

# Toward A More Declarative InstCombine

*Siddharth Bhat, Léo Stefanescu, George Rennie, John Regehr, Tobias Grosser*

(**x** y : BV **w**), x + y = y + x



```
(x' y' w' : BV 65), w' <= 64 ->  
let x := x' & (1 <<< w' - 1)  
let y := y' & (1 <<< w' - 1)  
x + y = y + x
```



# Léo



# George



# Tobias



# John



Sid

Will My  
Visa Arrive?



# What's InstCombine?

## LLVM's Amazing Peephole Optimizer!

$$X + X \quad \rightarrow \quad X \lll 1$$

$$-A + -B \quad \rightarrow \quad - (A + B)$$

$$-A + B \quad \rightarrow \quad B - A$$

... and thousands more!

# Is InstCombine Code Common (in lines of code)?

Vectorize | 20k ←

InstCombine | 25k ←

**InstCombine is larger than Vectorize?!**



# Declarative Code Hidden in Plain Sight C++

```
// X + X --> X << 1
if (LHS == RHS) {
    auto *Shl = BinaryOperator::CreateShl(LHS, ConstantInt::get(Ty, 1));
    Shl->setHasNoSignedWrap(I.hasNoSignedWrap());
    Shl->setHasNoUnsignedWrap(I.hasNoUnsignedWrap());
    return Shl;
}

Value *A, *B;
if (match(LHS, m_Neg(m_Value(A)))) {
    // -A + -B --> -(A + B)
    if (match(RHS, m_Neg(m_Value(B))))
        return BinaryOperator::CreateNeg(Builder.CreateAdd(A, B));

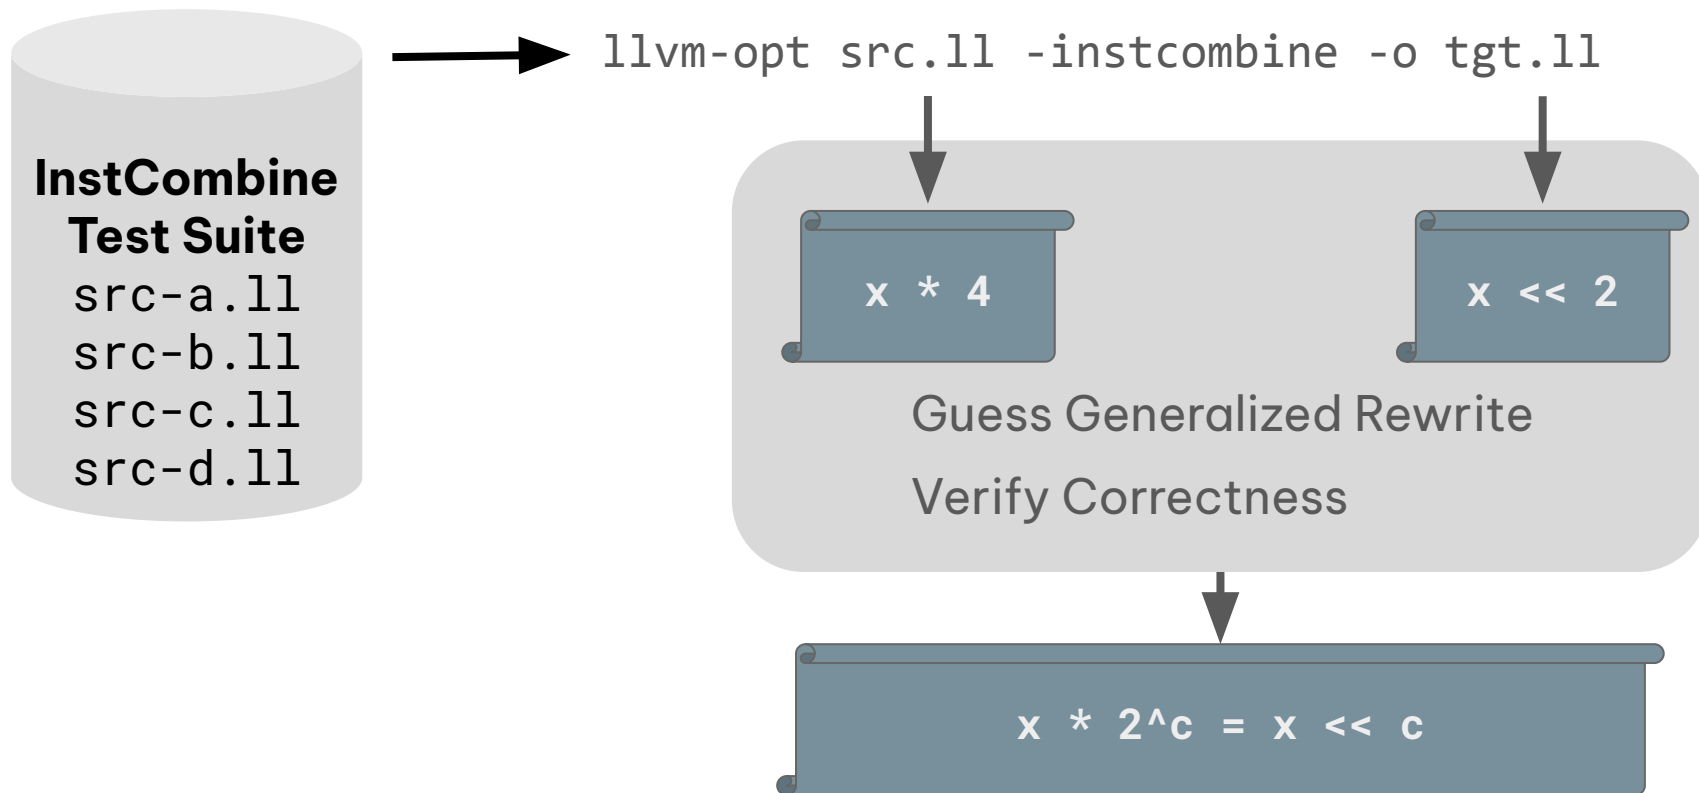
    // -A + B --> B - A
    auto *Sub = BinaryOperator::CreateSub(RHS, A);
    auto *OB0 = cast<OverflowingBinaryOperator>(LHS);
    Sub->setHasNoSignedWrap(I.hasNoSignedWrap() && OB0->hasNoSignedWrap());
}
```

InstCombine yearns  
to be declarative!

It facilitates

- **verification**
- maintenance
- optimization

# Idealized Pipeline For Extracting Declarative Rewrites



# LLVM Has “Verify” Already: Alive!

```
define i32 @src(i32) {  
    %r = udiv i32 %0, 8192  
    ret i32 %r  
}
```

Transformation seems to be correct!

```
define i32 @tgt(i32) {  
    %r = lshr i32 %0, 13  
    ret i32 %r  
}
```

# LLVM Has “Verify” Already: Alive! – It’s used Daily!

[InstCombine] | Fold `fma x,`  
only if the `and` (#100106)



XChy authored 3



dtcxzyw authc



zjaffal authored on Jun 23

Fixes [#113123](#)  
Alive proof: <https://alive2.llvm.org/ce/z/MFKNXs>

`main` (#113164)

1 parent [67ff5ba](#) cc

Alive2 proof (P1  
`smt-to`):  
<https://alive2.llvm.org/ce/z/MFKNXs>

`main` (#100106)

1 parent [eb9bf18](#)

[InstCombine] fold (`zext`  
|| `(a != c && b =`  
== `(b != c)`) (#94915)

resolves [#92966](#)

alive proof  
<https://alive2.llvm.org/ce/z/MFKNXs>

`main` (#94915) · llvmlorg

1 parent [3ae6755](#) commit f7f

[InstCombine] Extend Fold of Zero-  
extended Bit Test (#102100)

mskamp authored on Aug 21 · 51/56 · Verified

Previously, `(zext (icmp ne (and X, (1 << ShAmt)), 0))` has only been folded if the bit width of X and the result were equal. Use a `trunc` or `zext` instruction to also support other bit widths.

This is a follow-up to commit [533190a](#), which introduced a regression: `(zext (icmp ne (and (lshr X ShAmt) 1) 0))` is not folded any longer to `(zext/trunc (and (lshr X ShAmt) 1))` since the commit introduced the fold of `(icmp ne (and (lshr X ShAmt) 1) 0)` to `(icmp ne (and X (1 << ShAmt)) 0)`. The change introduced by this commit restores this fold.

Alive proof: <https://alive2.llvm.org/ce/z/MFKNXs>

Relates to issue [#86813](#) and pull request [#101838](#).

`main` (#102100)

1 parent [4f07508](#) commit 170a21e

# LLVM Has “Verify” Already: Alive!

```
define i32 @src(i32) {  
    %r = udiv i32 %0, 8192  
    ret i32 %r  
}
```

Transformation seems to be correct!

```
define i32 @tgt(i32) {  
    %r = lshr i32 %0, 13  
    ret i32 %r  
}
```

# LLVM Has “Verify” Already: Alive!

```
define i32 @src(i32) {  
    %r = udiv i32 %0, 1  
    ret i32 %r  
}
```

```
define i32 @tgt(i32) {  
    %r = lshr i32 %0, 13  
    ret i32 %r  
}
```

# LLVM Has “Verify” Already: Alive!

```
define i32 @src(i32) {  
    %r = udiv i32 %0, 1  
    ret i32 %r  
}
```

```
define i32 @tgt(i32) {  
    %r = lshr i32 %0, 13  
    ret i32 %r  
}
```

Transformation doesn't verify!

ERROR: Value mismatch

Example:

```
i32 %#0 = #x00000001 (1)
```

Source:

```
i32 %r = #x00000001 (1)
```

Target:

```
i32 %r = #x00000000 (0)
```

```
Source value: #x00000001 (1)
```

```
Target value: #x00000000 (0)
```

# LLVM Has “Verify” Already: Alive!

```
define i32 @src(i32) {  
  %r = udiv i32 %0, 1  
  ret i32 %r  
}
```

```
define i32 @tgt(i32) {  
  %r = lshr i32 %0, 13  
  ret i32 %r  
}
```

Transformation doesn't verify!

ERROR: Value mismatch

Example:

i32 %#0 = #x00000001 (1)

Source:

i32 %r = #x00000001 (1)

Target:

i32 %r = #x00000000 (0)

Source value: #x00000001 (1)

Target value: #x00000000 (0)

# LLVM Has “Verify” Already: Alive!

```
define i32 @src(i32) {  
  %r = udiv i32 %0, 1  
  ret i32 %r  
}
```

```
define i32 @tgt(i32) {  
  %r = lshr i32 %0, 13  
  ret i32 %r  
}
```

Transformation doesn't verify!

ERROR: Value mismatch

Example:

i32 %#0 = #x00000001 (1)

Source:

i32 %r = #x00000001 (1)

