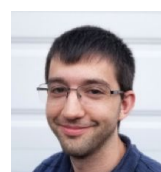
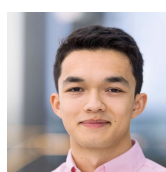
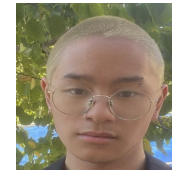
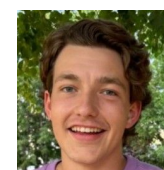
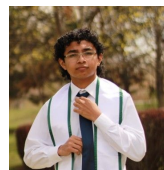
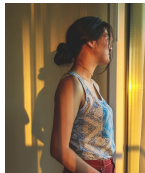


An abstract painting with vibrant colors like red, blue, yellow, and green, featuring thick brushstrokes and a central dark circular shape.

Relational Equality Saturation: *E-graphs meet Query Engines*

Zachary Tatlock (UW Allen School) @ NWDS 2026-03-13





Kernighan's Law:

- Debugging is twice as hard as writing the code in the first place.

Kernighan's Law:

- Debugging is twice as hard as writing the code in the first place.
- Therefore, if you write the code as cleverly as possible, you are, by definition, not smart enough to debug it.

Kernighan's Law:

- Debugging is twice as hard as writing the code in the first place.
- Therefore, if you write the code as cleverly as possible, you are, by definition, not smart enough to debug it.

Engineer complex systems
from **simple** and **modular** pieces!

PL Problems

1. Phase Ordering
2. Syntactic Brittleness
3. ...

$$(a * 2) / 2 \Rightarrow a$$

$$(a * 2) / 2 \Rightarrow a$$

REWRITE!

$$(a * 2) / 2 \Rightarrow a$$

REWRITE!

Useful

$$(x * y) / z = x * (y / z)$$

$$x / x = 1$$

$$x * 1 = x$$

$$(a * 2) / 2 \Rightarrow a$$

REWRITE!

Useful

$$(x * y) / z = x * (y / z)$$

$$x / x = 1$$

$$x * 1 = x$$

Less Useful

$$x * 2 = x \ll 1$$

$$x * y = y * x$$

$$x = x * 1$$

$$(a * 2) / 2$$

“happy path”

$$(x * y) / z = x * (y / z)$$

$$x / x = 1$$

$$x * 1 = x$$

$$(a * 2) / 2 \Rightarrow a * (2 / 2)$$

“happy path”

$$(x * y) / z = x * (y / z)$$

$$x / x = 1$$

$$x * 1 = x$$

$$(a * 2) / 2 \Rightarrow a * (2 / 2) \Rightarrow a * 1$$

“happy path”

$$(x * y) / z = x * (y / z)$$

$$x / x = 1$$

$$x * 1 = x$$

$$(a * 2) / 2 \Rightarrow a * (2 / 2) \Rightarrow a * 1 \Rightarrow a$$

“happy path”

$$(x * y) / z = x * (y / z)$$

$$x / x = 1$$

$$x * 1 = x$$

$$(a * 2) / 2 \Rightarrow a * (2 / 2) \Rightarrow a * 1 \Rightarrow a$$

“happy path”

$$(x * y) / z = x * (y / z)$$

$$x / x = 1$$

$$x * 1 = x$$



$(a * 2) / 2$

Pitfalls

$x * 2 = x \ll 1$

$x * y = y * x$

$x = x * 1$

$$(a * 2) / 2 \Rightarrow (a \ll 1) / 2$$

Pitfalls

$$x * 2 = x \ll 1$$

$$x * y = y * x$$

$$x = x * 1$$

$(a * 2) / 2 \Rightarrow (a \ll 1) / 2$  order

Pitfalls

$x * 2 = x \ll 1$

$x * y = y * x$

$x = x * 1$

$(a * 2) / 2 \Rightarrow (a \ll 1) / 2$  order

$(a * 2) / 2$

Pitfalls

$x * 2 = x \ll 1$

$x * y = y * x$

$x = x * 1$

$(a * 2) / 2 \Rightarrow (a \ll 1) / 2$  order

$(a * 2) / 2 \Rightarrow (2 * a) / 2$

Pitfalls

$x * 2 = x \ll 1$

$x * y = y * x$

$x = x * 1$

$(a * 2) / 2 \Rightarrow (a \ll 1) / 2$  order

$(a * 2) / 2 \Rightarrow (2 * a) / 2 \Rightarrow (a * 2) / 2$


Pitfalls

$x * 2 = x \ll 1$

$x * y = y * x$

$x = x * 1$

$(a * 2) / 2 \Rightarrow (a \ll 1) / 2$  order

$(a * 2) / 2 \Rightarrow (2 * a) / 2 \Rightarrow (a * 2) / 2$ 

diverge


Pitfalls

$x * 2 = x \ll 1$

$x * y = y * x$

$x = x * 1$

$(a * 2) / 2 \Rightarrow (a \ll 1) / 2$  order

$(a * 2) / 2 \Rightarrow (2 * a) / 2 \Rightarrow (a * 2) / 2$ 

diverge

a


Pitfalls

$x * 2 = x \ll 1$

$x * y = y * x$

$x = x * 1$

$(a * 2) / 2 \Rightarrow (a \ll 1) / 2$  order

$(a * 2) / 2 \Rightarrow (2 * a) / 2 \Rightarrow (a * 2) / 2$ 

diverge

$a \Rightarrow a * 1$


Pitfalls

$x * 2 = x \ll 1$

$x * y = y * x$

$x = x * 1$

$(a * 2) / 2 \Rightarrow (a \ll 1) / 2$  order

$(a * 2) / 2 \Rightarrow (2 * a) / 2 \Rightarrow (a * 2) / 2$ 

diverge

$a \Rightarrow a * 1 \Rightarrow a * 1 * 1$


Pitfalls

$x * 2 = x \ll 1$

$x * y = y * x$

$x = x * 1$

$(a * 2) / 2 \Rightarrow (a \ll 1) / 2$  order

$(a * 2) / 2 \Rightarrow (2 * a) / 2 \Rightarrow (a * 2) / 2$ 

diverge

$a \Rightarrow a * 1 \Rightarrow a * 1 * 1 \Rightarrow \dots$  infinite size



Pitfalls

$x * 2 = x \ll 1$

$x * y = y * x$

$x = x * 1$

$(a * 2) / 2 \Rightarrow (a \ll 1) / 2$  order

$(a * 2) / 2 \Rightarrow (2 * a) / 2 \Rightarrow (a * 2) / 2$ 


diverge

$a \Rightarrow a * 1 \Rightarrow a * 1 * 1 \Rightarrow \dots$  infinite size

Critical for other inputs!

Pitfalls

$$x * 2 = x \ll 1$$

$$x * y = y * x$$

$$x = x * 1$$

$$(a * 2) / 2 \Rightarrow a$$

Which rewrite? When?

Useful

$$(x * y) / z = x * (y / z)$$

$$x / x = 1$$

$$x * 1 = x$$

Less Useful

$$x * 2 = x \ll 1$$

$$x * y = y * x$$

$$x = x * 1$$

PL Problems

1. Phase Ordering 🤔
2. Syntactic Brittleness
3. ...

```
def foo(x):  
    y = abs(x)  
    return sqrt(y)
```

```
def foo(x):
```

```
    y = abs(x)
```

```
    return sqrt(y)
```

Safe? 🤔

```
def foo(x):
```

```
    y = abs(x)
```

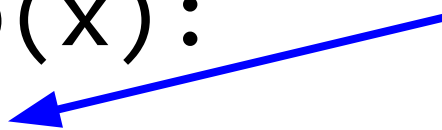
```
    return sqrt(y)
```

Safe? 🤔

Abstract Interpretation!

```
def foo(x):
```

$x \in [-\infty, \infty]$



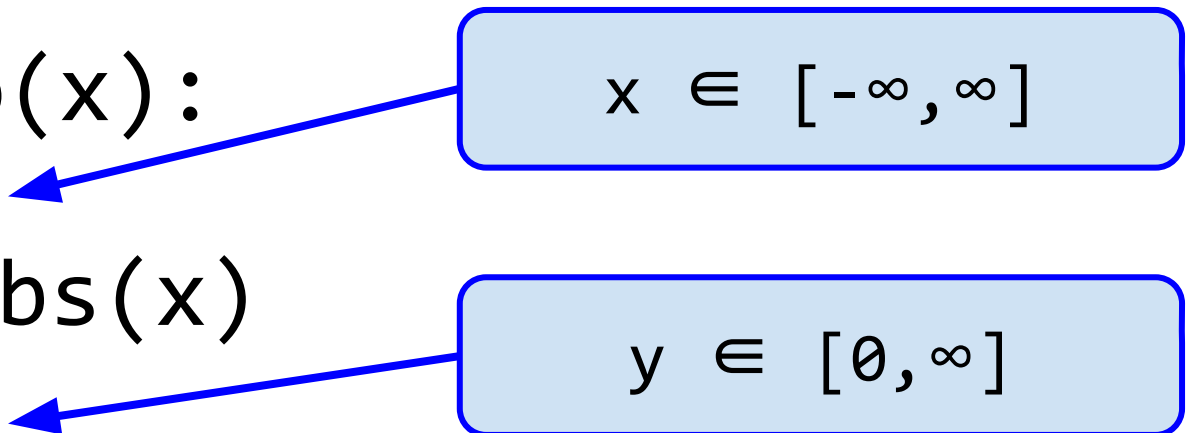
```
    y = abs(x)
```

```
    return sqrt(y)
```

Abstract Interpretation!

```
def foo(x):
```

$x \in [-\infty, \infty]$



```
    y = abs(x)
```

$y \in [0, \infty]$

```
    return sqrt(y)
```

Abstract Interpretation!

```
def foo(x):
```

$x \in [-\infty, \infty]$

```
    y = abs(x)
```

$y \in [0, \infty]$

```
    return sqrt(y)
```



Abstract Interpretation!

```
def foo(x):
```

$x \in [-\infty, \infty]$

```
    y = abs(x)
```

$y \in [0, \infty]$

```
    return sqrt(y)
```



BUT: abstraction can hide crucial facts!

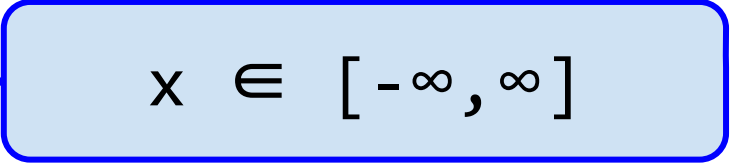
```
def foo(x):
```

```
    y = x - x
```

```
    return sqrt(y)
```

```
def foo(x):
```

$x \in [-\infty, \infty]$

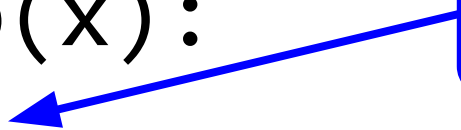


```
    y = x - x
```

```
    return sqrt(y)
```


```
def foo(x):
```

$x \in [-\infty, \infty]$



```
    y = x - x
```

$y \in [-\infty, \infty]$



```
    return sqrt(y)
```

```
def foo(x):
```

$x \in [-\infty, \infty]$

```
    y = x - x
```

$y \in [-\infty, \infty]$

```
    return sqrt(y)
```



```
def foo(x):
```

$$x \in [-\infty, \infty]$$

```
    y = x - x
```

$$y \in [-\infty, \infty]$$

```
    return sqrt(y)
```



OK, just preprocess? $x - x \Rightarrow 0$

```
def foo(x):
```

```
    y = x * x
```

```
    return sqrt(y)
```

```
def foo(x):
```

$x \in [-\infty, \infty]$

```
    y = x * x
```

$y \in [-\infty, \infty]$

```
    return sqrt(y)
```



UGH! Chasing our own tails!

PL Problems

1. Phase Ordering 🤔
2. Syntactic Brittleness 🤔
3. ...

PL Problems

1. Phase Ordering 🤔

2. Syntactic Brittleness 🤔

3. ...

Keep simple + modular!

But make robust?

PL Problems

1. Phase Ordering 🤔
2. Syntactic Brittleness 🤔
3. ...

Keep simple + modular!

But make robust?

DB techniques will
come to our rescue!

$$(a * 2) / 2 \Rightarrow a$$

Which rewrite? When?

Useful

$$(x * y) / z = x * (y / z)$$

$$x / x = 1$$

$$x * 1 = x$$

Less Useful

$$x * 2 = x \ll 1$$

$$x * y = y * x$$

$$x = x * 1$$

$$(a * 2) / 2 \Rightarrow a$$

Which rewrite? When?

Equality Saturation

Try applying all the rules in every order!

$$(a * 2) / 2 \Rightarrow a$$

Which rewrite? When?

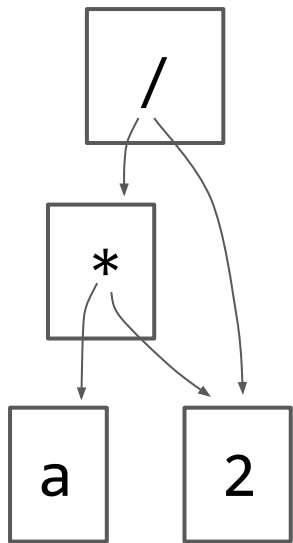
Equality Saturation

Try applying **all** the rules in **every** order?!

E-graphs

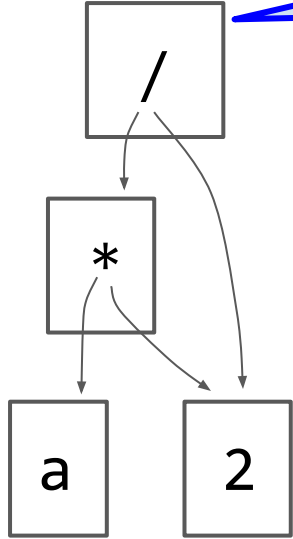
- Data structure from Greg Nelson's PhD thesis (1980)
- Used for congruence closure (Downey, Sethi, Tarjan 1980)
 - Intuition: union-find (Tarjan 1975) but function-aware
- Key for equality and uninterpreted funcs (EUF) theory in SMT
 - Intuition: the “glue” that connects other theories to SAT
- Historically: “baked in” to SMT solvers, no general libraries 😞

E-graphs

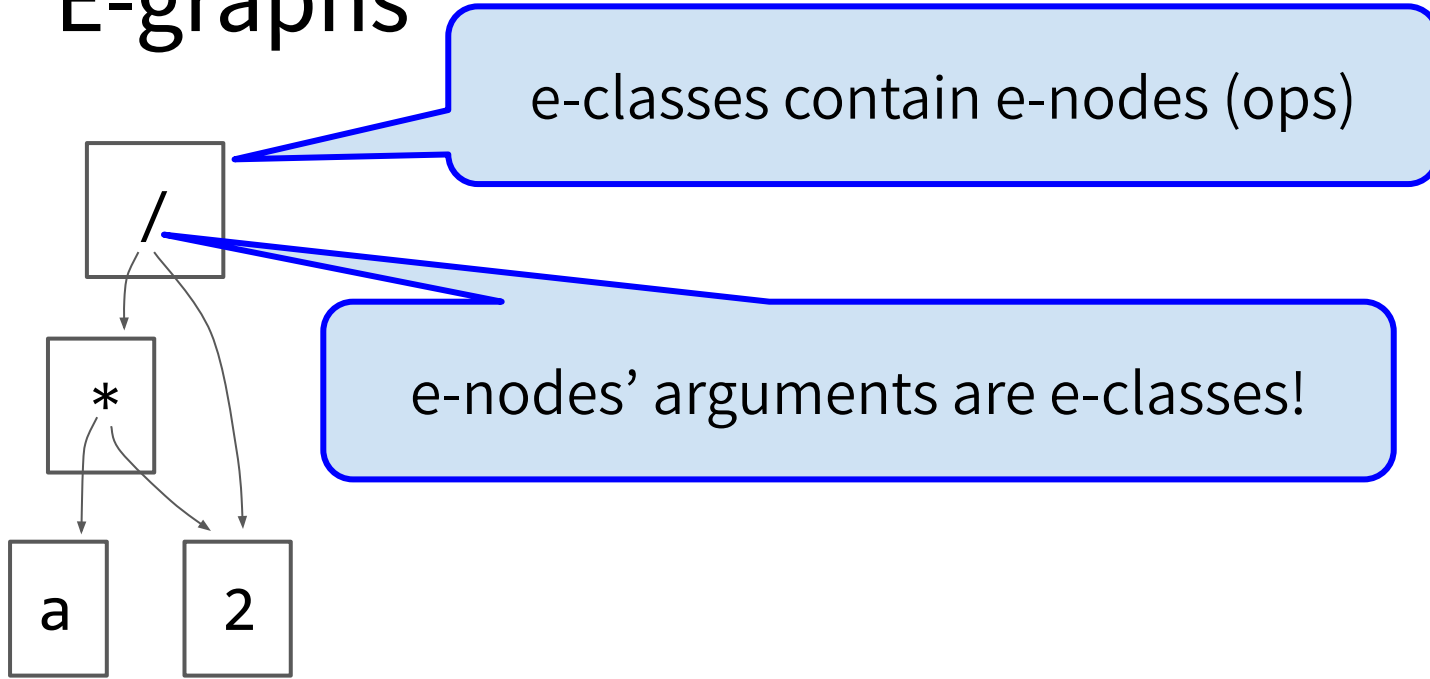


E-graphs

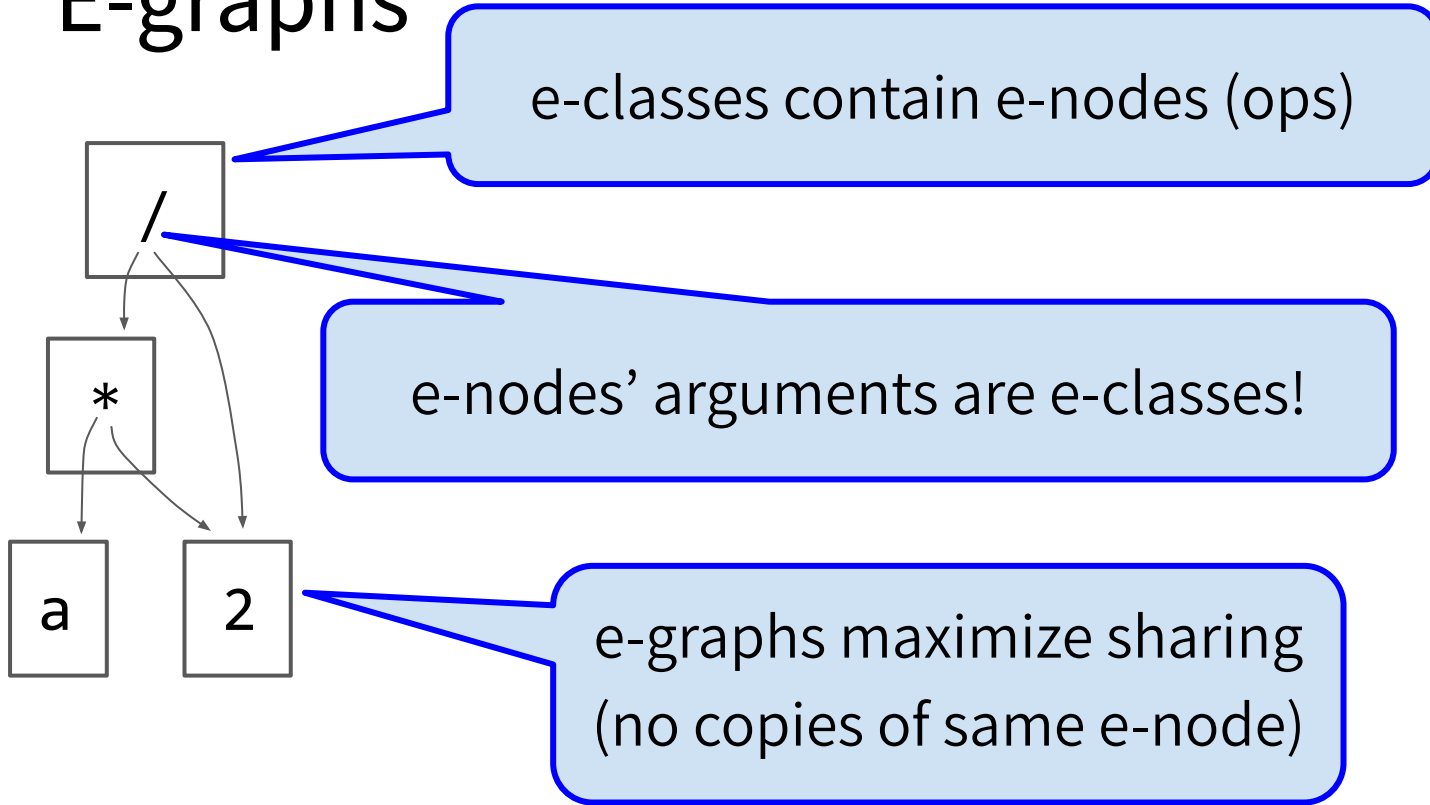
e-classes contain e-nodes (ops)



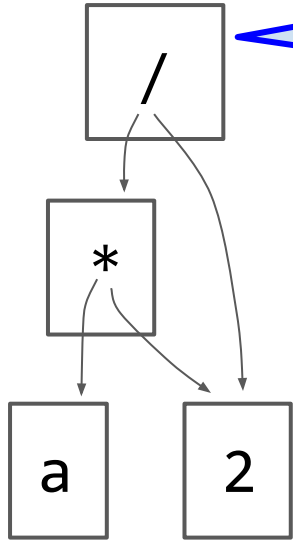
E-graphs



E-graphs

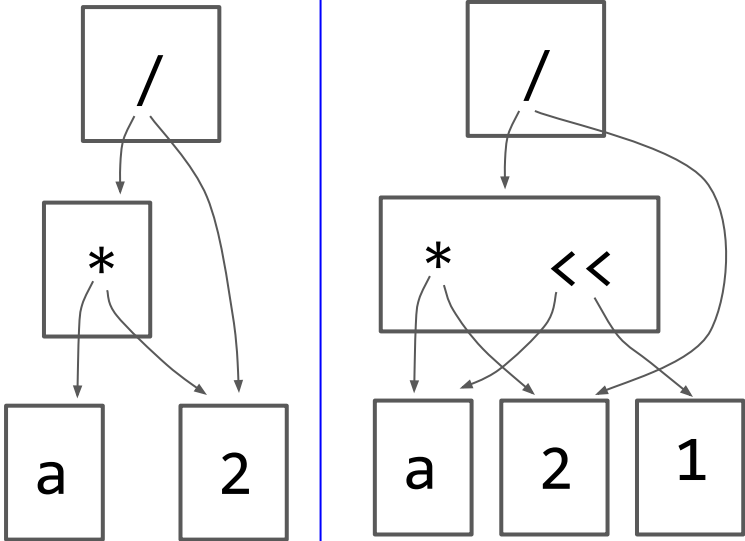


E-graphs



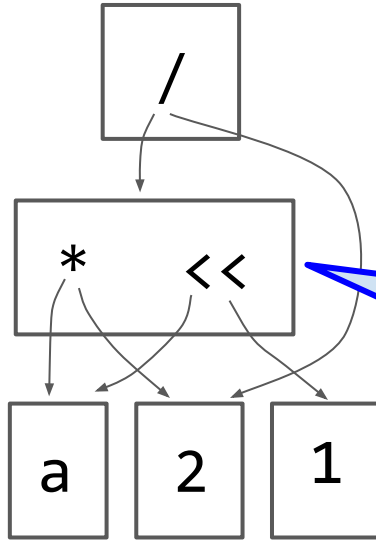
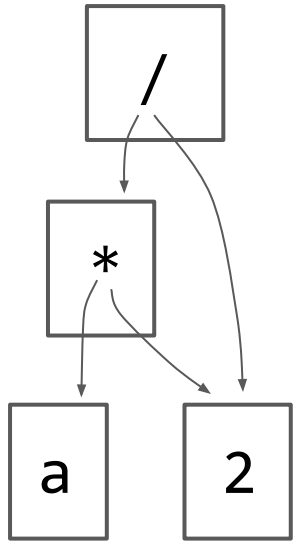
This e-classes represents
 $(a * 2) / 2$

