

# Thinking Fast and Correct: Automated Rewriting of Numerical Code through Compiler Augmentation



**Siyuan Brant Qian**<sup>1</sup> Vimarsh Sathia<sup>1</sup> Jan Hückelheim<sup>2</sup> Paul Hovland<sup>2</sup> William S. Moses<sup>1</sup>

siyuanq4@illinois.edu

<sup>1</sup> University of Illinois Urbana-Champaign, USA

<sup>2</sup> Argonne National Laboratory, USA

CGO 2026

Feb. 3, 2026

Optimizing numerical code is challenging at three levels:

**Context, Scope, and Scale.**

**Profiling + Compiler Optimization** are the key.



# Context: Where You Change Matters

---

```
double sum = 0.0;
for (double x : xs) {
    double val = sigmoid(exp(x));
    sum += val;
}
```



# Context: Where You Change Matters

---

```
double sum = 0.0;
for (double x : xs) {
    double val = sigmoid(exp(x));
    sum += val;
}
```

```
double sum = 0.0;
for (double x : xs) {
    float val = sigmoidf(expf(x));
    sum += val;
}
```



Lower Activation:  
1.6x speedup  
**error  $\approx 1e-12$  (safe!)**



# Context: Where You Change Matters

```
double sum = 0.0;
for (double x : xs) {
    double val = sigmoid(exp(x));
    sum += val;
}
```

```
double sum = 0.0;
for (double x : xs) {
    float val = sigmoidf(expf(x));
    sum += val;
}
```

```
float sum = 0.0f;
for (double x : xs) {
    float val = sigmoidf(expf(x));
    sum += val;
}
```

Lower Activation:  
1.6x speedup  
**error  $\approx 1e-12$  (safe!)**

Lower Reduction:  
**error  $\approx 1e-4$  (danger!)**



# Context: Where You Change Matters

```
double sum = 0.0;
for (double x : xs) {
    double val = sigmoid(exp(x));
    sum += val;
}
```

```
double sum = 0.0;
for (double x : xs) {
    float val = sigmoidf(expf(x));
    sum += val;
}
```

```
float sum = 0.0f;
for (double x : xs) {
    float val = sigmoidf(expf(x));
    sum += val;
}
```

Profiling just 10 iterations:  
Reduction is  $10^4 \times$  more sensitive\*!

Lower Activation:  
1.6x speedup  
**error  $\approx 1e-12$  (safe!)**

Lower Reduction:  
**error  $\approx 1e-4$  (danger!)**

\* ADAPT metric (Menon et al. '18).



# Context: *Structured Critical Value Hypothesis*

```
double sum = 0.0;
for (double x : xs) {
    double val = sigmoid(exp(x));
    sum += val;
}
```

```
double sum = 0.0;
for (double x : xs) {
    float val = sigmoidf(expf(x));
    sum += val;
}
```

```
float sum = 0.0f;
for (double x : xs) {
    float val = sigmoidf(expf(x));
    sum += val;
}
```

Profiling just 10 iterations:  
Reduction is  $10^4 \times$  more sensitive\*!

Accuracy-critical values are often **structural** (e.g., reduction, cancellation, one-hot, thresholds).  
A **small surrogate profile** reveals key information.



# Scope: Precision Isn't the Only Knob

---

FP32 (Original Order)

⚡ Fast

✗ Wrong

$$(1e8 + 1) - 1e8 \rightarrow 1.00000000e8 - 1e8 \rightarrow 0$$

+1 rounded away



# Scope: Precision Isn't the Only Knob

FP32 (Original Order)

$$(1e8 + 1) - 1e8 \rightarrow$$

$$1.00000000e8$$

+1 rounded away

⚡ Fast

✗ Wrong

$$- 1e8 \rightarrow 0$$

FP64 (Same Order)

$$(1e8 + 1) - 1e8 \rightarrow$$

$$1.000000001e8$$

🐢 Costly

✓ Correct

$$- 1e8 \rightarrow 1$$

Use FP64  
(Tradeoff!)



# Scope: Precision Isn't the Only Knob

FP32 (Original Order)

⚡ Fast

✗ Wrong

$$(1e8 + 1) - 1e8 \rightarrow 1.0000000e8 - 1e8 \rightarrow 0$$

+1 rounded away

FP64 (Same Order)

🐢 Costly

✓ Correct

$$(1e8 + 1) - 1e8 \rightarrow 1.00000001e8 - 1e8 \rightarrow 1$$

Use FP64  
(Tradeoff!)

FP32 (Reassociated)

⚡ Fast

✓ Correct

$$(1e8 - 1e8) + 1 \rightarrow 0 + 1 \rightarrow 1$$

Algebraic Rewrite  
(circumvents tradeoff  
but requires context)



# Scale: The Math Is in the Mess

---

```
double foo(double x) {  
    double y = 0.0;  
    if (x > 0.0)  
        y = pow(x, 3);  
    return y;  
}
```

Math behind control flow,  
memory I/O,  
function calls, ...;  
existing source/binary-level  
tools fail.



# Scale: The Math Is in the Mess

---

```
double foo(double x) {  
    double y = 0.0;  
    if (x > 0.0)  
        y = pow(x, 3);  
    return y;  
}
```

Math behind **control flow**,  
**memory I/O**,  
**function calls**, ...;  
existing source/binary-level  
tools fail.



A production compiler makes math  
*explicit* through optimizations  
(**SimplifyCFG**, **mem2reg**, **inlining**, ...)



# Scale: The Math Is in the Mess

```
double foo(double x) {  
    double y = 0.0;  
    if (x > 0.0)  
        y = pow(x, 3);  
    return y;  
}
```

Math behind control flow,  
memory I/O,  
function calls, ...;  
existing source/binary-level  
tools fail.

```
define double @foo(double %x) {  
entry:  
    %p = call double @powi(  
        double %x, i32 3)  
    %r = call double @max(  
        double 0.0, double %p)  
    ret double %r  
}
```

A production compiler makes math  
*explicit* through optimizations  
(SimplifyCFG, mem2reg, inlining, ...)



