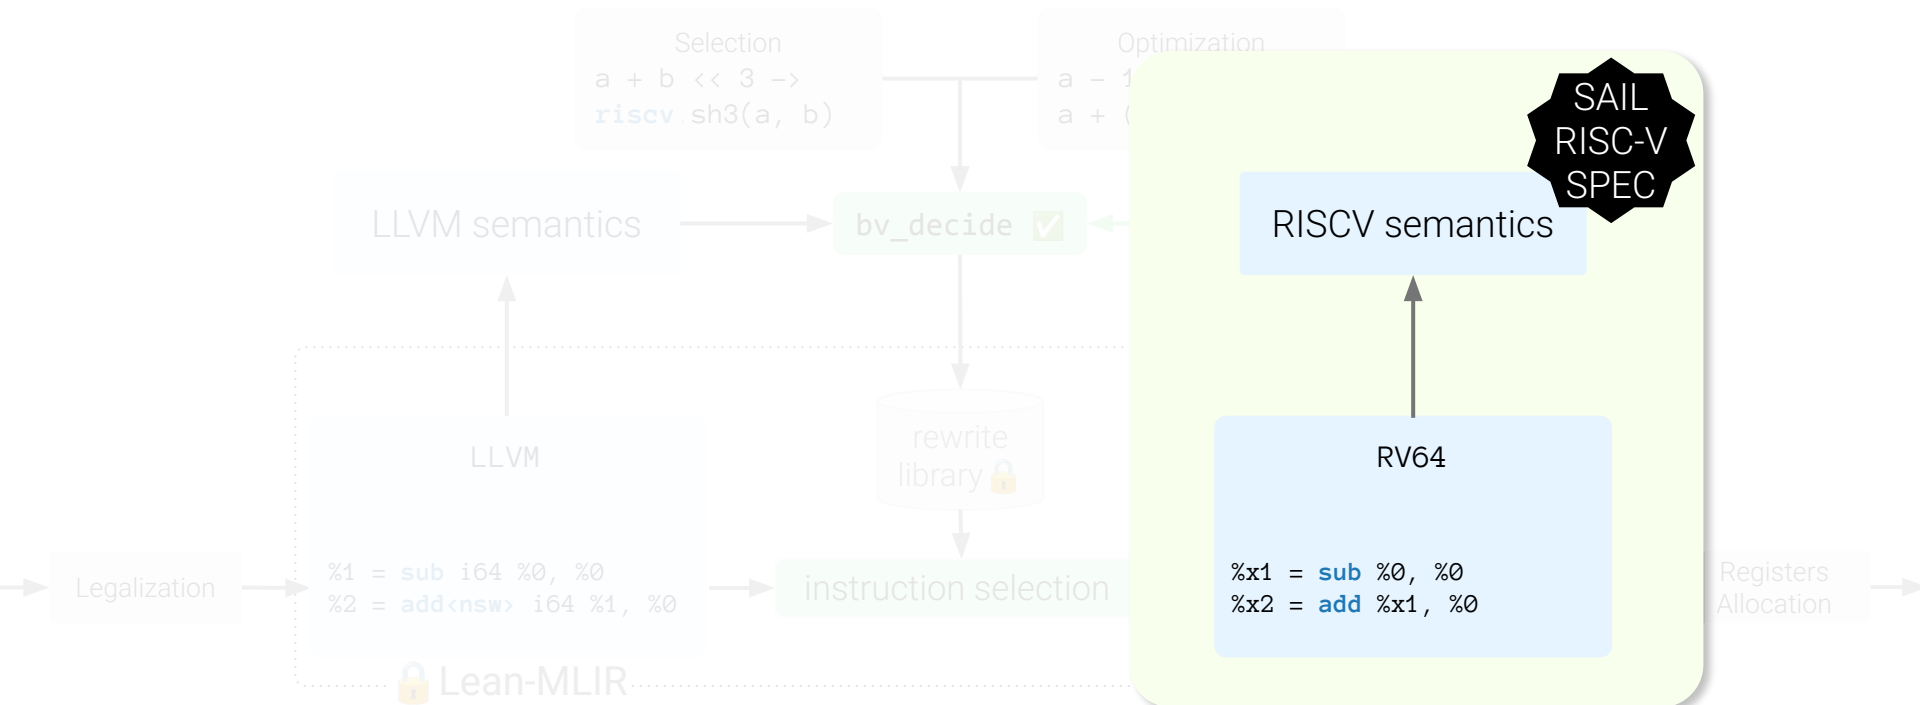
# Certified *assembly* output



# Obtaining *automation* and *trust*



# RISC-V semantics at the core of our work



# Bringing ISA dialects in MLIR

## 'x86' Dialect

### • Operations

- [x86.amx.tile\\_lo](#)
- [x86.amx.tile\\_mu](#)
- [x86.amx.tile\\_mu](#)
- [x86.amx.tile\\_st](#)
- [x86.amx.tile\\_ze](#)
- [x86.avx.best\\_to](#)

## 'xegpu' Dialect

The XeGPU dialect that models Intel GPU's ISA

The XeGPU dialect closely models a subset of the Xe GPU's ISA, providing an abstract performance GEMM code generation. It serves as a bridge dialect in the MLIR graduation working with MLIR memref and vector types, and complements the Arith, Math, Vector. XeGPU operations are introduced for special Xe instructions not modeled by the LLVM

## 'nvvm' Dialect

The NVVM dialect is MLIR's LLVM-IR-based, NVIDIA-specific backend dialect. It models NVVM intrinsics and public ISA functionality and introduces (e.g., global, shared, and cluster memory). While a NVVM op usually maps to a single Basic dialect to target Arm SME. other attributes so that a single NVVM

## 'ArmSME' Dialect

Basic dialect to target Arm SME.

## 'arm\_sve' Dialect

Basic dialect to target Arm SVE architectures

This dialect contains the definitions necessary to target specific Arm SVE scalable vector operations.

## ASTER : Assembly Tooling and Representations

MLIR C++ tool for programmable and highly-controllable assembly production on AMD GPUs.

Matrix  
[lg.matmul](#)  
-to-end

# Bringing ISA dialects in MLIR

## 'x86' Dialect

### • Operations

- [x86.amx.tile\\_lo](#)
- [x86.amx.tile\\_mu](#)
- [x86.amx.tile\\_mu](#)
- [x86.amx.tile\\_st](#)
- [x86.amx.tile\\_ze](#)
- [x86.amx.tile\\_ze](#)
- [x86.amx.tile\\_ze](#)

## 'xegpu' Dialect

The XeGPU dialect that models Intel GPU's ISA

The XeGPU dialect closely models a subset of the Xe GPU's ISA, providing an abstract performance GEMM code generation. It serves as a bridge dialect in the MLIR gradient working with MLIR memref and vector types, and complements the Arith, Math, Vector, and X86 dialects. XeGPU operations are introduced for special Xe instructions not modeled by the LLVM

## 'nvvm' Dialect

The NVVM dialect is MLIR's LLVM-IR-based, NVIDIA-specific backend dialect. It models NVVM intrinsics and public ISA functionality and introduces new operations (e.g., global, shared, and cluster memory). While a NVVM op usually maps to a single LLVM op, it also introduces other attributes so that a single NVVM op can be mapped to multiple LLVM ops.

## 'ArmSME' Dialect

Basic dialect to target Arm SME.

## 'arm\_sve' Dialect

Basic dialect to target Arm SVE architectures

## ASTER : Assembly Tool

MLIR C++ tool for programmable and highly-controlled hardware

## A Multi-level Compiler Backend for Accelerated Micro-kernels Targeting RISC-V ISA Extensions



# *Authoritative* ISA semantics for our RISC-V dialect



**RISC-V ISA**

# *Authoritative* ISA semantics for our RISC-V dialect

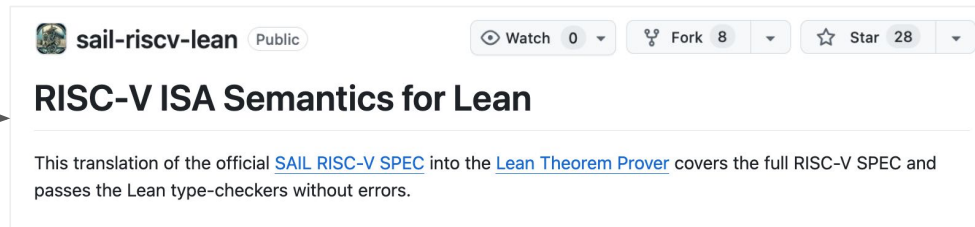



 **sail-riscv** Public

Watch 43 Fork 256 Star 686

## Formal Specification of the RISC-V ISA

This repository contains a formal specification of the RISC-V architecture, written in [Sail](#). It has been adopted by RISC-V International.



 **sail-riscv-lean** Public

Watch 0 Fork 8 Star 28

## RISC-V ISA Semantics for Lean

This translation of the official [SAIL RISC-V SPEC](#) into the [Lean Theorem Prover](#) covers the full RISC-V SPEC and passes the Lean type-checkers without errors.



