




eqsat



E-Graphs as a Persistent Compiler Abstraction



Jules Merckx
Ghent University



UNIVERSITY OF
CAMBRIDGE

Alexandre Lopoukhine
University of Cambridge



Jianyi Cheng
University of Edinburgh




Sam Coward
University College London


Bjorn De Sutter
Ghent University

Tobias Grosser
University of Cambridge

Code is transformed at different levels of abstraction

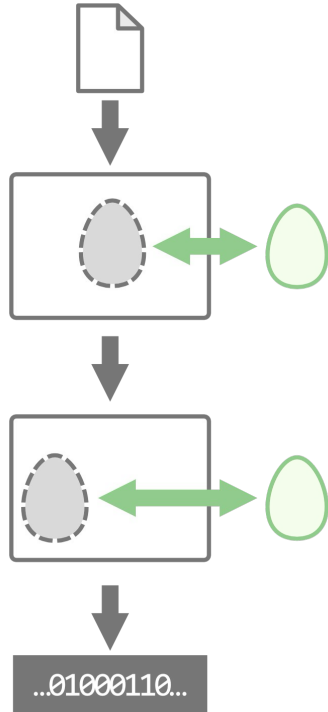


Egg is extensible but not integrated




Egg

Cranelift offers tight compiler integration



Egg




Cranelift


- Not at other levels of abstraction

+ Equality saturation is integrated


We aim to combine extensibility and integration



Egg

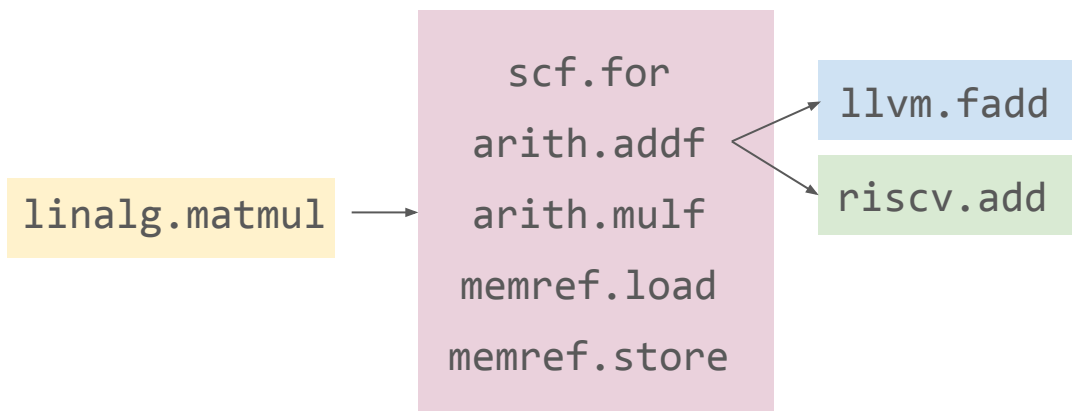


Cranelift



Our Work

MLIR: A User-Extensible Compiler Framework



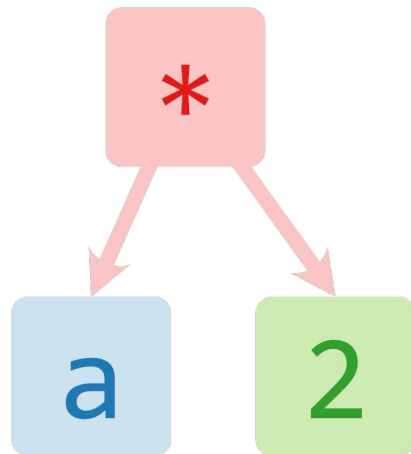
Start with arbitrary MLIR IR

```
func.func @f(%a : i64) -> i64 {  
  %two = arith.constant 2 : i64
```

```
  %res = arith.muli %a , %two : i64
```


```
  func.return %res : i64
```

```
}
```




Embed e-classes as MLIR operations

```
func.func @f(%a : i64) -> i64 {  
  %two = arith.constant 2 : i64  
  
  %c_two = eqsat.eclass %two : i64  
  %c_a = eqsat.eclass %a : i64  
  %res = arith.muli %c_a, %c_two : i64  
  
  %c_res = eqsat.eclass %res : i64  
  func.return %c_res : i64  
}
```