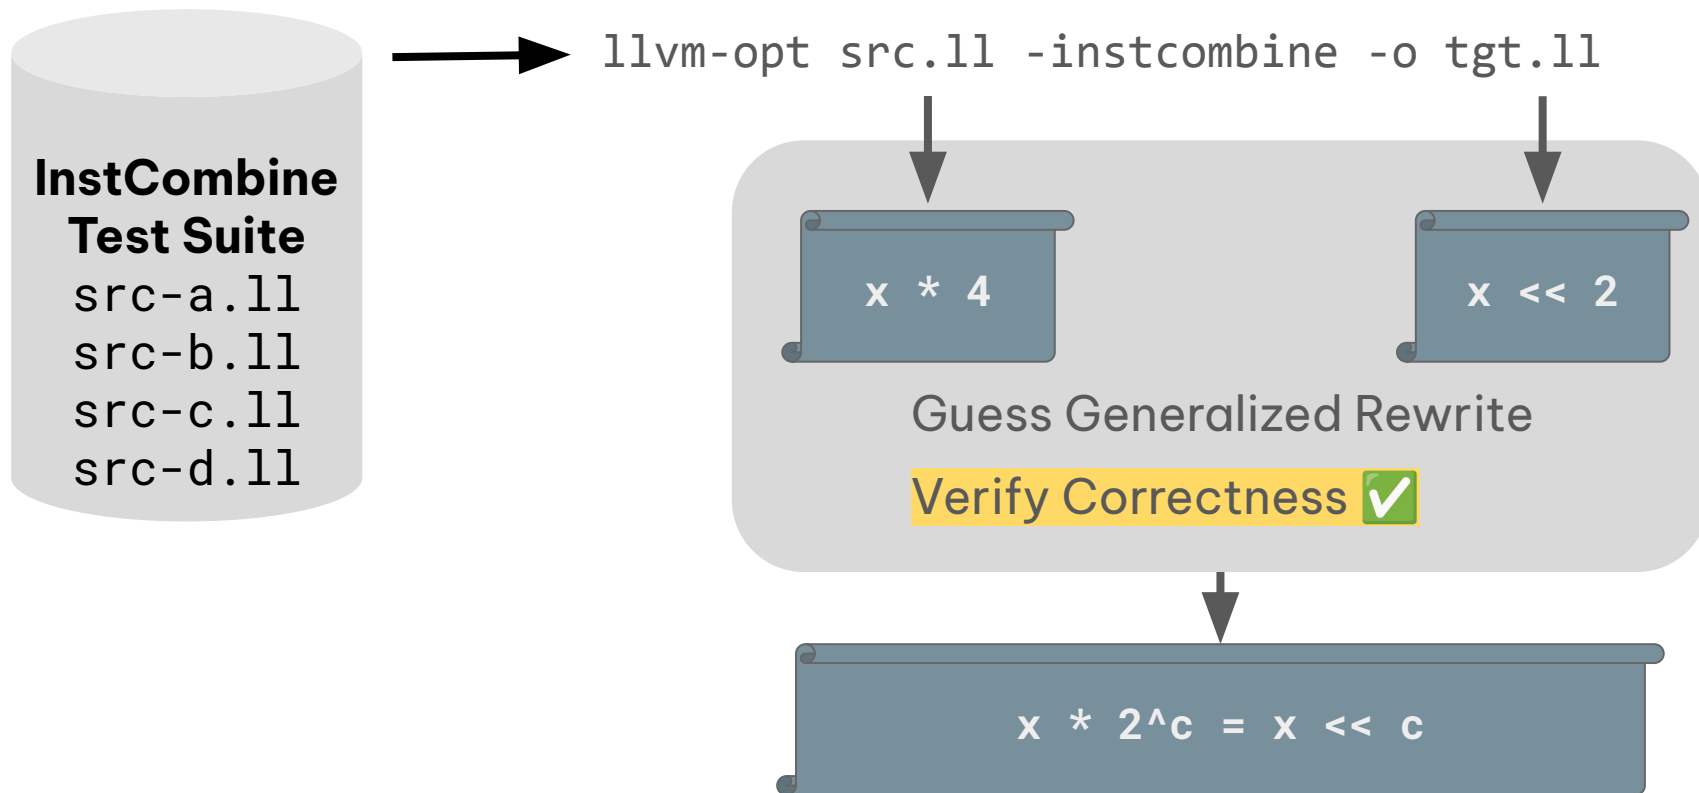
Target:

i32 %r = #x00000000 (0)

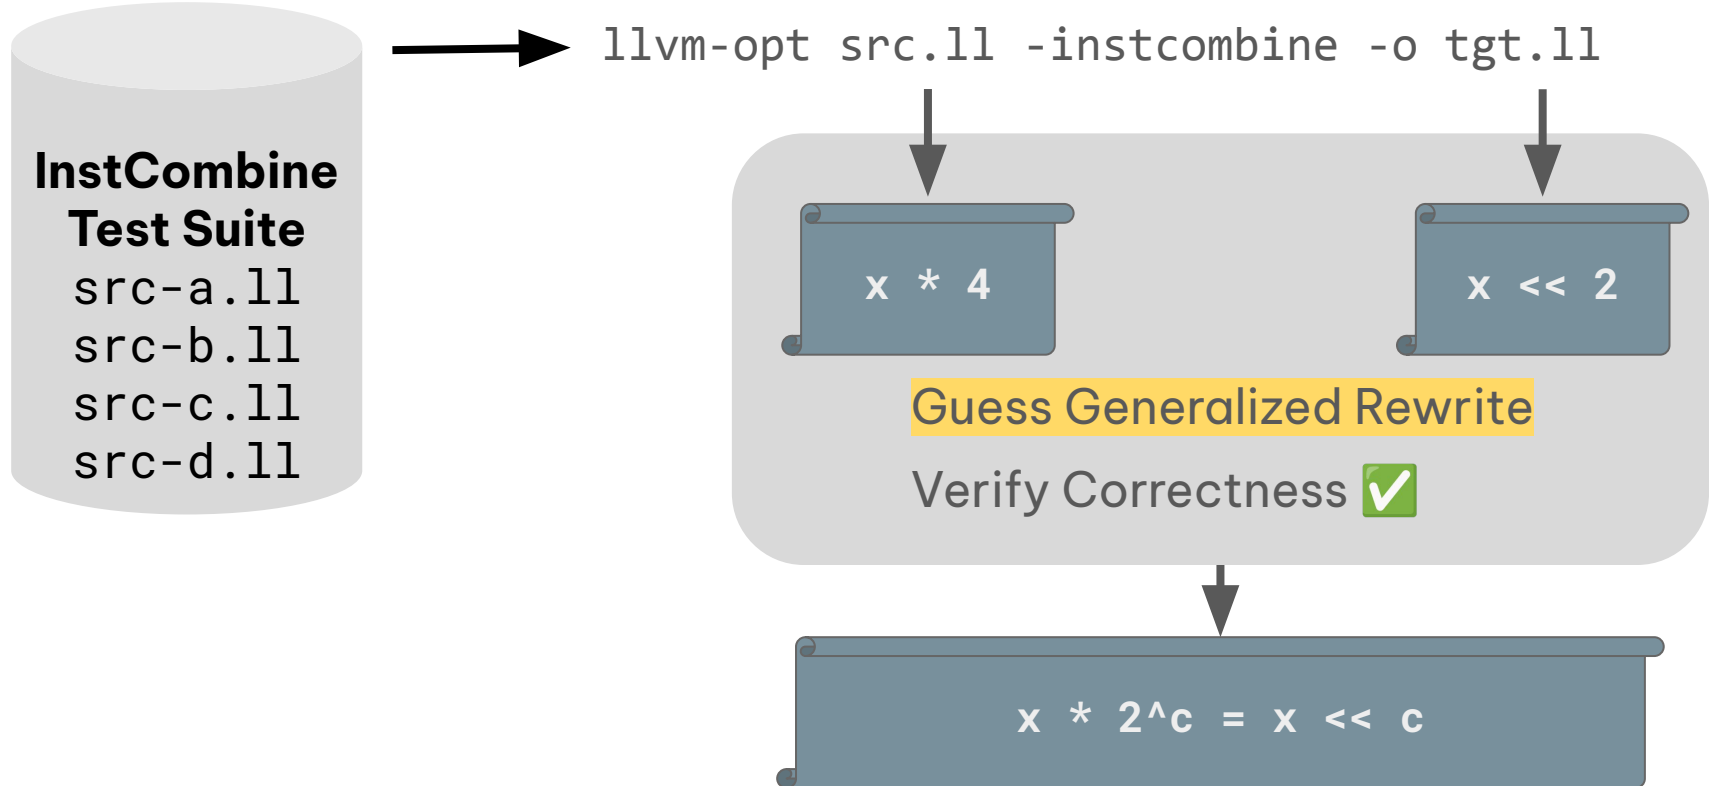
Source value: #x00000001 (1)

Target value: #x00000000 (0)

# Idealized Pipeline For Extracting Declarative Rewrites




# Idealized Pipeline For Extracting Declarative Rewrites



# How To Guess Good Width Generic Rewrites?

$(x\ y : BV\ 3),\ x + y = y + x$

$(x\ y : BV\ w),\ x + y = y + x$

 "Alive, is this rewrite true?"

## Problem 1: Constant Generalization

$(x : BV\ 5),\ (x \ll 4) \gg 4 = x \& 1$

$(x : BV\ w),\ (x \ll (w - 1)) \gg (w - 1) = x \& 1$

 How to come up with  $(w - 1)$ ?

## Problem 2: Precondition Generation

$(x : BV\ 3),\ (x.zext\ 4).sext\ 7 = x.zext\ 7$

$(x : BV\ w3),\ w3 < w4 \leq w7 \Rightarrow$

$(x.zext\ w4).xsext\ w7 = x.sext\ w7$

 How to come up with  $(w3 < w4 \leq w7)$ ?

## Hydra: Generalizing Peephole Optimizations with Program Synthesis

MANASIJ MUKHERJEE, University of Utah, USA

JOHN REGEHR, University of Utah, USA

Optimizing compilers rely on peephole optimizations to simplify combinations of instructions and remove redundant instructions. Typically, a new peephole optimization is added when a compiler developer notices an optimization opportunity—a collection of dependent instructions that can be improved—and manually derives a more general rewrite rule that optimizes not only the original code, but also other, similar collections of instructions. In this paper, we present Hydra, a tool that automates the process of generalizing peephole optimizations using a collection of techniques centered on program synthesis. One of the most important problems we have solved is finding a version of each optimization that is independent of the bitwidths of the optimization's inputs (when this version exists). We show that Hydra can generalize 75% of the ungeneralized missed peephole optimizations that LLVM developers have posted to the LLVM project's issue tracker. All of Hydra's generalized peephole optimizations have been formally verified, and furthermore we can automatically turn them into C++ code that is suitable for inclusion in an LLVM pass.

CCS Concepts: • **Software and its engineering** → **Translator writing systems and compiler generators**; *Source code generation*.

Additional Key Words and Phrases: program synthesis, generalization, superoptimization, llvm, alive2, souper, hydra, peephole optimization

### ACM Reference Format:

Manasij Mukherjee and John Regehr. 2024. Hydra: Generalizing Peephole Optimizations with Program Synthesis. *Proc. ACM Program. Lang.* 8, OOPSLA1, Article 120 (April 2024), 29 pages. <https://doi.org/10.1145/3649837>

## Hydra Generates Potentially Likely Candidates

# Manasij

A portrait of Manasij Mukherjee, a man with dark hair and glasses, looking slightly to the right. He is wearing a dark blue shirt.

## Hydra: Generalizing Peephole Optimizations with Program Synthesis

MANASIJ MUKHERJEE, University of Utah, USA

JOHN REGEHR, University of Utah, USA

Optimizing compilers rely on peephole optimizations to simplify combinations of instructions and remove redundant instructions. Typically, a new peephole optimization is added when a compiler developer notices an optimization opportunity—a collection of dependent instructions that can be improved—and manually derives a more general rewrite rule that optimizes not only the original code, but also other, similar collections of instructions. In this paper, we present Hydra, a tool that automates the process of generalizing peephole optimizations using a collection of techniques centered on program synthesis. One of the most important problems we have solved is finding a version of each optimization that is independent of the bitwidths of the optimization’s inputs (when this version exists). We show that Hydra can generalize 75% of the ungeneralized missed peephole optimizations that LLVM developers have posted to the LLVM project’s issue tracker. All of Hydra’s generalized peephole optimizations have been formally verified, and furthermore we can automatically turn them into C++ code that is suitable for inclusion in an LLVM pass.