# Formal Specification of the RISC-V ISA

This repository contains a formal specification of the RISC-V architecture, written in [Sail](#). It has been adopted by RISC-V International.

```
function clause execute REM (rs2, rs1, rd, is_unsigned) = {  
  
  let rs1_bits = X(rs1);  
  let rs2_bits = X(rs2);  
  
  let rs1_int = if is_unsigned then unsigned(rs1_bits) else signed(rs1_bits);  
  let rs2_int = if is_unsigned then unsigned(rs2_bits) else signed(rs2_bits);  
  let remainder = if rs2_int == 0 then rs1_int else rem_round_zero(rs1_int, rs2_int);  
  
  X(rd) = to_bits_truncate (remainder);  
  
  RETIRE_SUCCESS  
}
```



# Formal Specification of the RISC-V ISA

This repository contains a formal specification of the RISC-V architecture, written in [Sail](#). It has been adopted by RISC-V International.

```
function clause execute REM (rs2, rs1, rd, is_unsigned) = {  
  
  let rs1_bits =          X(rs1);  
  let rs2_bits =          X(rs2);  
  
  let rs1_int  = if is_unsigned      then unsigned(rs1_bits)      else signed(rs1_bits);  
  let rs2_int  = if is_unsigned      then unsigned(rs2_bits)      else signed(rs2_bits);  
  let remainder = if rs2_int == 0      then rs1_int else rem_round_zero(rs1_int, rs2_int);  
  
  X(rd) =      to_bits_truncate      (remainder);  
  
  RETIRE_SUCCESS  
}
```



# RISC-V ISA Semantics for Lean

This translation of the official [SAIL RISC-V SPEC](#) into the [Lean Theorem Prover](#) covers the full RISC-V SPEC and passes the Lean type-checkers without errors.

```
def execute_REM (rs2 : regidx) (rs1 : regidx) (rd : regidx) (is_unsigned : Bool) :=  
  
  let rs1_bits ← do (rX_bits rs1)  
  let rs2_bits ← do (rX_bits rs2)  
  
  let rs1_int := if (is_unsigned : Bool) then (BitVec.toNatInt rs1_bits) else (BitVec.toInt rs1_bits)  
  
  let rs2_int := if (is_unsigned : Bool) then (BitVec.toNatInt rs2_bits) else (BitVec.toInt rs2_bits)  
  
  let remainder := if ((rs2_int == 0) : Bool) then rs1_int else (Int.tmod rs1_int rs2_int)  
  
  (wX_bits rd (to_bits_truncate (1 := 64) remainder))  
  
  (pure RETIRE_SUCCESS)
```

# Authoritative ISA semantics for our RISC-V dialect

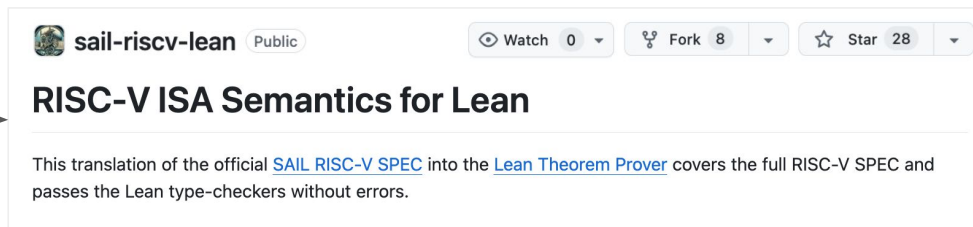


sail-riscv Public Watch 43 Fork 256 Star 686

## Formal Specification of the RISC-V ISA

This repository contains a formal specification of the RISC-V architecture, written in [Sail](#). It has been adopted by RISC-V International.

```
function clause execute REM
(rs2, rs1, rd, is_unsigned) = {
  let rs1_bits = X(rs1);
  let rs2_bits = X(rs2);
  let rs1_int = if is_unsigned then unsigned(rs1_bits)
                else signed(rs1_bits);
  let rs2_int = if is_unsigned then unsigned(rs2_bits)
                else signed(rs2_bits);
  let remainder = if rs2_int == 0 then rs1_int
                  else rem_round_zero(rs1_int, rs2_int);
  X(rd) = to_bits_truncate(remainder);
  RETIRE_SUCCESS
}
```



sail-riscv-lean Public Watch 0 Fork 8 Star 28

## RISC-V ISA Semantics for Lean

This translation of the official [SAIL RISC-V SPEC](#) into the [Lean Theorem Prover](#) covers the full RISC-V SPEC and passes the Lean type-checkers without errors.

```
def execute_REM (rs2 : regidx) (rs1 : regidx) (rd : regidx) (is_unsigned : Bool) :=
  let rs1_bits ← do (rX_bits rs1)
  let rs2_bits ← do (rX_bits rs2)
  let rs1_int := if (is_unsigned : Bool) then (BitVec.toNatInt rs1_bits)
                  else (BitVec.toInt rs1_bits)
  let rs2_int := if (is_unsigned : Bool) then (BitVec.toNatInt rs2_bits)
                  else (BitVec.toInt rs2_bits)
  let remainder := if ((rs2_int == 0) : Bool) then rs1_int
                    else (Int.tmod rs1_int rs2_int)
  (wX_bits rd (to_bits_truncate (1 := 64) remainder))
  (pure RETIRE_SUCCESS)
```

# Bitblastable, *authoritative* ISA semantics

SAIL model

```
def execute_REM (rs2 : regidx)
  (rs1 : regidx) (rd : regidx)
  (is_unsigned : Bool) :
  SailM ExecutionResult :=
let rs1_bits ← do (rX_bits rs1)
let rs2_bits ← do (rX_bits rs2)
let rs1_int := if (is_unsigned : Bool)
  then (BitVec.toNatInt rs1_bits) else (BitVec.toInt rs1_bits)
let rs2_int := if
  (is_unsigned : Bool)
  then (BitVec.toNatInt rs2_bits) else (BitVec.toInt rs2_bits)
let remainder := if ((rs2_int == 0) : Bool) then rs1_int else (Int.tmod rs1_int rs2_int)
(wX_bits rd (to_bits_truncate (1 := 64) remainder)) (pure RETIRE_SUCCESS)
```

# Bitblastable, *authoritative* ISA semantics

SAIL model

```
def execute_REM (rs2 : regidx)
  (rs1 : regidx) (rd : regidx)
  (is_unsigned : Bool) :
  SailM ExecutionResult :=
  let rs1_bits ← do (rX_bits rs1)
  let rs2_bits ← do (rX_bits rs2)
  let rs1_int := if (is_unsigned : Bool)
    then (BitVec.toNatInt rs1_bits) else (BitVec.toInt rs1_bits)
  let rs2_int := if
    (is_unsigned : Bool)
    then (BitVec.toNatInt rs2_bits) else (BitVec.toInt rs2_bits)
  let remainder := if ((rs2_int == 0) : Bool) then rs1_int else (Int.tmod rs1_int rs2_int)
  (wX_bits rd (to_bits_truncate (1 := 64) remainder)) (pure RETIRE_SUCCESS)
```

# Bitblastable, *authoritative* ISA semantics

SAIL model

```
def execute_REM (rs2 : regidx)
  (rs1 : regidx) (rd : regidx)
  (is_unsigned : Bool) :
  SailM ExecutionResult :=
  let rs1_bits ← do (rX_bits rs1)
  let rs2_bits ← do (rX_bits rs2)
  let rs1_int := if (is_unsigned : Bool)
    then (BitVec.toNatInt rs1_bits) else (BitVec.toInt rs1_bits)
  let rs2_int := if
    (is_unsigned : Bool)
    then (BitVec.toNatInt rs2_bits) else (BitVec.toInt rs2_bits)
  let remainder := if ((rs2_int == 0) : Bool) then rs1_int else (Int.tmod rs1_int rs2_int)
  (wX_bits rd (to_bits_truncate (1 := 64) remainder)) (pure RETIRE_SUCCESS)
```

Bitblastable  
semantics

```
def rem (rs2_val : BitVec 64)
  (rs1_val : BitVec 64) : BitVec 64 :=
  rs1_val.srem rs2_val
```

# Bitblastable, *authoritative* ISA semantics

SAIL model

```
def execute_REM (rs2 : regidx)
  (rs1 : regidx) (rd : regidx)
  (is_unsigned : Bool) :
  SailM ExecutionResult :=
  let rs1_bits ← do (rX_bits rs1)
  let rs2_bits ← do (rX_bits rs2)
  let rs1_int := if (is_unsigned : Bool)
    then (BitVec.toNatInt rs1_bits) else (BitVec.toInt rs1_bits)
  let rs2_int := if
    (is_unsigned : Bool)
    then (BitVec.toNatInt rs2_bits) else (BitVec.toInt rs2_bits)
  let remainder := if ((rs2_int == 0) : Bool) then rs1_int else (Int.tmod rs1_int rs2_int)
  (wX_bits rd (to_bits_truncate (1 := 64) remainder)) (pure RETIRE_SUCCESS)
```