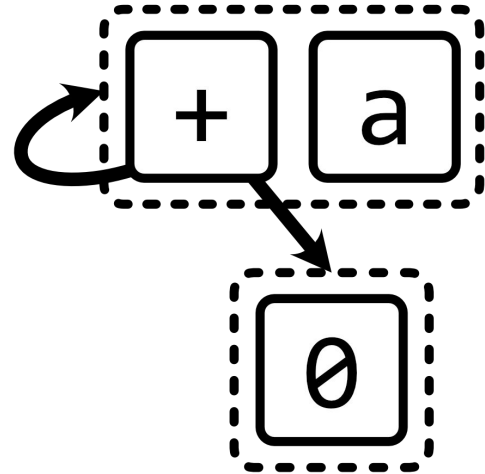
New e-nodes are introduced as operands to e-classes

```
func.func @f(%a : i64) -> i64 {  
  %two = arith.constant 2 : i64  
  %one = arith.constant 1  
  %c_one = eqsat.eclass %one  
  %c_two = eqsat.eclass %two : i64  
  %c_a = eqsat.eclass %a : i64  
  %res = arith.muli %c_a, %c_two : i64  
  %res1 = arith.shli %c_a, %c_one  
  %c_res = eqsat.eclass %res, %res1 : i64  
  func.return %c_res : i64  
}
```