# Scale: The Math Is in the Mess

```
double foo(double x) {  
    double y = 0.0;  
    if (x > 0.0)  
        y = pow(x, 3);  
    return y;  
}
```

Math behind control flow,  
memory I/O,  
function calls, ...;  
existing source/binary-level  
tools fail.

```
define double @foo(double %x) {  
entry:  
    %p = call double @powi(  
        double %x, i32 3)  
    %r = call double @max(  
        double 0.0, double %p)  
    ret double %r  
}
```

A production compiler makes math  
*explicit* through optimizations  
(SimplifyCFG, mem2reg, inlining, ...)

Aha, caller wants  $\max(0, x^3)$ !



Can we **automate** numerical rewriting techniques in compilers?

How much performance are we **leaving on the table** due to *suboptimal* choices of precision and expressions?



# Our Answer – Poseidon

---

- The first framework that automates numerical rewrites within a production compiler



# Our Answer – Poseidon

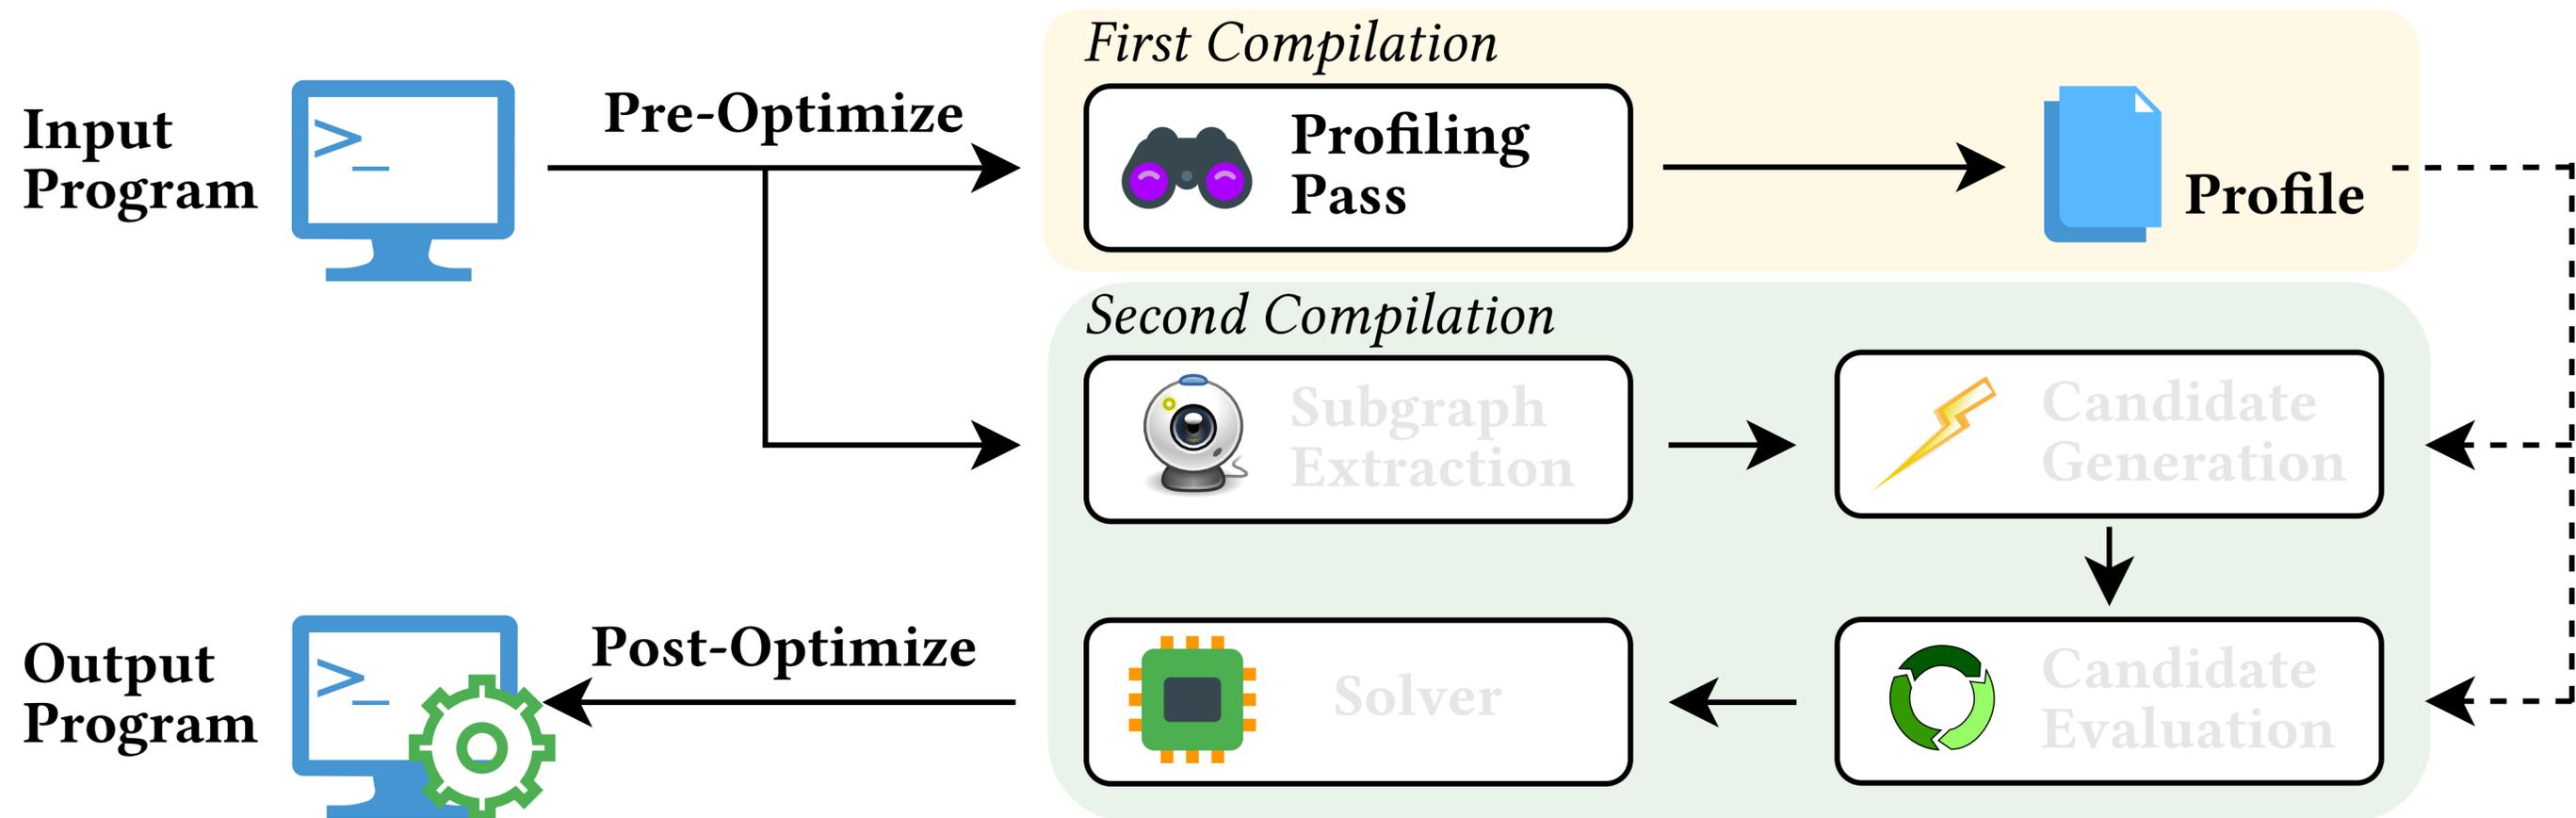
---

- The first framework that automates numerical rewrites within a production compiler
- **First Compile:** Instrument user program to collect numerical context