E-graphs: applying rewrite rules



$$x * 2 \rightarrow x \ll 1$$

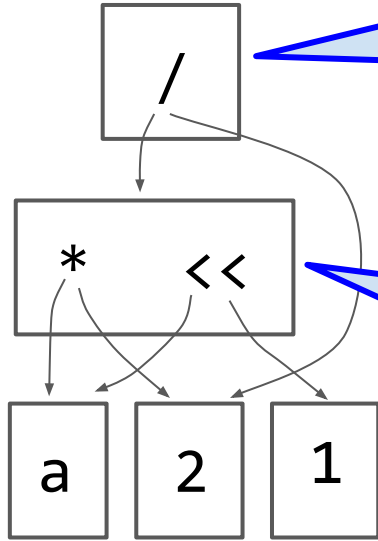
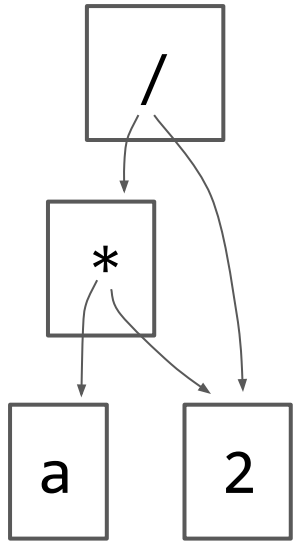
E-graphs: applying rewrite rules



This e-class *represents*
 $a * 2$ and $a \ll 1$

$$x * 2 \rightarrow x \ll 1$$

E-graphs: applying rewrite rules

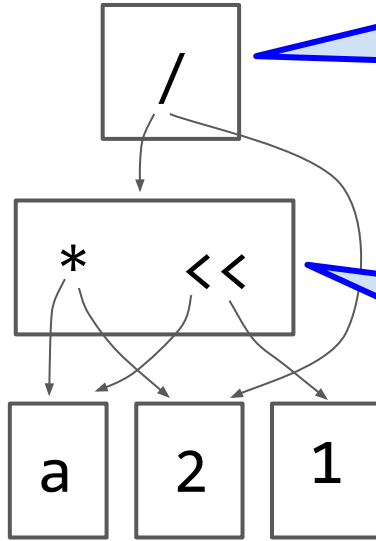
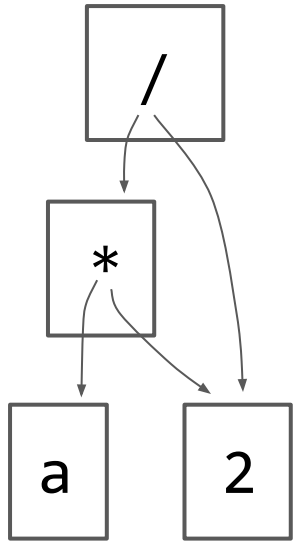


This e-class represents $(a*2)/2$ **and** $(a\ll 1)/2$

This e-class represents $a*2$ **and** $a\ll 1$

$x * 2 \rightarrow x \ll 1$

E-graphs: applying rewrite rules

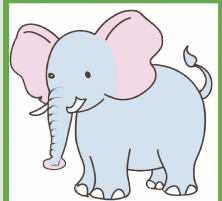


$x * 2 \rightarrow x \ll 1$

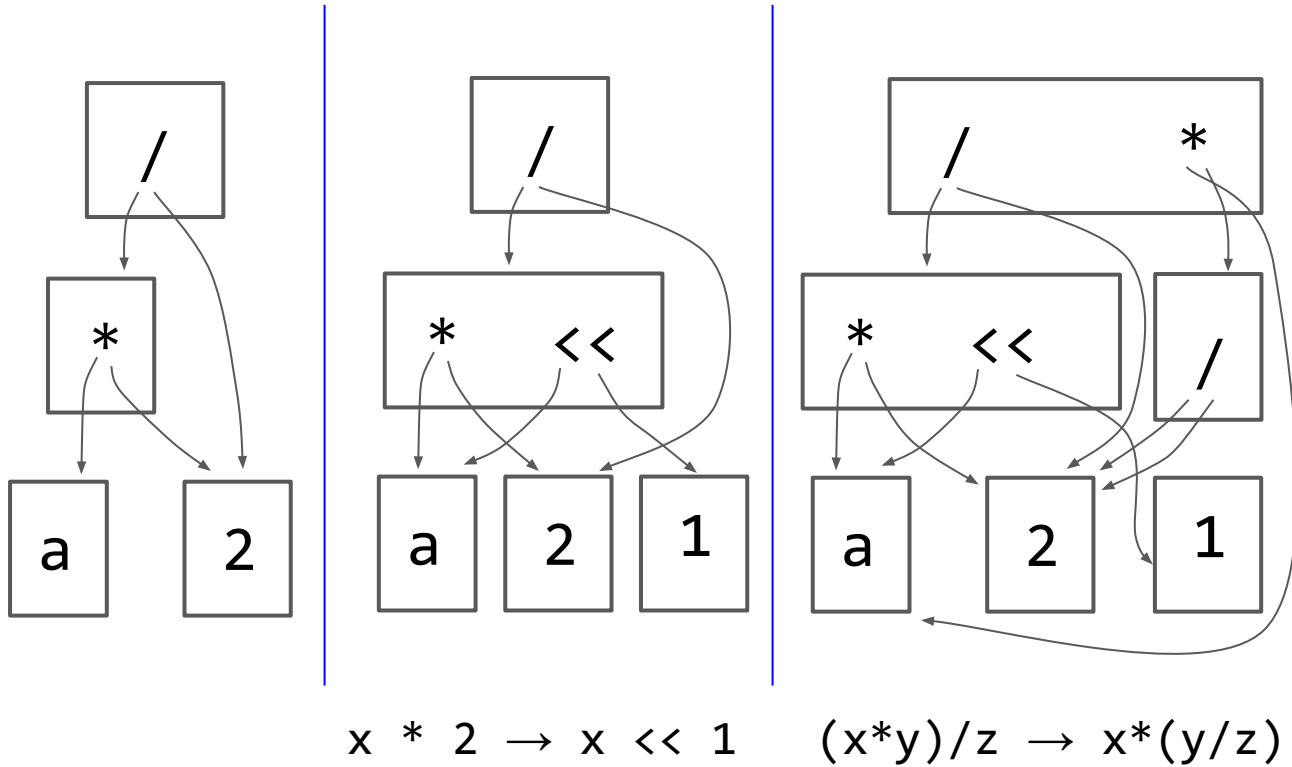
This e-class *represents*
 $(a * 2) / 2$ **and** $(a \ll 1) / 2$

This e-class *represents*
 $a * 2$ **and** $a \ll 1$

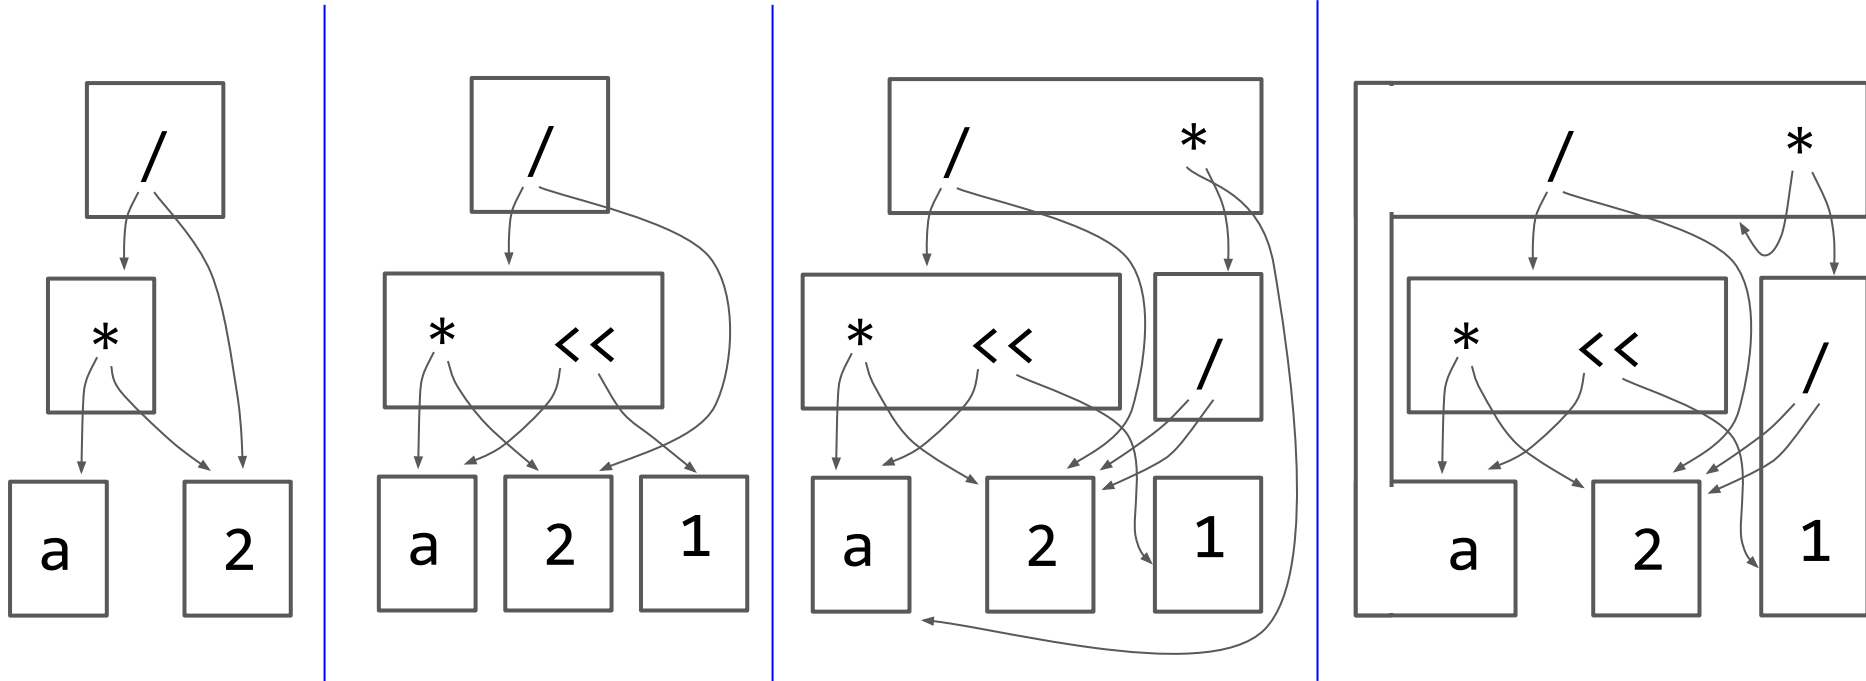
E-graphs never forget.
Rewrites don't lose info!



E-graphs: applying rewrite rules



E-graphs: applying rewrite rules



$$x * 2 \rightarrow x \ll 1$$

$$(x*y)/z \rightarrow x*(y/z)$$

$$x / x \rightarrow 1$$

$$x * 1 \rightarrow x$$

E-graphs: compact representation

Rewrites can **shrink** e-graphs!

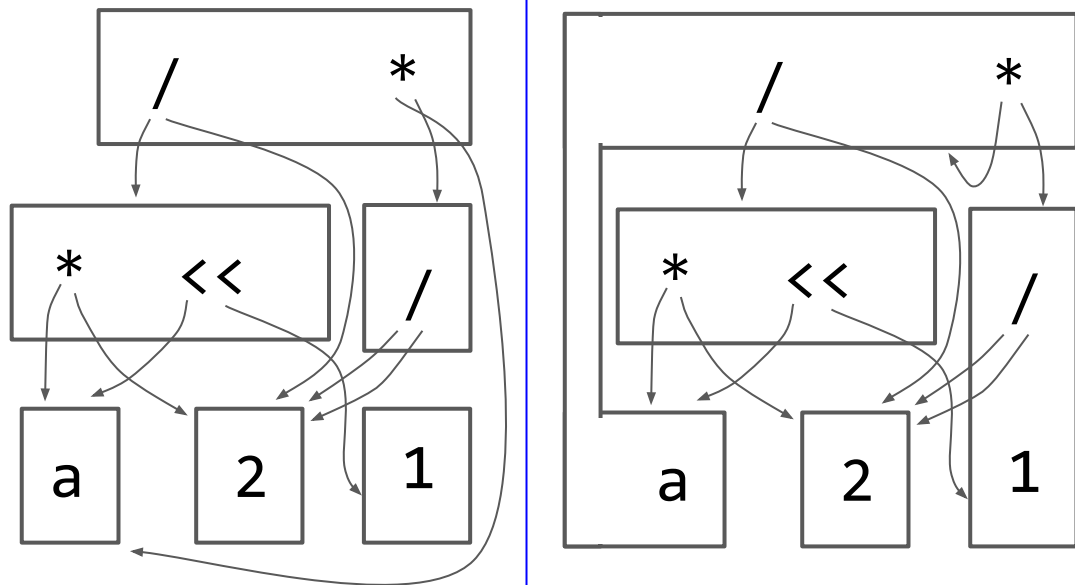
- $6 \rightarrow 5$ eclasses

E-graphs can represent ∞ terms

- $a, a * 1, a * 1 * 1, \dots$

E-graphs can “*saturate*”

- learn all derivable eqs 



$$x / x \rightarrow 1$$

$$x * 1 \rightarrow x$$

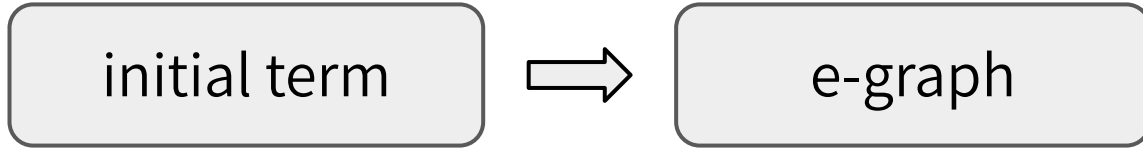
Equality Saturation

- Technique first used in Denali (Joshi, Nelson, Randall 2002)
 - Optimizes straight-line assembly kernels for the DEC Alpha
- Extended to loops in Peggy [POPL 2009]
 - Coined term “Equality Saturation”
 - Coinductive stream operators for algebraic loop rewrites
 - Used Rete algo from expert systems