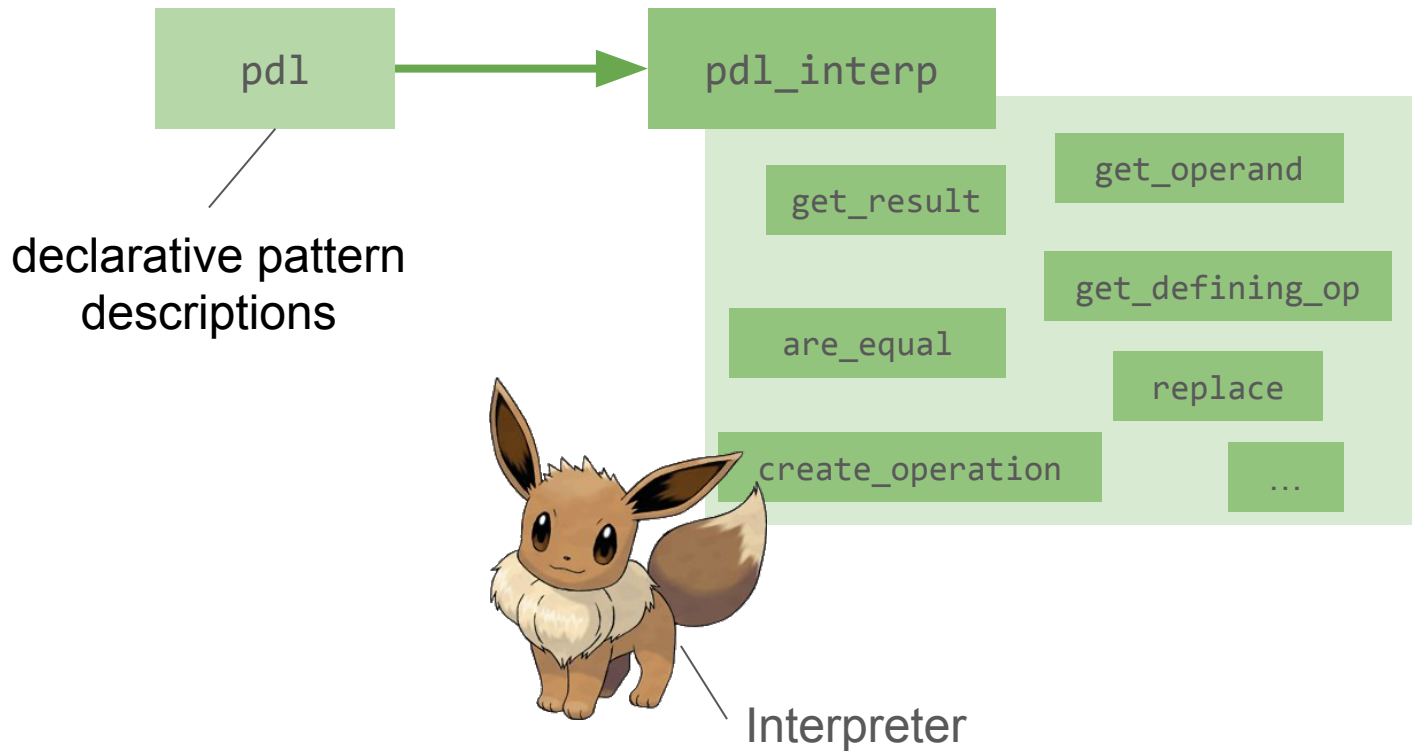


Cycles in graph regions

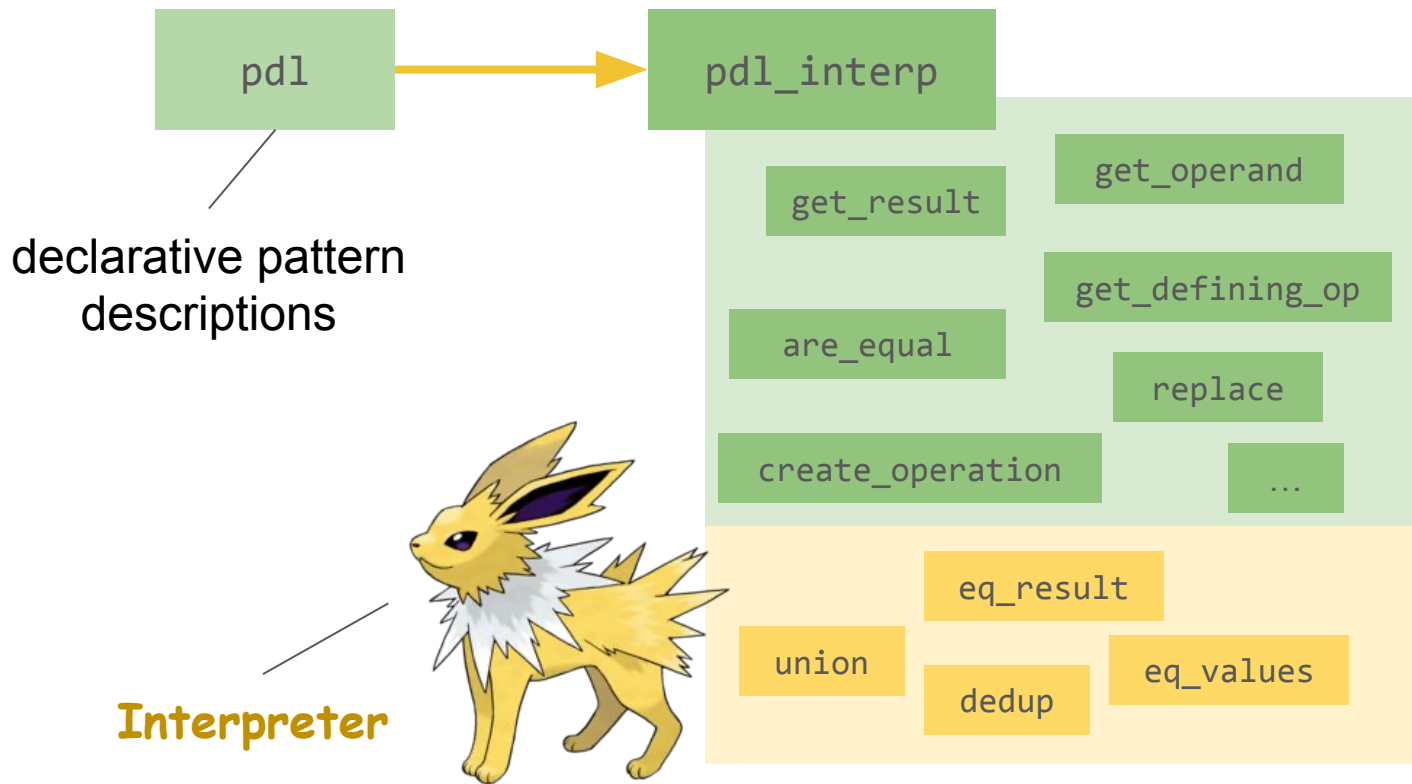
```
eqsat.egraph {  
    %c_zero = eqsat.eclass %zero  
    %sum = arith.addi %c_a, %c_zero  
    %c_a = eqsat.eclass %sum, %a  
}
```



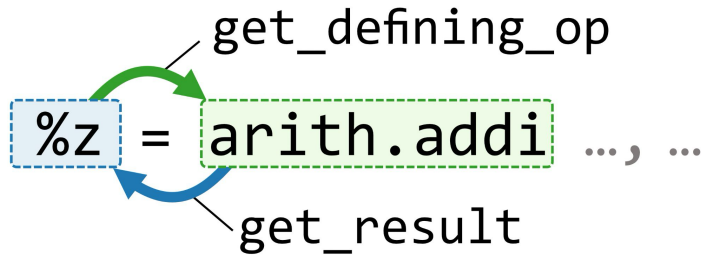
PDL: MLIR's declarative rewrite representation