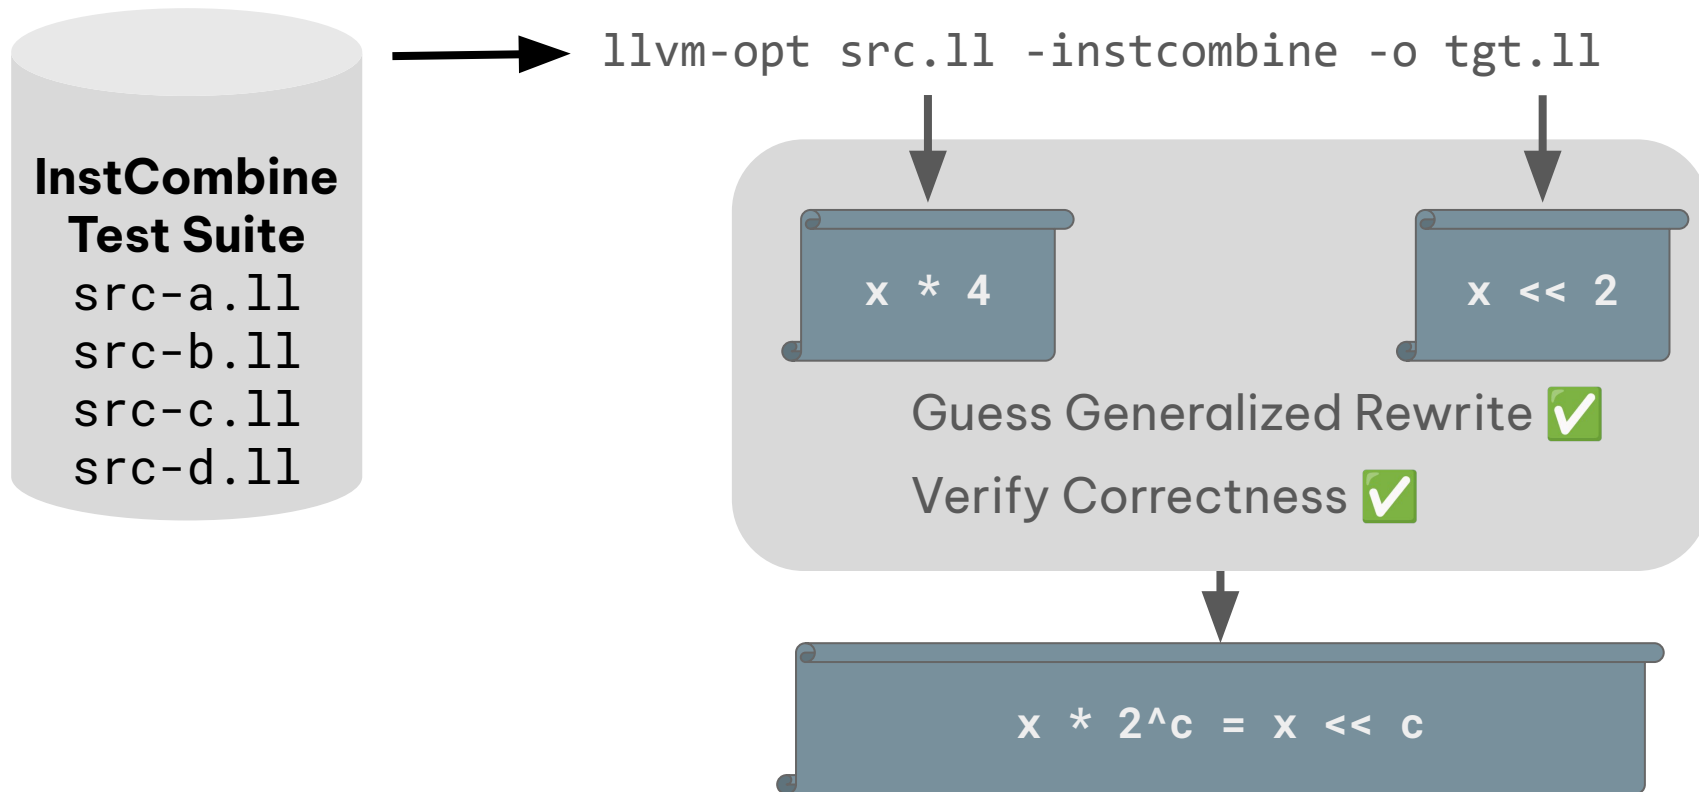
CCS Concepts: • **Software and its engineering** → **Translator writing systems and compiler generators**; *Source code generation.*

Additional Key Words and Phrases: program synthesis, generalization, superoptimization, llvm, alive2, souper, hydra, peephole optimization

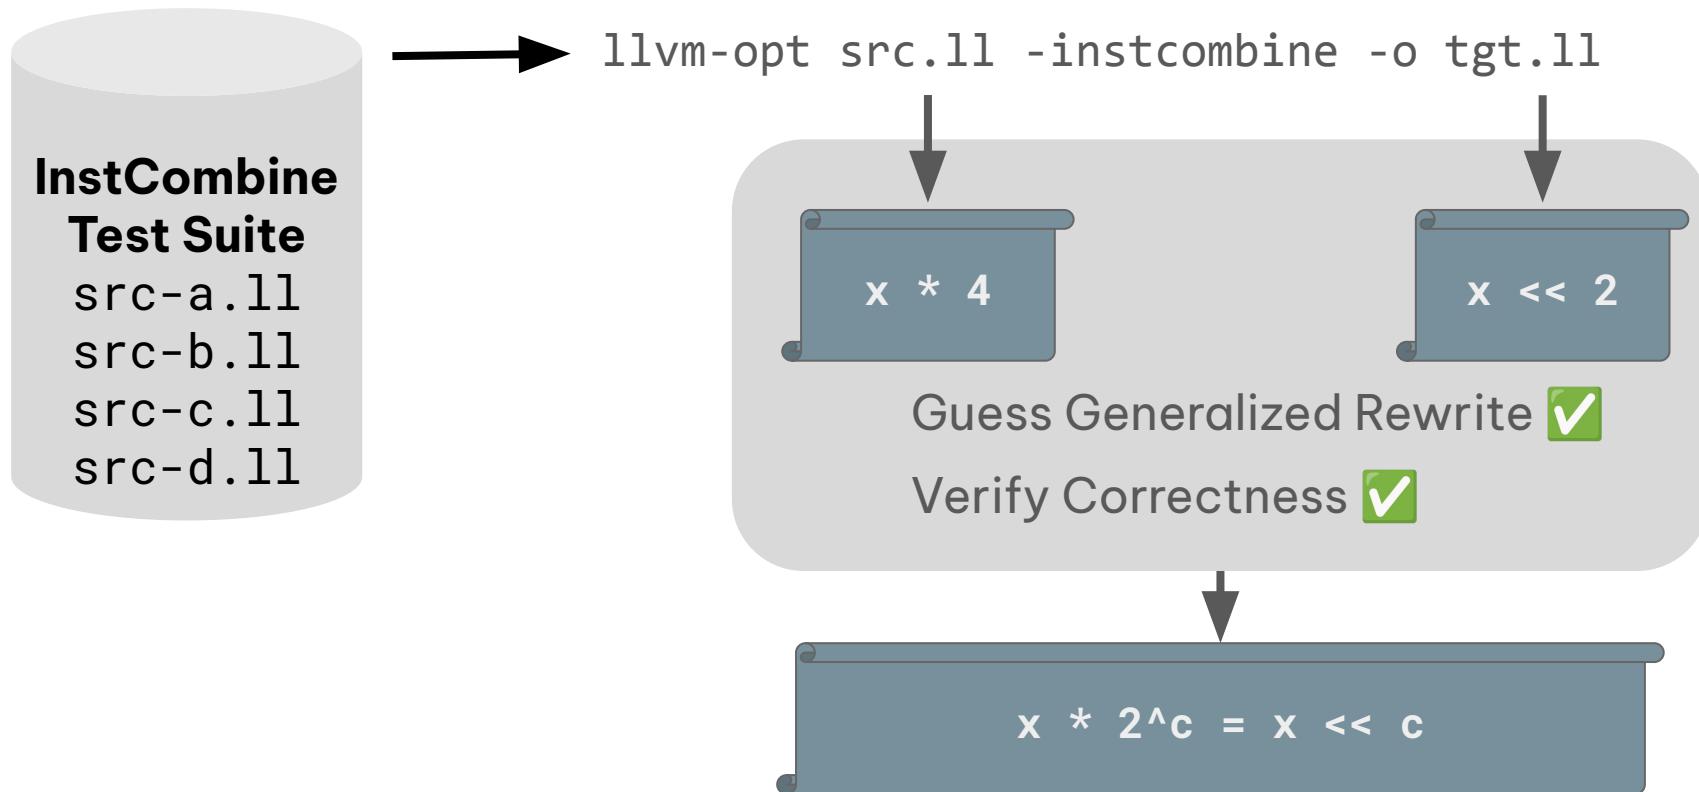
### ACM Reference Format:

Manasij Mukherjee and John Regehr. 2024. Hydra: Generalizing Peephole Optimizations with Program Synthesis. *Proc. ACM Program. Lang.* 8, OOPSLA1, Article 120 (April 2024), 29 pages. <https://doi.org/10.1145/3649837>

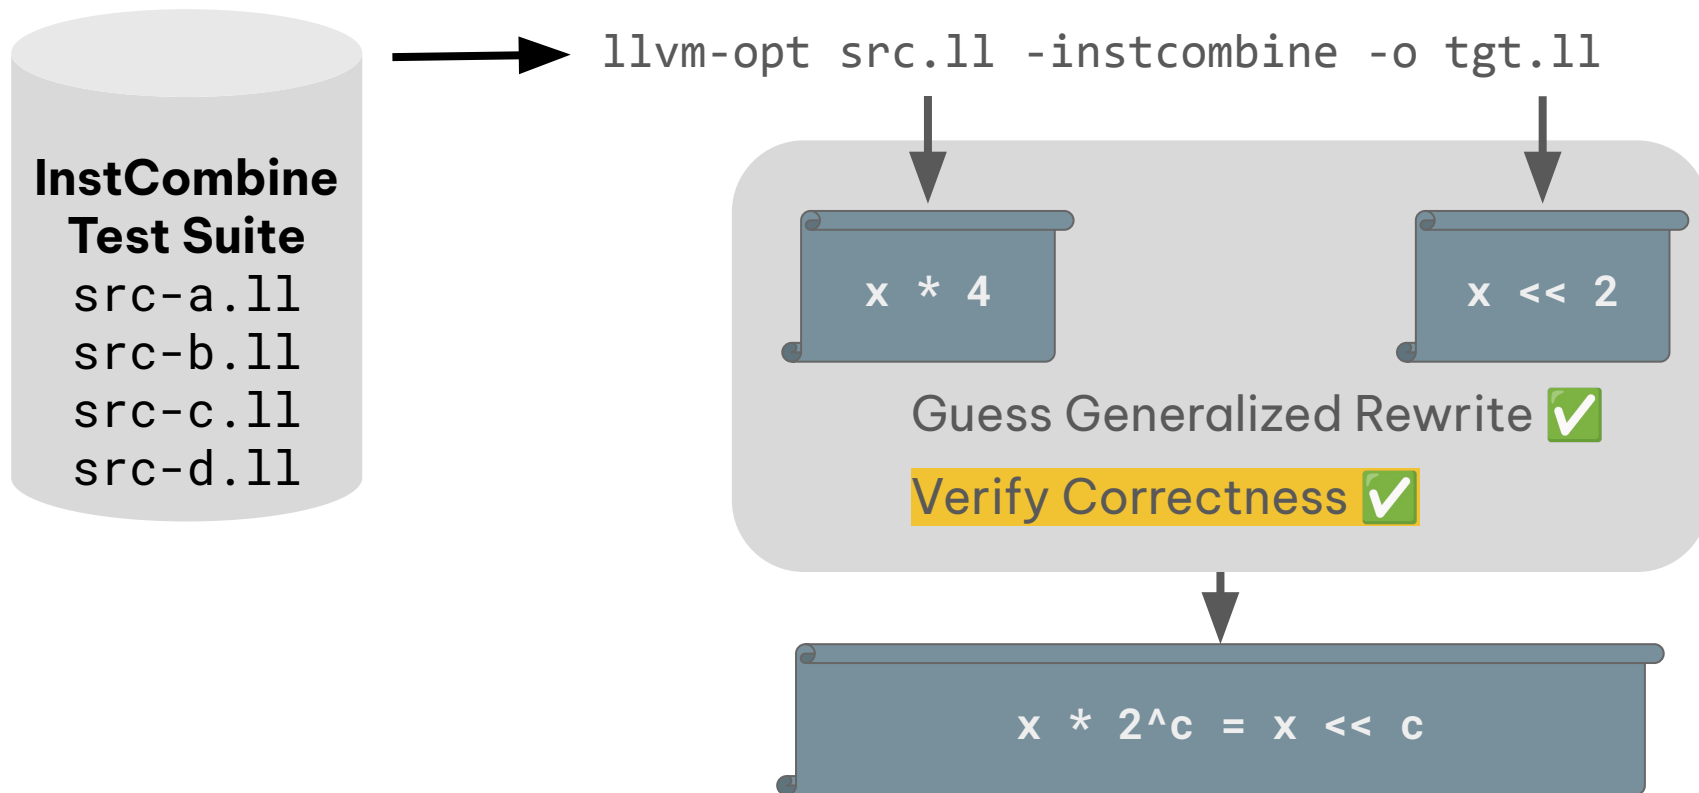
# Idealized Pipeline For Extracting Declarative Rewrites



# Did Sid Think This Is A Lightning Talk By Accident? 🤨



# Idealized Pipeline For Extracting Declarative Rewrites



# Verifying The Correctness Of Candidates

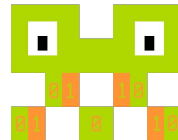
$(x\ y : \text{BV } w),\ x + y = y + x$

"Alive, is this rewrite true?"



$(x\ y : \text{BV } 1),\ x + y = y + x$

"Bitwuzla, is this rewrite true?"



$\forall (x\ y : \text{BV } 1),\ x + y = y + x$

fixed size bitvectors

theory problem to decide true/false

## About Bitwuzla

Bitwuzla is a Satisfiability Modulo Theories (SMT) solver for the theories of fixed-size bit-vectors, floating-point arithmetic, arrays and uninterpreted functions and their combinations.

# Verifying The Correctness Of Candidates

$(x\ y : \text{BV } w), x + y = y + x$

"Alive, is this rewrite true?"



$(x\ y : \text{BV } 1), x + y = y + x$

"Bitwuzla, is this rewrite true?"

$(x\ y : \text{BV } 2), x + y = y + x$

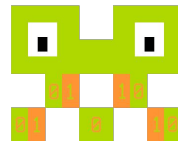
"Bitwuzla, is this rewrite true?"