Bitblastable  
semantics

```
def rem (rs2_val : BitVec 64)
  (rs1_val : BitVec 64) : BitVec 64 :=
  rs1_val.srem rs2_val
```

↓  
bv\_decide ✓

# Bitblastable, *authoritative* ISA semantics

SAIL model



Bitblastable semantics

```
def execute_REM (rs2 : regidx)
  (rs1 : regidx) (rd : regidx)
  (is_unsigned : Bool) :
  SailM ExecutionResult :=
  let rs1_bits ← do (rX_bits rs1)
  let rs2_bits ← do (rX_bits rs2)
  let rs1_int := if (is_unsigned : Bool)
    then (BitVec.toNatInt rs1_bits) else (BitVec.toInt rs1_bits)
  let rs2_int := if
    (is_unsigned : Bool)
    then (BitVec.toNatInt rs2_bits) else (BitVec.toInt rs2_bits)
  let remainder := if ((rs2_int == 0) : Bool) then rs1_int else (Int.tmod rs1_int rs2_int)
  (wX_bits rd (to_bits_truncate (1 := 64) remainder)) (pure RETIRE_SUCCESS)
```

```
def rem (rs2_val : BitVec 64)
  (rs1_val : BitVec 64) : BitVec 64 :=
  rs1_val.srem rs2_val
```

↓  
bv\_decide ✓

# ISA: lowering LLVM IR's arithmetic fragment

lui	srai	sraw	srl	srai
auipc	slliw	slli	sra	sub
addi	srliw	slti	addw	and
andi	sraiw	sltiu	subw	or
ori	slt	slli	sllw	xor
xori	sltu	srli	srlw	sll
addiw	add	fence	load	store

RISC-V **Base**

mul	mulw	mulh	mulhu	mulhsu
divuw	div	divu	rem	remu
remuw	divw	remw		

RISC-V **M**

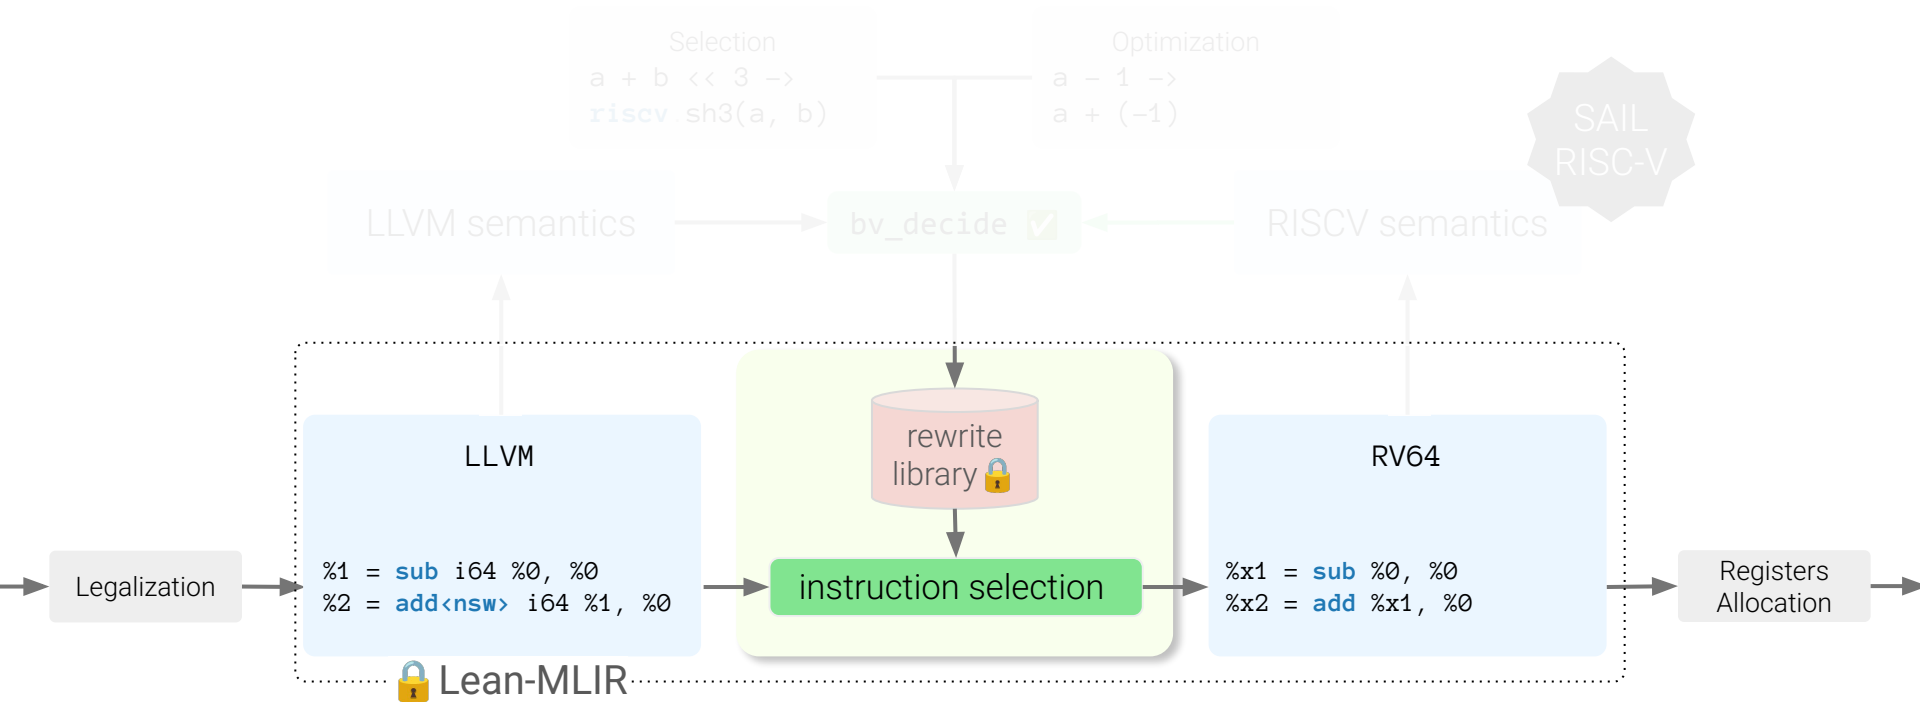
addw	sh1add	sh2add	sh3add	sh1adduw
sh3adduw	slliw	rol	ror	rolw
sextb	sexth	zexth	rori	roriw
orn	xnor	max	maxu	min
bclr	bclri	bext	bexti	binvi
bseti	czeroeqz	czeronez	orcb	rev8
clzw	ctz	ctzw	cpop	cpopw
sh2adduw	rorw	andn	minu	bset
clz	clmul	clmulh	clmulr	brev8
pack	packw	zip	unzip	andn
xperm4	xperm8			

RISC-V **B**

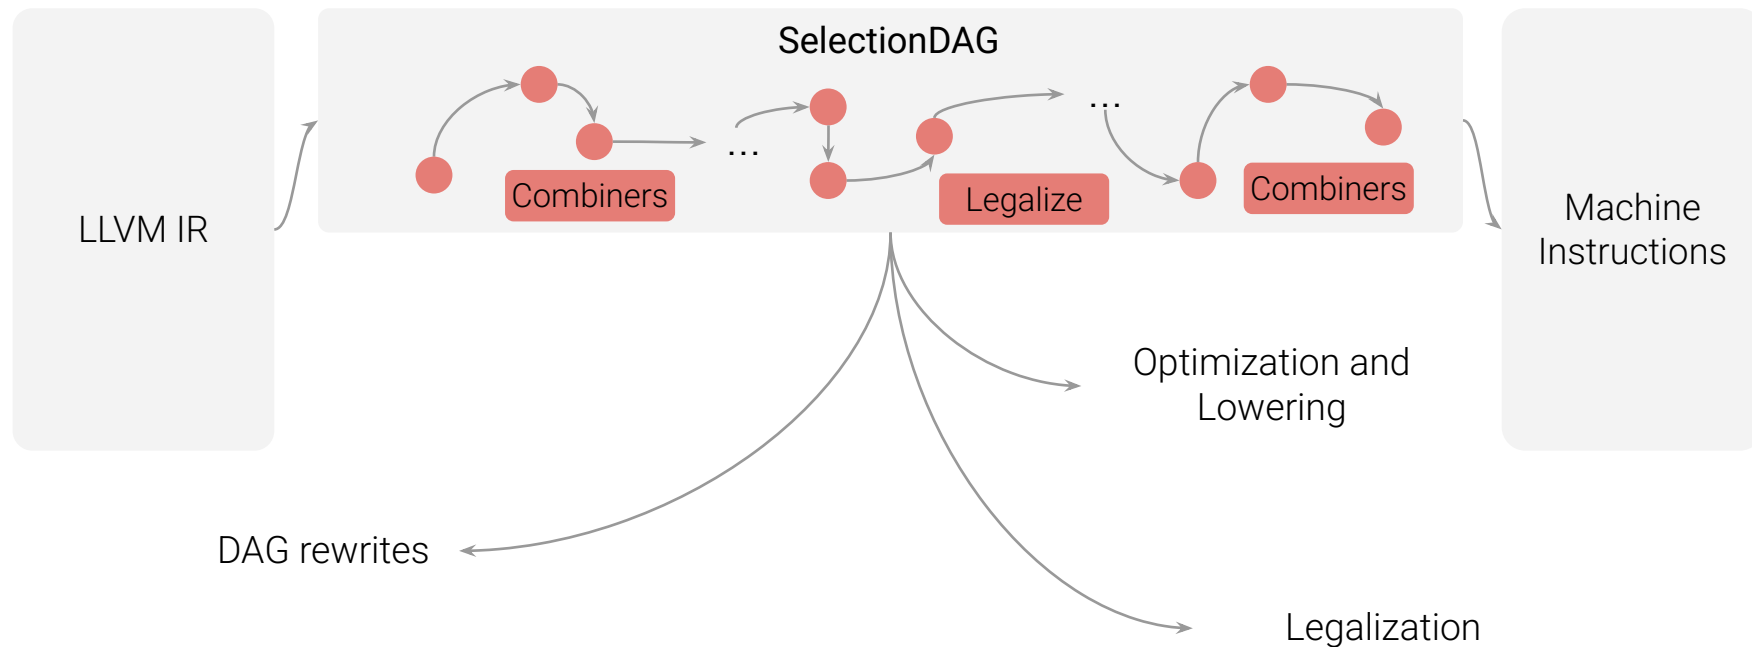
li	mv	not	neg	negw
zextb	seqz	snez	sltz	sgtz
sextw				