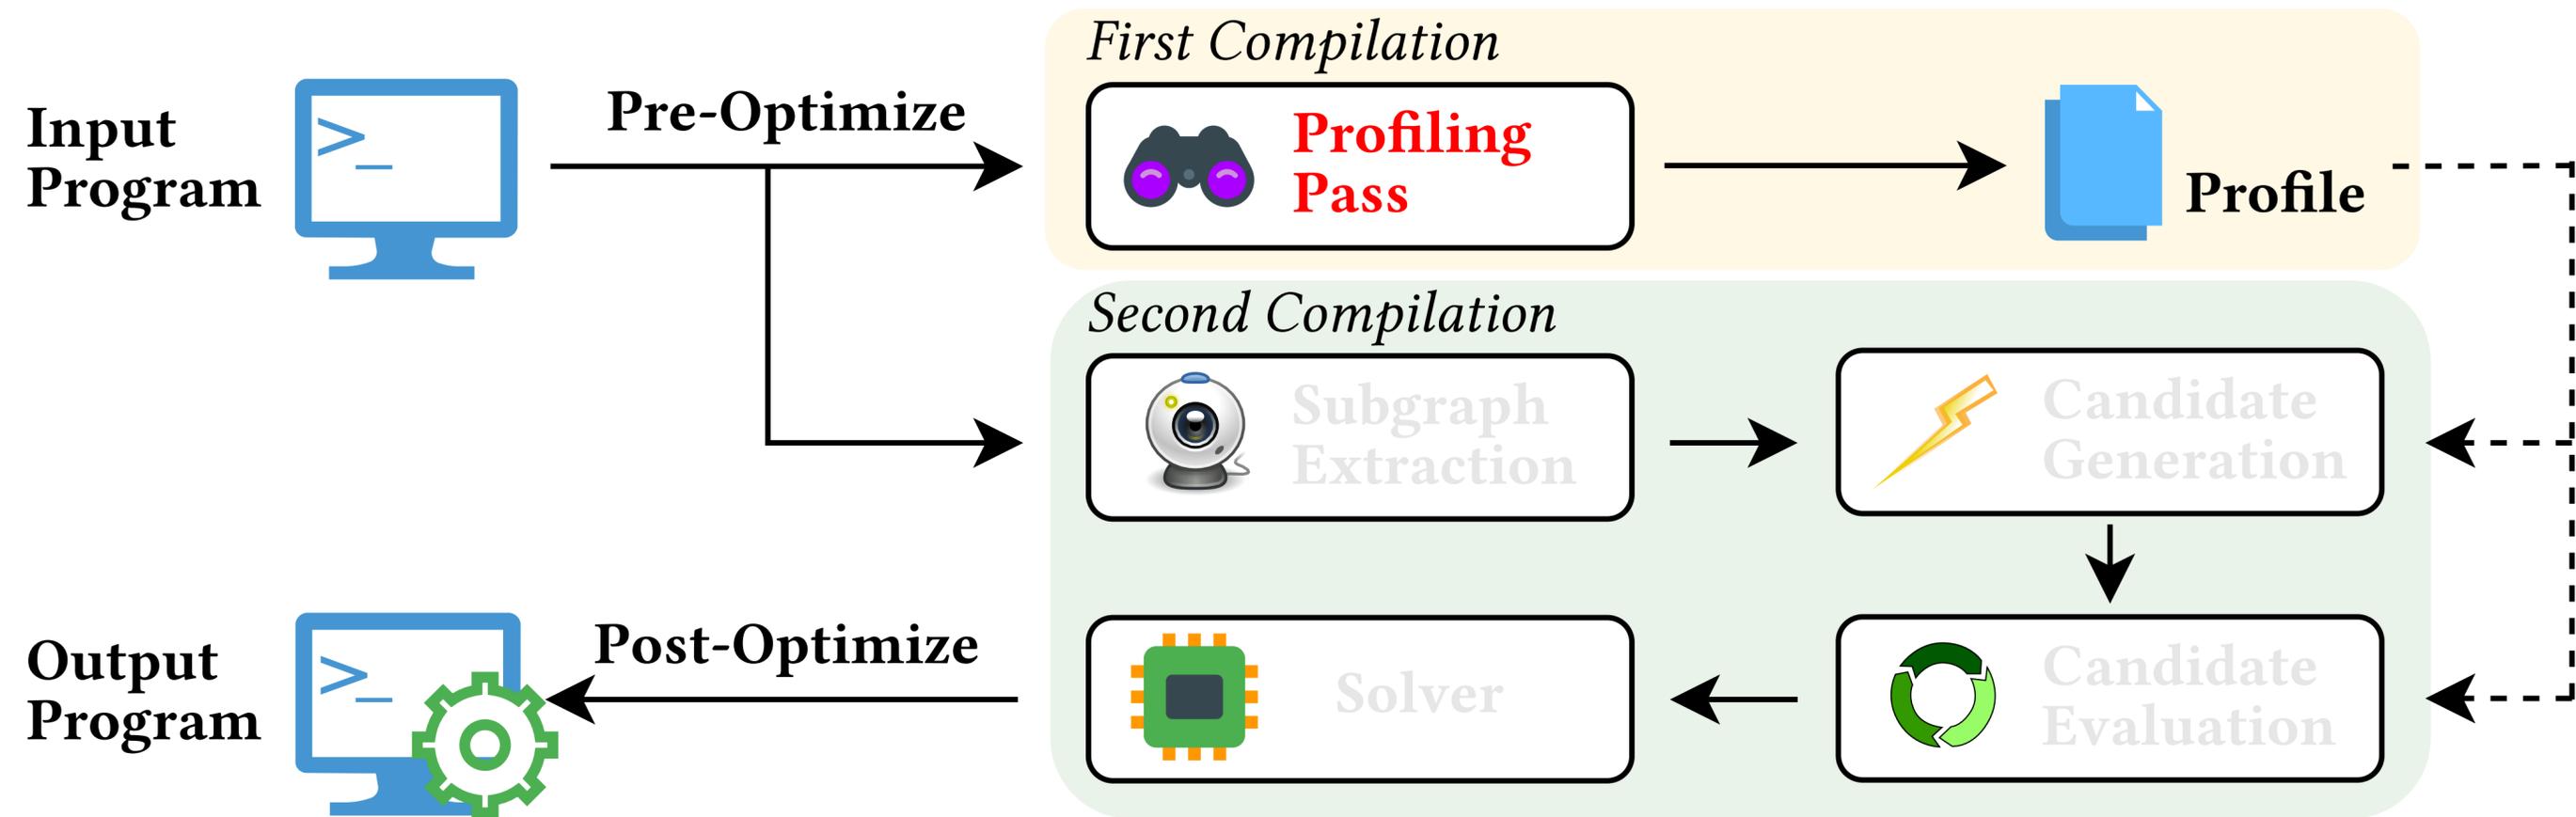
# Our Answer – Poseidon

- The first framework that automates numerical rewrites within a production compiler
- **Second Compile:** Perform *full-application scale* numerical rewrites