$(x\ y : \text{BV } 3), x + y = y + x$

"Bitwuzla, is this rewrite true?"

...

All widths from 1 to 64



# Verifying The Correctness Of Candidates Is Exponential!

```
(x : BV w3), w3 < w4 <= w7 =>  
  (x.zext w4).xsezt w7 = x.sext w7
```

"Alive, is this rewrite true?"

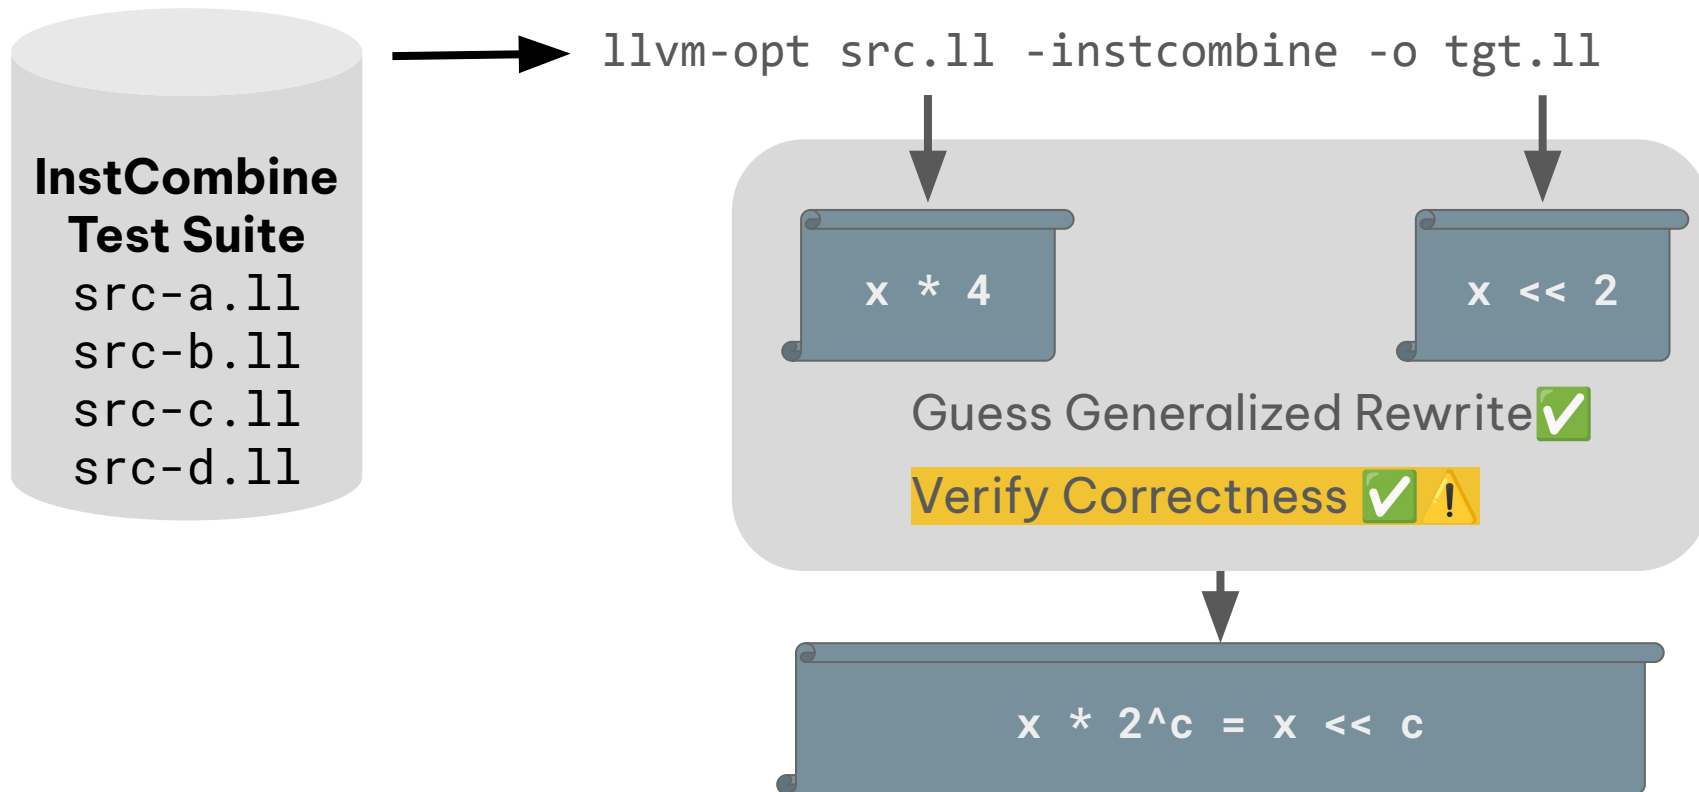


w3 : 64 choices  
w4 : 64 choices  
w7 : 64 choices

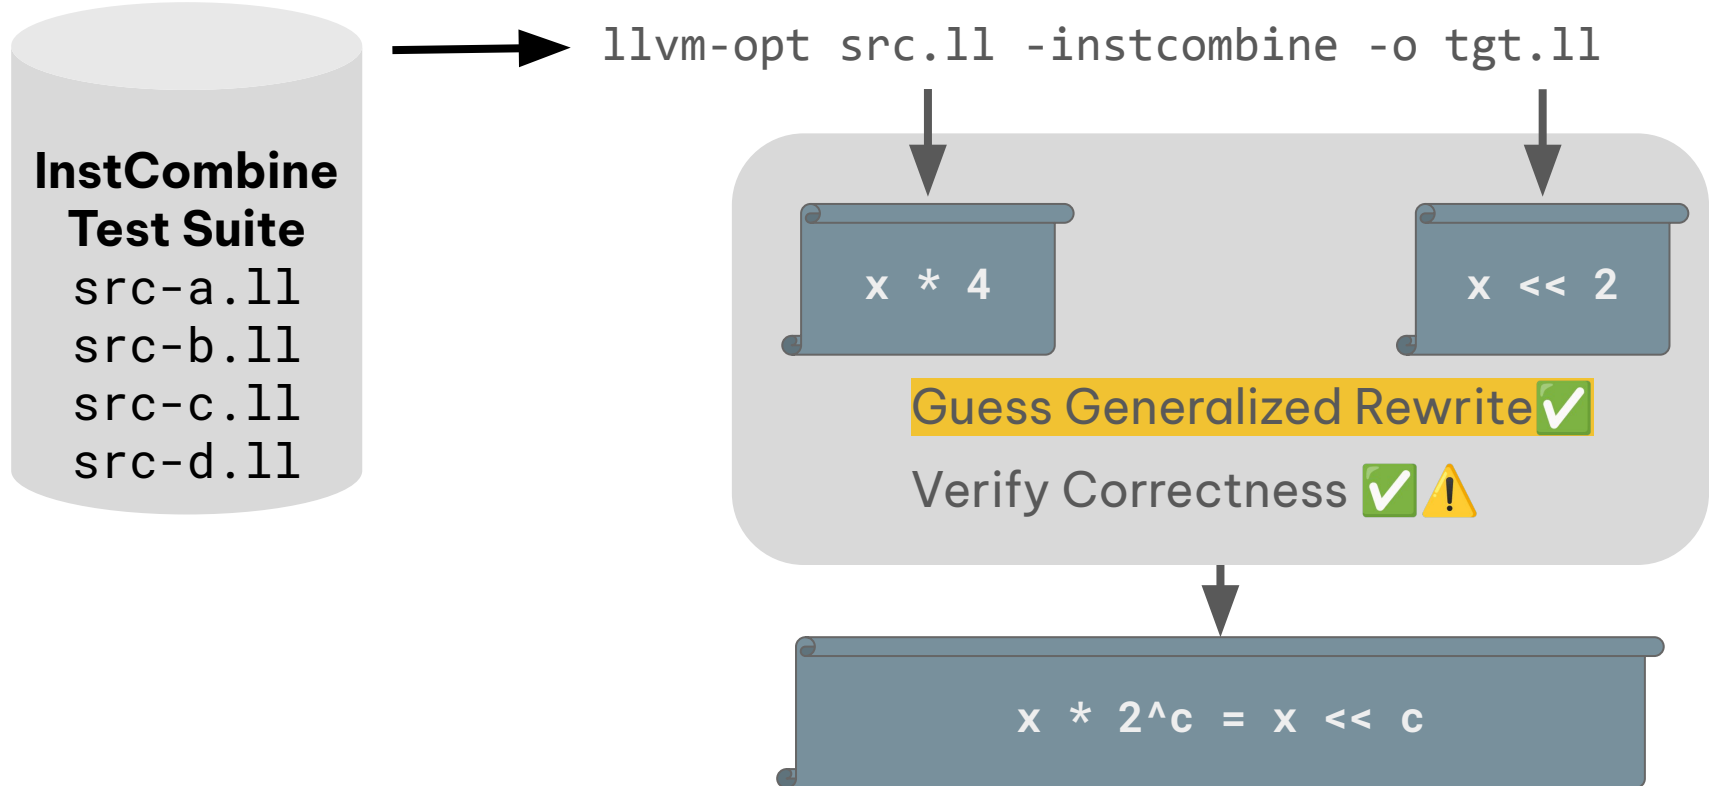
262144 combinations 🤖  
Alive needs **18.5 hours** to prove this!

**We Need An Algorithm That's Much Better Than Exhaustive Enumeration!**

# Idealized Pipeline For Extracting Declarative Rewrites



# Idealized Pipeline For Extracting Declarative Rewrites



# Manasij



## Hydra: Generalizing Peephole Optimizations with Program Synthesis

MANASIJ MUKHERJEE, University of Utah, USA

JOHN REGEHR, University of Utah, USA

Optimizing compilers rely on peephole optimizations to simplify combinations of instructions and remove redundant instructions. Typically, a new peephole optimization is added when a compiler developer notices an optimization opportunity—a collection of dependent instructions that can be improved—and manually derives a more general rewrite rule that optimizes not only the original code, but also other, similar collections of instructions. In this paper, we present Hydra, a tool that automates the process of generalizing peephole optimizations using a collection of techniques centered on program synthesis. One of the most important problems we have solved is finding a version of each optimization that is independent of the bitwidths of the optimization’s inputs (when this version exists). We show that Hydra can generalize 75% of the ungeneralized missed peephole optimizations that LLVM developers have posted to the LLVM project’s issue tracker. All of Hydra’s generalized peephole optimizations have been formally verified, and furthermore we can automatically turn them into C++ code that is suitable for inclusion in an LLVM pass.

CCS Concepts: • **Software and its engineering** → **Translator writing systems and compiler generators**; *Source code generation.*