RISC-V **pseudo**

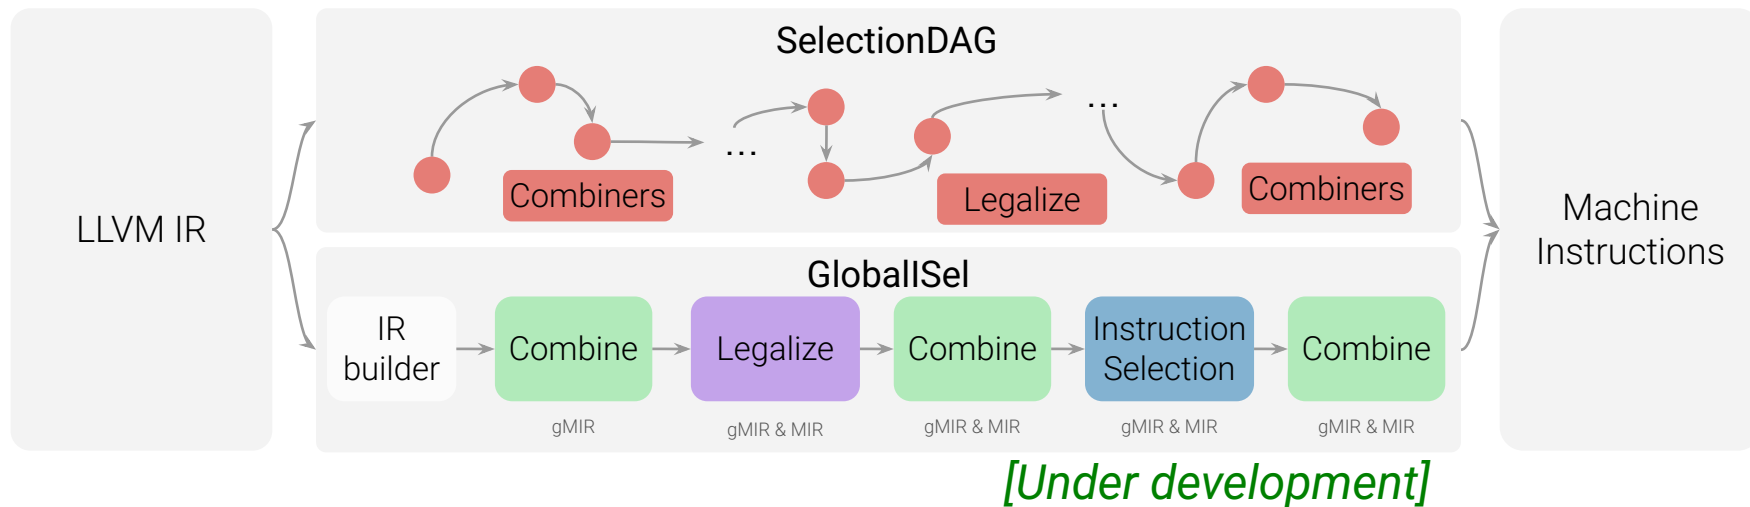
# A progressive *lowering* that mirrors LLVM