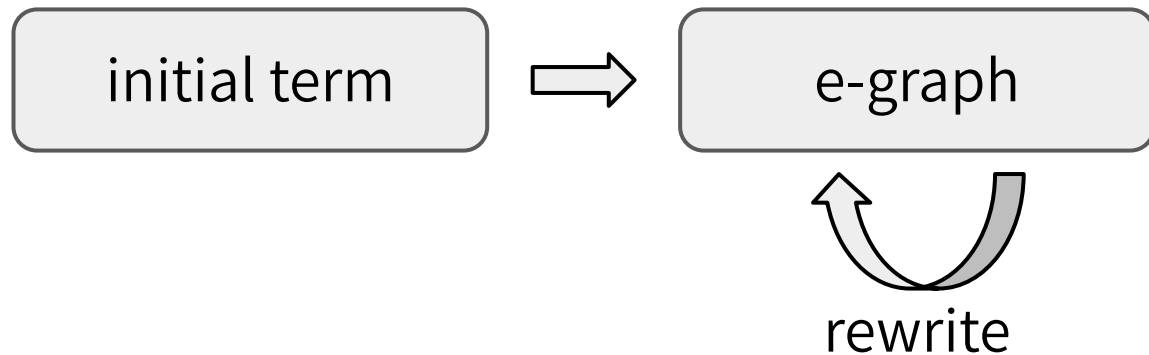
Equality Saturation

initial term

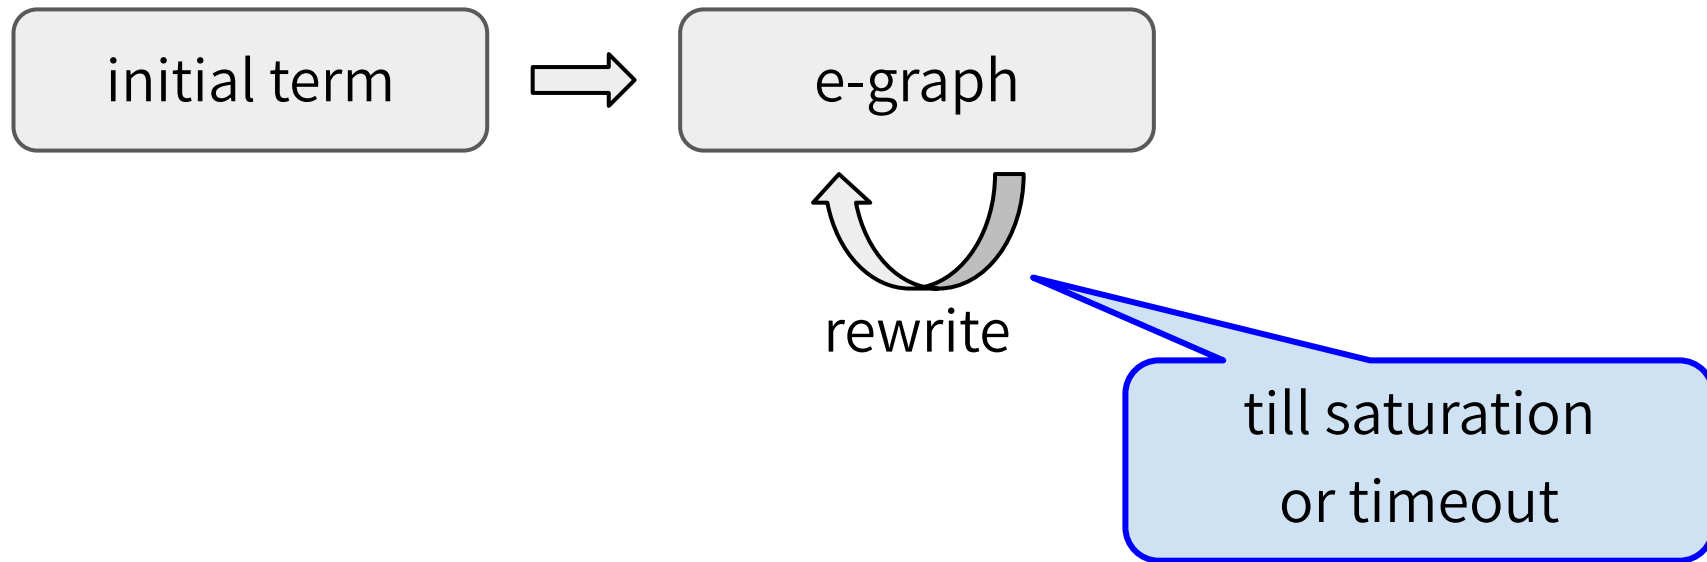
Equality Saturation



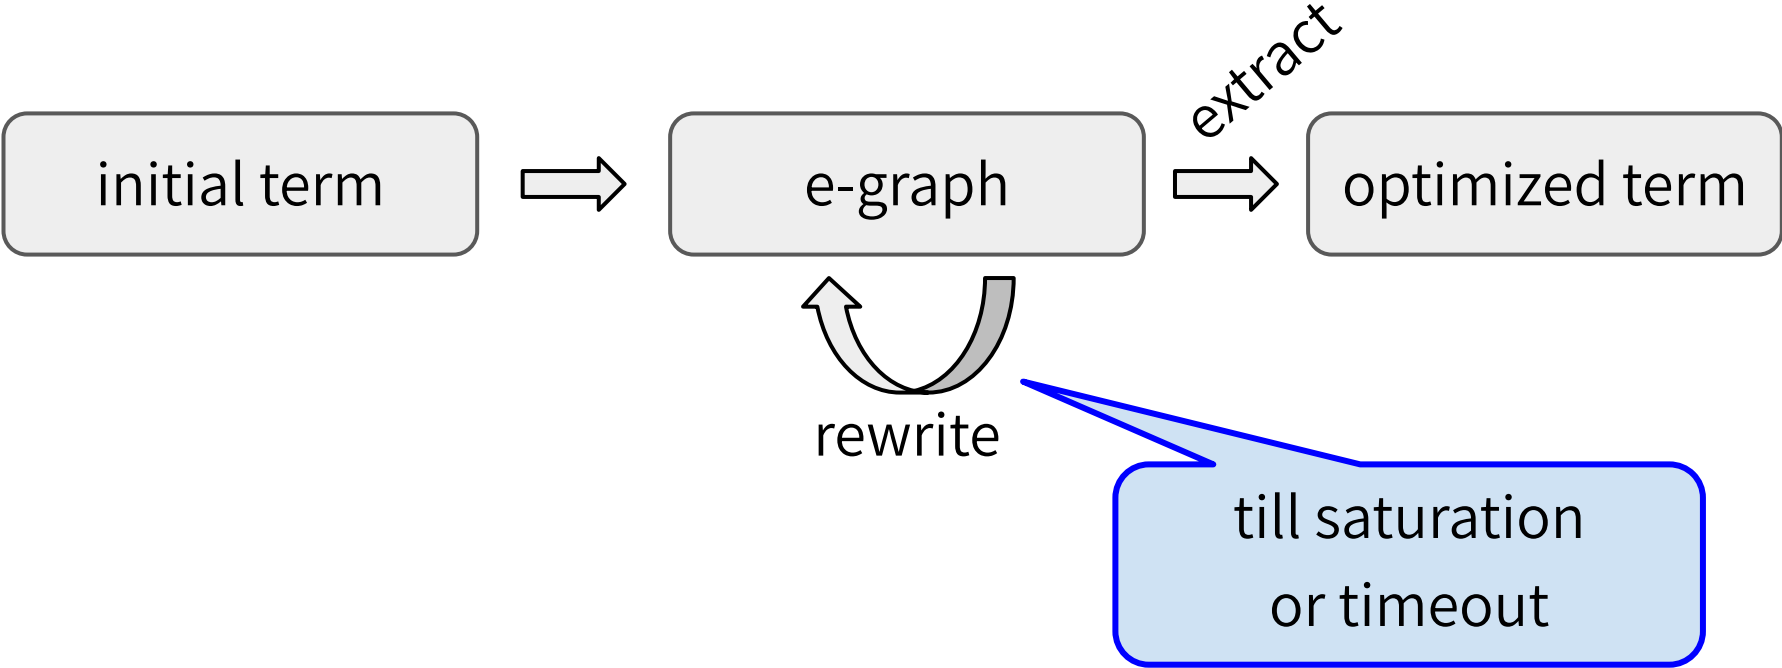
Equality Saturation



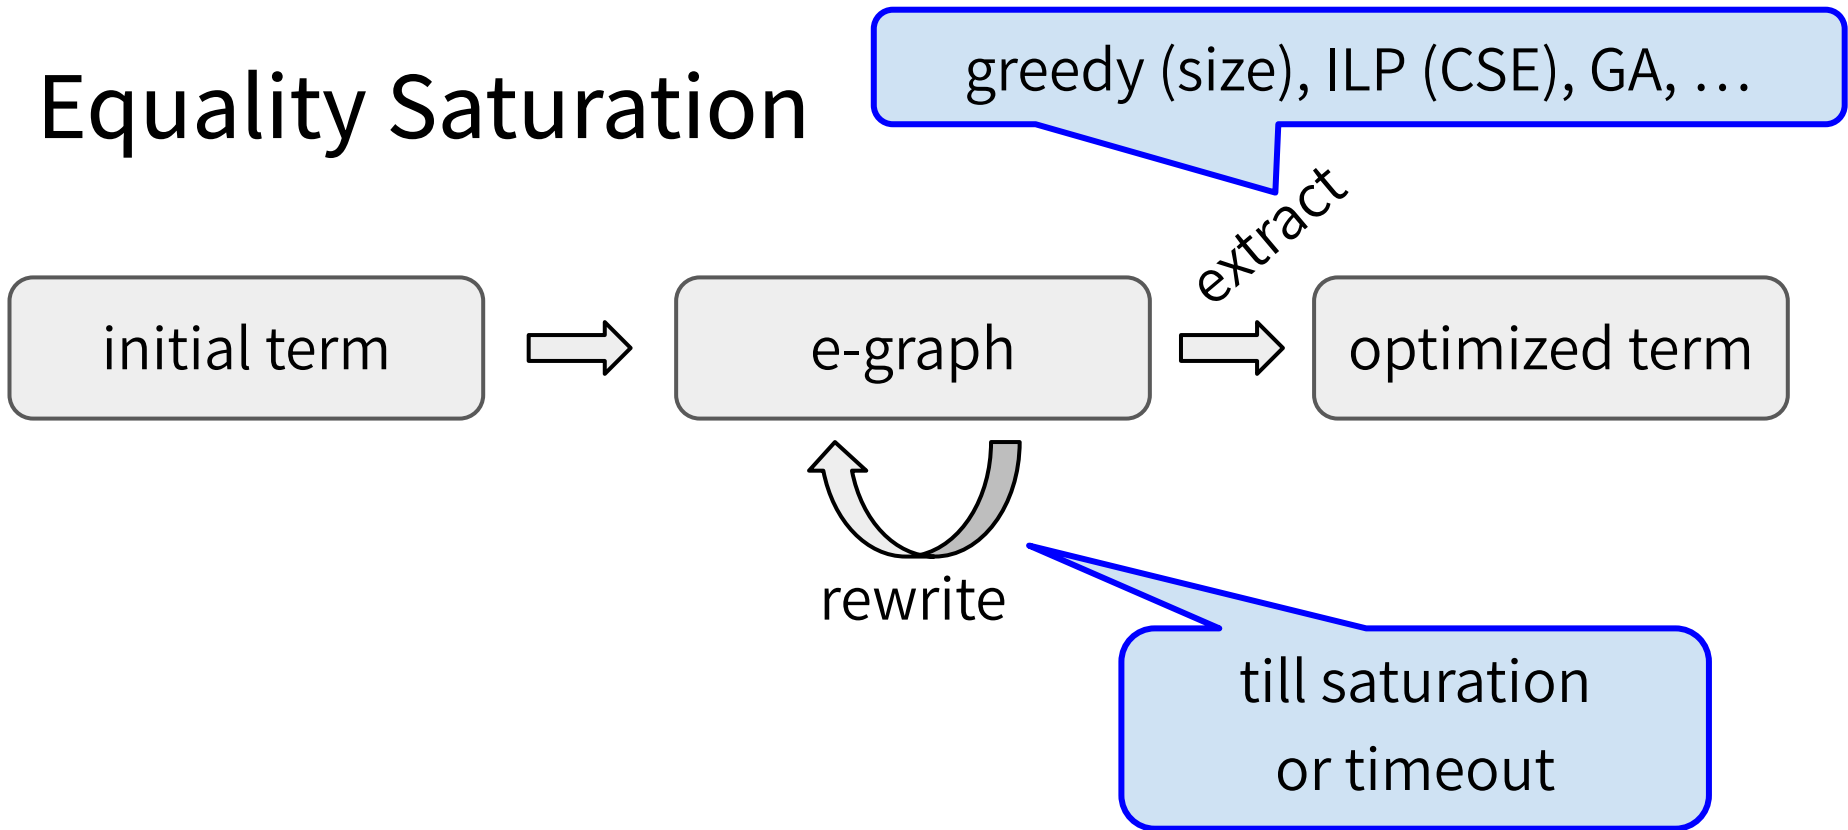
Equality Saturation



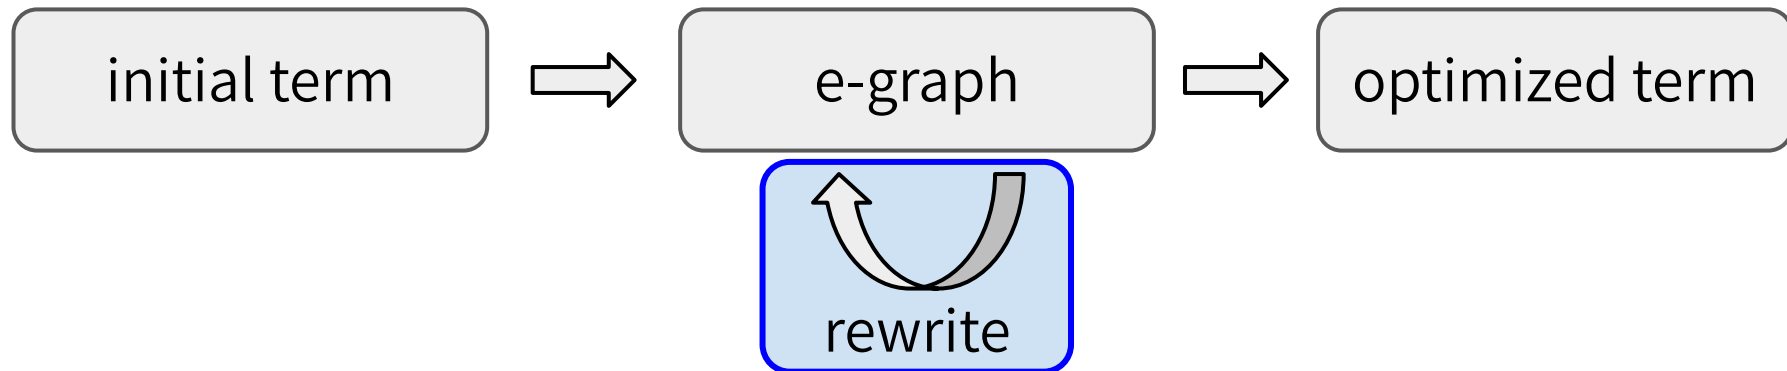
Equality Saturation



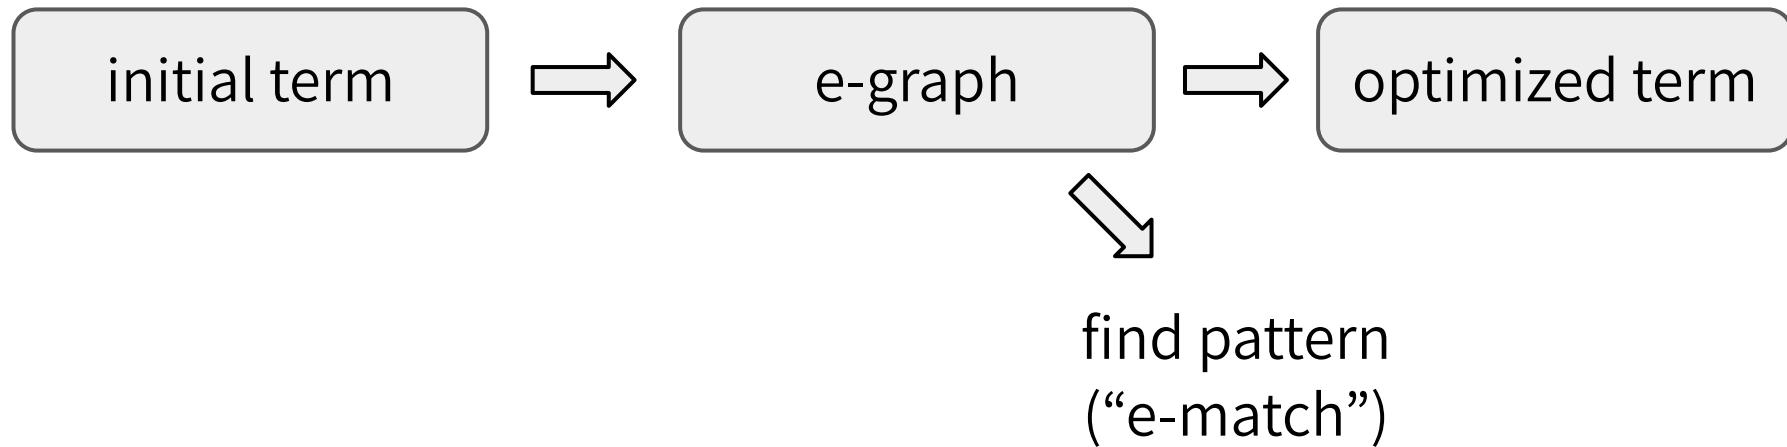
Equality Saturation



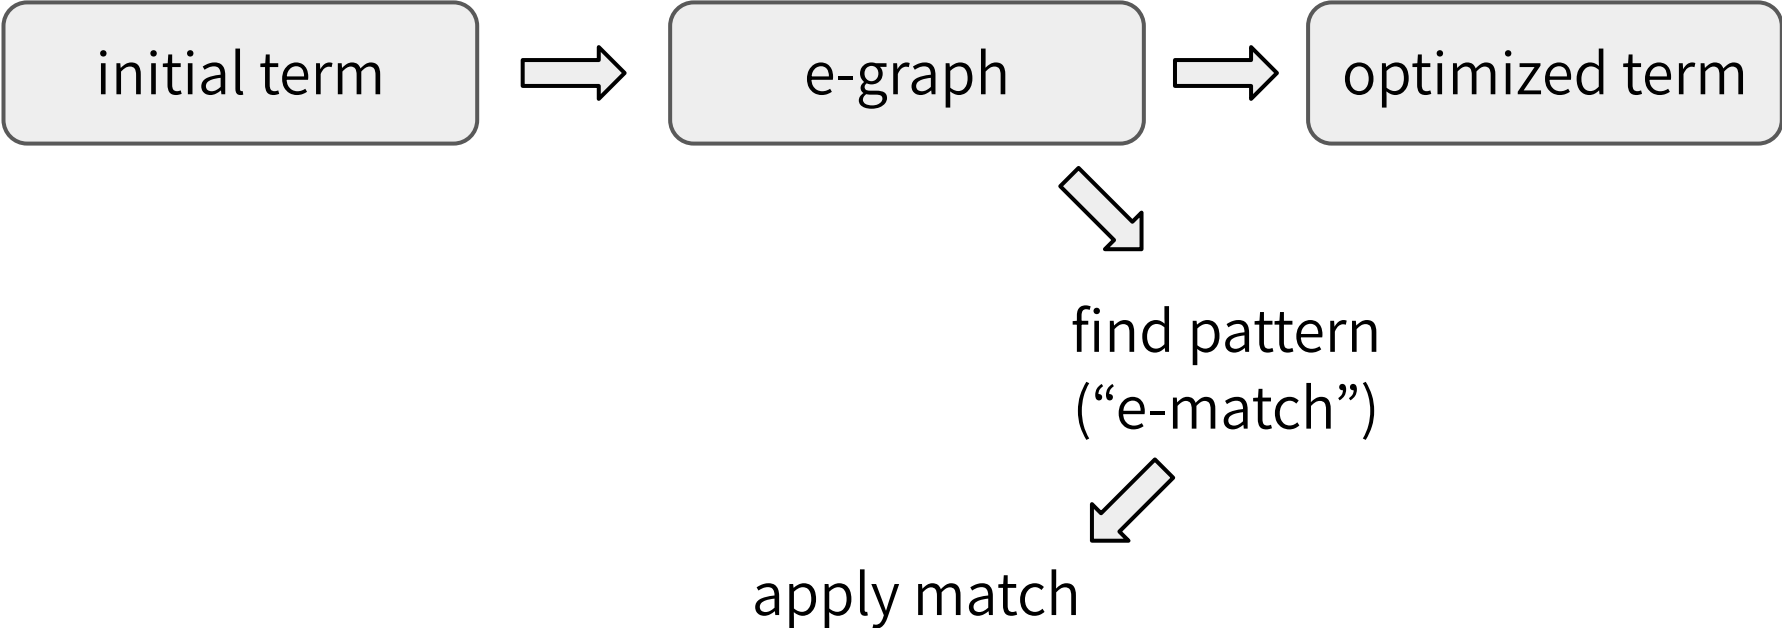
Equality Saturation



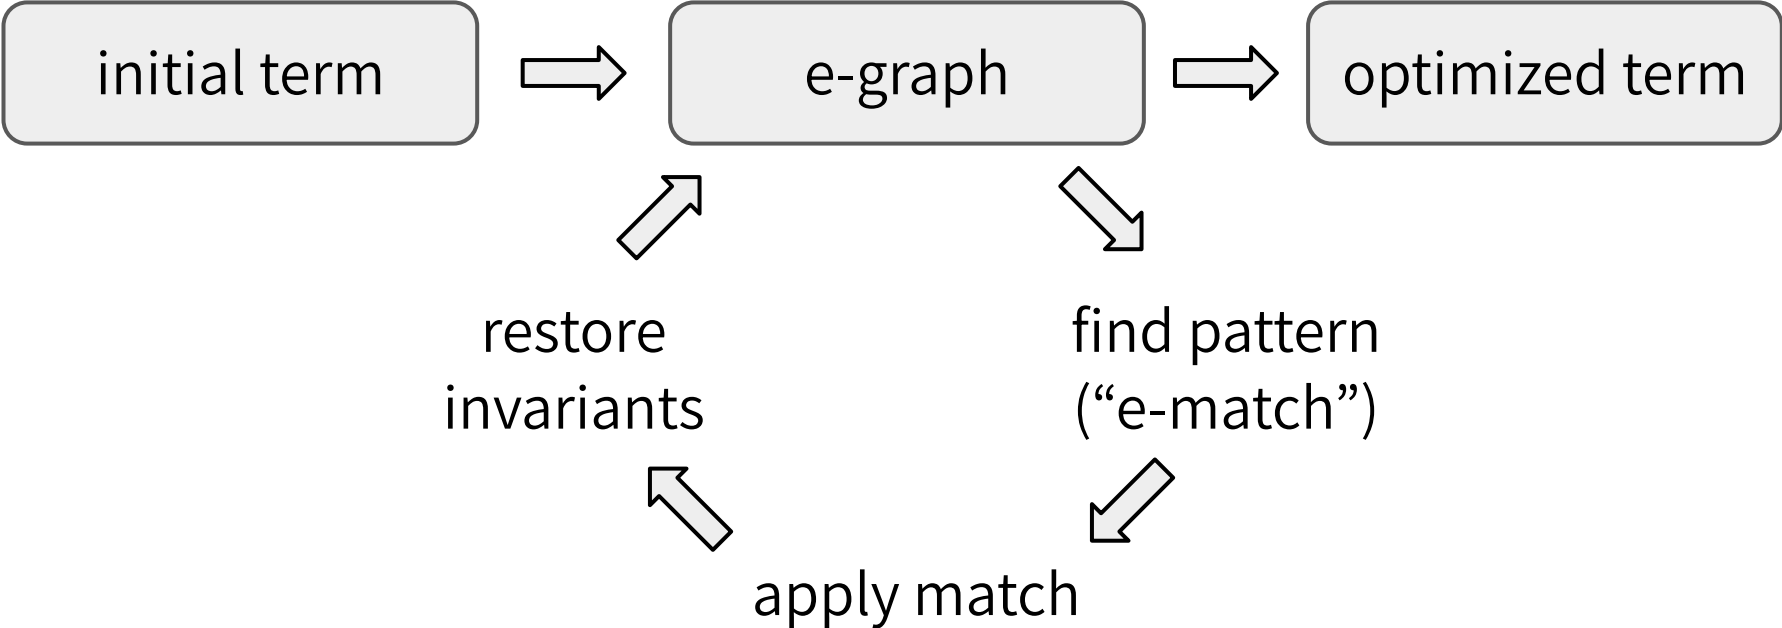
Equality Saturation



Equality Saturation



Equality Saturation



Equality Saturation



restore
invariants

find pattern
("e-match")



apply match

congruence

$$a = b \Rightarrow f(a) = f(b)$$

Equality Saturation



restore
invariants



find pattern
("e-match")



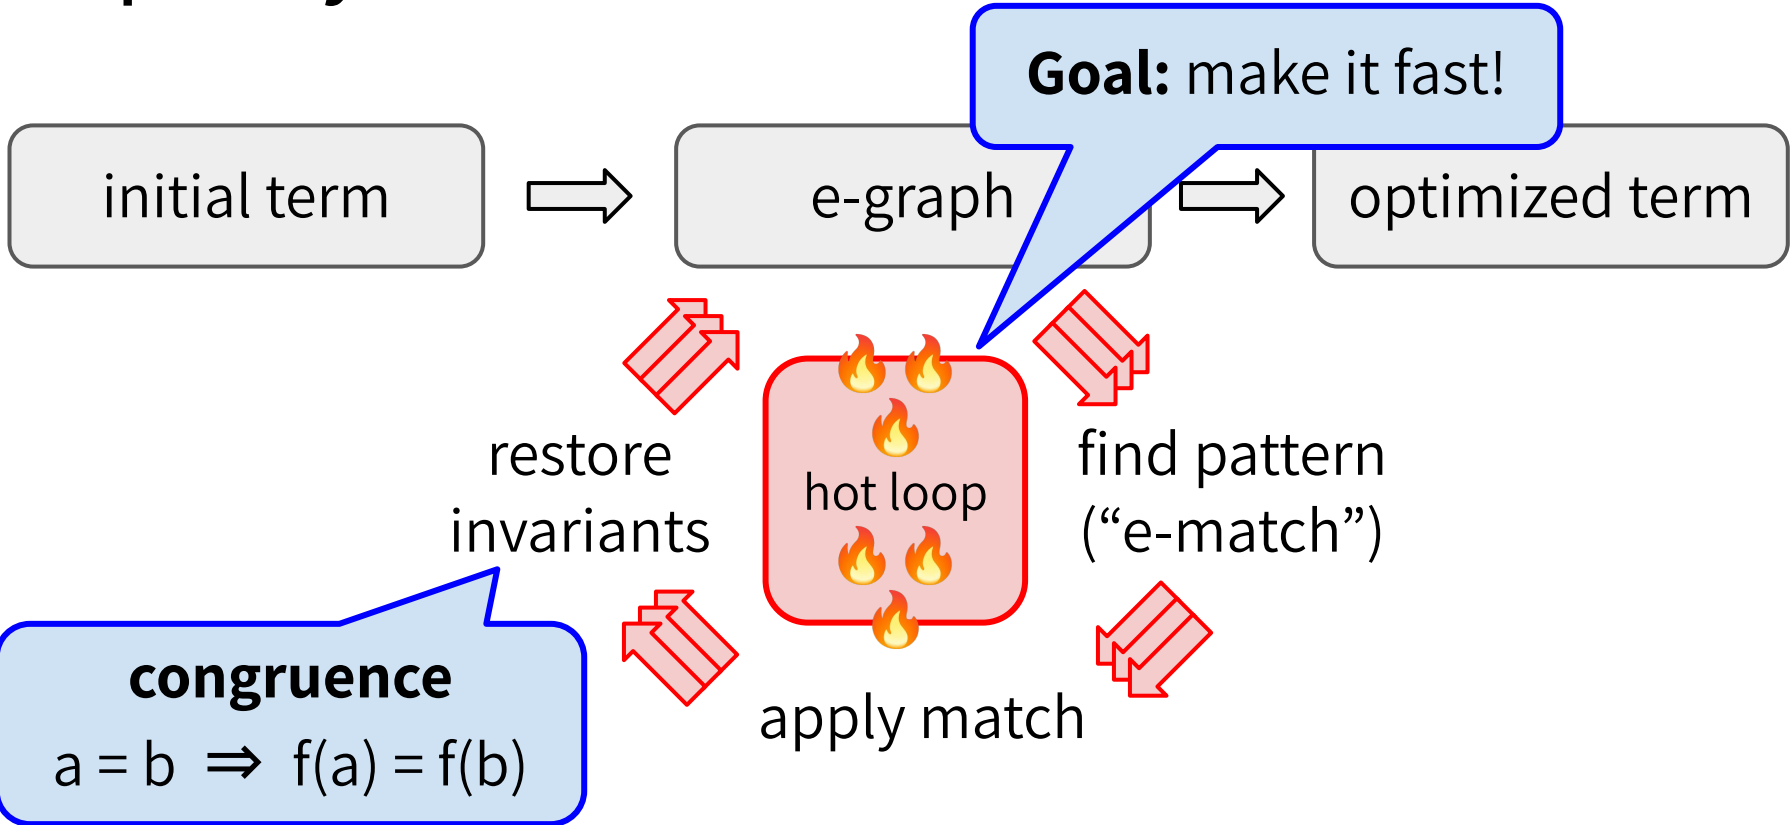
apply match



congruence

$$a = b \Rightarrow f(a) = f(b)$$

Equality Saturation



egg: Fast and Flexible EqSat [POPL 2021]

```
def equality_saturation(expr, rewrites):  
    egraph = initial_egraph(expr)
```

```
    while not egraph.is_sat:  
        for rw in rewrites:
```

```
            for (subst, ec) in rewrites:
```

```
                ec2 = egraph.add(rw.rhs.subst(subst))
```

```
                egraph.merge(ec, ec2)
```

```
    return egraph.extract_best()
```

batch reads

batch writes
(invariants broken)

invariants restored
once per iteration

```
def equality_saturation(expr, rewrites):  
    egraph = initial_egraph(expr)
```

```
    while not egraph.is_saturated_or_timeout():  
        matches = []
```

```
        for rw in rewrites:
```

```
            for (subst, ec) in egraph.ematch(rw.lhs):  
                matches.append((rw, subst, ec))
```

```
            for (rw, subst, ec) in matches:  
                ec2 = egraph.add(rw.rhs.subst(subst))  
                egraph.merge(ec, ec2)
```

```
            egraph.rebuild()
```