We implement e-matching with small modifications



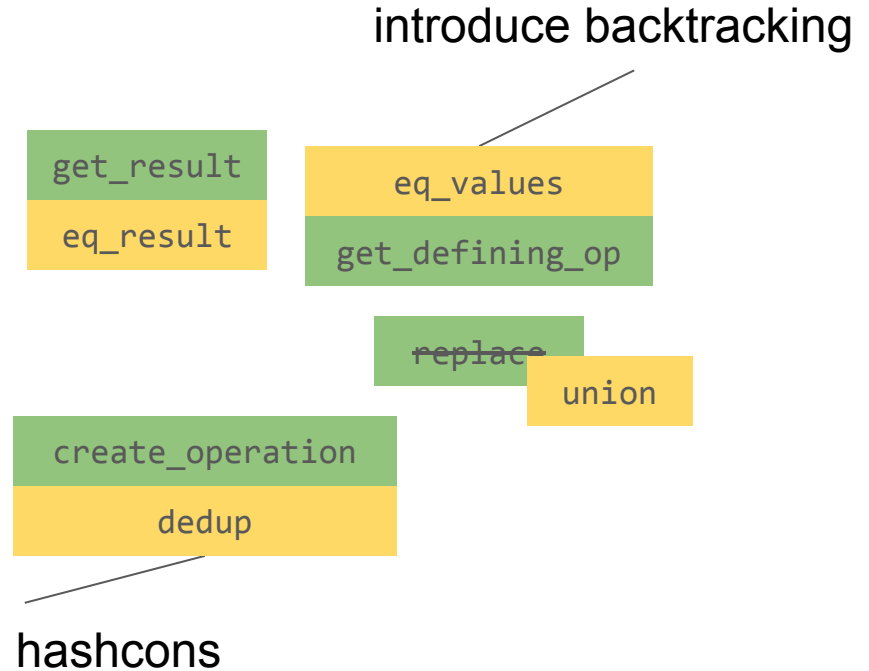
We have to take into account the use-def indirection