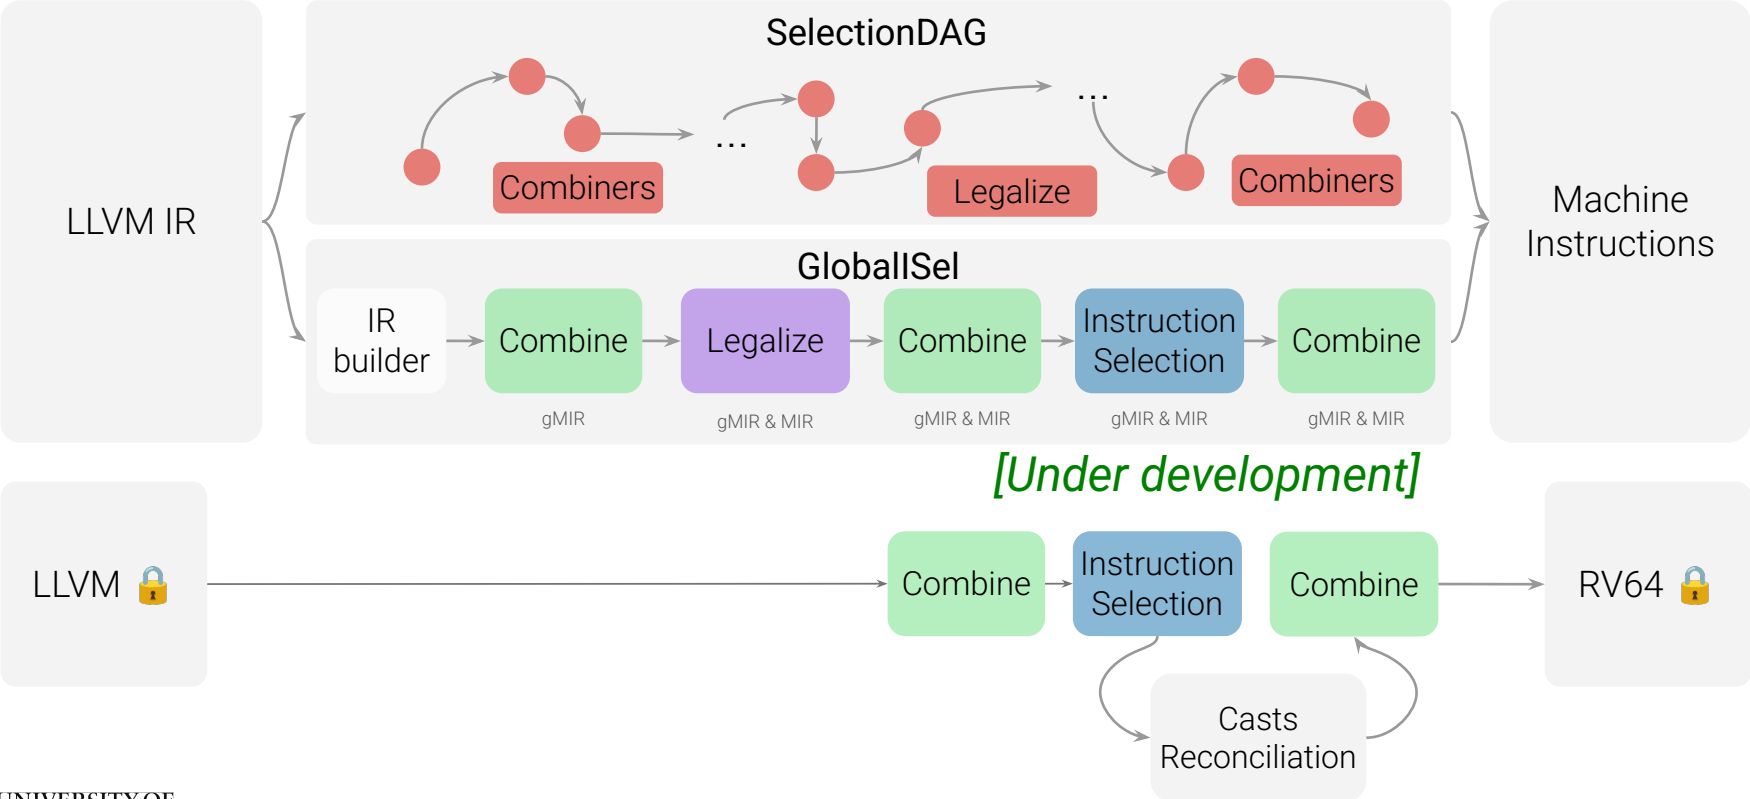
# LLVM's *instruction selectors*



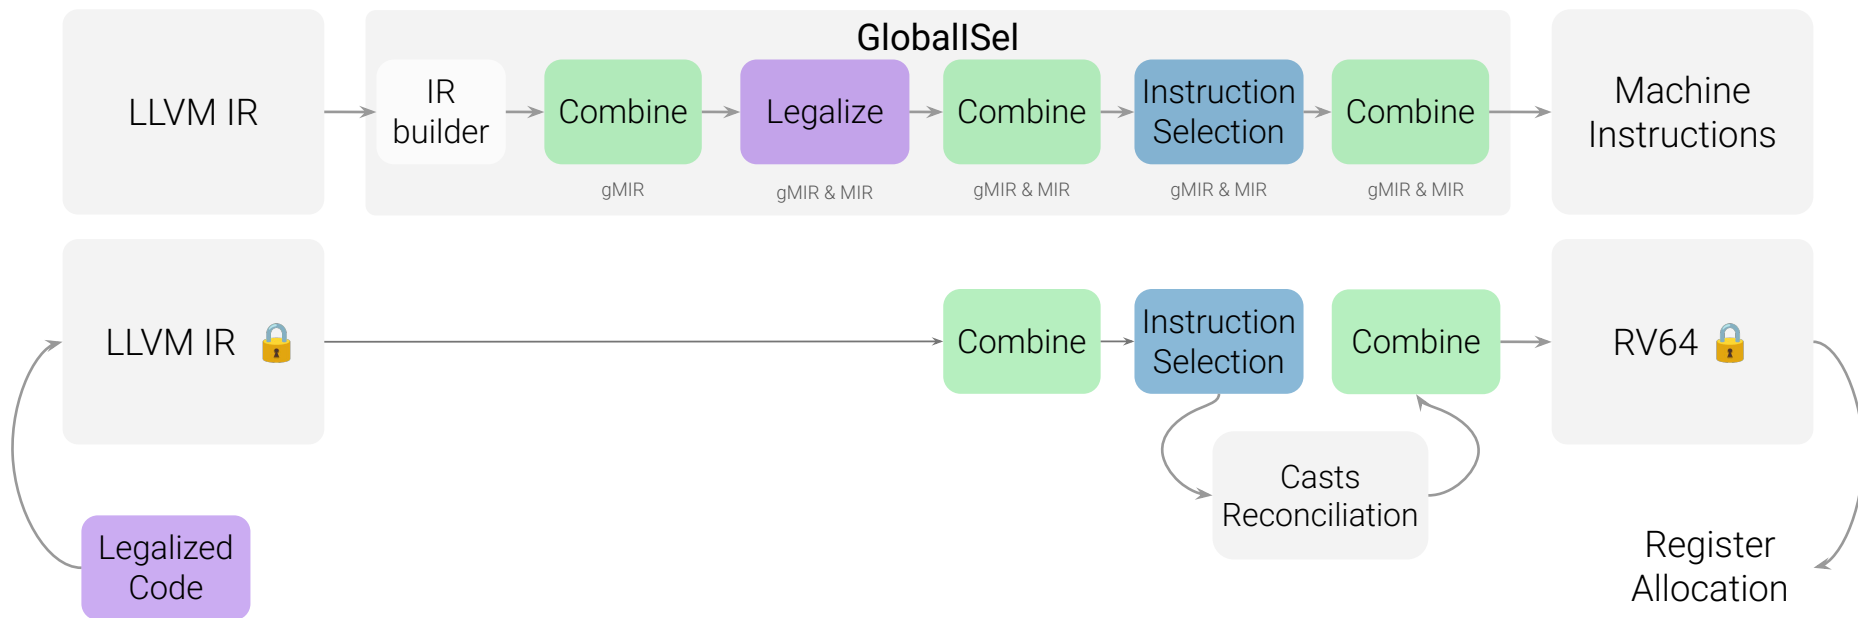
# LLVM's *instruction selectors*



# LLVM's *instruction selectors* and our Lean implementation



# GlobalSel vs. our Lean implementation



# LLVM has poison semantics

LLVM 

```
def and_llvm := [LV] {  
  ^entry (%lhs: i1, %rhs: i1 ):  
  %1 = llvm.and %lhs, %rhs : i1  
  llvm.return %1 : i1  
}]
```


Poison or  
concrete



# RISC-V semantics are concrete

LLVM 

```
def and_llvm := [LV] {  
  ^entry (%lhs: i1, %rhs: i1 ):  
  %1 = llvm.and %lhs, %rhs : i1  
  llvm.return %1 : i1  
}]
```

RV64 

```
def and_riscv := [LV] {  
  ^entry (%lhs: i1, %rhs: i1):  
  %lhsr = cast (%lhs) : (i1) -> (reg)  
  %rhrs = cast (%rhs) : (i1) -> (reg)  
  %0 = RV64.and %lhsr, %rhrs : reg  
  %1 = cast (%0) : reg -> (i1)  
  llvm.return %1 : i1  
}]
```




# We need to prove *refinement*



LLVM 

```
def and_llvm := [LV] {  
  ^entry (%lhs: i1, %rhs: i1):  
    %1 = llvm.and %lhs, %rhs : i1  
    llvm.return %1 : i1  
}]
```

RV64 

```
def and_riscv := [LV] {  
  ^entry (%lhs: i1, %rhs: i1):  
    %lhsr = cast (%lhs) : (i1) -> (reg)  
    %rhrs = cast (%rhs) : (i1) -> (reg)  
    %0 = RV64.and %lhsr, %rhrs : reg  
    %1 = cast (%0) : (reg) -> (i1)  
    llvm.return %1 : i1  
}]
```

```
def llvm_and_lower_riscv : LLVMRewriteRefine 1 [Ty.llvm (.bv 1), Ty.llvm (.bv 1)] where  
  lhs:= and_llvm  
  rhs:= and_riscv  
  correct := by simp_lowering <;>; bv_decide
```

# Refinement verification



$\forall (x\ y : \llbracket \text{Ty.llvm} (\text{MTy.bv } 1) \rrbracket),$   
 $\text{LLVM.and } x\ y \sqsubseteq \text{castriscvToLLVM } (\text{RV64.and } (\text{castLLVMToriscv } y) (\text{castLLVMToriscv } x))$



$\forall (x\ y : \llbracket \text{Ty.llvm} (\text{MTy.bv } 1) \rrbracket), (\text{do let } x' \leftarrow x; \text{ let } y' \leftarrow y$   
 $\text{PoisonOr.value } (x' \ \&\& \ y')) \sqsubseteq$   
 $\text{PoisonOr.value } (\text{BitVec.signExtend } 1$   
 $\quad (\text{BitVec.signExtend } 64 \ (x.\text{toOption.getD } 0\#1) \ \&\&$   
 $\quad \text{BitVec.signExtend } 64 \ (y.\text{toOption.getD } 0\#1)))$



$\forall (x\ y : \text{BitVec } 1),$   
 $x \ \&\& \ y = \text{BitVec.signExtend } 1 \ (\text{BitVec.signExtend } 64 \ x \ \&\& \ \text{BitVec.signExtend } 64 \ y)$

→ **bv\_decide**

# 158 lowering patterns



```
^bb0(%x : i64, %y : i64):  
  %3 = llvm.add %x, %y  
    <{overflow<none>}> : i64  
  llvm.return(%3) : i64
```



```
^bb0(%x : i64, %y : i64):  
  %0 = cast(%reg1) : i64 -> reg  
  %1 = cast(%reg2) : i64 -> reg  
  %2 = RV64.add %0, %1 : reg  
  %3 = cast(%2) : reg -> i64  
  llvm.return(%3) : i64
```

# With *pseudoinstructions*



```
^bb0(%x : i64, %y : i64):  
  %3 = llvm.add %x, %y  
    <{overflow<none>}> : i64  
  llvm.return(%3) : i64
```

```
^bb0(%x : i64, %y : i64):  
  %0 = cast(%reg1) : i64 -> reg  
  %1 = cast(%reg2) : i64 -> reg  
  %2 = RV64.add %0, %1 : reg  
  %3 = cast(%2) : reg -> i64  
  llvm.return(%3) : i64
```

```
^bb0(%x : i64, %y : i64):  
  %1 = llvm.icmp.eq %x, %y : i1  
  llvm.return(%1) : i1
```

```
xor a0, a0, a1  
seqz a0, a0  
ret
```