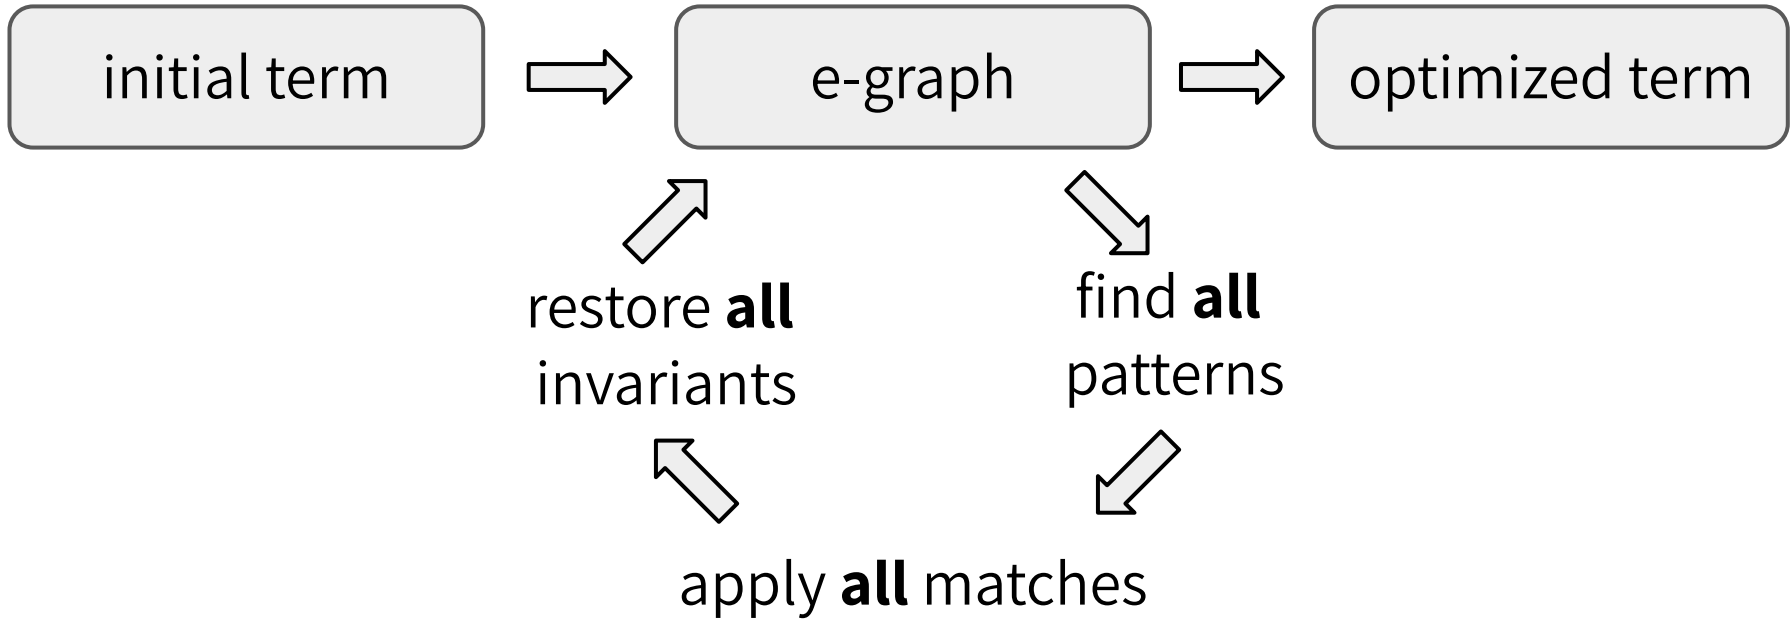
```
    return egraph.extract_best()
```