Cannot generalize on multiple widths

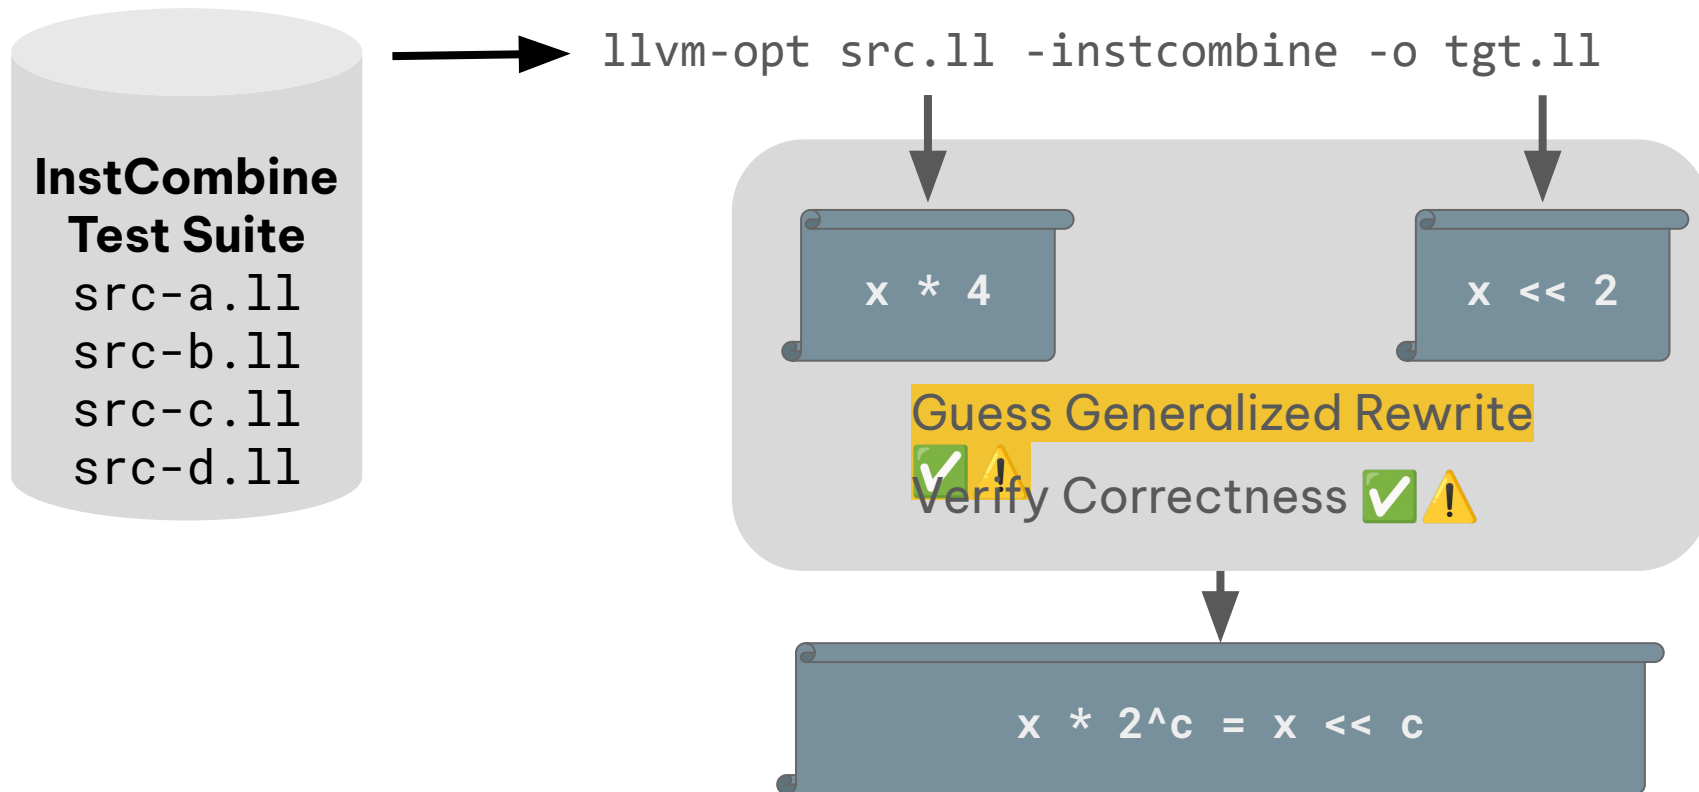
Can we have fewer human-tuned heuristics?

3649837

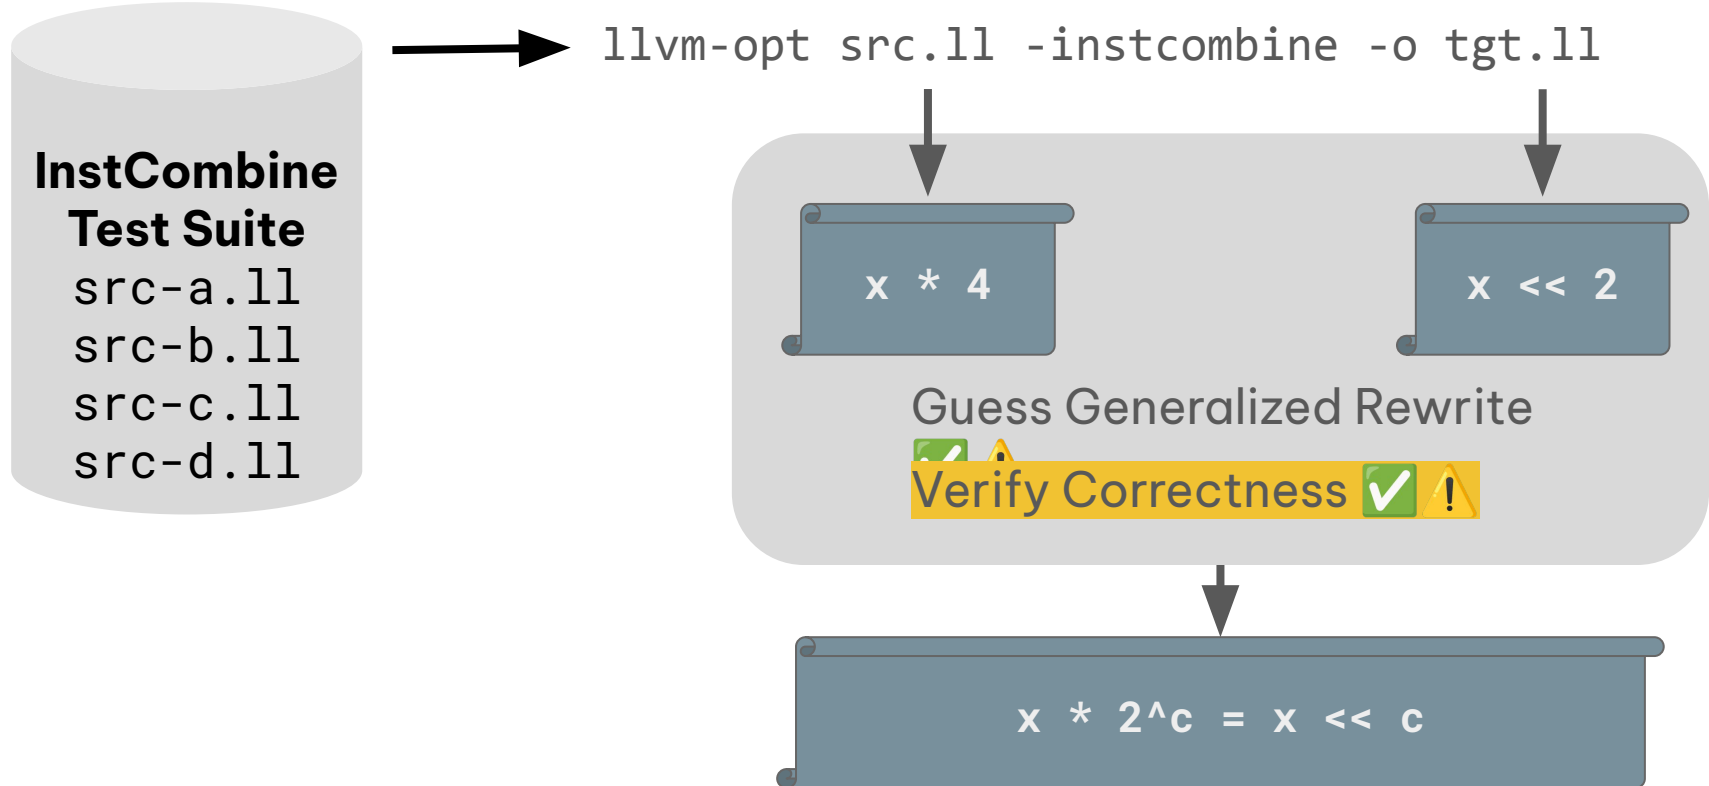
iper,

ram  
145/

# Idealized Pipeline For Extracting Declarative Rewrites



# Idealized Pipeline For Extracting Declarative Rewrites



# Verifying The Correctness Of Candidates

$(x\ y : \text{BV } w), x + y = y + x$

"Alive, is this rewrite true?"



$(x\ y : \text{BV } 1), x + y = y + x$

"Bitwuzla, is this rewrite true?"

$(x\ y : \text{BV } 2), x + y = y + x$

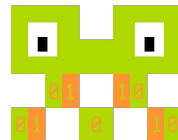
"Bitwuzla, is this rewrite true?"

$(x\ y : \text{BV } 3), x + y = y + x$

"Bitwuzla, is this rewrite true?"

...

All widths from 1 to 64



# Toward A Faster Algorithm

$(x\ y : BV\ w),\ x + y = y + x$



$(x\ y : BV\ 1),\ x + y = y + x$

Bitwuzla, is this rewrite true?"

$x\ y : BV\ 2),\ x + y = y + x$

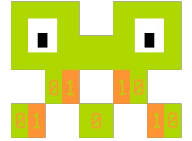
Bitwuzla, is this rewrite true?"

$x\ y : BV\ 3),\ x + y = y + x$

Bitwuzla, is this rewrite true?"

...

All widths from 1 to 64



# Toward A Faster Algorithm: Can We Eliminate Looping?

O Aristotle, can we have bitwuzla loop on the widths?



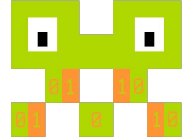
$y + x$   
Bitwuzla, is this rewrite true?"

$x + y : BV(2), x + y = y + x$   
Bitwuzla, is this rewrite true?"

$x + y : BV(3), x + y = y + x$   
Bitwuzla, is this rewrite true?"

...

All widths from 1 to 64



# Toward A Faster Algorithm: Can We Eliminate Looping?

But Plato, Bitwuzla only knows about bitvectors, not natural numbers!



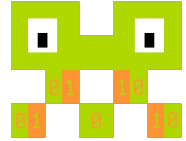
$y + x$   
Bitwuzla, is this rewrite true?"

$x + y : BV(2), x + y = y + x$   
Bitwuzla, is this rewrite true?"

$x + y : BV(3), x + y = y + x$   
Bitwuzla, is this rewrite true?"

...

All widths from 1 to 64



# Toward A Faster Algorithm: Can We Eliminate Looping?

Aha, but what is a natural number, if not a unsigned bitvector?



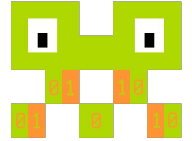
$y + x$   
Bitwuzla, is this rewrite true?"

$x + y : BV(2), x + y = y + x$   
Bitwuzla, is this rewrite true?"

$x + y : BV(3), x + y = y + x$   
Bitwuzla, is this rewrite true?"

...

All widths from 1 to 64




# Toward A Faster Algorithm: Widths As Bitvectors

$x$	: BV 3	-	-	$x_2$	$x_1$	$x_0$	
$w_3\text{mask}$	: BV 5	0	0	1	1	1	( $w_3\text{mask} := 1 \lll 3 - 1$ )
$x'$	: BV 5	*	*	$x_2$	$x_1$	$x_0$	
$x' \ \& \ w_3\text{mask}$	: BV 5	0	0	$x_2$	$x_1$	$x_0$	

Widths Like '3' Can Be Encoded By **Bitvectors** Like ' $w_3\text{mask}$ '!

# Toward A Faster Algorithm: Widths As Bitvectors

$(x\ y : BV\ w),\ x + y = y + x$  

$(x\ y : BV\ 1),\ x + y = y + x$

"Bitwuzla, is this rewrite true?"

$(x\ y : BV\ 2),\ x + y = y + x$

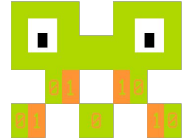
"Bitwuzla, is this rewrite true?"

$(x\ y : BV\ 3),\ x + y = y + x$

"Bitwuzla, is this rewrite true?"

...

upto width 64



# Toward A Faster Algorithm: Widths As Bitvectors

$(x\ y : BV\ w),\ x + y = y + x$   $\longrightarrow$

$(x\ y : BV\ 1),\ x + y = y + x$

"Bitwuzla, is this rewrite true?"

$(x\ y : BV\ 2),\ x + y = y + x$

"Bitwuzla, is this rewrite true?"

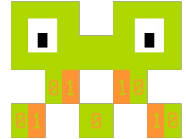
$(x\ y : BV\ 3),\ x + y = y + x$

"Bitwuzla, is this rewrite true?"

...

upto width 64

$(\quad : BV\ 65)$



# Toward A Faster Algorithm: Widths As Bitvectors

$(x\ y : \text{BV } w),\ x + y = y + x$   $\longrightarrow$



$(w' : \text{BV } 65),\ w' \leq 64 \rightarrow$

$(x\ y : \text{BV } 1),\ x + y = y + x$

"Bitwuzla, is this rewrite true?"

$(x\ y : \text{BV } 2),\ x + y = y + x$

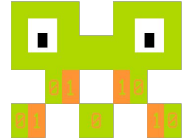
"Bitwuzla, is this rewrite true?"

$(x\ y : \text{BV } 3),\ x + y = y + x$

"Bitwuzla, is this rewrite true?"

...

upto width 64



# Toward A Faster Algorithm: Widths As Bitvectors

(`x` `y` : BV `w`), `x` + `y` = `y` + `x` →



(`x'` `w'` : BV 65), `w'` <= 64 ->  
let `x` := `x'`

(`x` `y` : BV 1), `x` + `y` = `y` + `x`

"Bitwuzla, is this rewrite true?"

(`x` `y` : BV 2), `x` + `y` = `y` + `x`

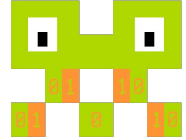
"Bitwuzla, is this rewrite true?"

(`x` `y` : BV 3), `x` + `y` = `y` + `x`

"Bitwuzla, is this rewrite true?"