# With *pseudoinstructions*



```
^bb0(%x : i64, %y : i64):  
  %3 = llvm.add %x, %y  
    <{overflow<none>}> : i64  
  llvm.return(%3) : i64
```

```
^bb0(%x : i64, %y : i64):  
  %0 = cast(%reg1) : i64 -> reg  
  %1 = cast(%reg2) : i64 -> reg  
  %2 = RV64.add %0, %1 : reg  
  %3 = cast(%2) : reg -> i64  
  llvm.return(%3) : i64
```

```
^bb0(%x : i64, %y : i64):  
  %1 = llvm.icmp.eq %x, %y : i1  
  llvm.return(%1) : i1
```

```
xor a0, a0, a1  
seqz a0, a0  
ret
```

```
def seqz_pseudo (rs1_val : BitVec 64) : BitVec 64 :=  
  RV64.sltiu 1 rs1_val
```

# With *pseudoinstructions*



```
^bb0(%x : i64, %y : i64):  
  %3 = llvm.add %x, %y  
    <{overflow<none>}> : i64  
  llvm.return(%3) : i64
```

```
^bb0(%x : i64, %y : i64):  
  %0 = cast(%reg1) : i64 -> reg  
  %1 = cast(%reg2) : i64 -> reg  
  %2 = RV64.add %0, %1 : reg  
  %3 = cast(%2) : reg -> i64  
  llvm.return(%3) : i64
```

```
^bb0(%x : i64, %y : i64):  
  %1 = llvm.icmp.eq %x, %y : i1  
  llvm.return(%1) : i1
```


```
^bb0(%x : i64, %y : i64):  
  %0 = cast(%x) : i64 -> reg  
  %1 = cast(%y) : i64 -> reg  
  %2 = RV64.xor %0, %1 : reg  
  %3 = RV64.sltiu %0, 1 : reg  
  %4 = cast(%3) : reg -> i64  
  llvm.return(%4) : i64
```

# Verifying *optimizations*



```
// Fold (x * 0) -> 0
def binop_right_to_zero: GICombineRule<
  (defs root:$dst),
  (match (G_MUL $dst, $lhs, 0:$zero)),
  (apply (GIReplaceReg $dst, $zero))
>;
```

# Verifying *optimizations*



```
// Fold (x * 0) -> 0
def binop_right_to_zero: GICombineRule<
  (defs root:$dst),
  (match (G_MUL $dst, $lhs, 0:$zero)),
  (apply (GIReplaceReg $dst, $zero))
>;
```

```
def binop_right_to_zero :
  LLVMRewriteRefine 64 [Ty.llvm (.bv 64)] where
  lhs := [LV| {
    ^entry (%0: i64):
      %1 = llvm.mlir.constant (0) : i64
      %2 = llvm.mul %0, %1 : i64
      llvm.return %2 : i64
  }]
  rhs := [LV| {
    ^entry (%0: i64):
      %1 = llvm.mlir.constant (0) : i64
      llvm.return %1 : i64
  }]
```


# Verifying *optimizations*



```
// Fold (x * 0) -> 0
def binop_right_to_zero: GICombineRule<
  (defs root:$dst),
  (match (G_MUL $dst, $lhs, 0:$zero)),
  (apply (GIReplaceReg $dst, $zero))
>;
```

```
def binop_right_to_zero :
  LLVMRewriteRefine 64 [Ty.llvm (.bv 64)] where
  lhs := [LV| {
    ^entry (%0: i64):
      %1 = llvm.mlir.constant (0) : i64
      %2 = llvm.mul %0, %1 : i64
      llvm.return %2 : i64
  }]
  rhs := [LV| {
    ^entry (%0: i64):
      %1 = llvm.mlir.constant (0) : i64
      llvm.return %1 : i64
  }]
```

# Verifying *optimizations*



```
// Fold (x * 0) -> 0
def binop_right_to_zero: GICombineRule<
  (defs root:$dst),
  (match (G_MUL $dst, $lhs, 0:$zero)),
  (apply (GIReplaceReg $dst, $zero))
>;
```


```
def binop_right_to_zero :
  LLVMRewriteRefine 64 [Ty.llvm (.bv 64)] where
  lhs := [LV| {
    ^entry (%0: i64):
      %1 = llvm.mlir.constant (0) : i64
      %2 = llvm.mul %0, %1 : i64
      llvm.return %2 : i64
  }]
  rhs := [LV| {
    ^entry (%0: i64):
      %1 = llvm.mlir.constant (0) : i64
      llvm.return %1 : i64
  }]
  correct := by simp_lowering <;>; bv_decide
```

# Extracting *optimizations*



```
// Fold (x * 0) -> 0
def binop_right_to_zero: GICombineRule<
  (defs root:$dst),
  (match (G_MUL $dst, $lhs, 0:$zero)),
  (apply (GIReplaceReg $dst, $zero))
>;
```

```
def binop_right_to_zero :
  LLVMRewriteRefine 64 [Ty.llvm (.bv 64)] where
  lhs := [LV| {
    ^entry (%0: i64):
      %1 = llvm.mlir.constant (0) : i64
      %2 = llvm.mul %0, %1 : i64
      llvm.return %2 : i64
  }]
  rhs := [LV| {
    ^entry (%0: i64):
      %1 = llvm.mlir.constant (0) : i64
      llvm.return %1 : i64
  }]
  correct := by simp_lowering <;>; bv_decide
```

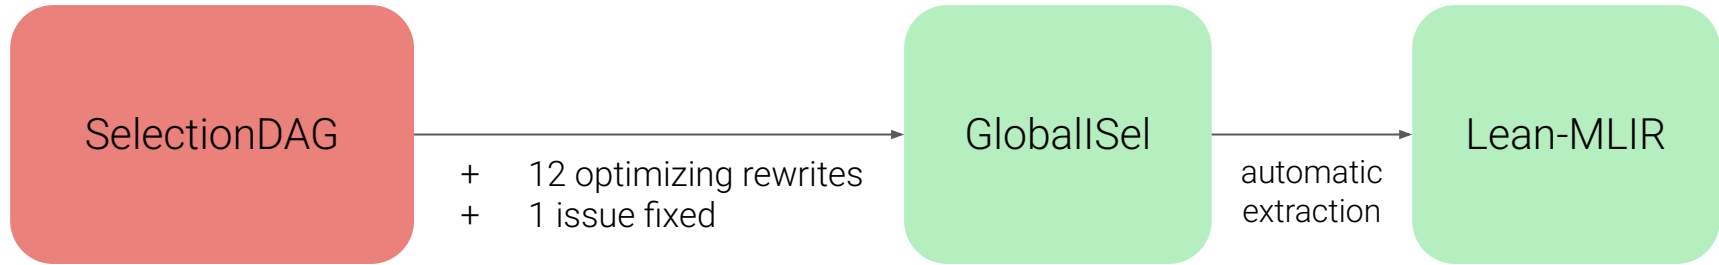


# Porting LLVM *optimizations*



Group	OO, PreLeg	PreLeg.	PostLeg.	#Rules	#Rules In Frag.	#Mechanized
trivial	✓	✓		6	4	4
optnone	✓			12	5	5
identity		✓	✓	20	8	8
cmp		✓		7	4	4
integer_reassoc		✓		13	13	13
cast		✓		11	9	9
cast_of_cast		✓		11	10	10
others		✓	✓	148	27	13
<b>Total</b>				<b>228</b>	<b>80</b>	<b>66</b>
<b>Mechanized [%]</b>						<b>82.5%</b>

# Improving *GlobalSel* improves verified instruction selection



A vertical stack of four GitHub-style approval comments. Each comment consists of a profile picture icon, a green checkmark, the user's name, and the text of the approval. The comments are as follows:

- arsenm approved these changes [on Jan 29](#)
- qcolombet approved these changes [on Feb 20](#)
- jayfoad approved these changes [last week](#)
- Pierre-vh linked a pull request that will close this issue [2 hours ago](#)