congruence

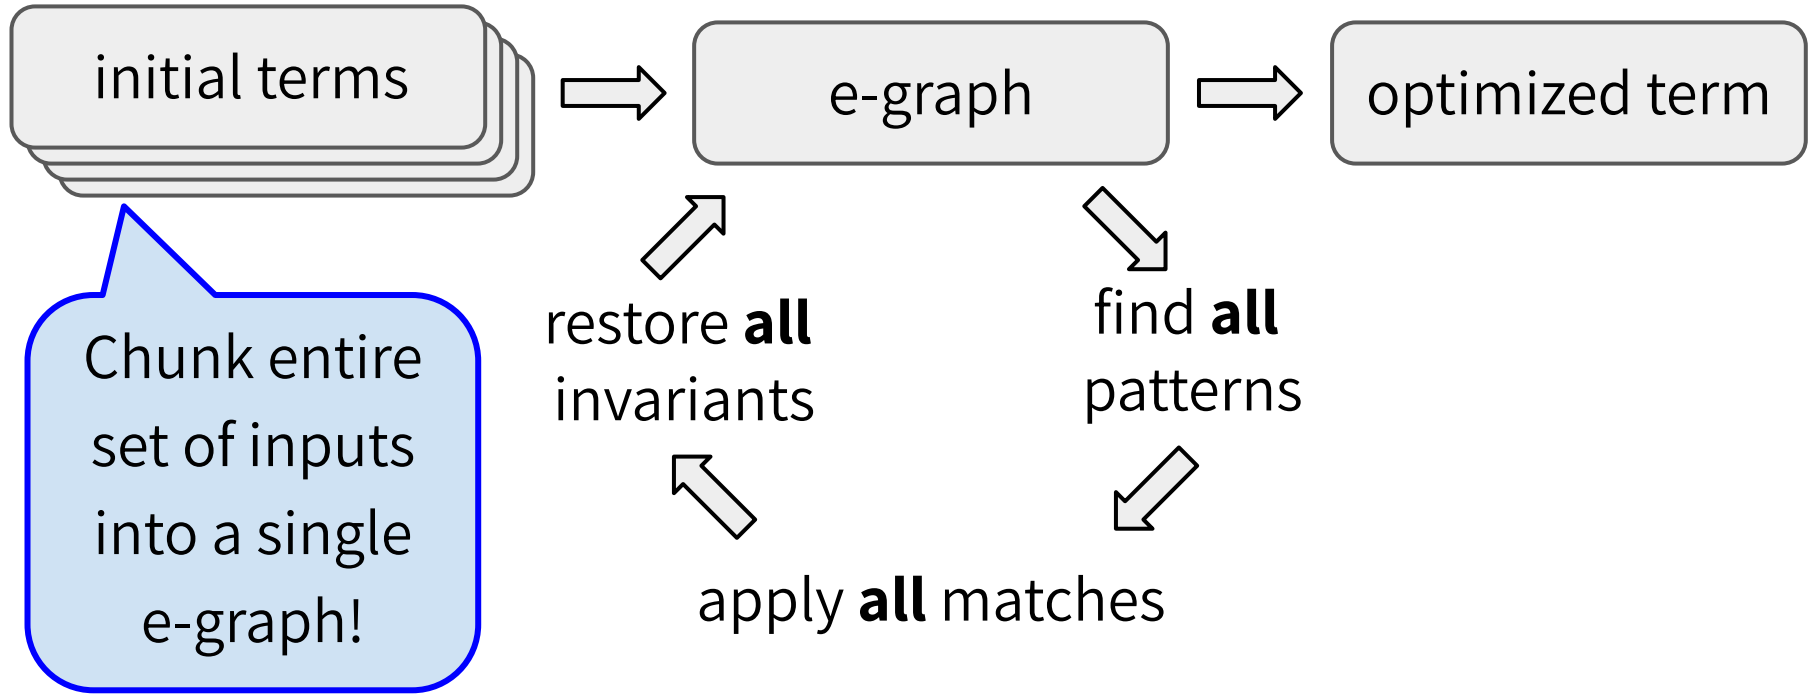
$a = b \Rightarrow f(a) = f(b)$

Adapted from Downey, Sethi, Tarjan 1980

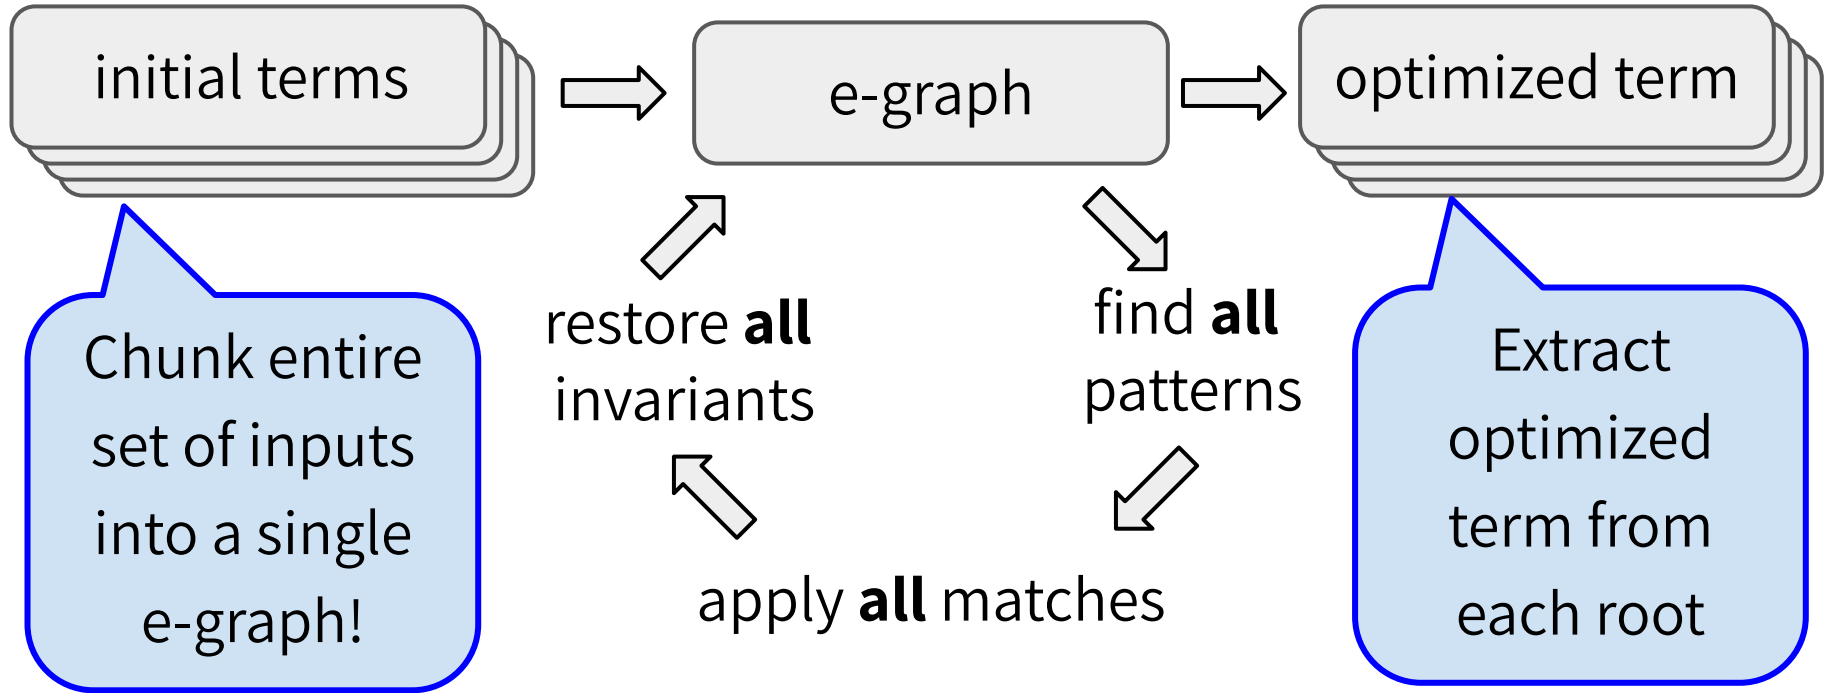
More amortization via *batching* in egg



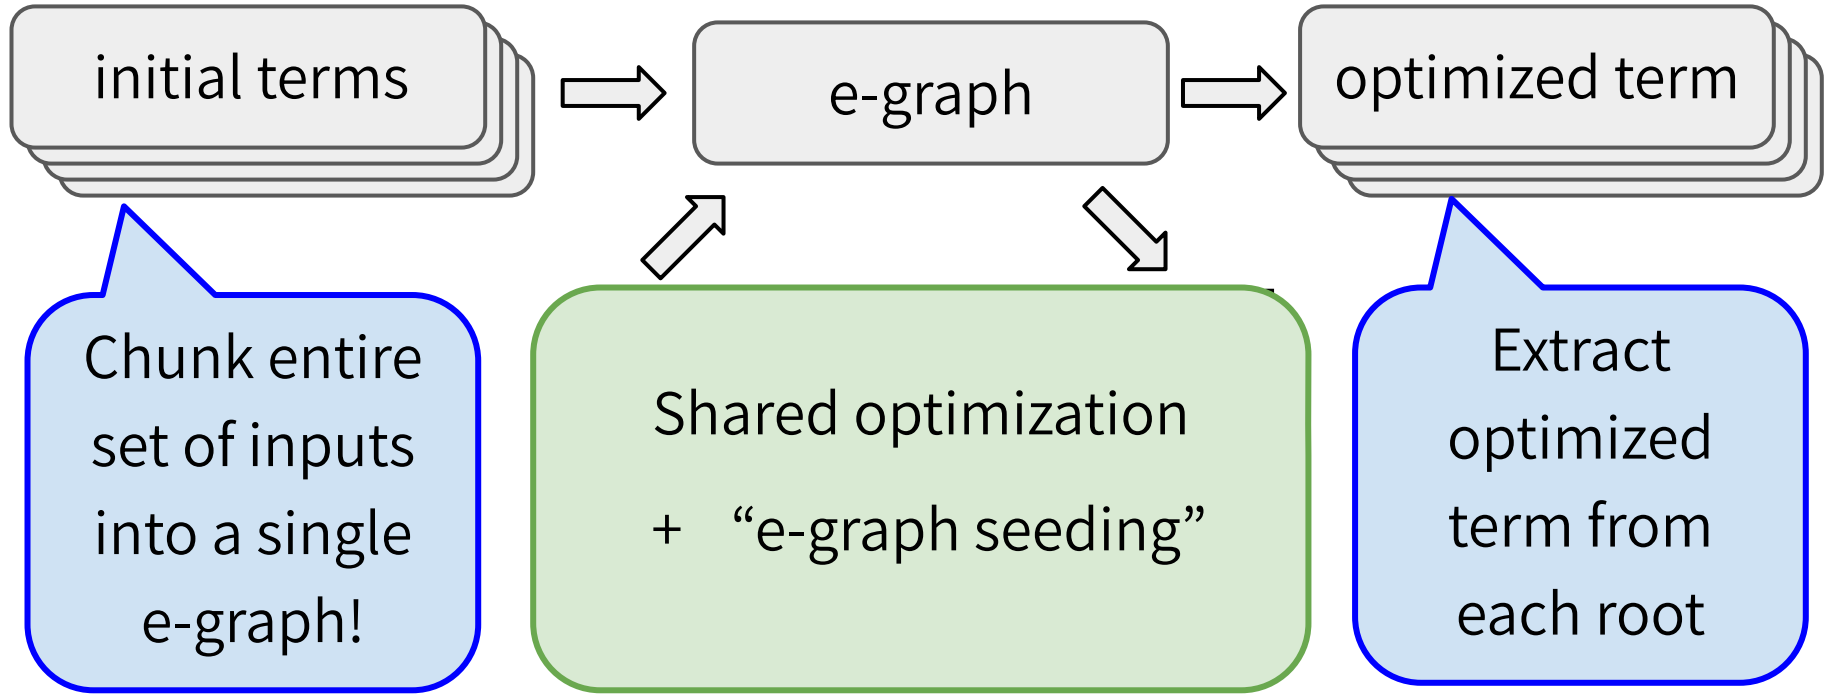
More amortization via *batching* in egg



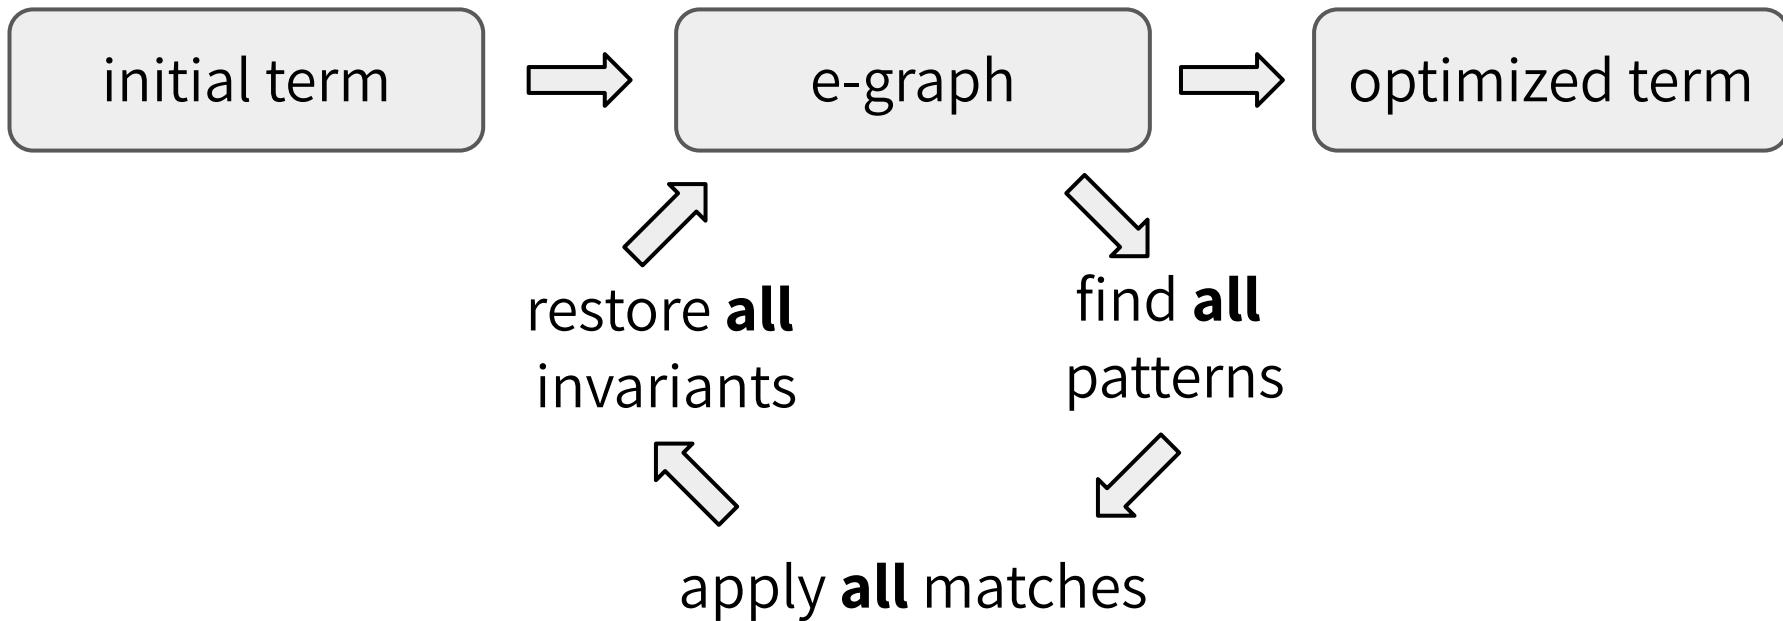
More amortization via *batching* in egg



More amortization via *batching* in egg

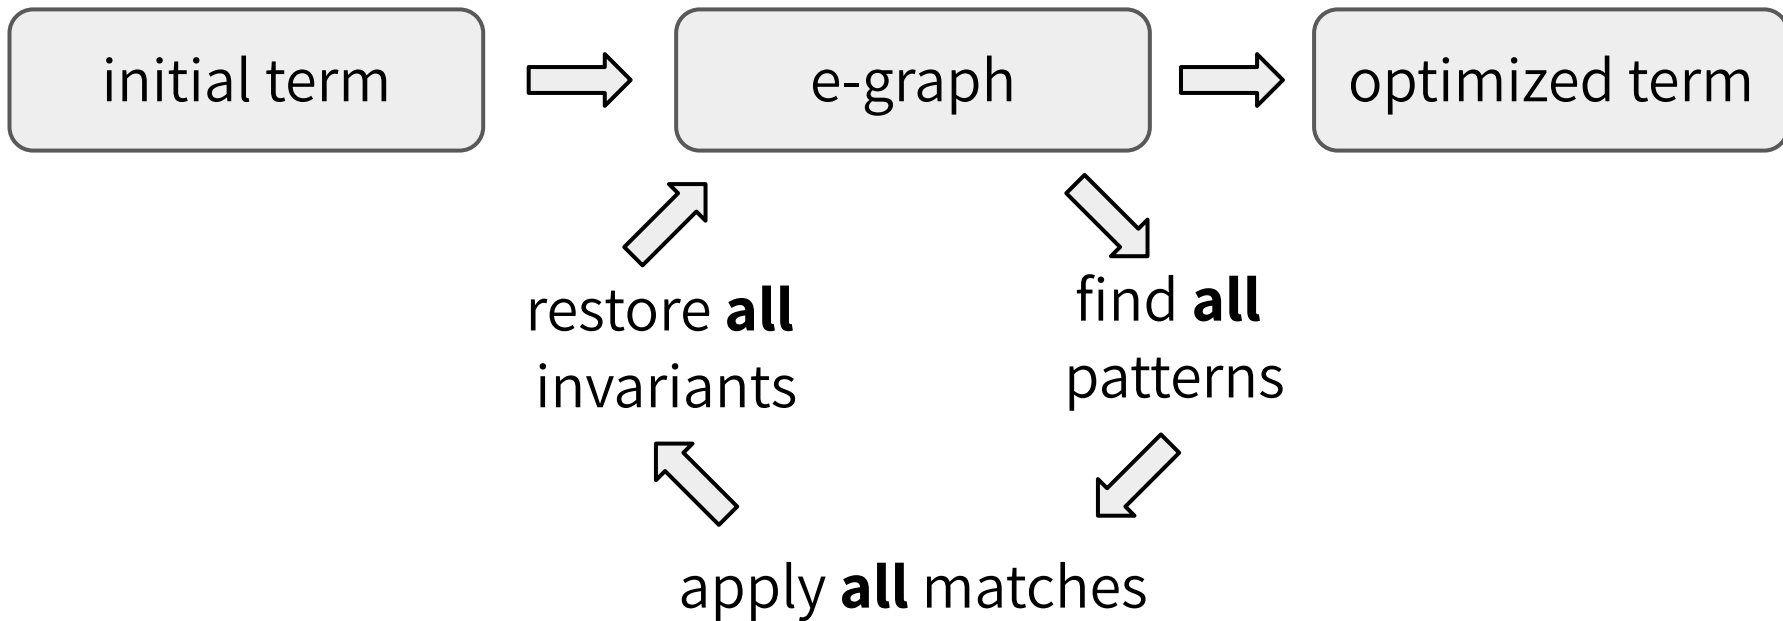


egg's Equality Saturation

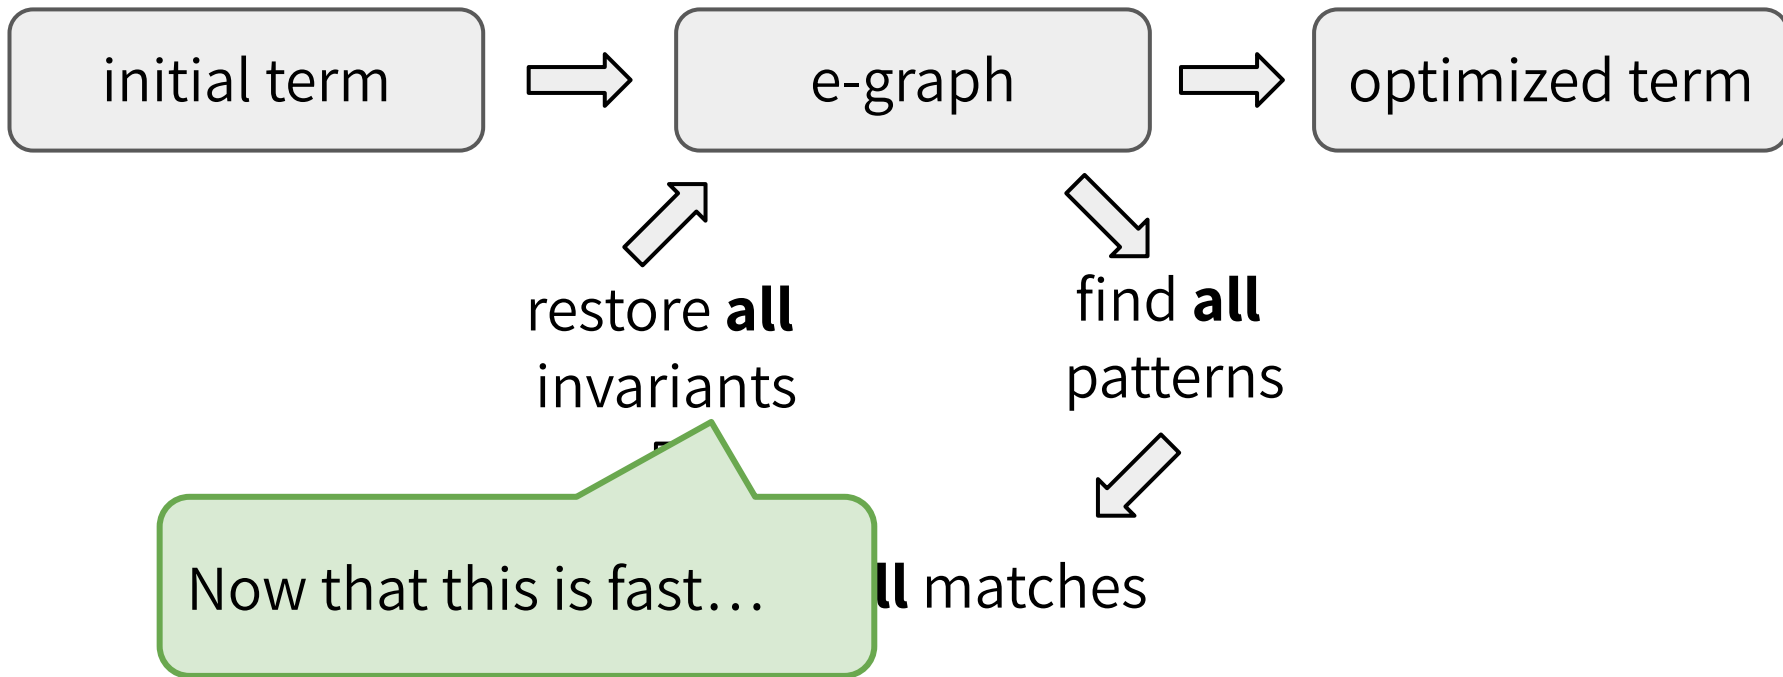


egglog : EqSat + Datalog

egg's Equality Saturation



egg's Equality Saturation



egg's Equality Saturation



restore **all**
invariants

find **all**
patterns

Now that this is fast...

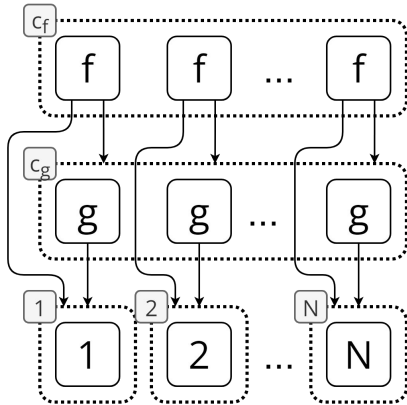
we bottleneck on matching 😓

E-matching: pattern matching over e-graphs

- *E-matching*: find substs from pattern variables to e-classes
- Substs guaranteed to be represented by the matched e-graph

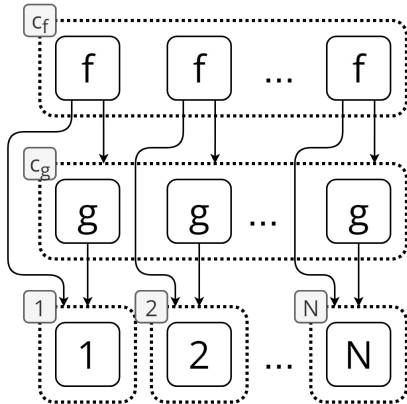
E-matching: pattern matching over e-graphs

- *E-matching*: find substs from pattern variables to e-classes
- Substs guaranteed to be represented by the matched e-graph



E-matching: pattern matching over e-graphs

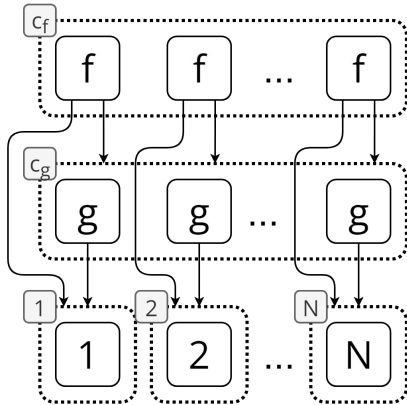
- *E-matching*: find substs from pattern variables to e-classes
- Substs guaranteed to be represented by the matched e-graph



$f(\alpha, g(\alpha))$ will match $f(1, g(1))$, $f(2, g(2))$, ..., $f(N, g(N))$, witnessed by $\{\alpha \mapsto 1\}$, $\{\alpha \mapsto 2\}$, ..., $\{\alpha \mapsto N\}$.

E-matching: pattern matching over e-graphs

- *E-matching*: find substs from pattern variables to e-classes
- Substs guaranteed to be represented by the matched e-graph

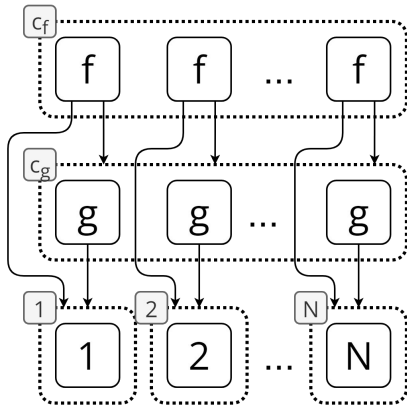


$f(\alpha, g(\alpha))$ will match $f(1, g(1))$, $f(2, g(2))$, ..., $f(N, g(N))$, witnessed by $\{\alpha \mapsto 1\}$, $\{\alpha \mapsto 2\}$, ..., $\{\alpha \mapsto N\}$.

$f(1, \alpha)$ will match $f(1, g(1))$, $f(1, g(2))$, ..., $f(1, g(N))$, witnessed by $\{\alpha \mapsto c_g\}$.

E-matching: pattern matching over e-graphs

- *E-matching*: find substs from pattern variables to e-classes
- Substs guaranteed to be represented by the matched e-graph
- NP-complete wrt to pattern size (Kozen 1977) 🤯

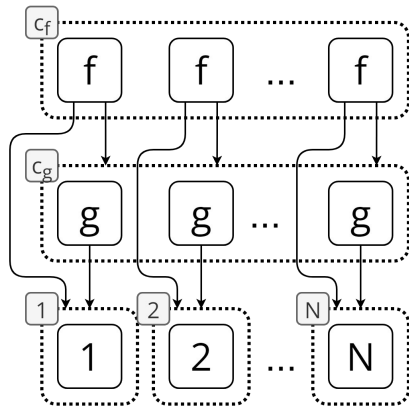


$f(\alpha, g(\alpha))$ will match $f(1, g(1))$, $f(2, g(2))$, ..., $f(N, g(N))$, witnessed by $\{\alpha \mapsto 1\}$, $\{\alpha \mapsto 2\}$, ..., $\{\alpha \mapsto N\}$.

$f(1, \alpha)$ will match $f(1, g(1))$, $f(1, g(2))$, ..., $f(1, g(N))$, witnessed by $\{\alpha \mapsto c_g\}$.

E-matching: pattern matching over e-graphs

- *E-matching*: find substs from pattern variables to e-classes
- Substs guaranteed to be represented by the matched e-graph
- NP-complete wrt to pattern size (Kozen 1977) 🤯

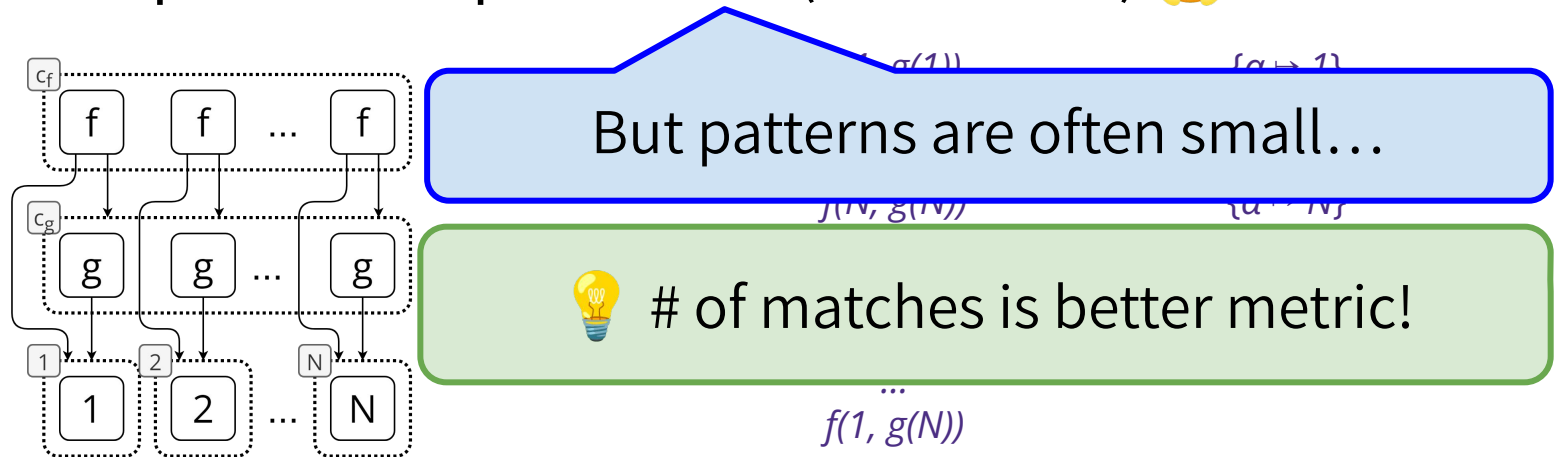


But patterns are often small...

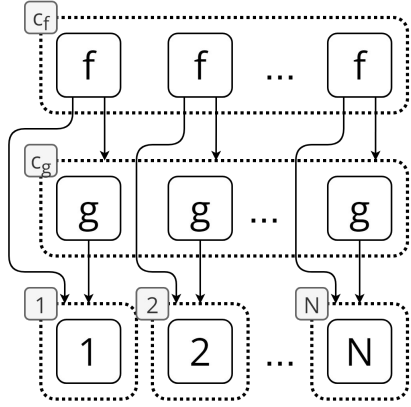
$f(1, \alpha)$ will match $f(1, g(1))$, $f(1, g(2))$, ..., $f(1, g(N))$, witnessed by $\{\alpha \mapsto c_g\}$.

E-matching: pattern matching over e-graphs

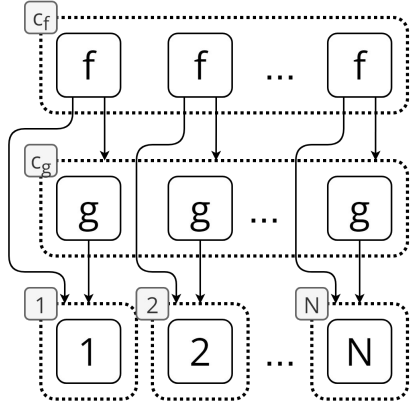
- *E-matching*: find substs from pattern variables to e-classes
- Substs guaranteed to be represented by the matched e-graph
- NP-complete wrt to pattern size (Kozen 1977) 🤯



Traditional e-matching via backtracking



Traditional e-matching via backtracking



$f(a, g(a))$

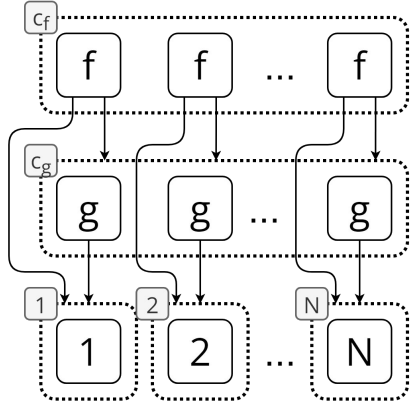


for e-class **c** **in** e-graph **E**:

Backtracking search $f(a, g(a))$



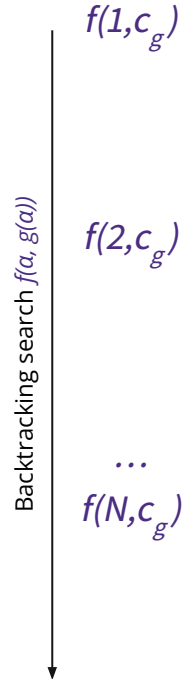
Traditional e-matching via backtracking



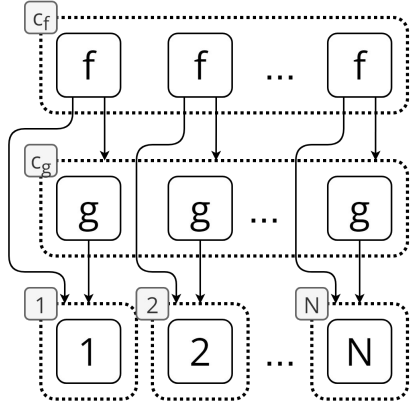
$f(a, g(a))$



```
for e-class c in e-graph E:  
  for f-node  $n_1$  in c:
```



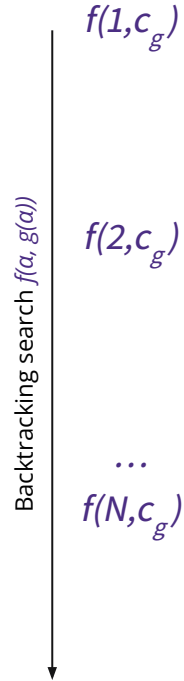
Traditional e-matching via backtracking



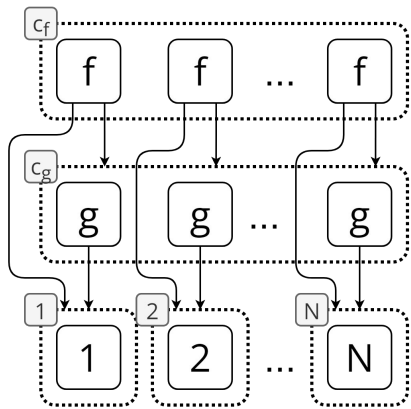
$f(\alpha, g(\alpha))$



```
for e-class c in e-graph E:  
  for f-node  $n_1$  in c:  
    subst = {root  $\mapsto$  c,  $\alpha \mapsto n_1.child_1$ }
```



Traditional e-matching via backtracking



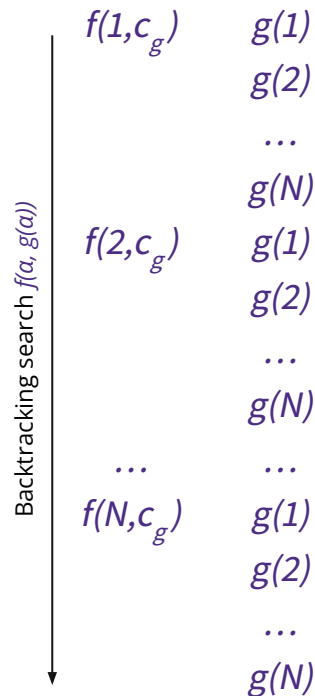
$f(\alpha, g(\alpha))$



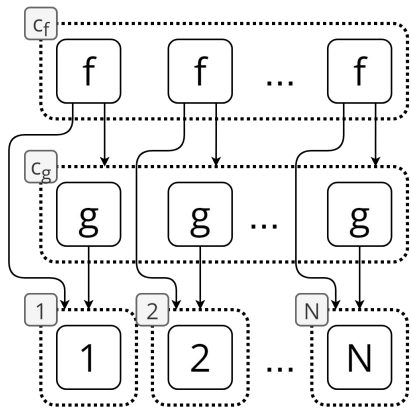
```

for e-class c in e-graph E:
  for f-node  $\mathbf{n}_1$  in c:
    subst = {root  $\mapsto$  c,  $\alpha \mapsto \mathbf{n}_1$ .child1}
    for g-node  $\mathbf{n}_2$  in  $\mathbf{n}_1$ .child2:

```



Traditional e-matching via backtracking

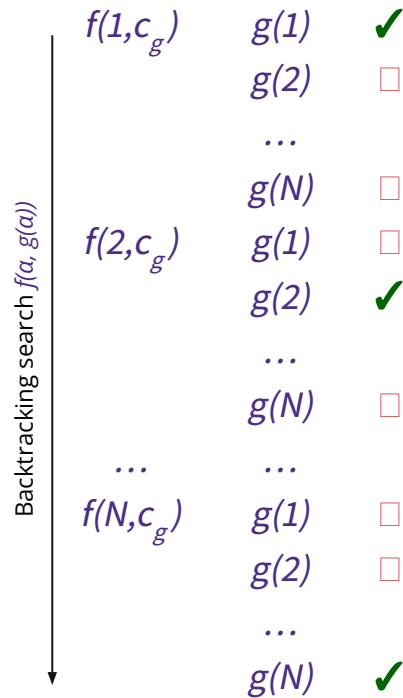


$f(\alpha, g(\alpha))$

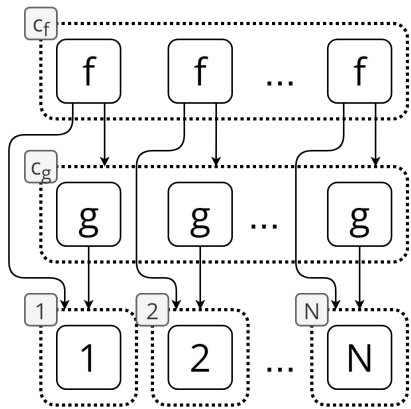


```

for e-class c in e-graph E:
  for f-node  $n_1$  in c:
    subst = {root  $\mapsto$  c,  $\alpha \mapsto n_1.child_1$ }
    for g-node  $n_2$  in  $n_1.child_2$ :
      if subst[ $\alpha$ ] =  $n_2.child_1$ :
  
```



Traditional e-matching via backtracking

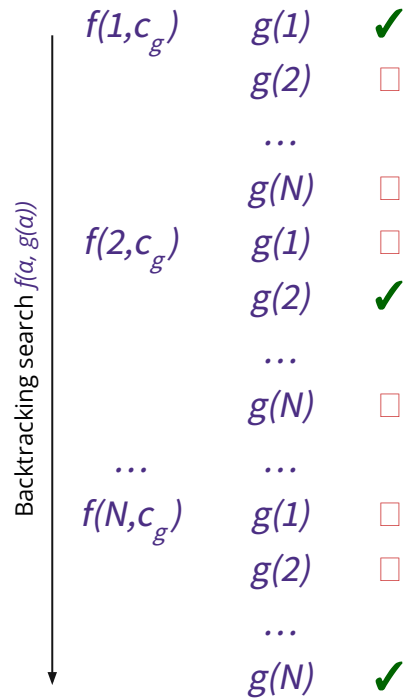


$f(\alpha, g(\alpha))$

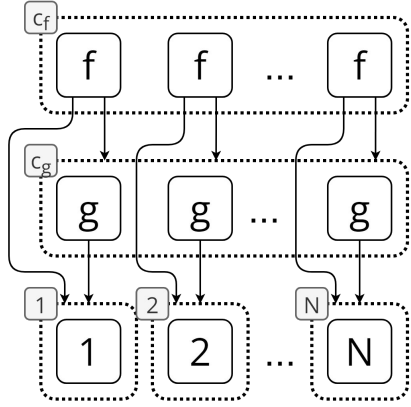


```

for e-class c in e-graph E:
  for f-node  $\mathbf{n}_1$  in c:
    subst = {root  $\mapsto$  c,  $\alpha \mapsto \mathbf{n}_1$ .child1}
    for g-node  $\mathbf{n}_2$  in  $\mathbf{n}_1$ .child2:
      if subst[ $\alpha$ ] =  $\mathbf{n}_2$ .child1:
        yield subst
  
```



Traditional e-matching via backtracking




$f(\alpha, g(\alpha))$



```

for e-class c in e-graph E:
  for f-node n1 in c:
    subst = {root ↦ c, α ↦ n1.child1}
    for g-node n2 in n1.child2:
      if subst[α] = n2.child1:
  
```

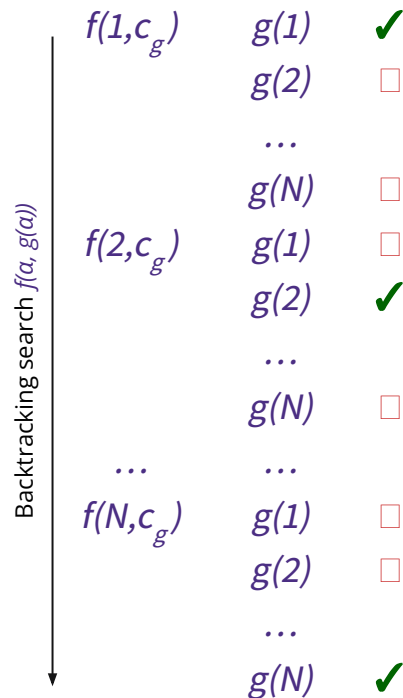
$O(N^2)$, yet at most $O(N)$ matches


Backtracking search $f(\alpha, g(\alpha))$

$f(1, c_g)$	$g(1)$	✓
	$g(2)$	□
	...	
	$g(N)$	□
$f(2, c_g)$	$g(1)$	□
	$g(2)$	✓
	...	
	$g(N)$	□
...	...	
$f(N, c_g)$	$g(1)$	□
	$g(2)$	□
	...	
	$g(N)$	✓

Traditional e-matching via backtracking

- Many optimizations in literature
 - custom VMs for “CSE”
 - specific patterns
 - mod-time analysis
- No data complexity bounds!