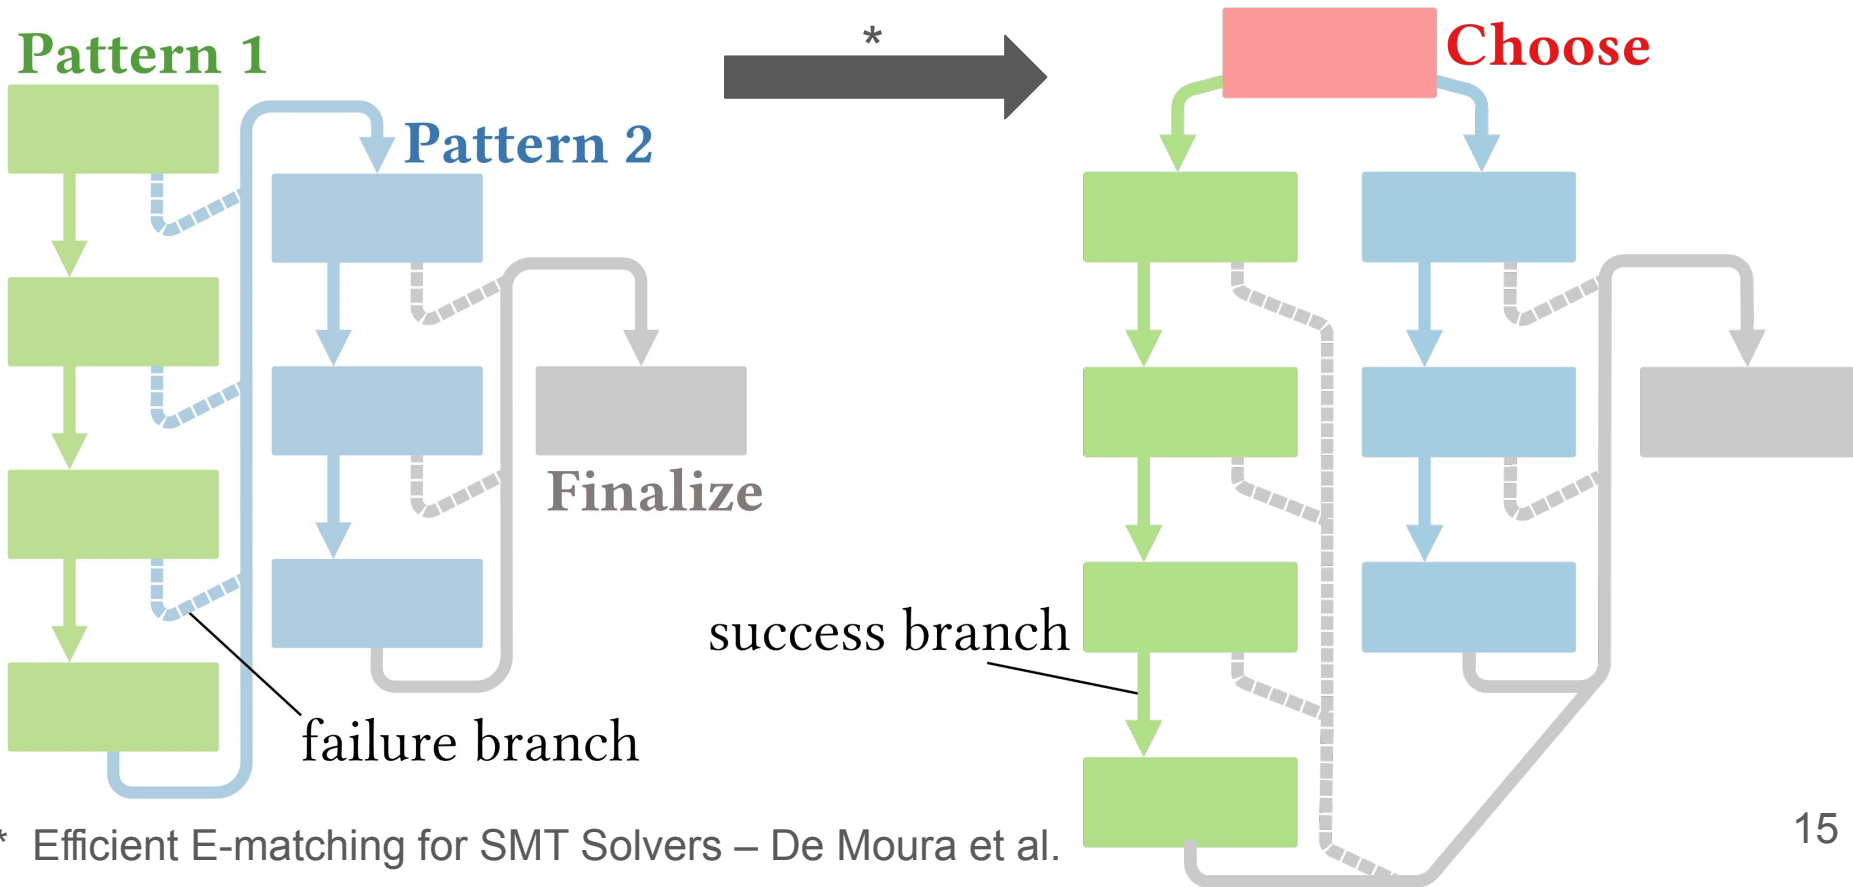
`get_result`

`get_defining_op`

We have to take into account the use-def indirection