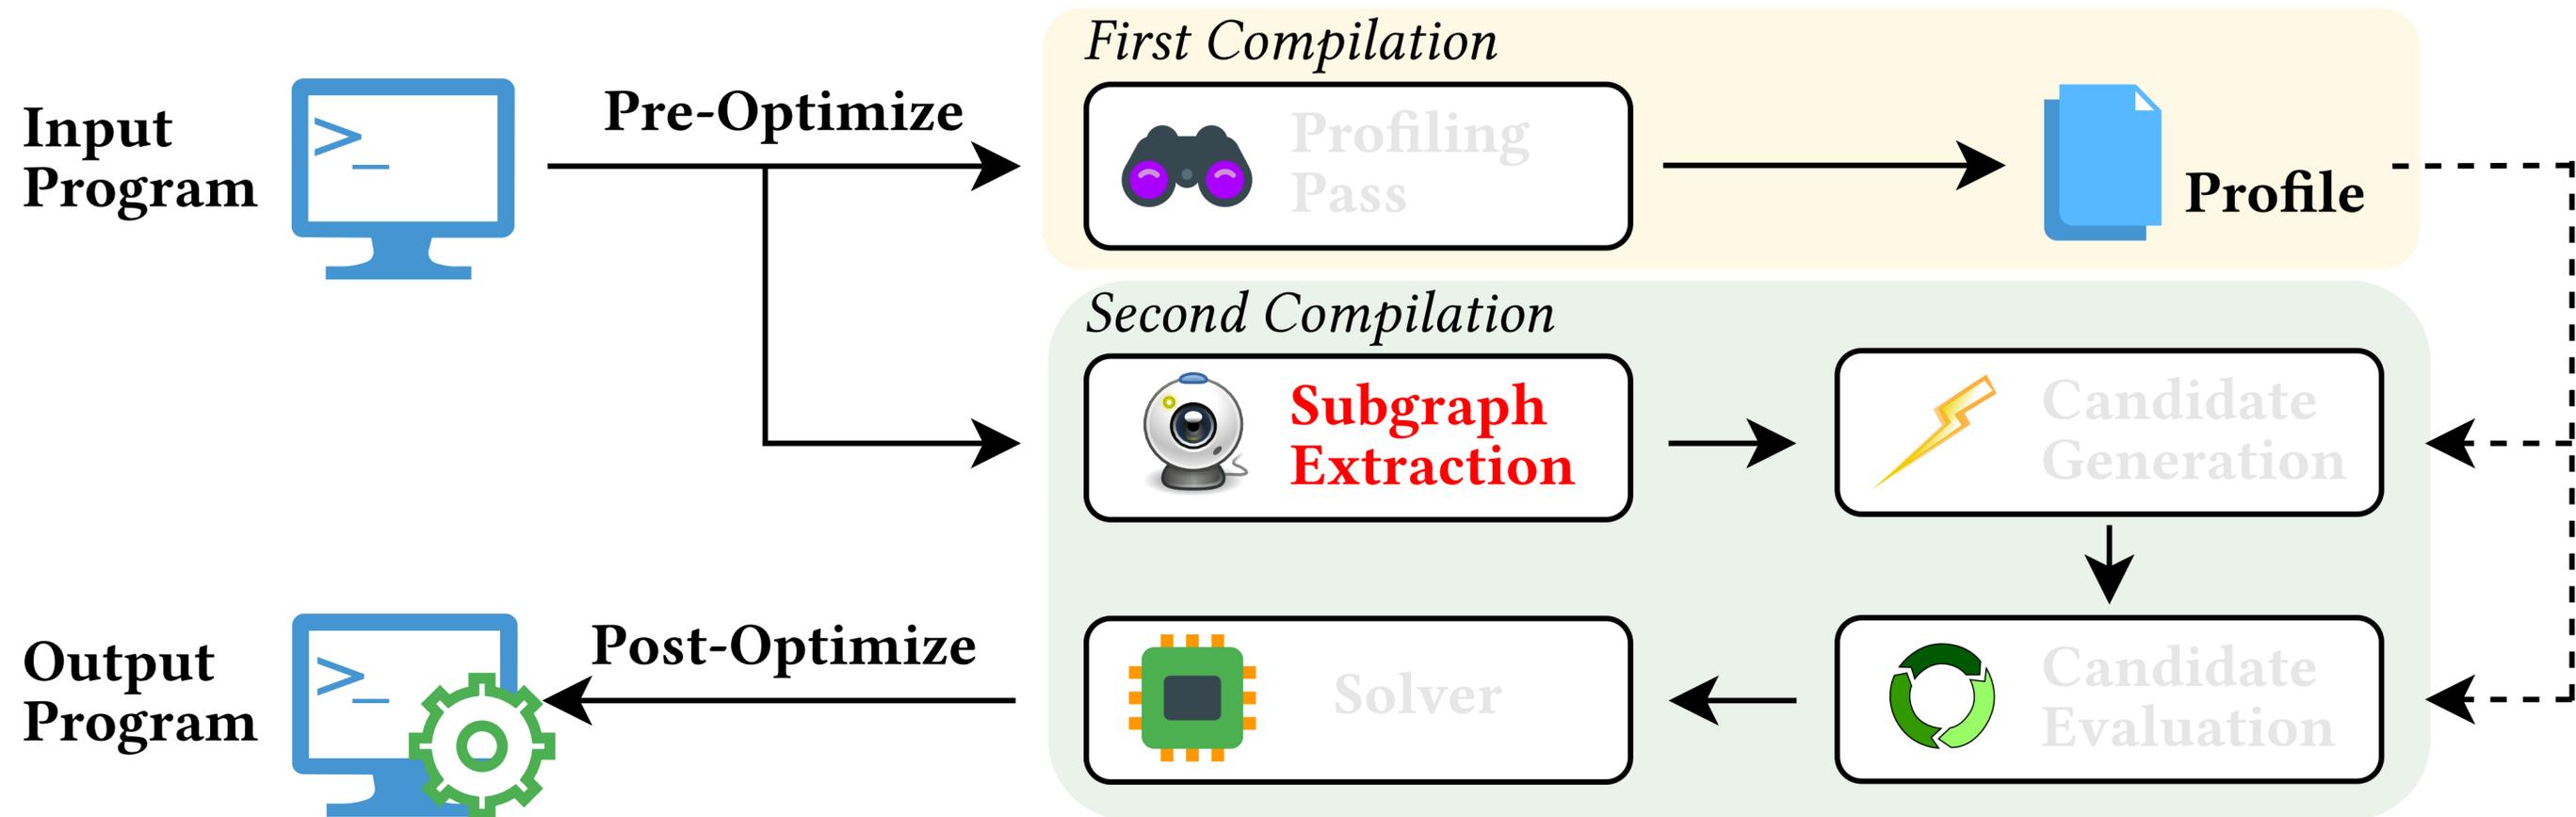
# Profiling Pass

- Perform **pre-optimizations** to expose underlying math
- Emit profiles: execution counts, running sums of values/gradients per instruction