Key insight: e-matching is a DB problem!

E-matching in e-graphs

Finding substitutions such that substituted terms are represented in an e-graph.

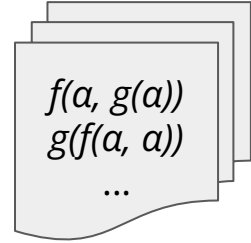
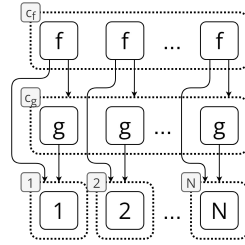
Conjunctive queries in DBs

Finding substitutions such that substituted atoms are present in a relational DB.

Relational e-matching

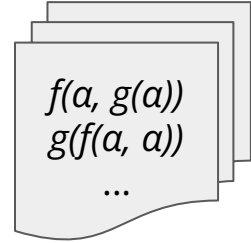
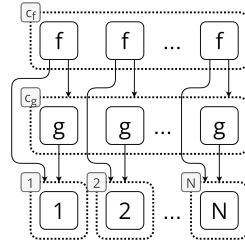
Relational e-matching

- Given e-graph + patterns



Relational e-matching

- Given e-graph + patterns
- Transform e-graph to tables

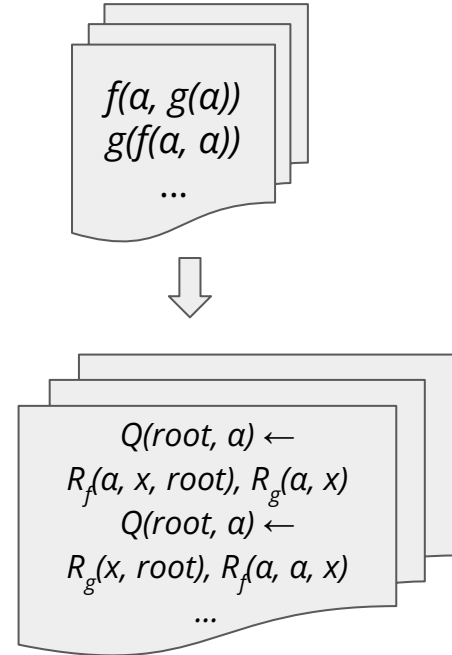
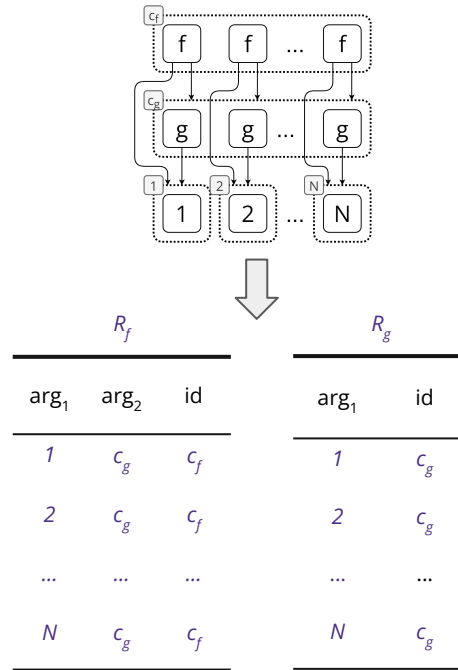


↓

R_f			R_g	
arg_1	arg_2	id	arg_1	id
1	c_g	c_f	1	c_g
2	c_g	c_f	2	c_g
...
N	c_g	c_f	N	c_g

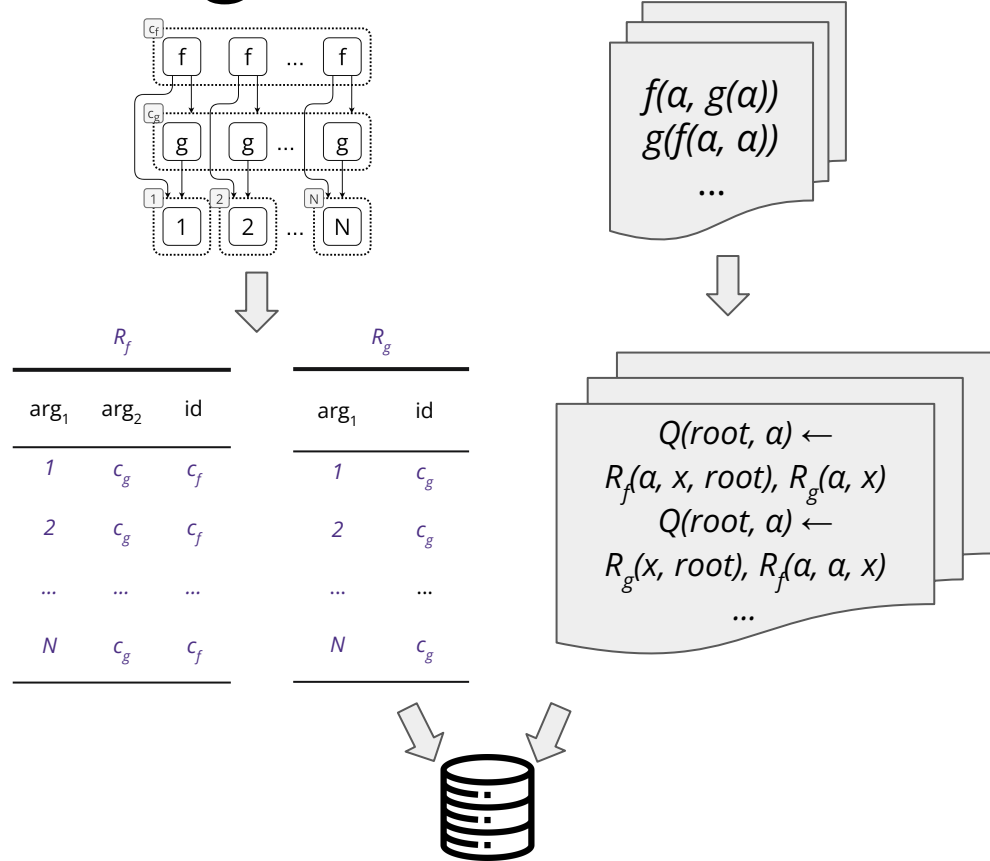
Relational e-matching

- Given e-graph + patterns
- Transform e-graph to tables
- Compile patterns to queries



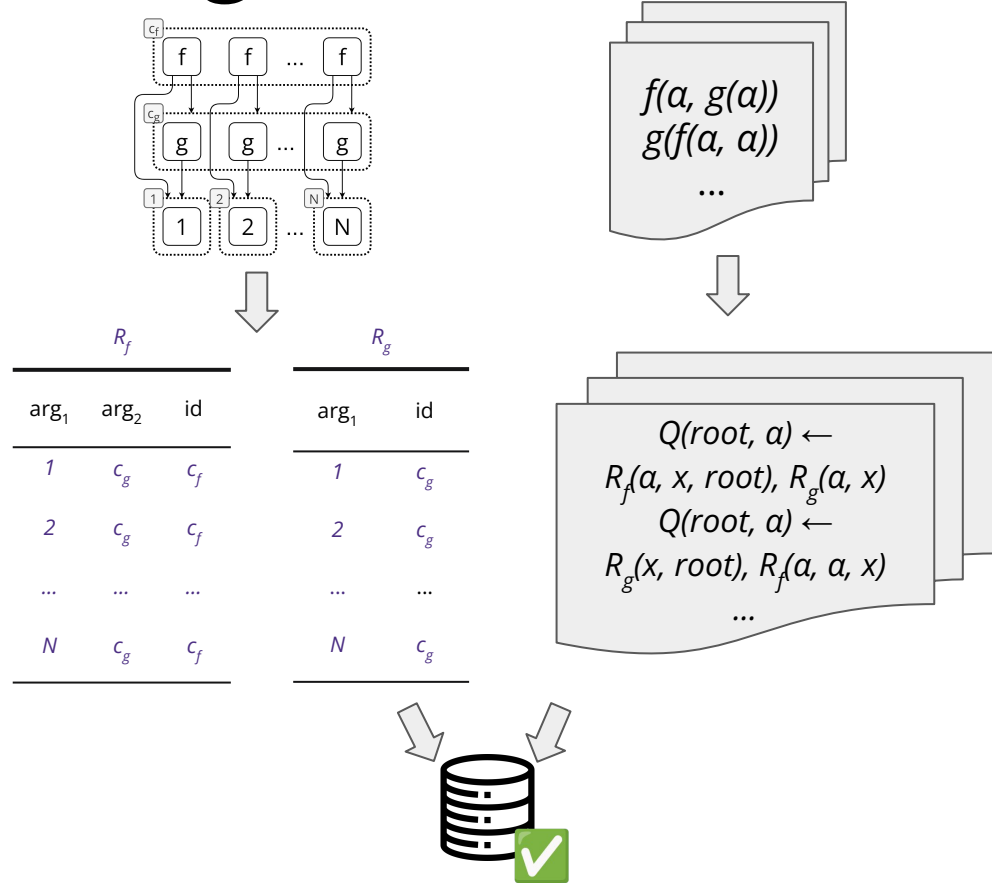
Relational e-matching

- Given e-graph + patterns
- Transform e-graph to tables
- Compile patterns to queries
- Use DB query engine to e-match!

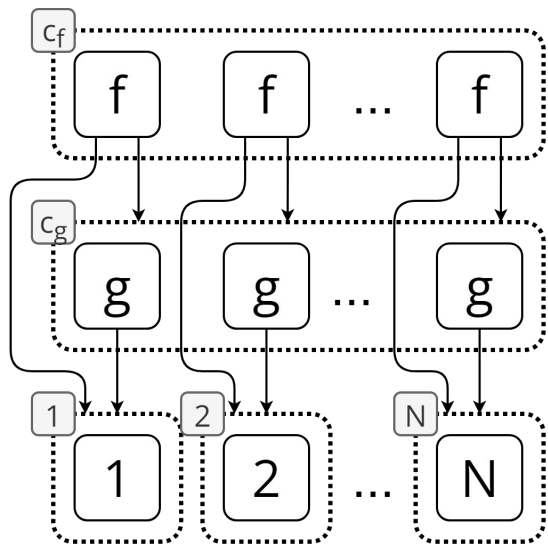


Relational e-matching

- Given e-graph + patterns
- Transform e-graph to tables
- Compile patterns to queries
- Use DB query engine to e-match!
- Derive bounds from DB theory!



E-graphs as tables (relational DBs)

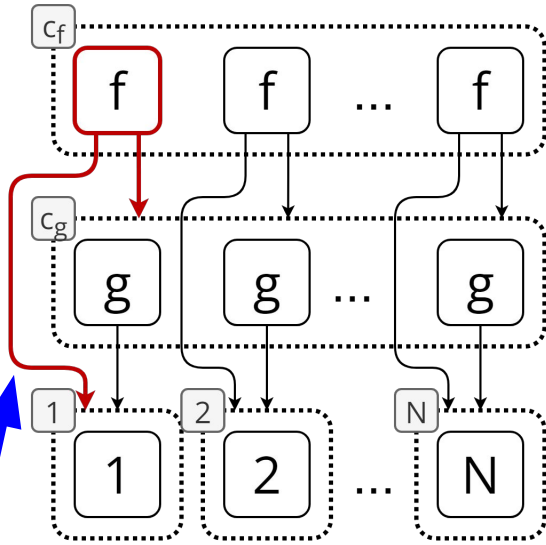


arg ₁	arg ₂	id
1	C_g	C_f
2	C_g	C_f
...
N	C_g	C_f

arg ₁	id
1	C_g
2	C_g
...	...
N	C_g

id
i

E-graphs as tables (relational DBs)



each e-node becomes a row

R_f		
arg ₁	arg ₂	id
1	C_g	C_f
2	C_g	C_f
...
N	C_g	C_f

R_g	
arg ₁	id
1	C_g
2	C_g
...	...
N	C_g

$R_{i=1\dots N}$	
id	
i	

E-match patterns as conjunctive queries

$f(\alpha, g(\alpha))$

E-match patterns as conjunctive queries

$f(\alpha, g(\alpha))$



$Q(\text{root}, \alpha) \leftarrow$

$R_f(\alpha, x, \text{root}),$

$R_g(\alpha, x)$

E-match patterns as conjunctive queries

$f(\alpha, g(\alpha))$



$Q(\text{root}, \alpha) \leftarrow$

$R_f(\alpha, x, \text{root}),$

$R_g(\alpha, x)$

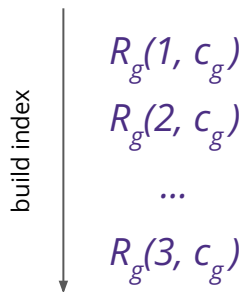


```
ind = {}
```

```
for ( $\alpha, x$ ) in  $R_{fg}$ : # build index  
  ind.insert(( $\alpha, x$ ))
```

build index

$R_g(1, c_g)$
 $R_g(2, c_g)$
...
 $R_g(3, c_g)$



E-match patterns as conjunctive queries

$f(\alpha, g(\alpha))$



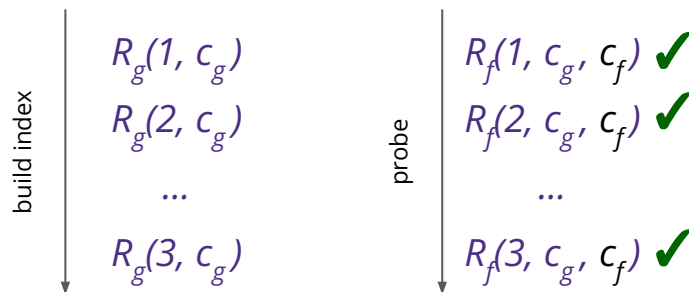
$Q(\text{root}, \alpha) \leftarrow$

$R_f(\alpha, x, \text{root}),$

$R_g(\alpha, x)$



```
ind = {}  
for ( $\alpha, x$ ) in  $R_g$ : # build index  
    ind.insert(( $\alpha, x$ ))  
for ( $\alpha, x, \text{root}$ ) in  $R_f$ : # probe  
    if ( $\alpha, x$ ) in ind:  
        yield {root  $\mapsto$  root,  $\alpha \mapsto \alpha$ }
```



Why is relational e-matching faster?

$f(\alpha, g(\alpha))$

Enum all terms of shape $f(\alpha, g(\beta))$

Check if $\alpha = \beta$ only before yielding

$Q(\text{root}, \alpha) \leftarrow$

$R_f(\text{root}, \alpha, x), R_g(x, \alpha)$

Build indices on both α and x .

Only enum terms where constraints on both x and α are satisfied.

structural
constraints

equality
constraints

Data complexity results (see paper)

THEOREM 9. *Relational e-matching is worst-case optimal; that is, fix a pattern p , let $M(p, E)$ be the set of substitutions yielded by e-matching on an e-graph E with N e-nodes, relational e-matching runs in time $O(\max_E(|M(p, E)|))$.*

THEOREM 10. *Fix an e-graph E with N e-nodes that compiles to a database I , and a fix pattern p that compiles to conjunctive query $Q(\bar{X}) \leftarrow R_1(\bar{X}_1), \dots, R_m(\bar{X}_m)$. Relational e-matching p on E runs in time $O\left(\sqrt{|Q(I)| \times \prod_i |R_i|}\right) \leq O\left(\sqrt{|Q(I)| \times N^m}\right)$.*

New Capabilities: *Multi-patterns*

```
(rule (  
  (eq X (matmul a b))  
  (eq Y (matmul a c))  
)(  
  (eq F (matmul a (concat b c)))  
  (eq X (split1 F))  
  (eq Y (split2 F))  
))
```

New Capabilities: *Multi-patterns*

```
(rule (  
  (eq X (matmul a b))  
  (eq Y (matmul a c))  
)(  
  (eq F (matmul a (concat b c)))  
  (eq X (split1 F))  
  (eq Y (split2 F))  
))
```

search for multiple patterns
anywhere in the e-graph

New Capabilities: *Multi-patterns*

```
(rule (  
  (eq X (matmul a b))  
  (eq Y (matmul a c))  
)  
(  
  (eq F (matmul a (concat b c)))  
  (eq X (split1 F))  
  (eq Y (split2 F))  
)  
)
```

search for multiple patterns
anywhere in the e-graph

perform multiple merges,
each on separate e-classes!

New Capabilities: *Merge-only rules*

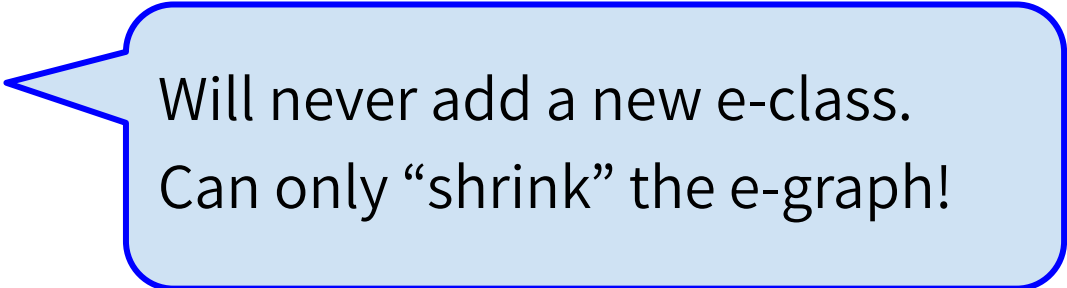
```
(rule (  
  (eq R (+ (+ a b) c))  
)(  
  (eq R (+ a (+ b c))))  
))
```

New Capabilities: *Merge-only rules*

```
(rule (  
  (eq R (+ (+ a b) c))  
)(  
  (eq R (+ a (+ b c))))  
))
```

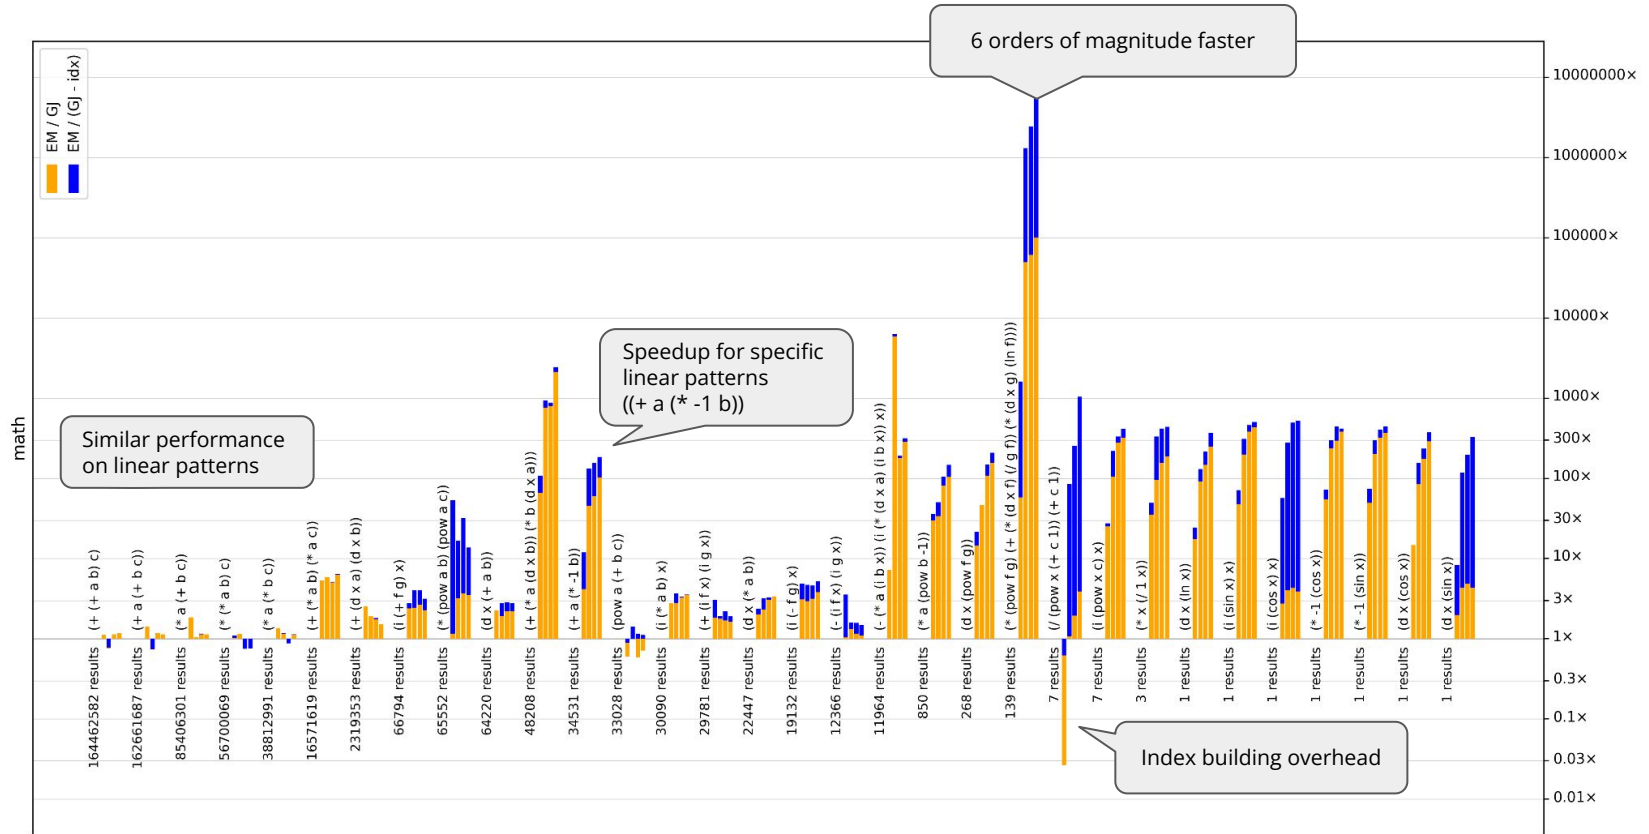
versus

```
(rule (  
  (eq R1 (+ (+ a b) c))  
  (eq R2 (+ a (+ b c))))  
)(  
  (eq R1 R2)  
))
```



Will never add a new e-class.
Can only “shrink” the e-graph!