Naive pdl_interp can be exponentially slow



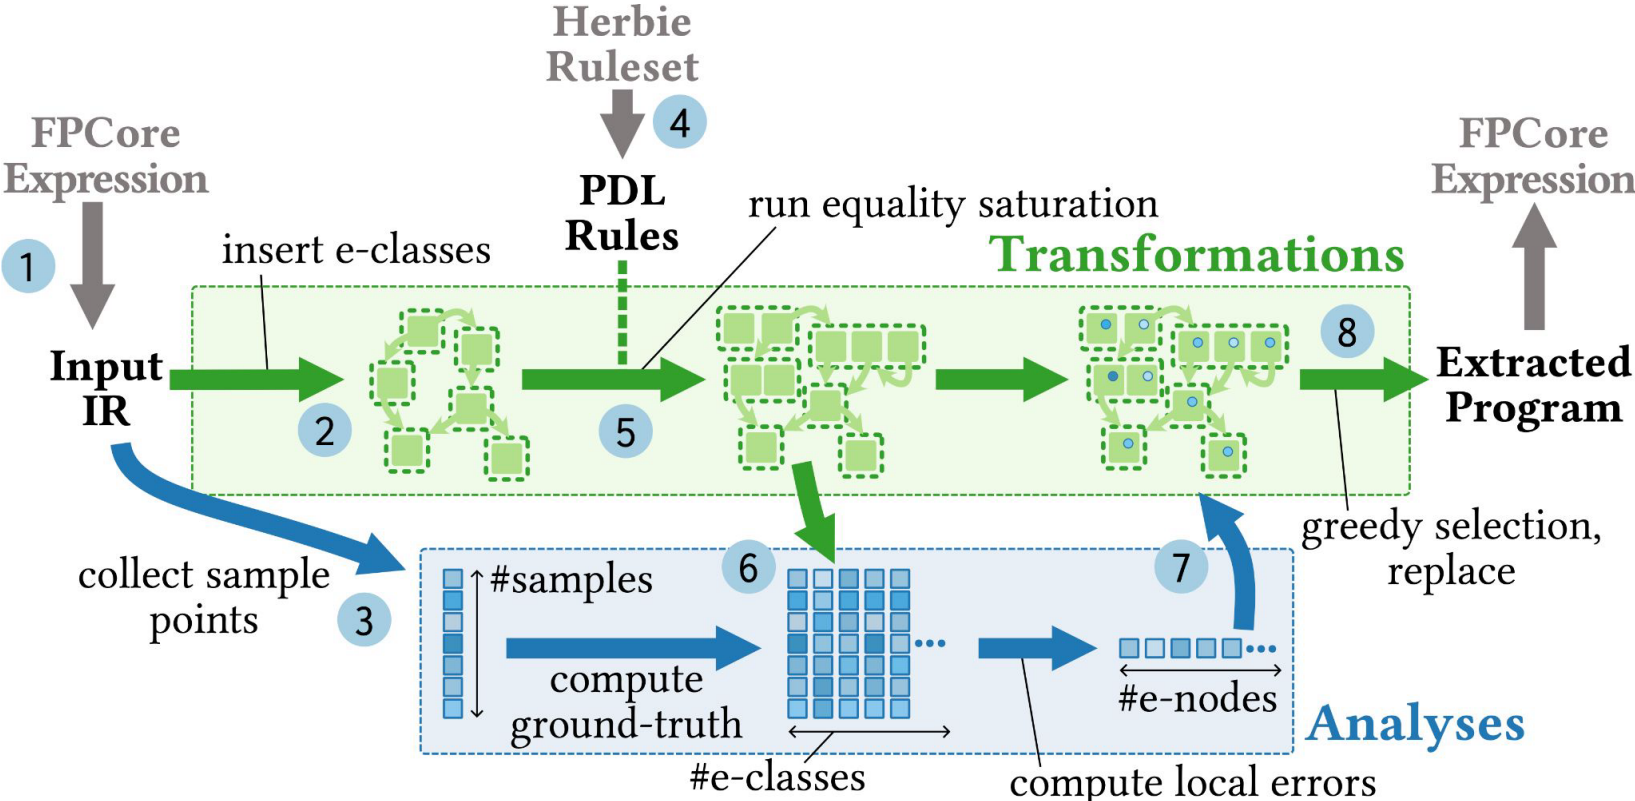
* Efficient E-matching for SMT Solvers – De Moura et al.