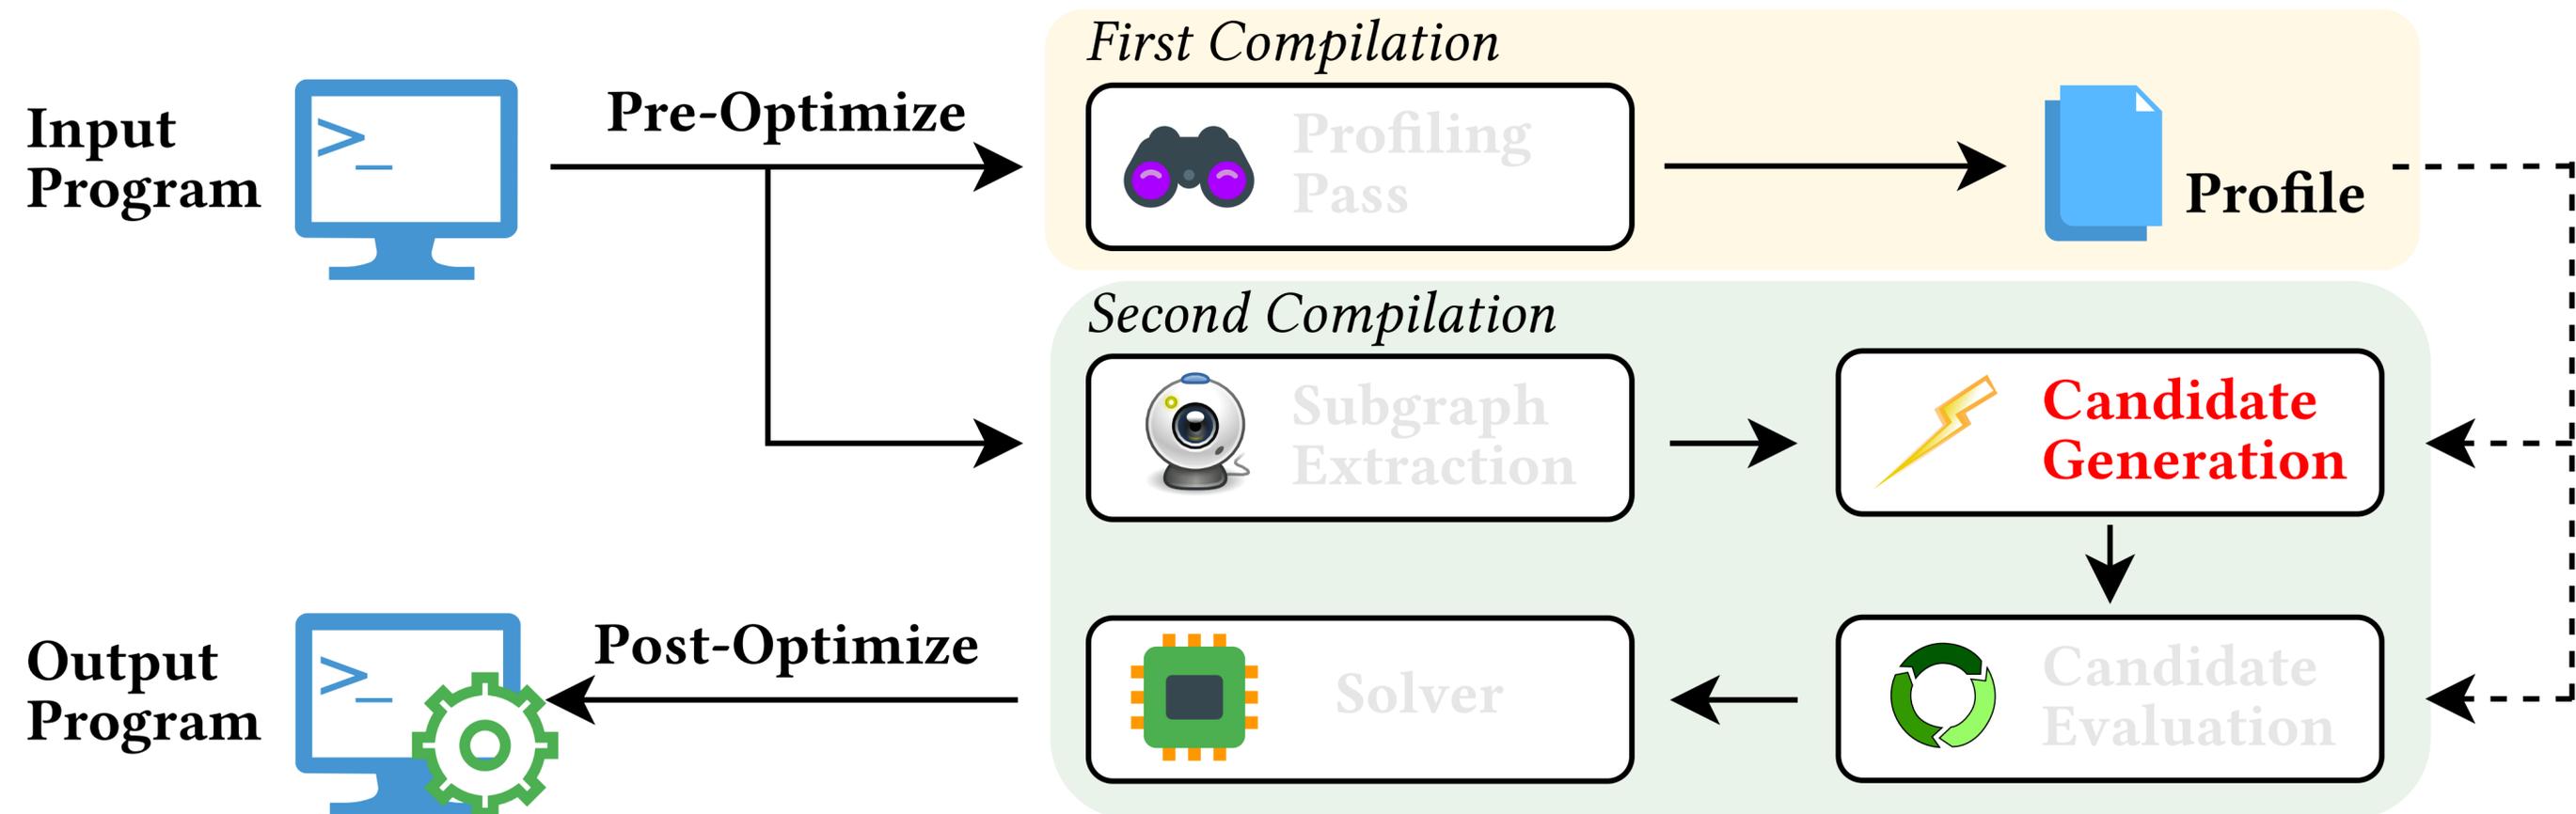
# Subgraph Extraction

- Similarly, perform **pre-optimizations** to expose underlying math
- **Rewrite Regions:** floating-point def-use subgraph (linear-time flood fill)