Relational e-matching : asymptotic speedup



Graph \rightarrow Tables \rightarrow Graph \rightarrow Tables \rightarrow ...

- Relational e-match prototype: repeated **graph** \leftrightarrow **DB** conversion
- **egglog** : fully embrace DB
 - Now EqSat “just” Datalog + congruence!
 - Need to keep fast invariant maintenance
 - Want to leverage Datalog strengths as well

Congruence: just a functional dependence!

Remember key e-graph invariant: $\mathbf{a = b \Rightarrow f(a) = f(b)}$

Congruence: just a functional dependence!

Remember key e-graph invariant: $a = b \Rightarrow f(a) = f(b)$

Rebuilding *still* amortizes maintaining this invariant in DB!

For terms, congruence violation manifests as “same args, diff class”

Congruence: just a functional dependence!

Remember key e-graph invariant: $a = b \Rightarrow f(a) = f(b)$

Rebuilding *still* amortizes maintaining this invariant in DB!

For terms, congruence violation manifests as “same args, diff class”

Foo

ArgL ArgR EC

A	B	X
A	B	Y

Congruence: just a functional dependence!

Remember key e-graph invariant: $a = b \Rightarrow f(a) = f(b)$

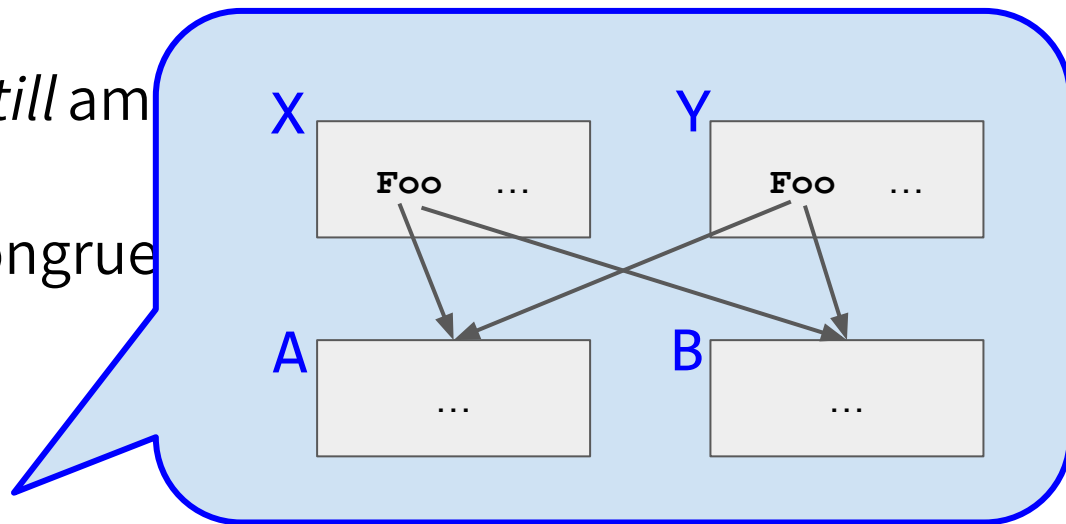
Rebuilding *still* am

For terms, congrue

Foo

ArgL ArgR EC

A	B	X
A	B	Y



DB!

args, diff class”

Congruence: just a functional dependence!

Remember key e-graph invariant: $a = b \Rightarrow f(a) = f(b)$

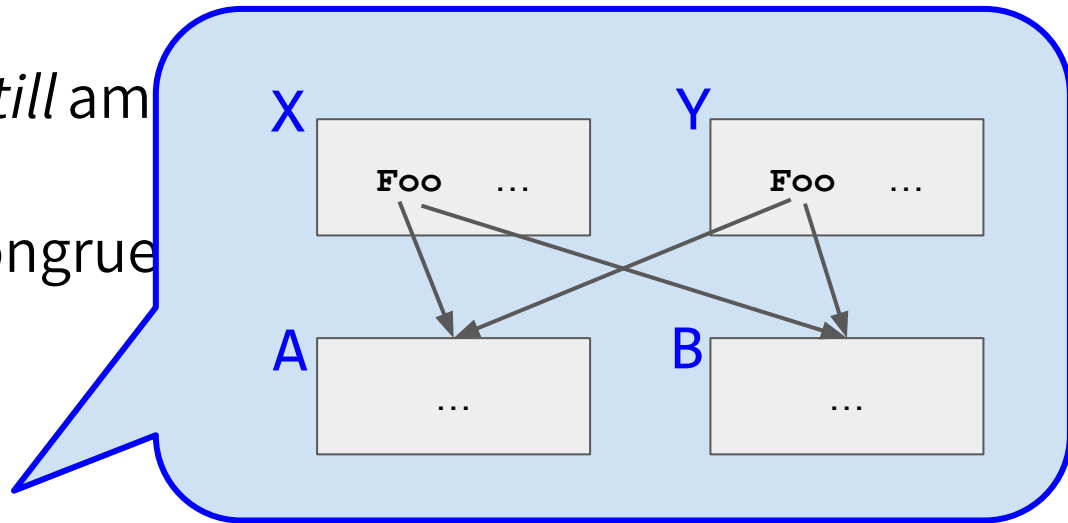
Rebuilding *still* am

For terms, congrue

DB!

args, diff class”

Foo		
ArgL	ArgR	EC
A	B	X
A	B	Y



In tables = FD violation!

Congruence: just a functional dependence!

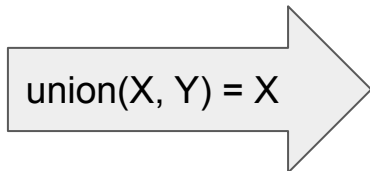
Remember key e-graph invariant: $a = b \Rightarrow f(a) = f(b)$

Rebuilding *still* amortizes maintaining this invariant in DB!

For terms, congruence violation manifests as “same args, diff class”

Foo

ArgL	ArgR	EC
A	B	X
A	B	Y



Foo

ArgL	ArgR	EC
A	B	X

Congruence: just a functional dependence!

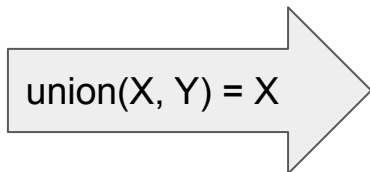
Remember key e-graph invariant: $a = b \Rightarrow f(a) = f(b)$

Rebuilding *still* amortizes maintaining this invariant in DB!

For terms, congruence violation manifests as “same args, diff class”

Foo

ArgL	ArgR	EC
A	B	X
A	B	Y



Foo

ArgL	ArgR	EC
A	B	X

May cause additional
congruence merges!

Congruence: just a functional dependence!

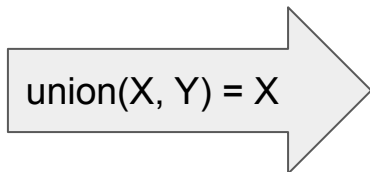
Remember key e-graph invariant: $a = b \Rightarrow f(a) = f(b)$

Rebuilding *still* amortizes maintaining this invariant in DB!

For terms, congruence violation manifests as “same args, diff class”

Foo

ArgL	ArgR	EC
A	B	X
A	B	Y



Foo

ArgL	ArgR	EC
A	B	X

May cause additional congruence merges!

Merge in the union find and keep rebuilding to fixpoint!

New Capabilities: *Incremental e-matching*

Just Datalog semi-naive eval!

$$Q = R \bowtie S \bowtie T$$

$$\Delta Q = \Delta R \bowtie S \bowtie T$$

$$U \quad R \bowtie \Delta S \bowtie T$$

$$U \quad R \bowtie S \bowtie \Delta T$$

PL Problems

1. Phase Ordering 🥲 ⇒ 🤩

2. Syntactic Brittleness 🥲

3. ...

Keep simple + modular!

But make robust?

DB techniques will
come to our rescue!

PL Problems

1. Phase Ordering 🥲 ⇒ 🤪

2. Syntactic Brittleness 🥲

3. ...

Keep simple + modular!

But make robust?

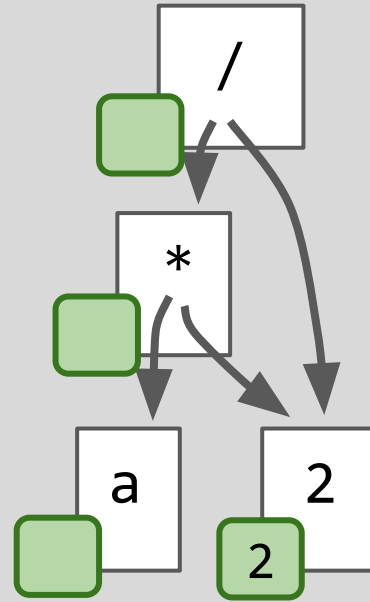
DB techniques will
come to our rescue!

Syntactic rewriting is not enough...

- How many rules do we need for constant folding?
 - $2 + 2 \rightarrow 4$, $3 + 4 \rightarrow 7$, $4 + 6 \rightarrow 10$, ... *a lot!*
- What about satisfying guards for conditional rules?
 - $x / x \rightarrow 1$ only ok if $x \neq 0$
- In general, many optimizations depend on analyses!
 - nullability, tensor shape, intervals, free variables, ...

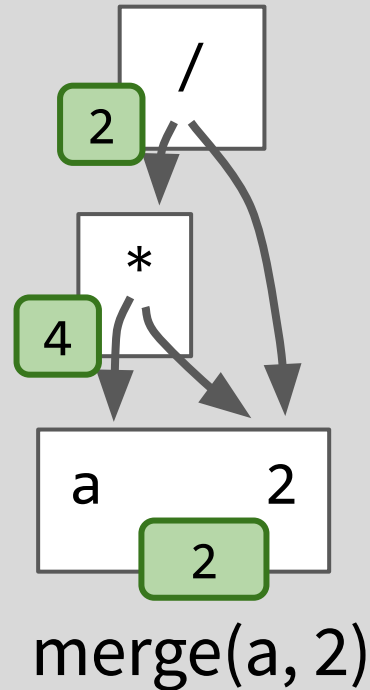
Constant folding

- Option<Number> “fact” per e-class
- Try to eval new e-nodes
- Option “or” on merge



Constant folding

- Option<Number> “fact” per e-class
- Try to eval new e-nodes
- Option “or” on merge
- Only propagates up!



Original egg *e-class analyses* too limited:

- Facts could only propagate from children to parents
 - Injectivity difficult for f where $f(x) = f(y) \rightarrow x = y$
- Analyses / rules implemented + executed separately
 - Only “one kind” of fact per domain (products clunky)
 - Custom Rust code to realize conditional rewrites
 - Difficult to give unified theory of EqSat! Termination?

Program analysis modulo equivalence

- Tightest summary over all equivalent represented terms!

To demonstrate an advantage of this approach, consider the following example, for $x \in [0, 1]$, $y \in [1, 2]$, where the following concrete-equivalences are discovered via rewriting:

$$1 - \frac{2y}{x+y} \in \left[-3, \frac{1}{3}\right] \quad (5)$$

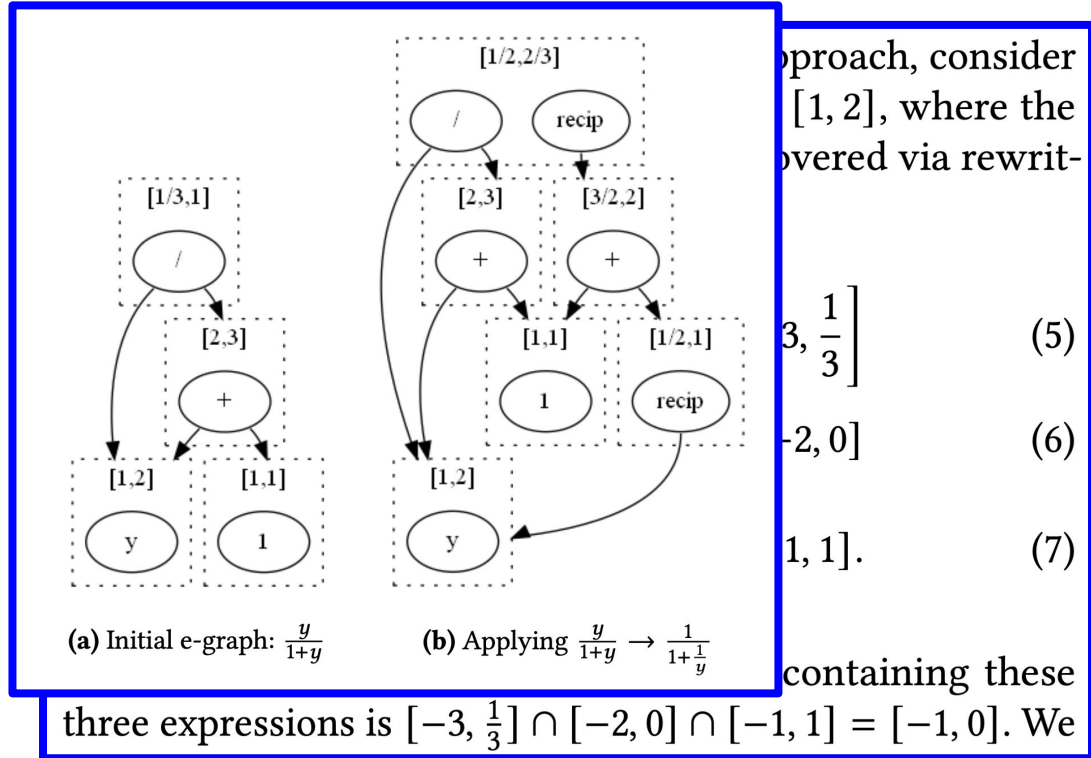
$$\cong \frac{x-y}{x+y} \in [-2, 0] \quad (6)$$

$$\cong \frac{2x}{x+y} - 1 \in [-1, 1]. \quad (7)$$

The interval associated with the e-class containing these three expressions is $[-3, \frac{1}{3}] \cap [-2, 0] \cap [-1, 1] = [-1, 0]$. We

Program analysis modulo equivalence

- Tightest summary over all equivalent represented terms!
- *Virtuous cycle*: facts enable rewrites, rewrites improve facts!




egglog: Analyses via FDs and “merge functions”

- **Key Idea** : table for function F has args \rightarrow e-class FD


egglog: Analyses via FDs and “merge functions”

- **Key Idea** : table for function F has args \rightarrow e-class FD
- For “terms” enforcing FD should restore congruence

egglog: Analyses via FDs and “merge functions”

- **Key Idea** : table for function F has args \rightarrow e-class FD
- For “terms” enforcing FD should restore congruence 
- But what about analysis facts? Can't just union...

egglog: Analyses via FDs and “merge functions”

- **Key Idea** : table for function F has args \rightarrow e-class FD
- For “terms” enforcing FD should restore congruence 
- But what about analysis facts? Can't just union...
 - Facts are also just functions!
 - **Allow user-specified FD repair!**

egglog: FDs and “merge functions”

```
(datatype Math
  (Num Rational)
  (Var String)
  (Const String)
  (Add Math Math)
  (Sub Math Math)
  (Mul Math Math)
  (Div Math Math)
  (Neg Math)
  (Sqrt Math)
  ...)
```

egglog: FDs and “merge functions”

```
(datatype Math
  (Num Rational)
  (Var String)
  (Const String)
  (Add Math Math)
  (Sub Math Math)
  (Mul Math Math)
  (Div Math Math)
  (Neg Math)
  (Sqrt Math)
  ...)
```

```
(function hi (Math) Rational :merge min)
(function lo (Math) Rational :merge max)
(relation non-zero (Math))
```

egglog: FDs and “merge functions”

```
(datatype Math
  (Num Rational)
  (Var String)
  (Const String)
  (Add Math Math)
  (Sub Math Math)
  (Mul Math Math)
  (Div Math Math)
  (Neg Math)
  (Sqrt Math)
  ...)
```

```
(function hi (Math) Rational :merge min)
(function lo (Math) Rational :merge max)
(relation non-zero (Math))
```

```
(rewrite (Add (Num a) (Num b))
         (Num (+ a b)))
(rewrite (Sub (Num a) (Num b))
         (Num (- a b)))
(rewrite (Mul (Num a) (Num b))
         (Num (* a b)))
(rewrite (Div (Num a) denom)
         (Num (/ a b))
         :when ((= denom (Num b))
                (non-zero denom)))
```

```
...
```

egglog: FDs and “merge functions”

```
(datatype Math
  (Num Rational)
  (Var String)
  (Const String)
  (Add Math Math)
  (Sub Math Math)
  (Mul Math Math)
  (Div Math Math)
  (Neg Math)
  (Sqrt Math)
  ...)
```

```
(function hi (Math) Rational :merge min)
(function lo (Math) Rational :merge max)
(relation non-zero (Math))
```

```
(rule ((= e (Add a b))
      (= la (lo a))
      (= lb (lo b)))
      ((set (lo e) (+ la lb))))

(rule ((= e (Add a b))
      (= ha (hi a))
      (= hb (hi b)))
      ((set (hi e) (+ ha hb))))
```

egglog: FDs and “merge functions”

```
(datatype Math
  (Num Rational)
  (Var String)
  (Const String)
  (Add Math Math)
  (Sub Math Math)
  (Mul Math Math)
  (Div Math Math)
  (Neg Math)
  (Sqrt Math)
  ...)
```

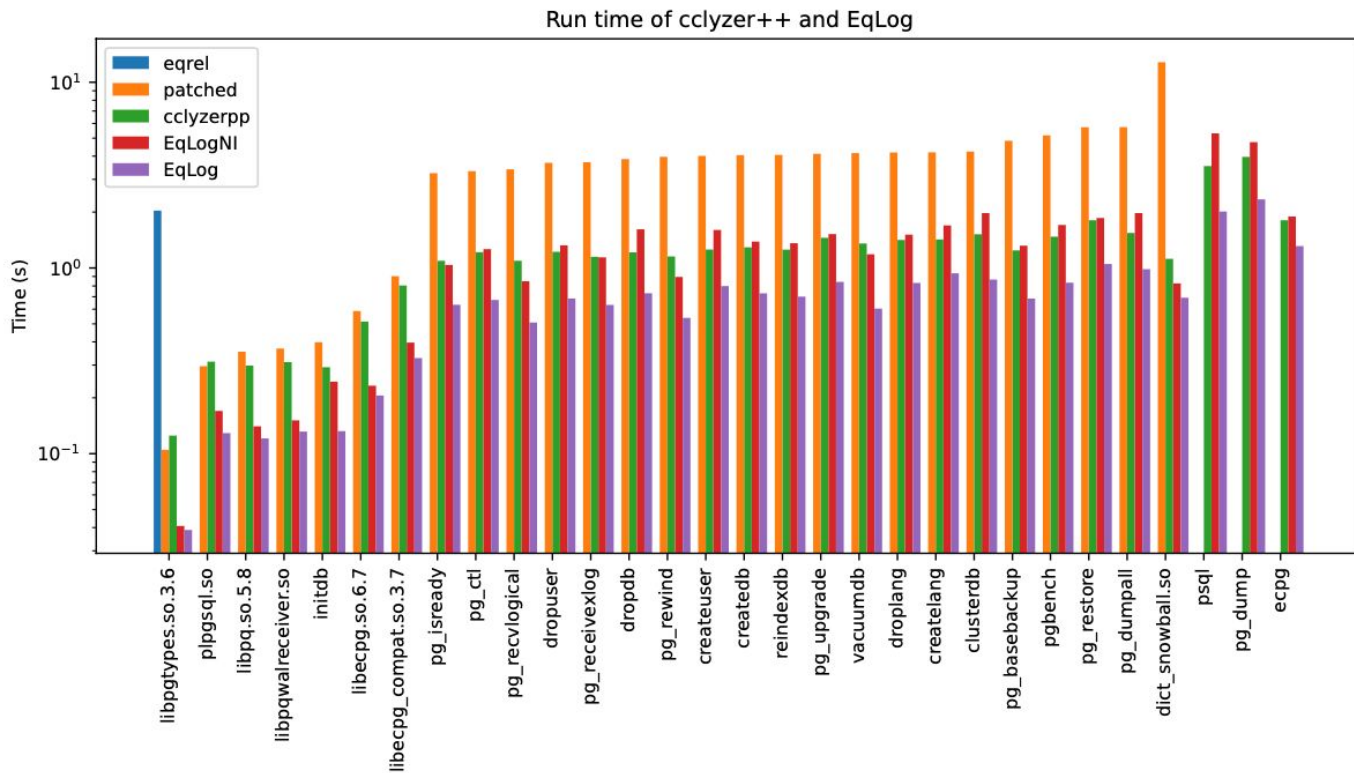
```
(function hi (Math) Rational :merge min)
(function lo (Math) Rational :merge max)
(relation non-zero (Math))
```

```
(rule ((= l (lo e))
      (> l r-zero))
      ((non-zero e)))

(rule ((= h (hi e))
      (< h r-zero))
      ((non-zero e)))

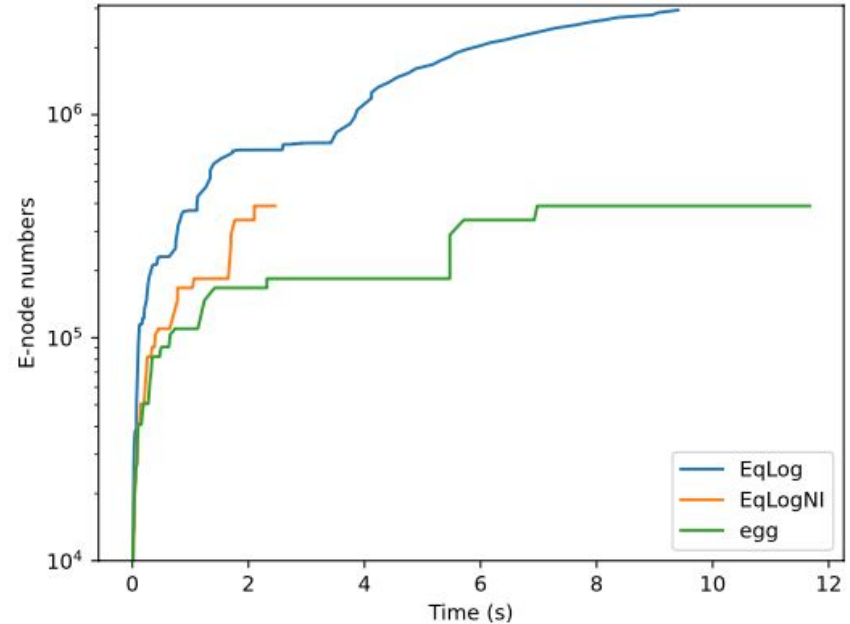
(rule ((= e (Num ve))
      (set (lo e) ve)
      (set (hi e) ve)))
```

egglog Pointer Analysis



Huge egglog Wins

- No graph \Leftrightarrow DB conversions
- Analyses now “just” rewrites
- Semi-naive evaluation \rightarrow incremental e-matching!
- Provides easy text format
- Key: flexible framework!!
- Support unified EqSat/DB theory



PL Problems

1. Phase Ordering 🥲 ⇒ 😄

Keep simple + modular!