Reuse lattices for e-class analyses

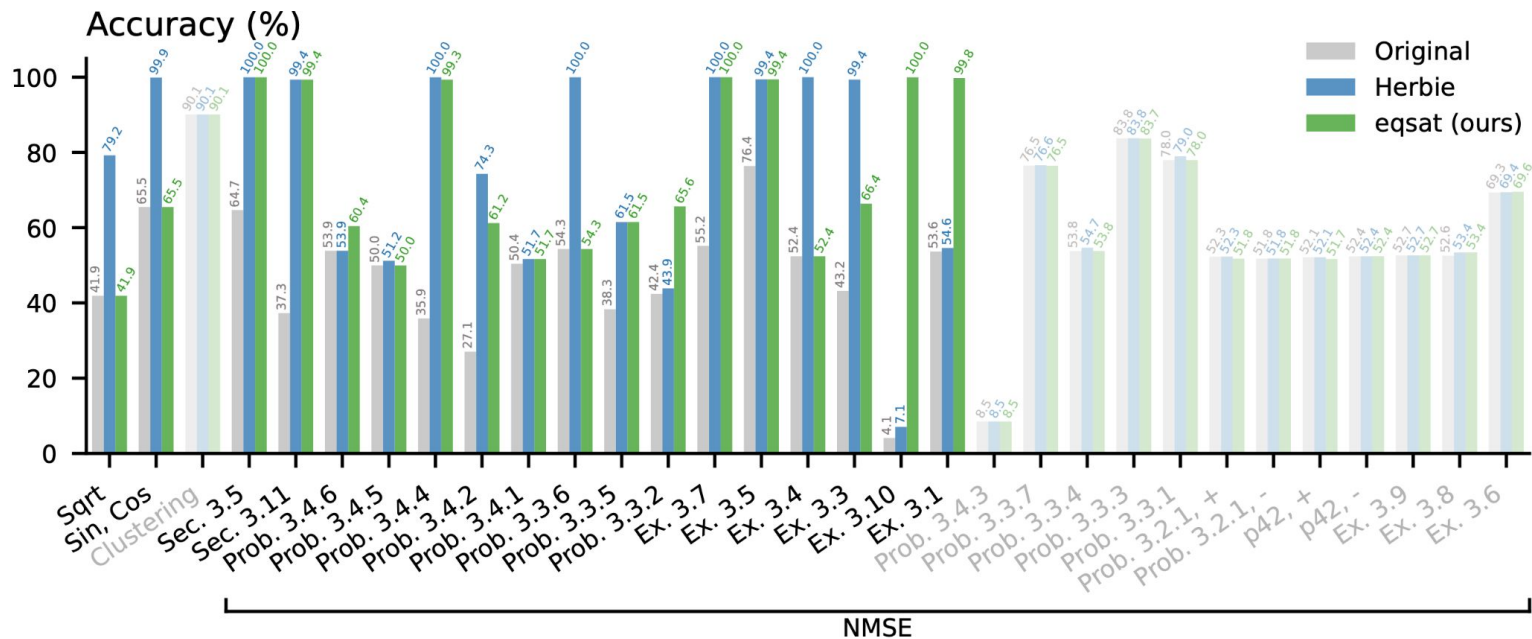
meet

In forward dataflow: join of eclass states is the meet of its operands


Proof of Concept: Herbie in MLIR



We match Herbie accuracy-wise



... But Much Slower*



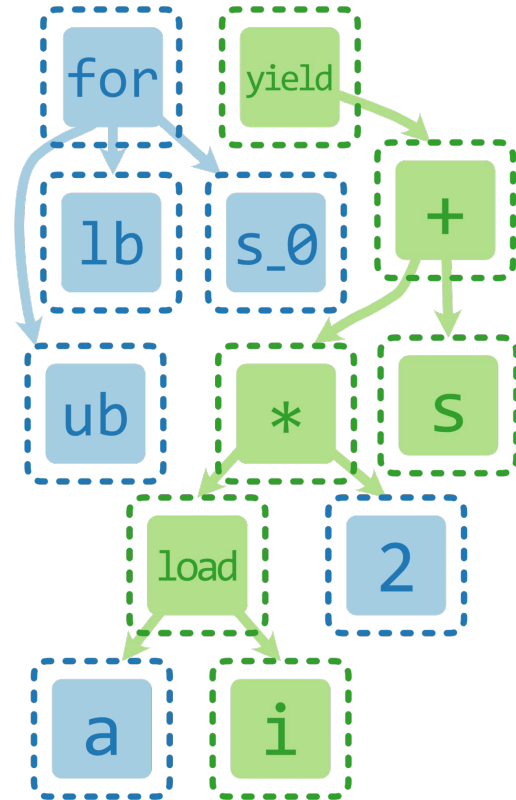
*Timing includes very inefficient arbitrary precision interval analysis

Control flow is kept intact

```
%s = scf.for %i = %lb to %ub
  iter_args(%s = %s_0) {
    %x = memref.load %a[%i]
    %term = arith.mulf %x, %two
    %s_new = arith.addf %s, %term
    scf.yield %s_new
  }
```

scoped hashcons?