...

upto width 64



# Toward A Faster Algorithm: Widths As Bitvectors

$(x\ y : BV\ w),\ x + y = y + x \longrightarrow$



$(x'\ w' : BV\ 65),\ w' \leq 64 \rightarrow$   
let  $x := x' \ \&\ (1 \lll w' - 1)$

$(x\ y : BV\ 1),\ x + y = y + x$

"Bitwuzla, is this rewrite true?"

$(x\ y : BV\ 2),\ x + y = y + x$

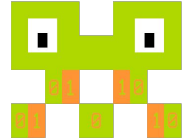
"Bitwuzla, is this rewrite true?"

$(x\ y : BV\ 3),\ x + y = y + x$

"Bitwuzla, is this rewrite true?"

...

upto width 64



# Toward A Faster Algorithm: Widths As Bitvectors

`(x y : BV w), x + y = y + x`  $\longrightarrow$

$\downarrow$  `x' masked with w is x!`

`(x' w' : BV 65), w' <= 64 ->`  
`let x := x' & (1 <<< w' - 1)`

`(x y : BV 1), x + y = y + x`

"Bitwuzla, is this rewrite true?"

`(x y : BV 2), x + y = y + x`

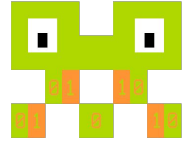
"Bitwuzla, is this rewrite true?"

`(x y : BV 3), x + y = y + x`

"Bitwuzla, is this rewrite true?"

...

upto width 64



# Toward A Faster Algorithm: Widths As Bitvectors

$(x \ y : BV \ w), \ x + y = y + x \longrightarrow$

$\downarrow$   
 $x'$  masked with  $w$  is  $x!$   
 $y'$  masked with  $w$  is  $y!$

$(x' \ y' \ w' : BV \ 65), \ w' \leq 64 \rightarrow$   
let  $x := x' \ \& \ (1 \lll w' - 1)$   
let  $y := y' \ \& \ (1 \lll w' - 1)$

$(x \ y : BV \ 1), \ x + y = y + x$

"Bitwuzla, is this rewrite true?"

$(x \ y : BV \ 2), \ x + y = y + x$

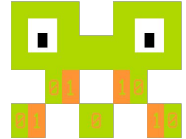
"Bitwuzla, is this rewrite true?"

$(x \ y : BV \ 3), \ x + y = y + x$

"Bitwuzla, is this rewrite true?"

...

upto width 64



# Toward A Faster Algorithm: Widths As Bitvectors

$(x\ y : \text{BV } w)$ ,  $x + y = y + x$   $\longrightarrow$

$\downarrow$   $x'$  masked with  $w$  **is**  $x!$   
 $y'$  masked with  $w$  **is**  $y!$

```
(x' y' w' : BV 65), w' <= 64 ->  
  let x := x' & (1 <<< w' - 1)  
  let y := y' & (1 <<< w' - 1)  
  x + y = y + x
```

$(x\ y : \text{BV } 1)$ ,  $x + y = y + x$

"Bitwuzla, is this rewrite true?"

$(x\ y : \text{BV } 2)$ ,  $x + y = y + x$

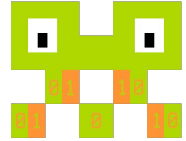
"Bitwuzla, is this rewrite true?"

$(x\ y : \text{BV } 3)$ ,  $x + y = y + x$

"Bitwuzla, is this rewrite true?"

...

upto width 64



# Toward A Faster Algorithm: Widths As Bitvectors

$(x\ y : BV\ w),\ x + y = y + x$   $\longrightarrow$

$\downarrow$   $x'$  masked with  $w$  **is**  $x!$   
 $y'$  masked with  $w$  **is**  $y!$

```
(x' y' w' : BV 65), w' <= 64 ->  
  let x := x' & (1 <<< w' - 1)  
  let y := y' & (1 <<< w' - 1)  
  x + y = y + x
```

"Bitwuzla, is this rewrite true?"

$(x\ y : BV\ 1),\ x + y = y + x$

"Bitwuzla, is this rewrite true?"

$(x\ y : BV\ 2),\ x + y = y + x$

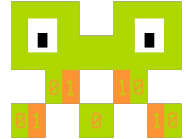
"Bitwuzla, is this rewrite true?"

$(x\ y : BV\ 3),\ x + y = y + x$


"Bitwuzla, is this rewrite true?"


...

upto width 64



# Upshot: One Solver Call For All Widths Simultaneously

$(x\ y : BV\ w),\ x + y = y + x$  

  $x'$  masked with  $w$  **is**  $x!$   
 $y'$  masked with  $w$  **is**  $y!$

```
(x' y' w' : BV 65), w' <= 64 ->  
  let x := x' & (1 <<< w' - 1)  
  let y := y' & (1 <<< w' - 1)  
  x + y = y + x
```

"Bitwuzla, is this rewrite true?"

$(x\ y : BV\ 1),\ x + y = y + x$

"Bitwuzla, is this rewrite true?"

$(x\ y : BV\ 2),\ x + y = y + x$

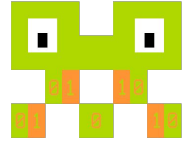
"Bitwuzla, is this rewrite true?"

$(x\ y : BV\ 3),\ x + y = y + x$

"Bitwuzla, is this rewrite true?"

...

upto width 64





One call for **all** widths, **simultaneously**



Bitwuzla is fast for one call!

# Upshot: One Solver Call For All Widths Simultaneously

$(x\ y : BV\ w),\ x + y = y + x$  

  $x'$  masked with  $w$  **is**  $x!$   
 $y'$  masked with  $w$  **is**  $y!$

```
(x' y' w' : BV 65), w' <= 64 ->  
  let x := x' & (1 <<< w' - 1)  
  let y := y' & (1 <<< w' - 1)  
  x + y = y + x
```

**1sec**

"Bitwuzla, is this rewrite true?"

$(x\ y : BV\ 1),\ x + y = y + x$

"Bitwuzla, is this rewrite true?"

$(x\ y : BV\ 2),\ x + y = y + x$

"Bitwuzla, is this rewrite true?"

$(x\ y : BV\ 3),\ x + y = y + x$

"Bitwuzla, is this rewrite true?"

...

upto width 64 **18.5h**



One call for **all** widths, **simultaneously**



Bitwuzla is fast for one call!

# Everything Can Be Translated From Many Widths to One

Multi-width PBV	Mono-width PBV $\llbracket \cdot \rrbracket$
<i>Variable Declarations</i> (preconditions asserted)	
$\text{Var}(x_w) : \mathbb{N}$	$(x_w : \text{BitVec } o)$
$\text{Var}(x, w) : \text{BitVec } e$	$(x : \text{BitVec } o), \quad \text{assert } (x \ \&\llbracket e \rrbracket^{\text{mask}} = x)$
<i>Width Expressions</i> (widths stored directly as bitvectors)	
$\text{Var}(x_w)$	$x_w$
$\text{const}(c)$	$c$
$\text{min}(u, v)$	$\text{ite}(\llbracket u \rrbracket <_u \llbracket v \rrbracket, \llbracket u \rrbracket, \llbracket v \rrbracket)$
$\text{wadd}(u, v)$	$\llbracket u \rrbracket + \llbracket v \rrbracket$
<i>Bitvector Expressions</i> (mask result to width via $\llbracket w \rrbracket^{\text{mask}} = (1 \ll \llbracket w \rrbracket) - 1$ )	
$\text{Var}(a, w)$	$a$ (pre-masked by variable declaration)
$\text{Const}(x, w)$	$x \ \&\llbracket w \rrbracket^{\text{mask}}$
$\text{add}(a, b, w)$	$(\llbracket a \rrbracket + \llbracket b \rrbracket) \ \&\llbracket w \rrbracket^{\text{mask}}$
$\text{nand}(a, b, w)$	$\sim (\llbracket a \rrbracket \ \&\llbracket b \rrbracket) \ \&\llbracket w \rrbracket^{\text{mask}}$
$\text{zext}(a, v, w)$	$\llbracket a \rrbracket \ \&\llbracket w \rrbracket^{\text{mask}}$
$\text{sxt}(a, v, w)$	<b>let</b> $s = \text{ite}(\text{msb}_v(\llbracket a \rrbracket), \sim 0, 0)$ (sign-extend fill) <b>let</b> $\text{ext} = s \ \&\sim \llbracket v \rrbracket^{\text{mask}}$ (bits above width $v$ ) $(\llbracket a \rrbracket \parallel \text{ext}) \ \&\llbracket w \rrbracket^{\text{mask}}$
$\text{append}(a, v, b, w)$	$(\llbracket a \rrbracket \ll \llbracket w \rrbracket) \parallel \llbracket b \rrbracket$
<i>Predicate Expressions</i>	
$\text{eq}(a, b, w)$	$\llbracket a \rrbracket = \llbracket b \rrbracket$ (inputs pre-masked)
$\text{ult}(a, b, w)$	$\llbracket a \rrbracket <_u \llbracket b \rrbracket$ (inputs pre-masked)
$\text{slt}(a, b, w)$	$\text{ite}(\text{msb}_w(\llbracket a \rrbracket) = \text{msb}_w(\llbracket b \rrbracket),$ $\text{msb}_w(\llbracket a \rrbracket) \oplus \text{ult}(\llbracket a \rrbracket, \llbracket b \rrbracket),$ $\text{msb}_w(\llbracket a \rrbracket))$
$\text{and}(p, q)$	$\llbracket p \rrbracket \wedge \llbracket q \rrbracket$
$\text{or}(p, q)$	$\llbracket p \rrbracket \vee \llbracket q \rrbracket$
$\text{weq}(u, v)$	$\llbracket u \rrbracket = \llbracket v \rrbracket$
$\text{wlt}(u, v)$	$\llbracket u \rrbracket <_u \llbracket v \rrbracket$

where  $\text{msb}_w(a) \triangleq (a \ \&\llbracket w \rrbracket \gg 1) \neq 0$ .

The  $\text{PBV}_1$  encoding is equisatisfiable with the  $\text{PBV}_n$  encoding:

**Proof!** 

**THEOREM 4.1.** *If  $p$  is a  $\text{PBV}_n$  formula, then  $\models p$  iff  $\models \llbracket p \rrbracket$ .*

**PROOF.** *Left to right direction.* Let  $\rho$  be a model of  $p$ . We begin by constructing a model  $\bar{\rho}$  for  $\llbracket p \rrbracket$ . In particular,  $\rho$  associates each width variable  $x_w^1, \dots, x_w^n$  with a natural number  $\rho(x_w^1), \dots, \rho(x_w^n)$ . We define  $\bar{\rho}$  to map the unique width variable  $o$  of  $\llbracket p \rrbracket$  to  $\max\{\rho(x_w^1), \dots, \rho(x_w^n)\}$ , and map the bitvector variables of  $\llbracket p \rrbracket$  as follows:

$$\bar{\rho}(x_w^i) = \rho(x_w^i) \text{ as a bitvector of size } \bar{\rho}(o)$$

$$\bar{\rho}(x_i) = \rho(x_i) \text{ zero-extended to size } \bar{\rho}(o).$$

We check that all the variables  $x_i$  of width  $w$  satisfy the condition  $\rho(x_i) \ \&\ (2^{\rho(w)} - 1) = \rho(x_i)$  since they are obtained via zero extension. We then prove that for a well-typed width expression  $w$ ,  $\rho(w)$  is equal to the bitvector  $\bar{\rho}(\llbracket w \rrbracket)$  interpreted as a natural number. This is a straightforward induction on the typing derivation.