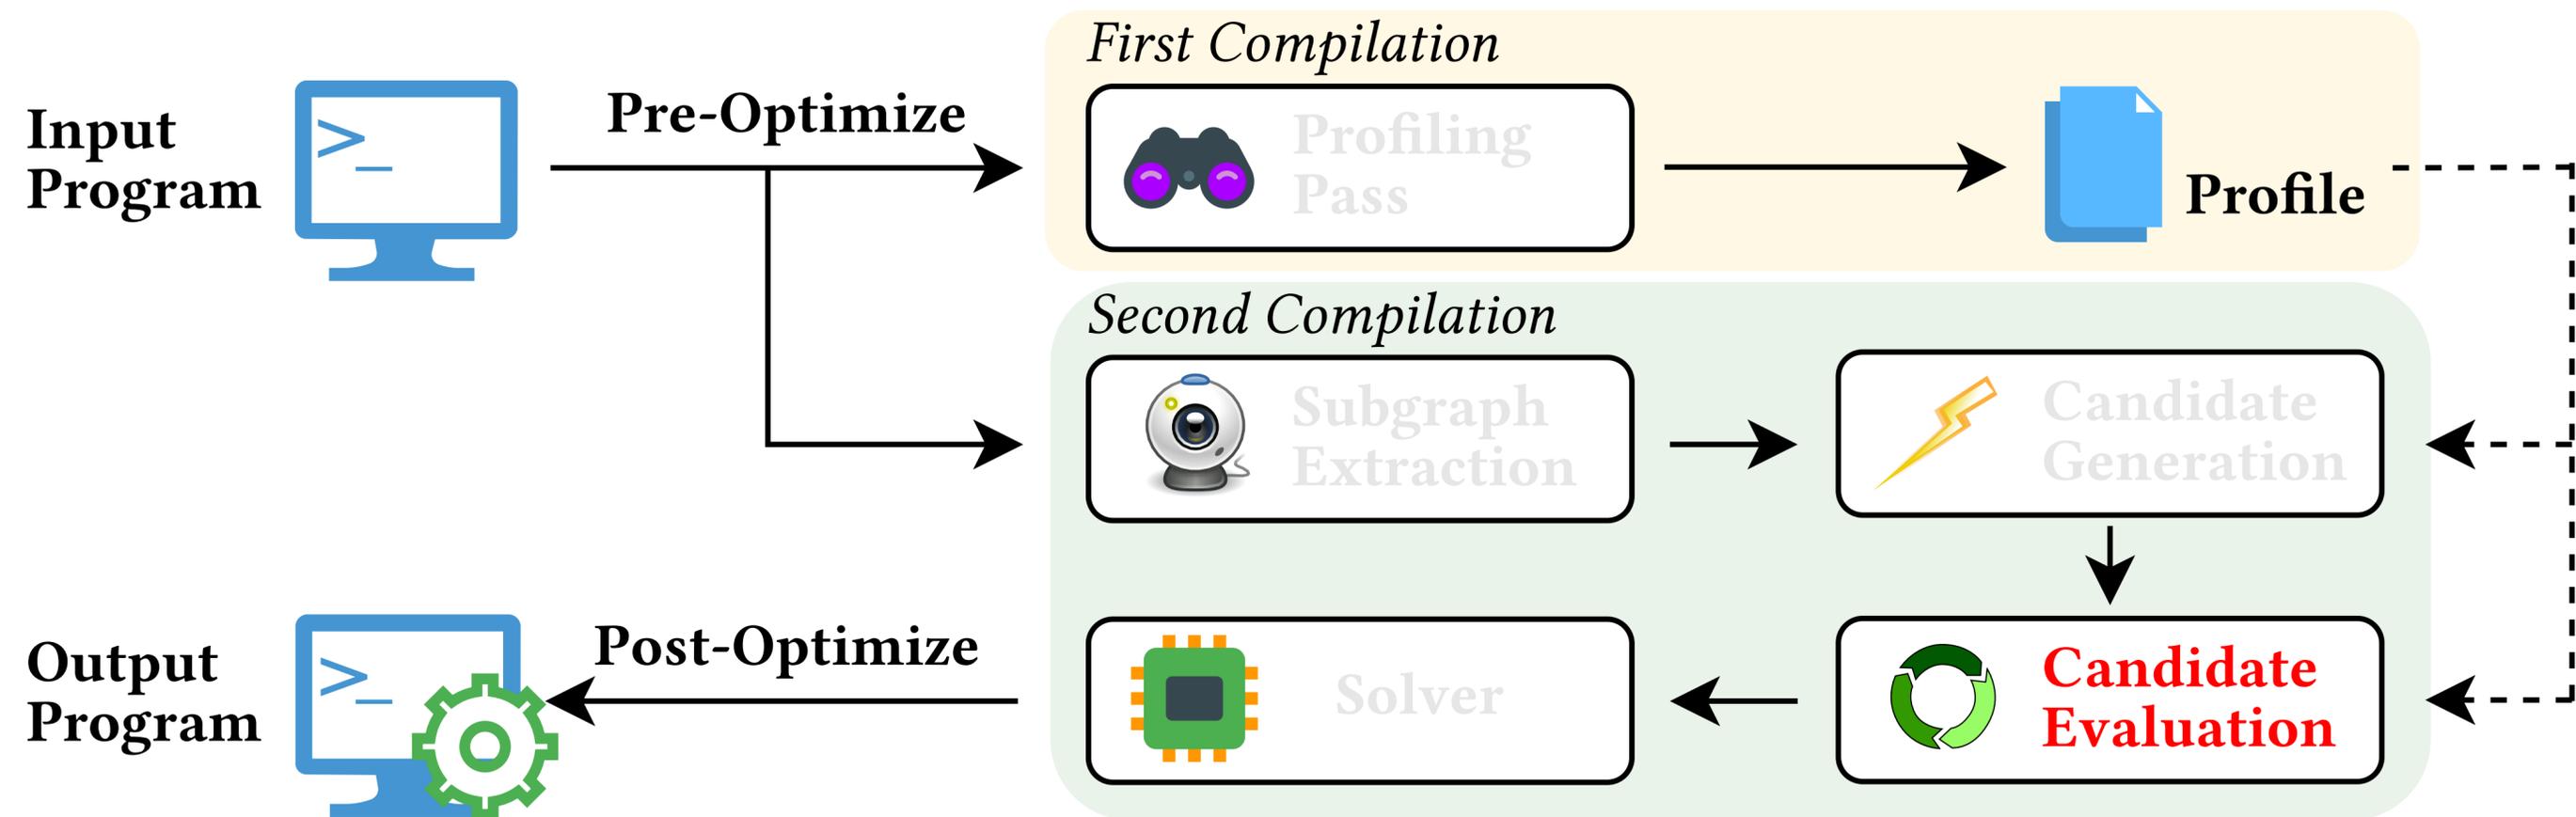
# Candidate Generation

- Algebraic rewrites from external tools (e.g., Herbie)
- Profile-guided precision changes
- Fully extensible!