Relation not visible in
regular e-graph

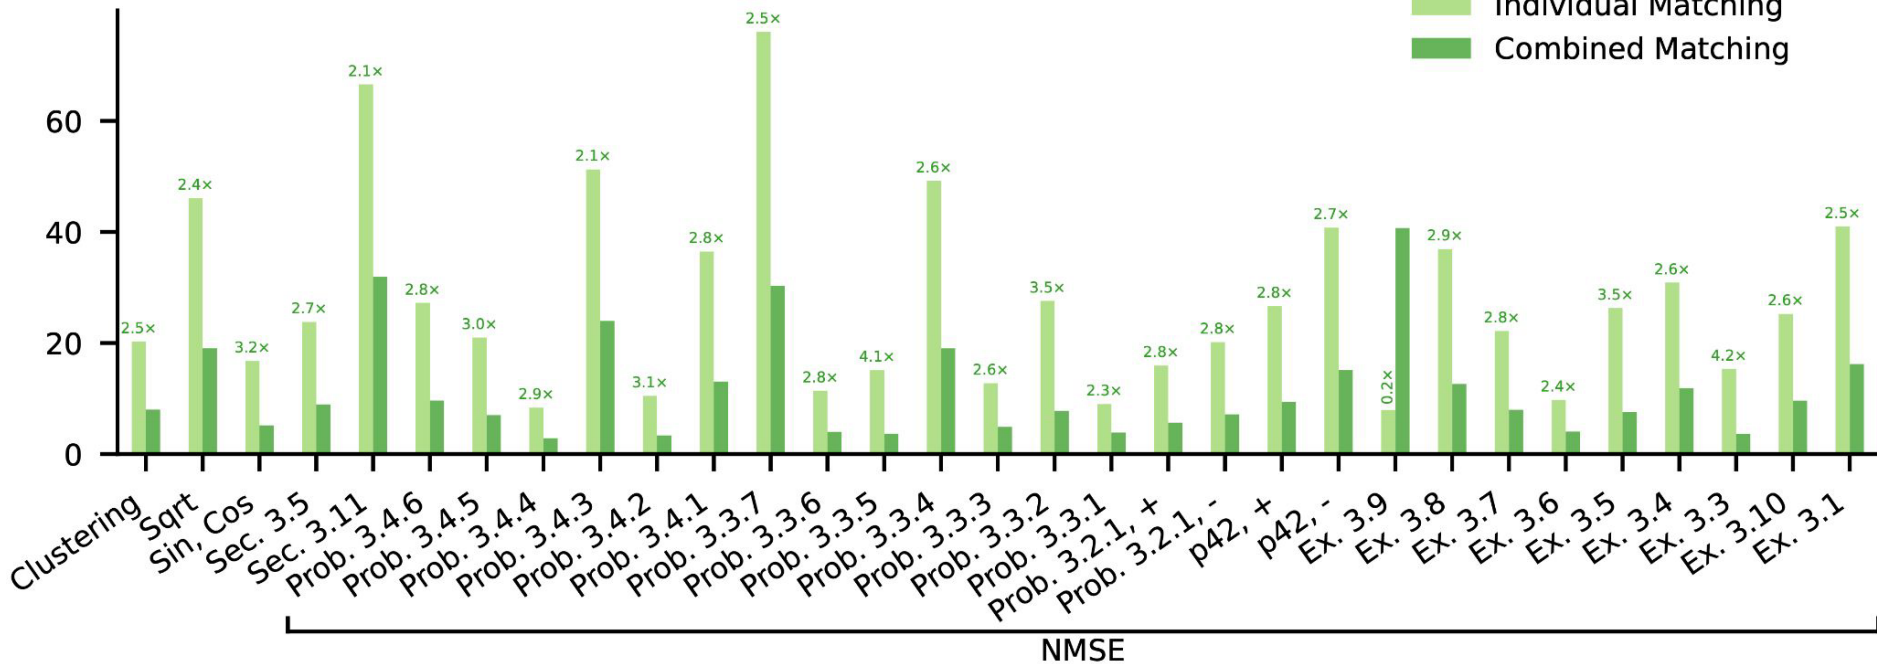


Things we want to work on

- Cranelift-style acyclic e-graphs for instruction selection
- Large scale tensor-program optimizations
- Control-flow rewrites
- “Partial” e-graphs (only put some terms in an e-class)
 - Automatic e-graph insertion
- E-matching optimizations

Time (s)

Individual Matching
Combined Matching



? Extending `pd1_interp` for region matching

? How to handle operations without results

```
memref.store %val, %A[%a, %b] : memref<8x?xi32, #layout>
```

? Which compiler passes benefit most

- Function call inlining
- Multi-level analyses
- ...

? Applications

- Herbie-like floating point optimization
- Multi-level, end-to-end compilation
- ...