# Improving *GlobalSel* improves verified instruction selection

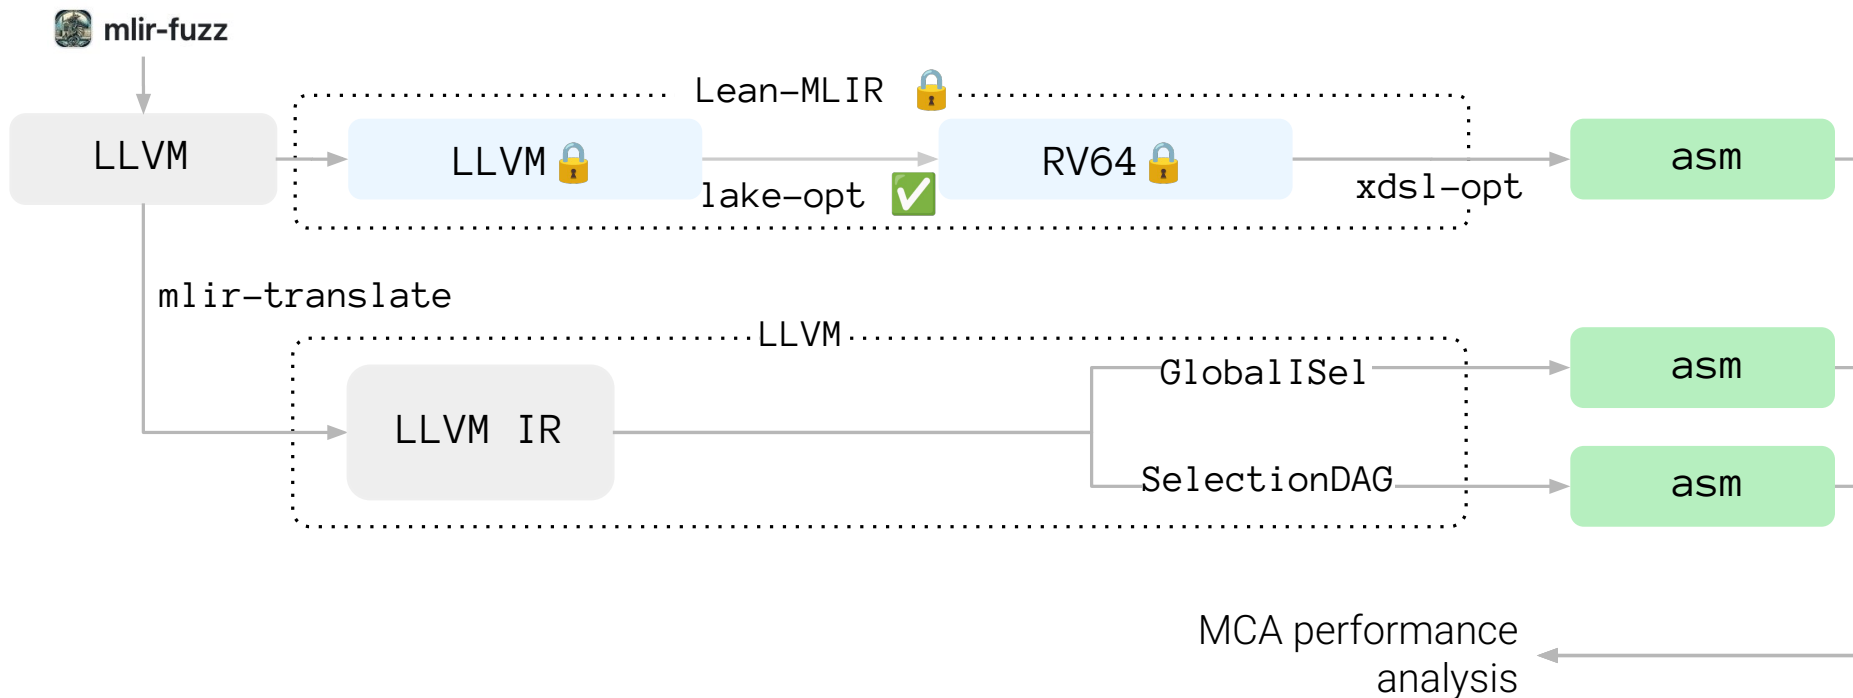
Select

**WE ARE DONE VERIFYING LLVM**

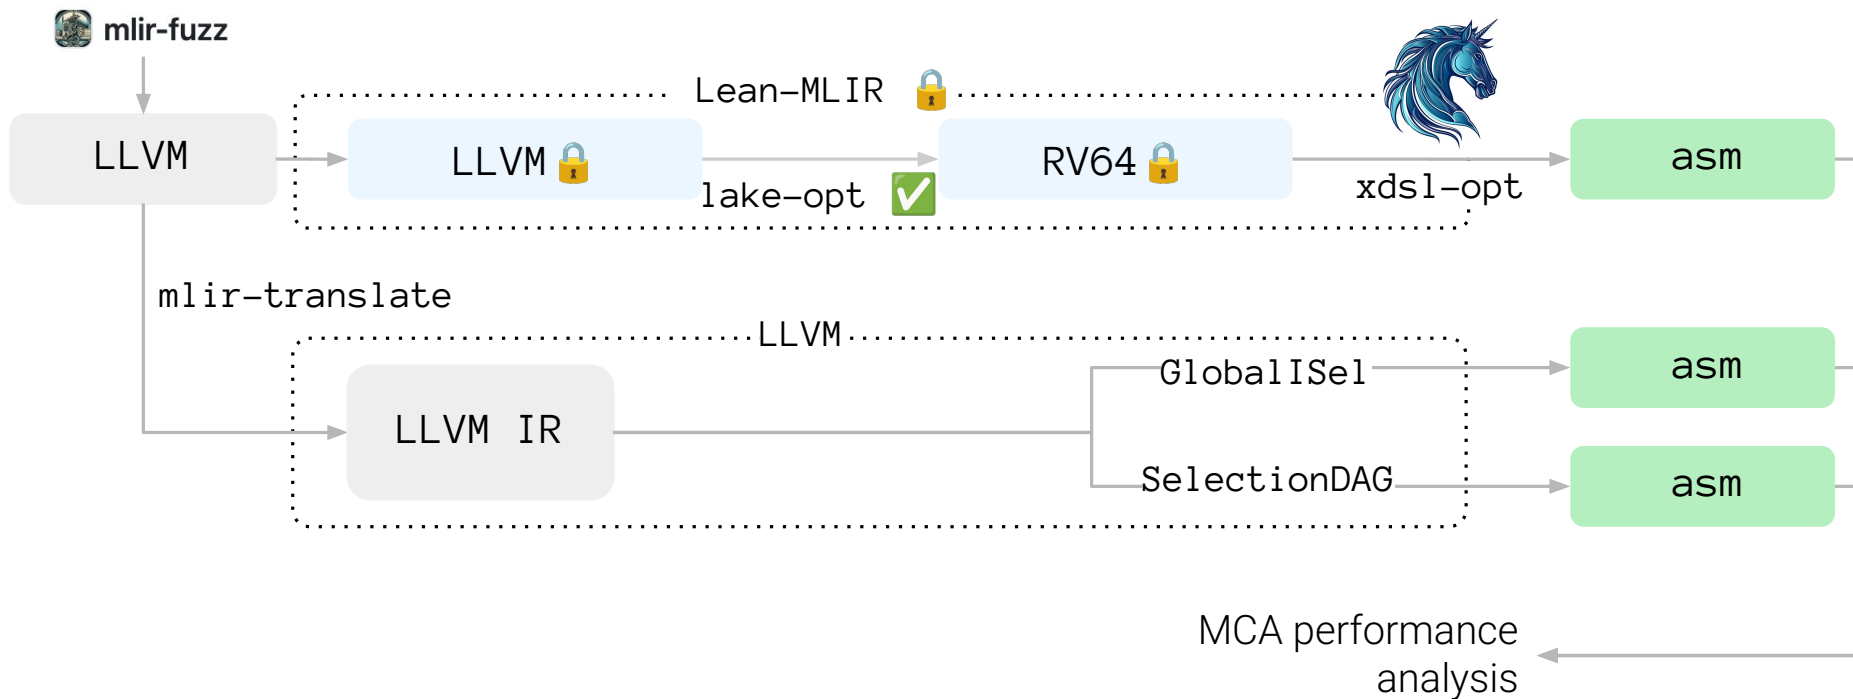
**...ARE WE?**



# Comparing with LLVM's *instruction selection*

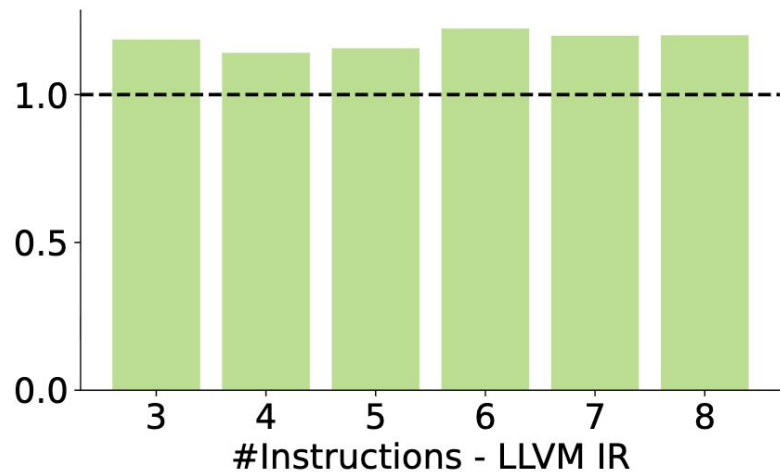


# Comparing with LLVM's *instruction selection*

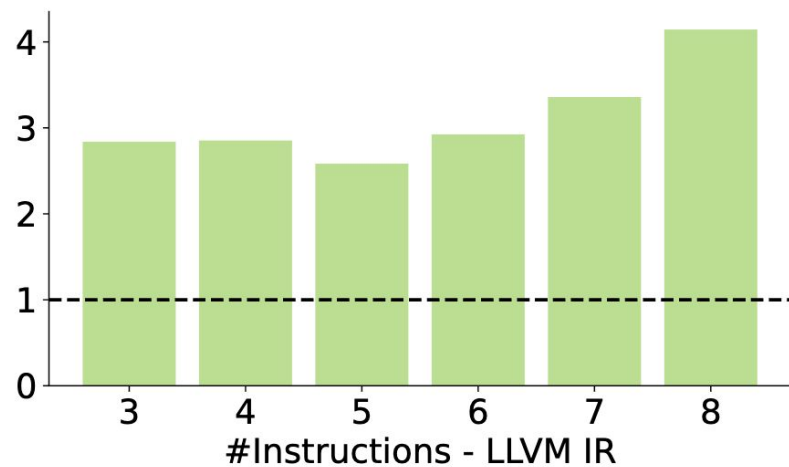


# MCA Estimate: #Cycles

LeanMLIR  
GlobalSel




LeanMLIR  
SelDAG




# MCA Estimate: #Cycles

LeanMLIR  
GlobalSel




LeanMLIR  
SelDAG




# MCA Estimate: #Instructions

LeanMLIR  
GlobalSel

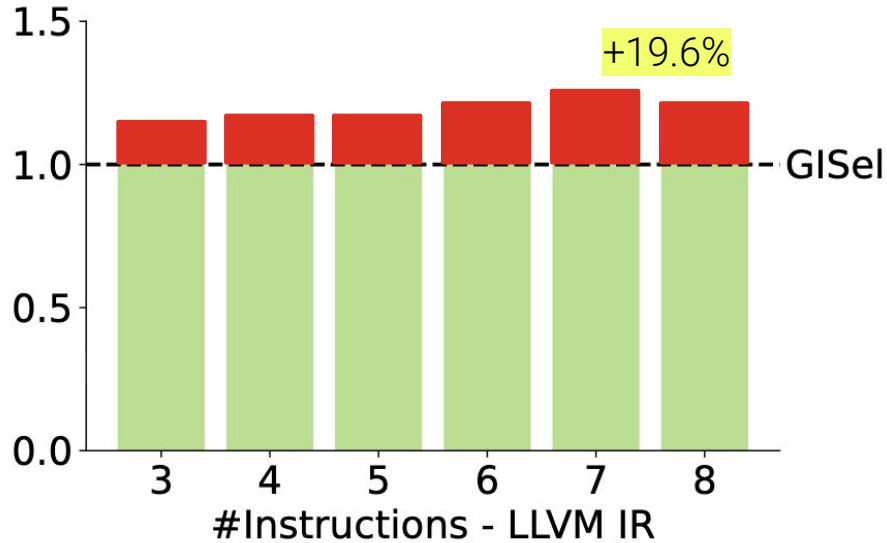