The crux of the proof is to relate the bitvector expressions: for any bitvector expression  $a$  of

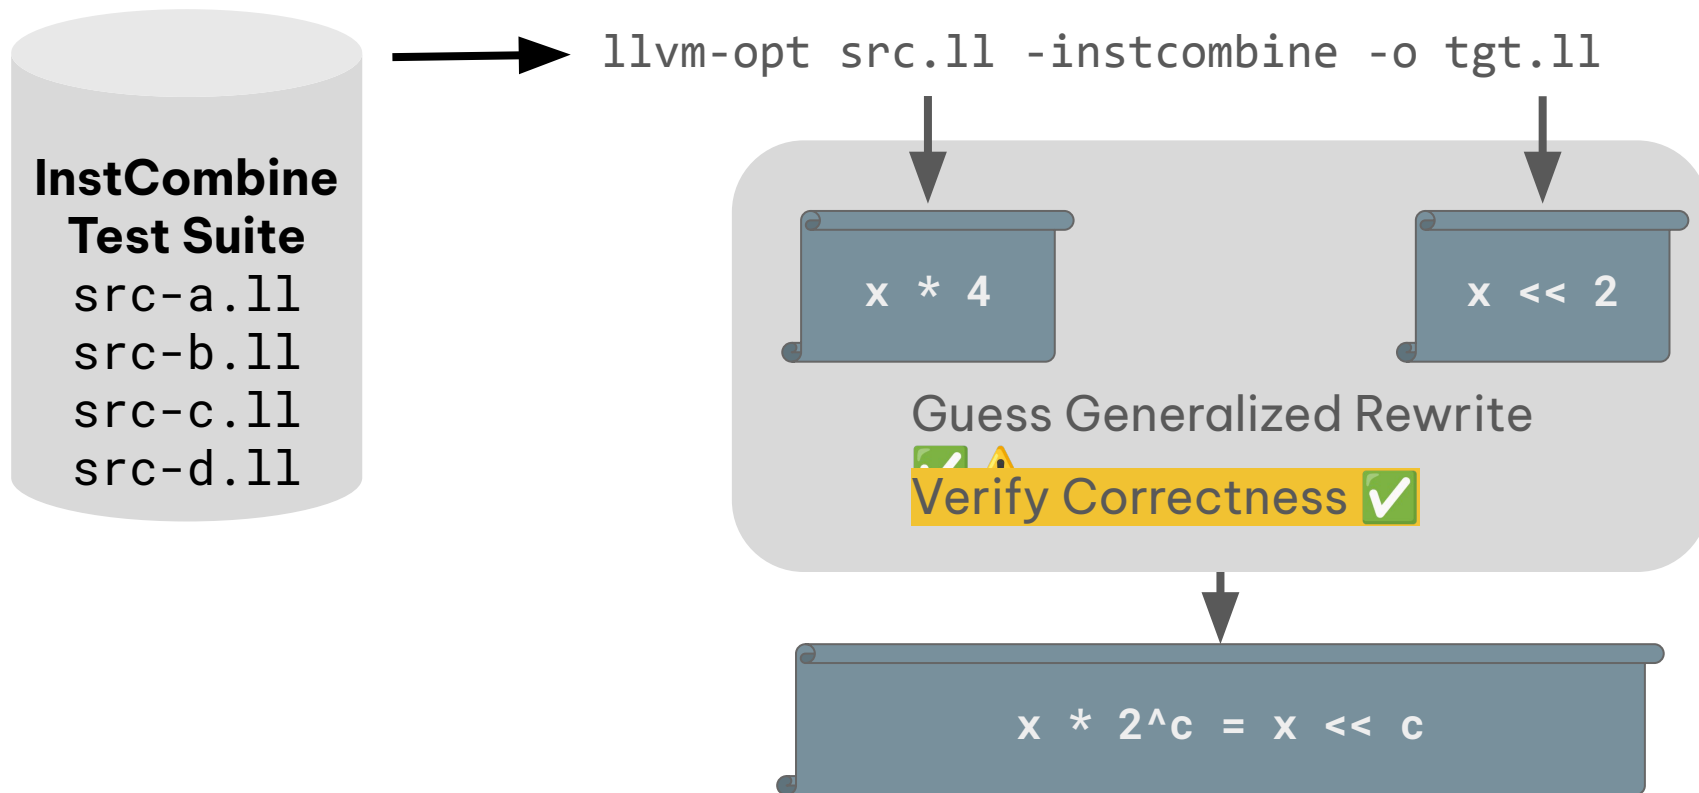


One call for **all** widths, **simultaneously**

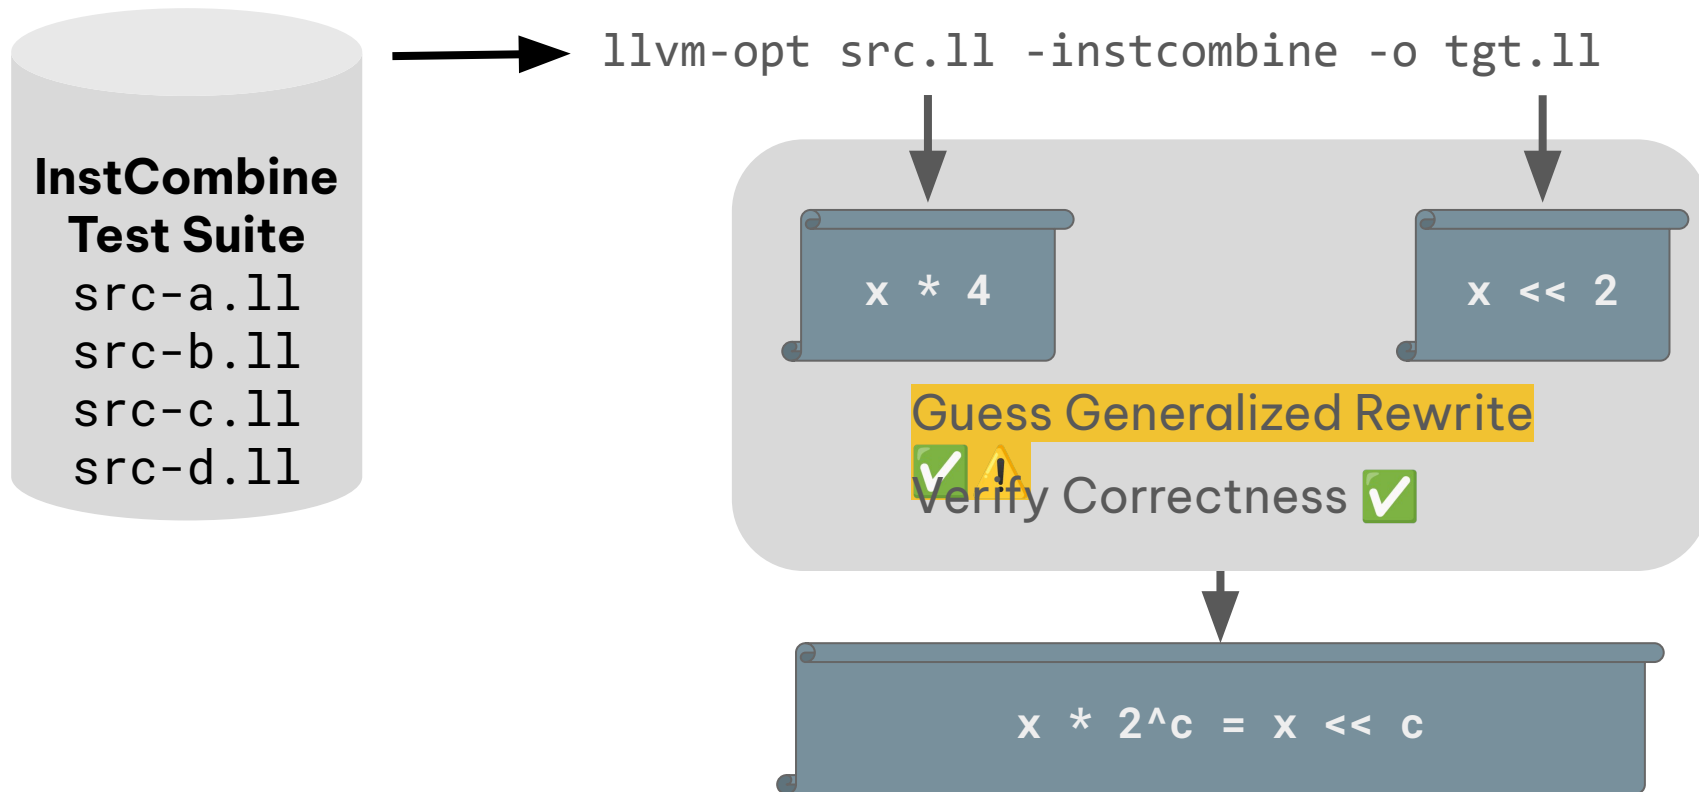


Bitwuzla is fast for one call!

# Idealized Pipeline For Extracting Declarative Rewrites



# Idealized Pipeline For Extracting Declarative Rewrites



# Guess Generalization

# Guess Generalization By Prayer To LLM



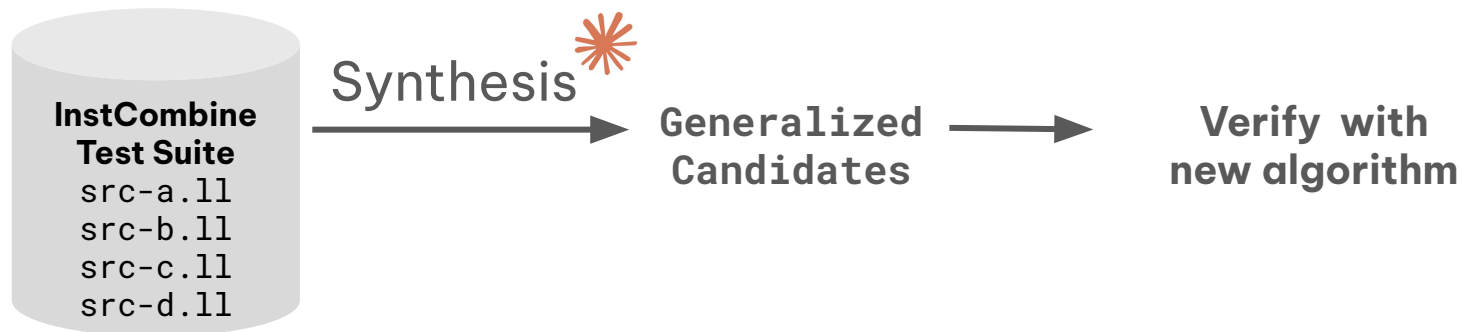
**John Regehr** 23:16

yeah, program synthesis as a research area is basically done-- our enumerative work is totally obsolete except in narrow circumstances

I think the verification side of synthesis is where we should focus all of our efforts now




# LLM-Driven Synthesis, Algorithmic Verification



Start with LLVM rewrite problems from lean-mlir (**1553** problems of the form LHS  $\rightarrow$  RHS)

Ask the agentic harness (claude code, Opus 4.6) to sort these by "clean rewrites"

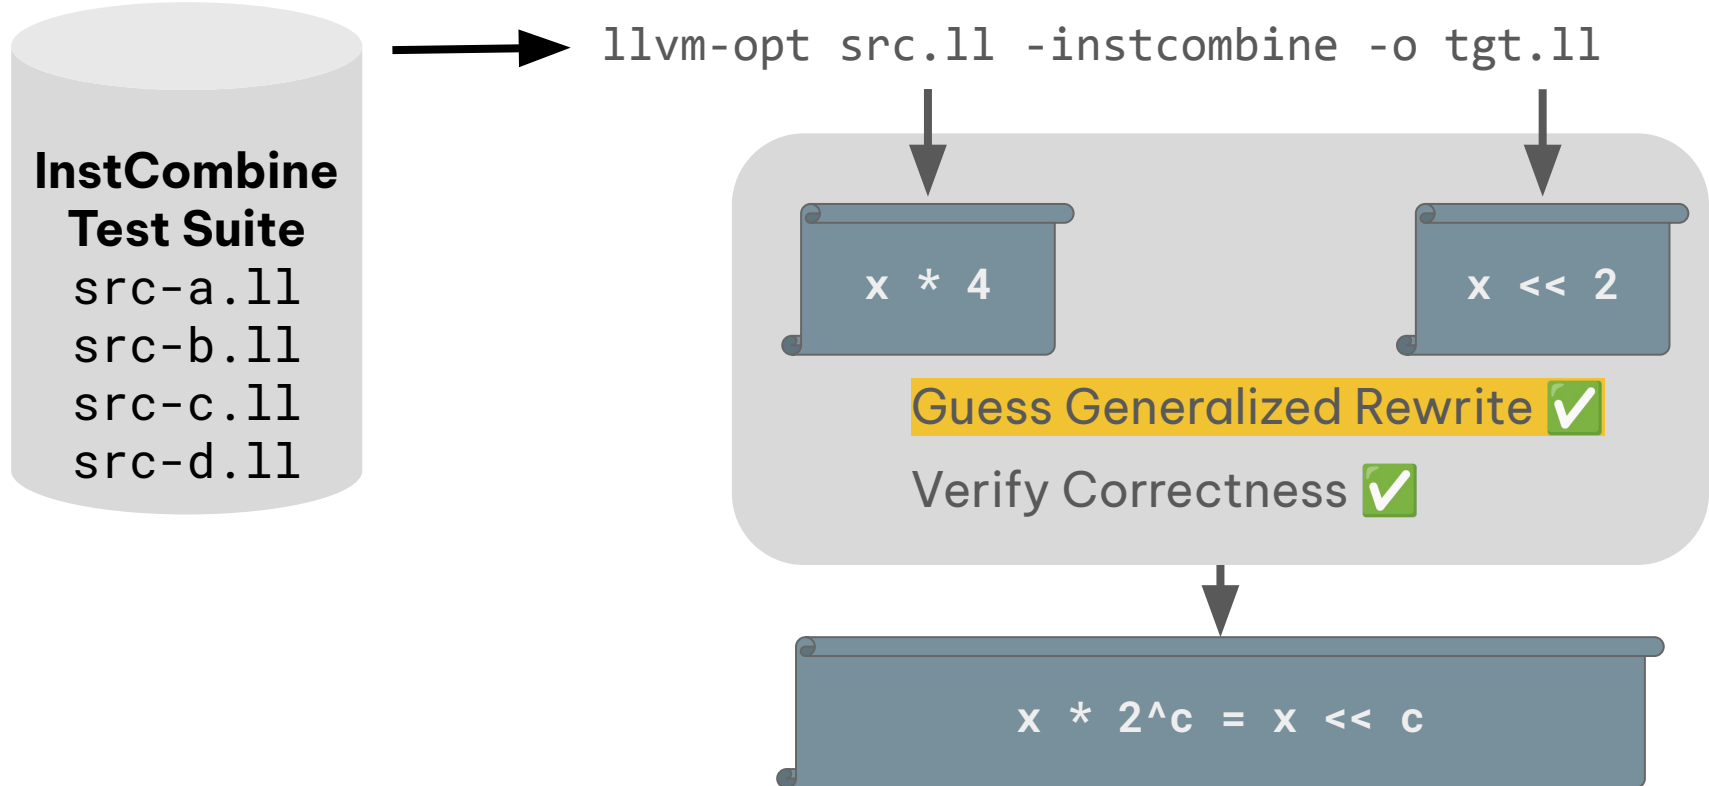
Ask the agentic harness to generalize these clean rewrites