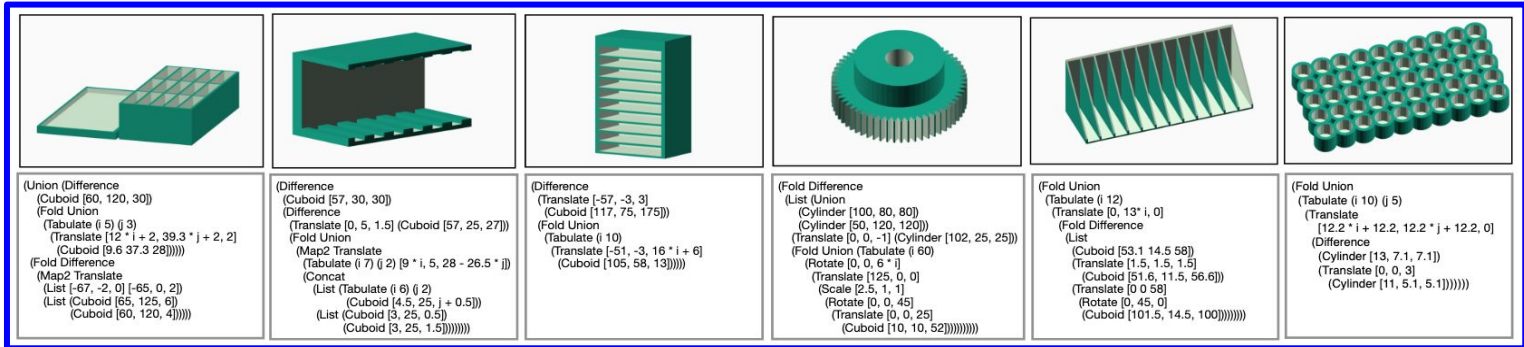
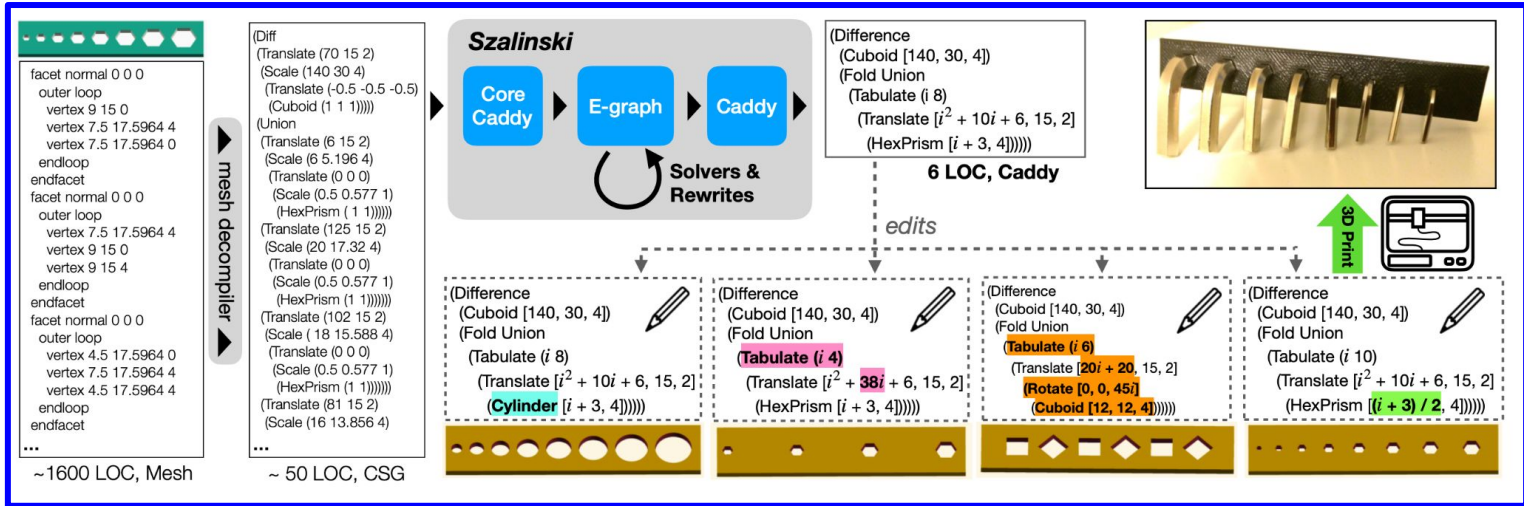
But make robust?

2. Syntactic Brittleness 🥲 ⇒ 😄

DB techniques will
come to our rescue!

3. ...

Decompile 3D CAD [ICFP 2018, PLDI 2020, POPL 2023]

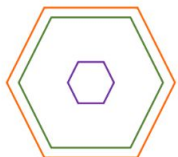


Decompile 3D CAD [ICFP 2018, PLDI 2020, POPL 2023]

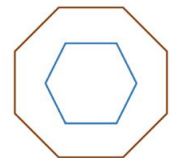
A Initial corpus (size 208)



```
(combine
  (xform
    (repeat (xform line ...) 6 [1 (2π / 6) 0 0])
    [2 0 0 0])
  (xform
    (repeat (xform line ...) 6 [1 (2π / 6) 0 0])
    [1 0 0 0]))
```



```
(combine (combine
  (xform
    (repeat (xform line ...) 6 [1 (2π / 6) 0 0])
    [2.25 0 0 0]))
  (xform
    (repeat (xform line ...) 6 [1 (2π / 6) 0 0])
    [2 0 0 0]))
  (xform
    (repeat (xform line ...) 6 [1 (2π / 6) 0 0])
    [0.5 0 0 0]))
```



```
(combine
  (xform
    (repeat (xform line ...) 8 [1 (2π / 8) 0 0])
    [2 0 0 0])
  (xform
    (repeat (xform line ...) 6 [1 (2π / 6) 0 0])
    [1 0 0 0]))
```

B Compressed corpus (size 72)

library:

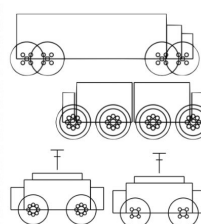
```
f0 = λx0 x1 -> // polygon with x1 sides scaled by x0
  xform
    (repeat (xform line ...) x1 [1 (2π / x1) 0 0])
    [x0 0 0 0]
```

refactored programs:

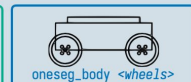
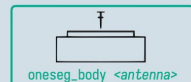
```
(combine
  (f0 6 2)
  (f0 6 1))
```

```
(combine (combine
  (f0 6 2.25)
  (f0 6 2))
  (f0 6 0.5))
```

library
learning

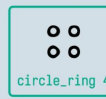


```
oneseg_body = λrest ->
  C (C (C (T (T (r_s 0 0) (xform_x 0)) (xform_x -8))
    (T (T (r_s 16 4.5) (M 1 0 0 2.25))
      (xform_x 0)))
    (T (T (r_s 0 0) (xform_x 0)) (xform_x 8)))
  (T (r_s 12 1) (M 1 0 0 5))) rest
```



one-segment vehicle body

```
circle_ring = λn -> repeat
  (T (T c (M 0.25 0 0 0))
    (M 1 0 0.53033 0.53033))
  n
  (M 1 ((2 * π) / x0) 0 0)
```



ring of circles

Compile \mathbb{R} to IEEE 754 [PLDI 2015, ARITH 2021, ASPLOS 2025]

$$\frac{(-b) + \sqrt{b \cdot b - 4 \cdot (a \cdot c)}}{2 \cdot a}$$

↓

if $b \leq -2.1714197031320663 \cdot 10^{+114}$:

$$-\frac{b}{a}$$

elif $b \leq 2.9809086538561536 \cdot 10^{-153}$:

$$\frac{\sqrt{\text{fma}(b, b, c \cdot (a - 4))} - b}{a \cdot 2}$$

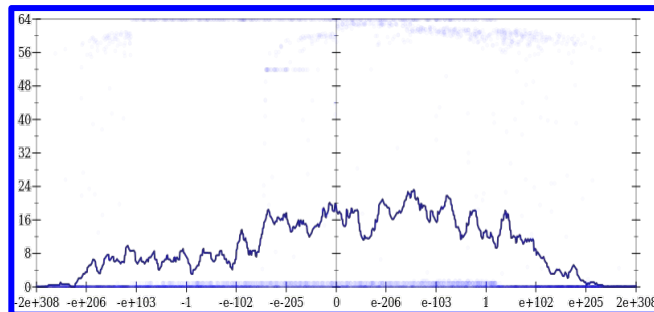
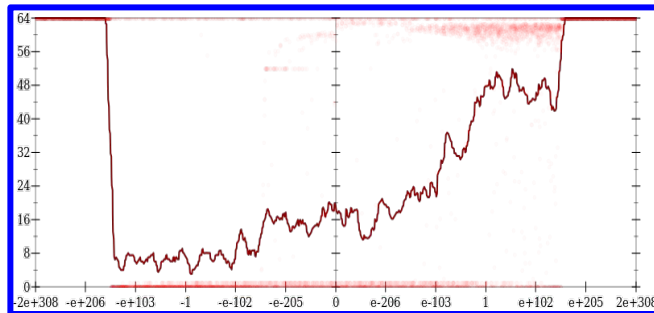
elif $b \leq 3.095118518558678 \cdot 10^{+20}$:

$$t_0 := 4 \cdot (a \cdot c)$$

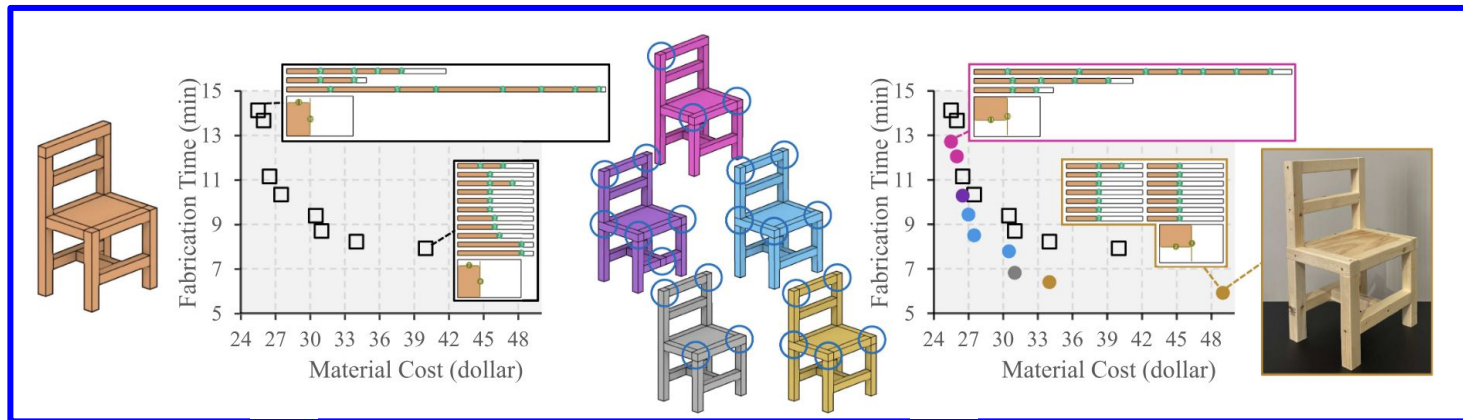
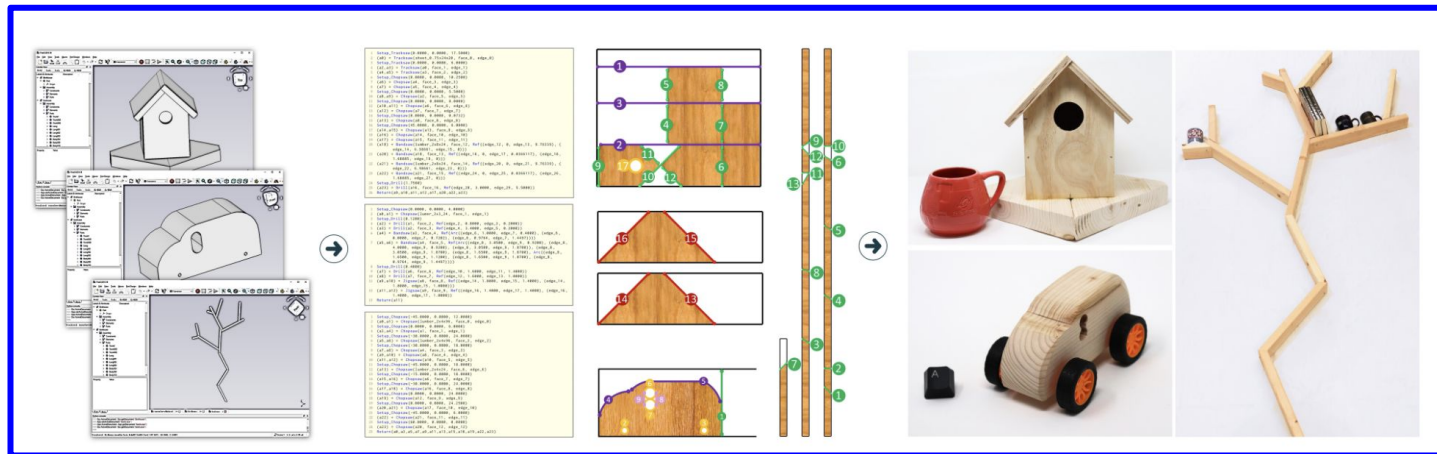
$$\frac{t_0 \cdot \frac{0.5}{a}}{(-b) - \sqrt{b \cdot b - t_0}}$$

else :

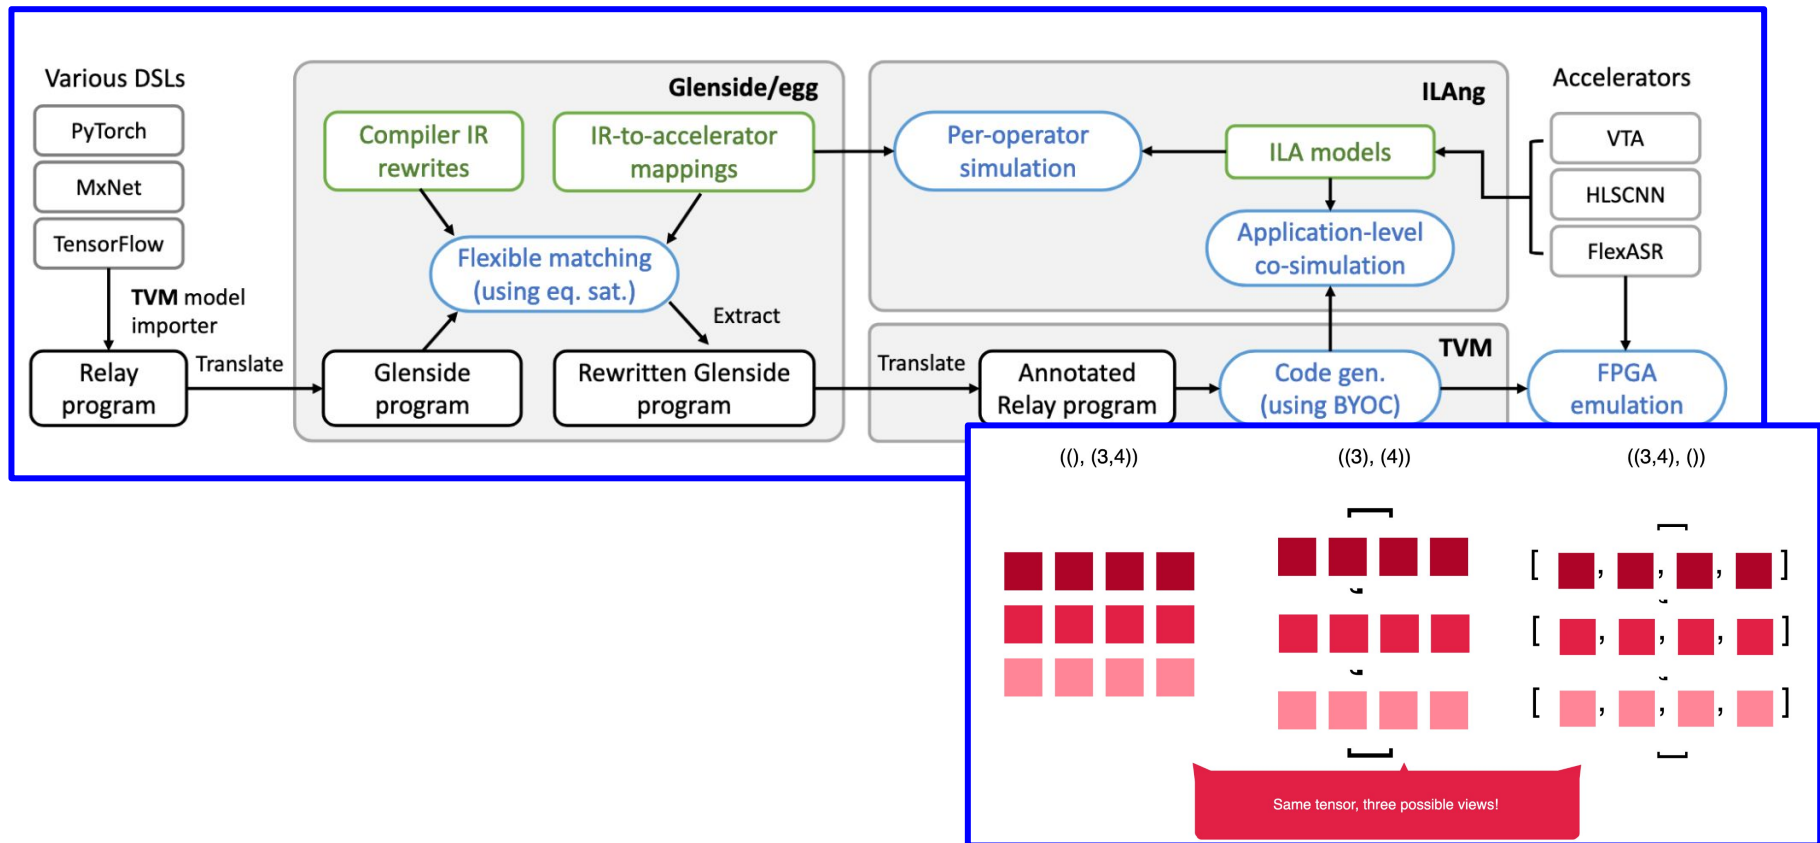
$$-\frac{c}{b}$$



Carpentry Compiler [SIGGRAPH ASIA 2019, TOG 2022]



Targeting HW Accelerators [MAPS 2021, TODAES 2023]



Many Other Applications!



Automatic Generation of Vectorizing Compilers for Customizable Digital Signal Processors

Samuel Thomas
sgt@cs.utexas.edu
The University of Texas at
Austin, TX, USA

Vectorization for Digital Signal Processors via Equality Saturation

Alexa VanHattum

Rachit Nigam
Cornell University
Ithaca, NY, USA

Vincent T. Lee
Facebook Reality Labs Research
Redmond, WA, USA

SPORES: SUM-PRODUCT OPTIMIZATION VIA RELATIONAL EQUALITY SATURATION FOR LARGE SCALE LINEAR ALGEBRA

Adrian Sampson
Cornell University
Ithaca, NY, USA

Yisu Remy Wang
University of Washington
Seattle, Washington
remywang@cs.washington.edu

Shana Hutchison
University of Washington
Seattle, Washington
shutchis@cs.washington.edu

Jonathan Lear
University of Washington
Seattle, Washington
jleang@cs.washington.edu

Bill Howe
University of Washington
Seattle, Washington
billhowe@cs.washington.edu

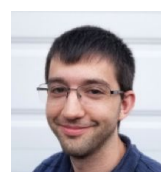
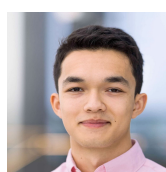
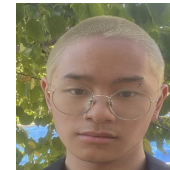
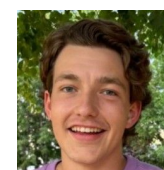
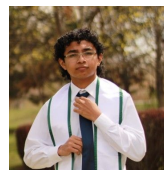
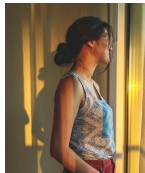
Dan Suciu
University of Washington
Seattle, Washington
suciu@cs.washington.edu

Automating Constraint-Aware Datapath Optimization using E-Graphs

Samuel Coward
Numerical Hardware Group
Intel Corporation
Email: samuel.coward@intel.com

George A. Constantinides
Electrical and Electronic Engineering
Imperial College London
Email: g.constantinides@imperial.ac.uk

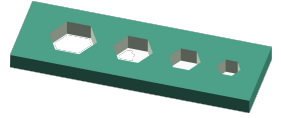
Theo Drane
Numerical Hardware Group
Intel Corporation
Email: theo.drane@intel.com



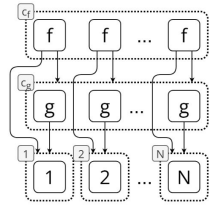
Thank You!

[egraphs-good.github.io](https://github.com/egraphs-good) • egraphs.org

1. Rule-based transformations everywhere



2. E-graphs and Equality Saturation (EqSat)



3. Adapting DB results: more speed and flexibility