LeanMLIR  
SelDAG




# MCA Estimate: #Instructions


LeanMLIR  
GlobalSel




LeanMLIR  
SelDAG




The same .asm output!



# What's next



Bitblastable  
Vector Semantics



Close the gap  
SelectionDAG → GlobalSel

## Verified Intermediate Representation

A verified implementation of the MLIR SSA-based datastructures.

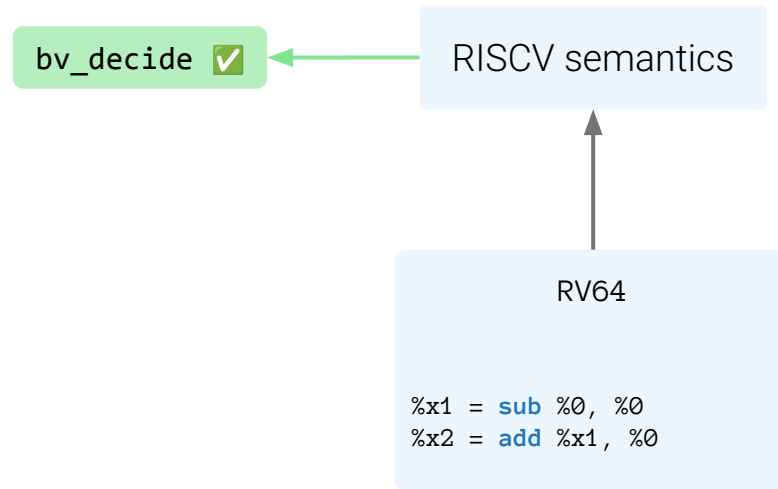
<https://github.com/opencompil/veir/>

Realistic use case  
and performance

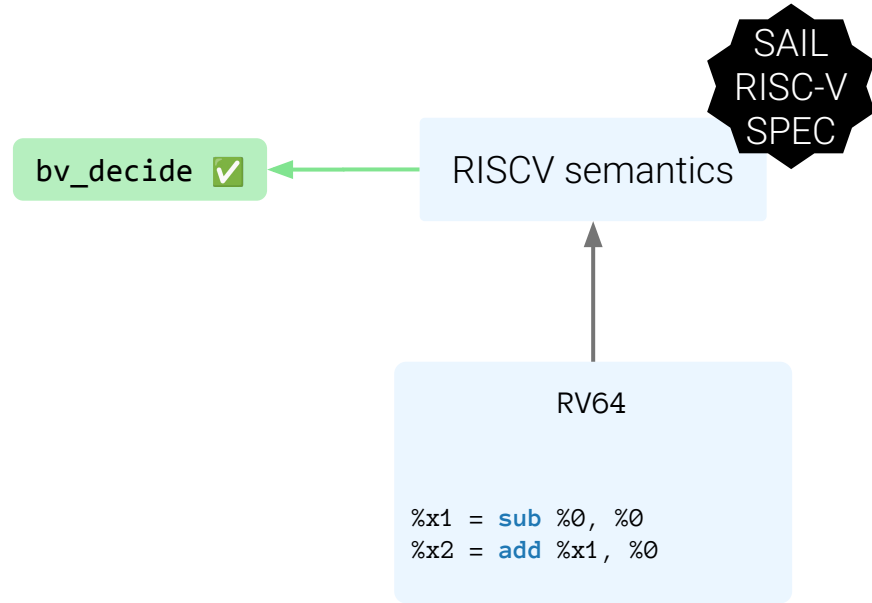
*Questions?*

bv\_decide ✓

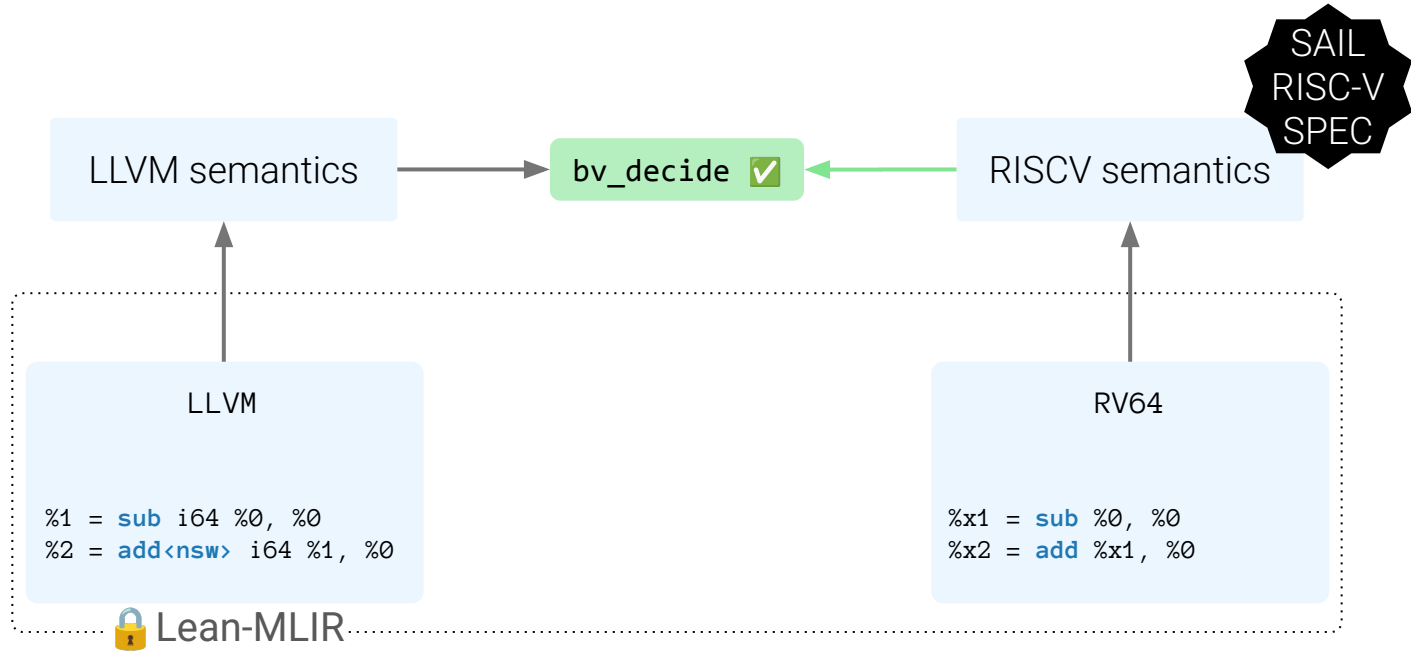
# Questions?



# Questions?



# Questions?



# Questions?