- + tool use of our verifier given to agent, agent uses tool repeatedly
- + ( fast  $\Rightarrow$  no bottleneck here, catches hallucinations quickly)

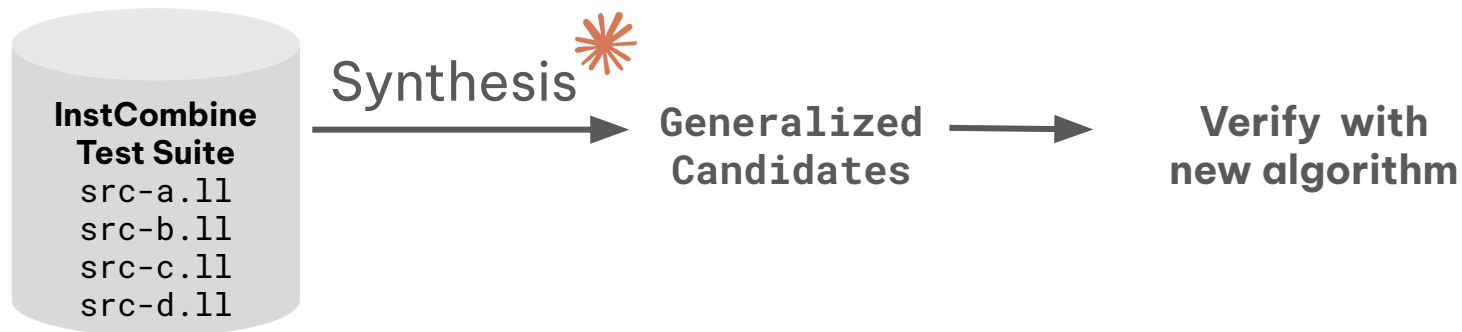
Produced 309 "candidate rewrites" after burning 1 week of CC subscription in 1 hr 

Check that 309 candidates are legal, giving **309** verified rewrites

# Idealized Pipeline For Extracting Declarative Rewrites



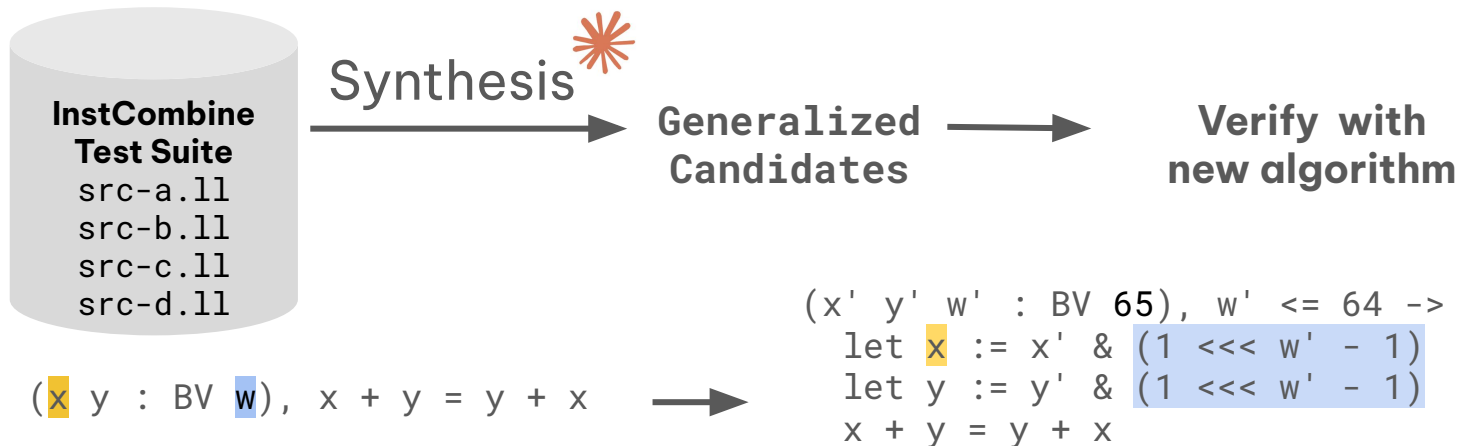
# What More Do We Need To Scale This Up?



- 1) Extracted Rewrites are in declarative SMT-LIB format.  
How to teach LLVM to consume and execute these as part of the rewriter?
- 2) Need to run this pipeline at scale, I ran out of credits for thinking.  
Given enough thinking tokens, maybe everything will be generalized?
- 3) Support for nsw/nuw, KnownBits, floating point, ...

# LLM-Driven Synthesis, Algorithmic Verification

## New Algorithms To Move Toward Declarative Compilers!



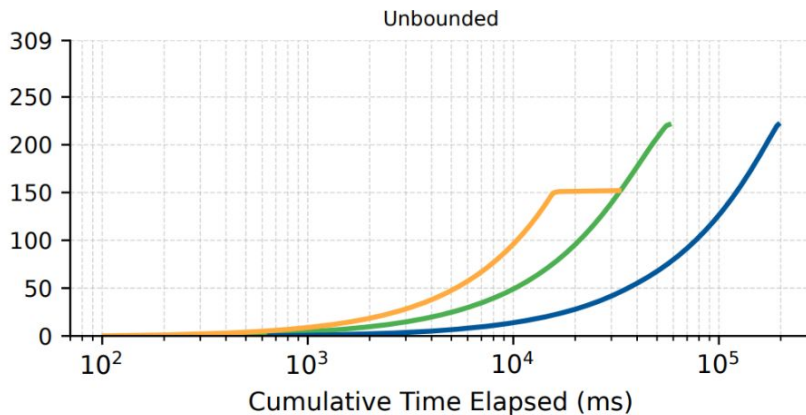
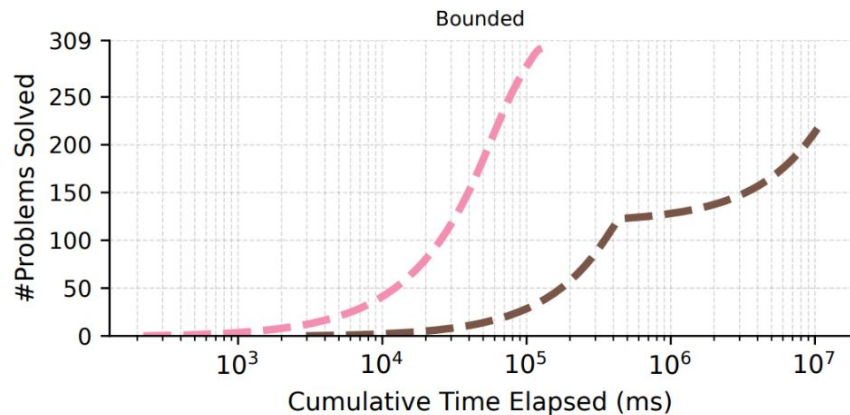
**Blase - Blaster for StrEams** [github.com/opencomp1/lean-mlir/tree/main/Blase](https://github.com/opencomp1/lean-mlir/tree/main/Blase)

A decision procedure that is sound and complete for parametric bitvector expressions with linear arithmetic, bitwise operations, zeroExtend, signExtend, and bitwidth constraints.

To use the bleeding edge of `Blase` your project, add the following to your `lakefile.toml`:

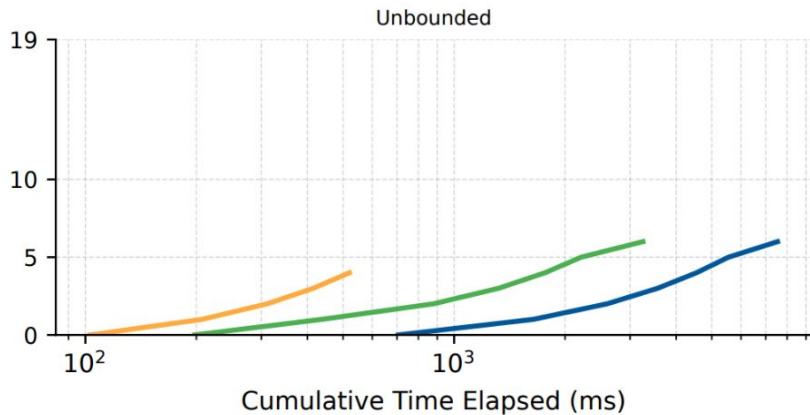
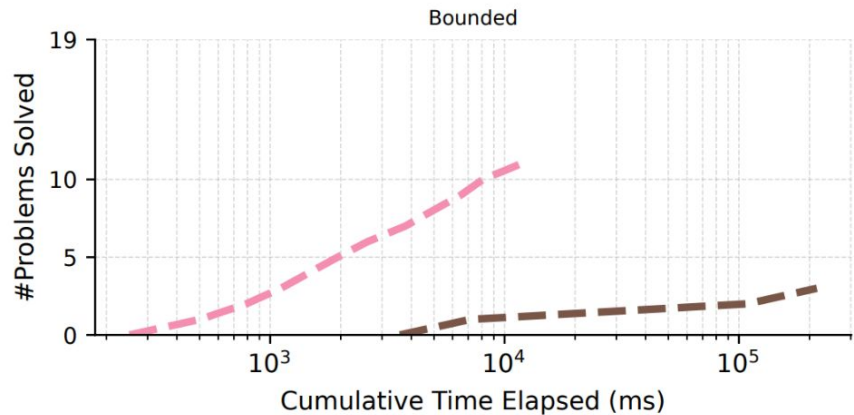
```
[[require]]  
name = "Blase"  
git = { url = "https://github.com/opencomp1/lean-mlir", subDir = "Blase/" }  
rev = "main"
```





— BMC-Naive (b=64)   
 - - BMC-Ours (b=64)   
 — Unbounded-Ours-k-ind   
 — Unbounded-Ours-ric3   
 — CVC5-berger

(b) InstCombine (309 multi-width problems). BMC-Ours saturates the dataset despite multiple widths; BMC-Naive collapses from enumeration blowup.



— BMC-Naive (b=64)   
 - - BMC-Ours (b=64)   
 — Unbounded-Ours-k-ind   
 — Unbounded-Ours-ric3   
 — CVC5-berger

(c) ROVER (19 problems, up to 12 symbolic widths). Naive enumeration is completely impractical; BMC-Ours and the unbounded solvers remain effective.



## But Why 32? Because Can Only Solve Concrete Width

```
define i32 @src(i32) {          (set-logic QF_UFBV)
  %r = udiv i32 %0, 8192
  ret i32 %r
}                                (define-fun src
                                ((x (_ BitVec 32)))
                                (_ BitVec 32)
                                (bvudiv x (_ bv32 8192)))
define i32 @tgt(i32) {        (define-fun tgt
                                ((x (_ BitVec 32)))
                                (_ BitVec 32)
                                (bvlshr x (_ bv32 32)))
  %r = lshr i32 %0, 13
  ret i32 %r
}
```



# BV SMT Solver

Transformation doesn't verify!

ERROR: Value mismatch

Example:

```
i32 %#0 = #x00000001 (1)
```

Source:

```
i32 %r = #x00000001 (1)
```

Target:

```
i32 %r = #x00000000 (0)
```

```
Source value: #x00000001 (1)
```

```
Target value: #x00000000 (0)
```