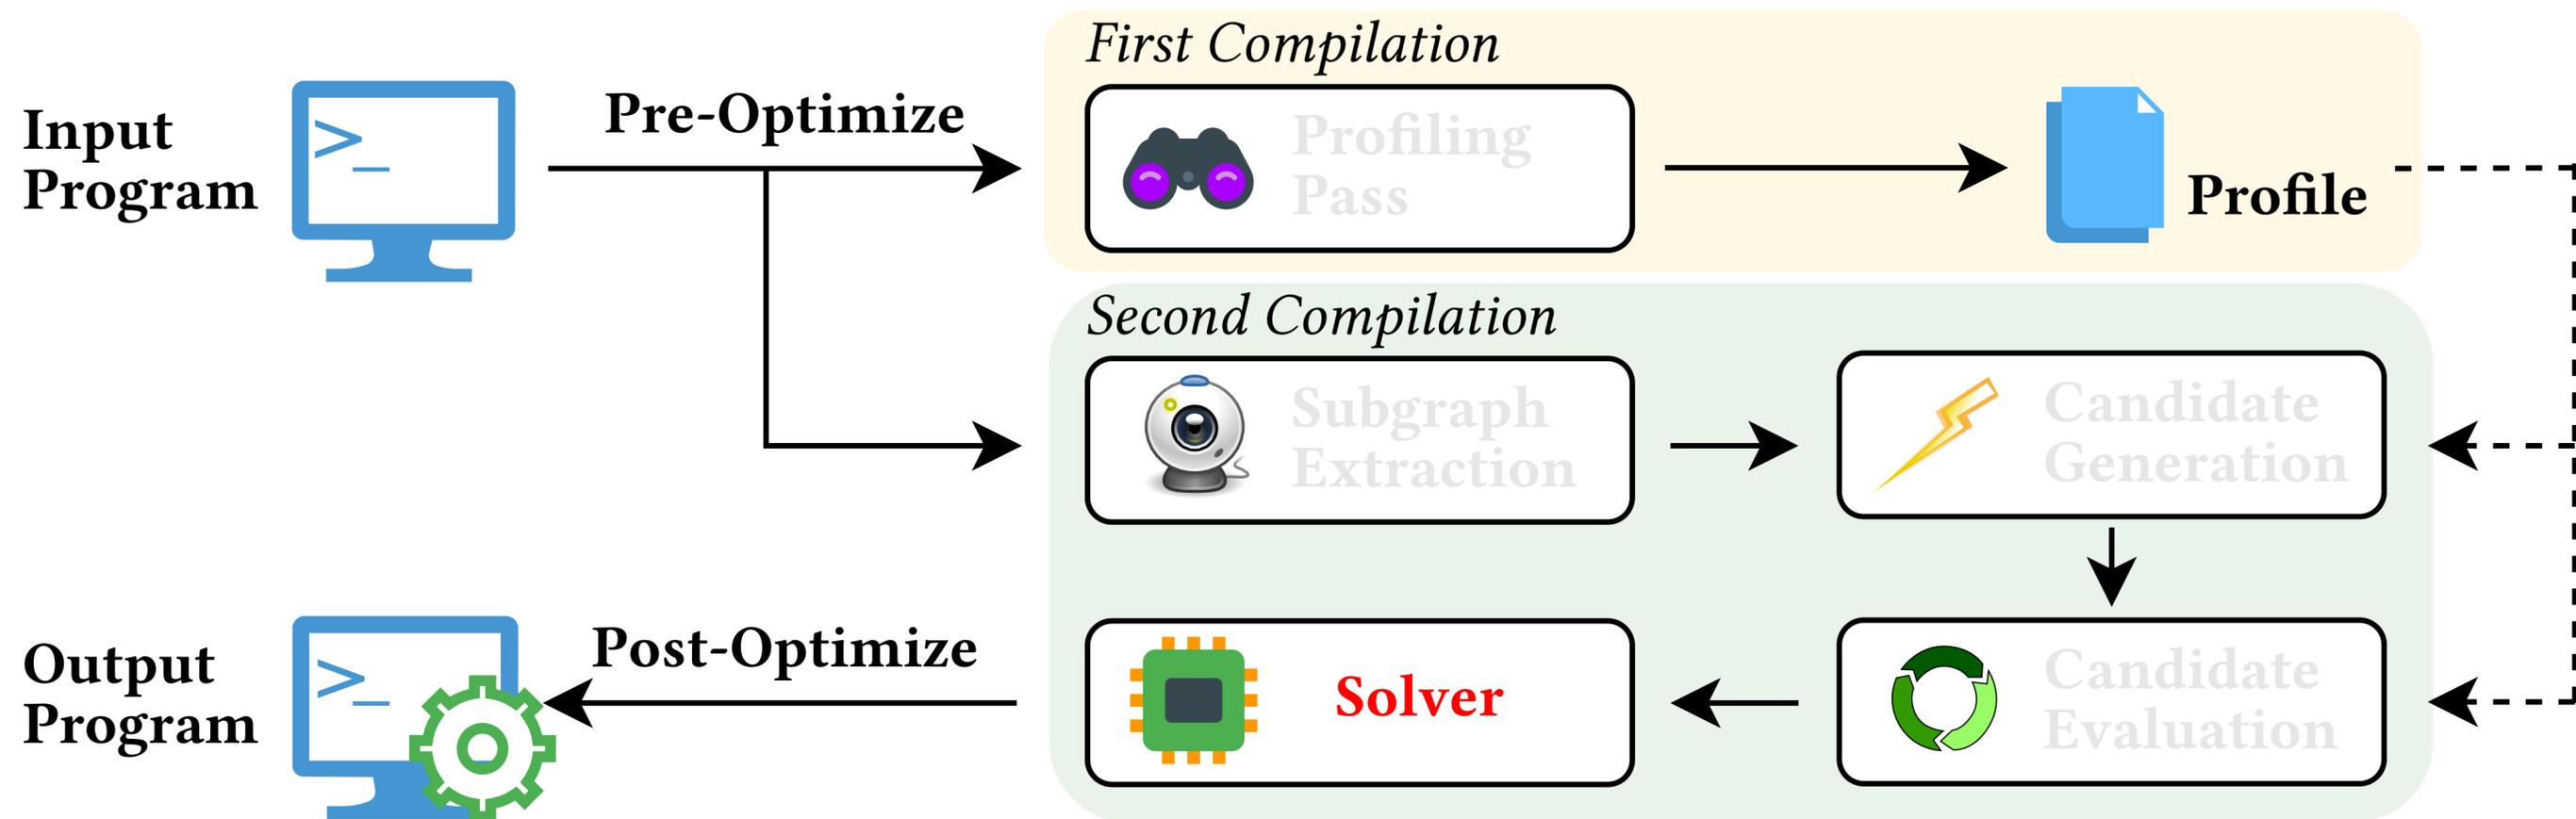
# Candidate Evaluation

- Goal: **predict** candidate's performance/accuracy impact using the profile
- Cost: Sum over instructions of  $(cost \times execution\ count)$
- Accuracy: Sum over instructions of  $(local\ error \times global\ sensitivity)$



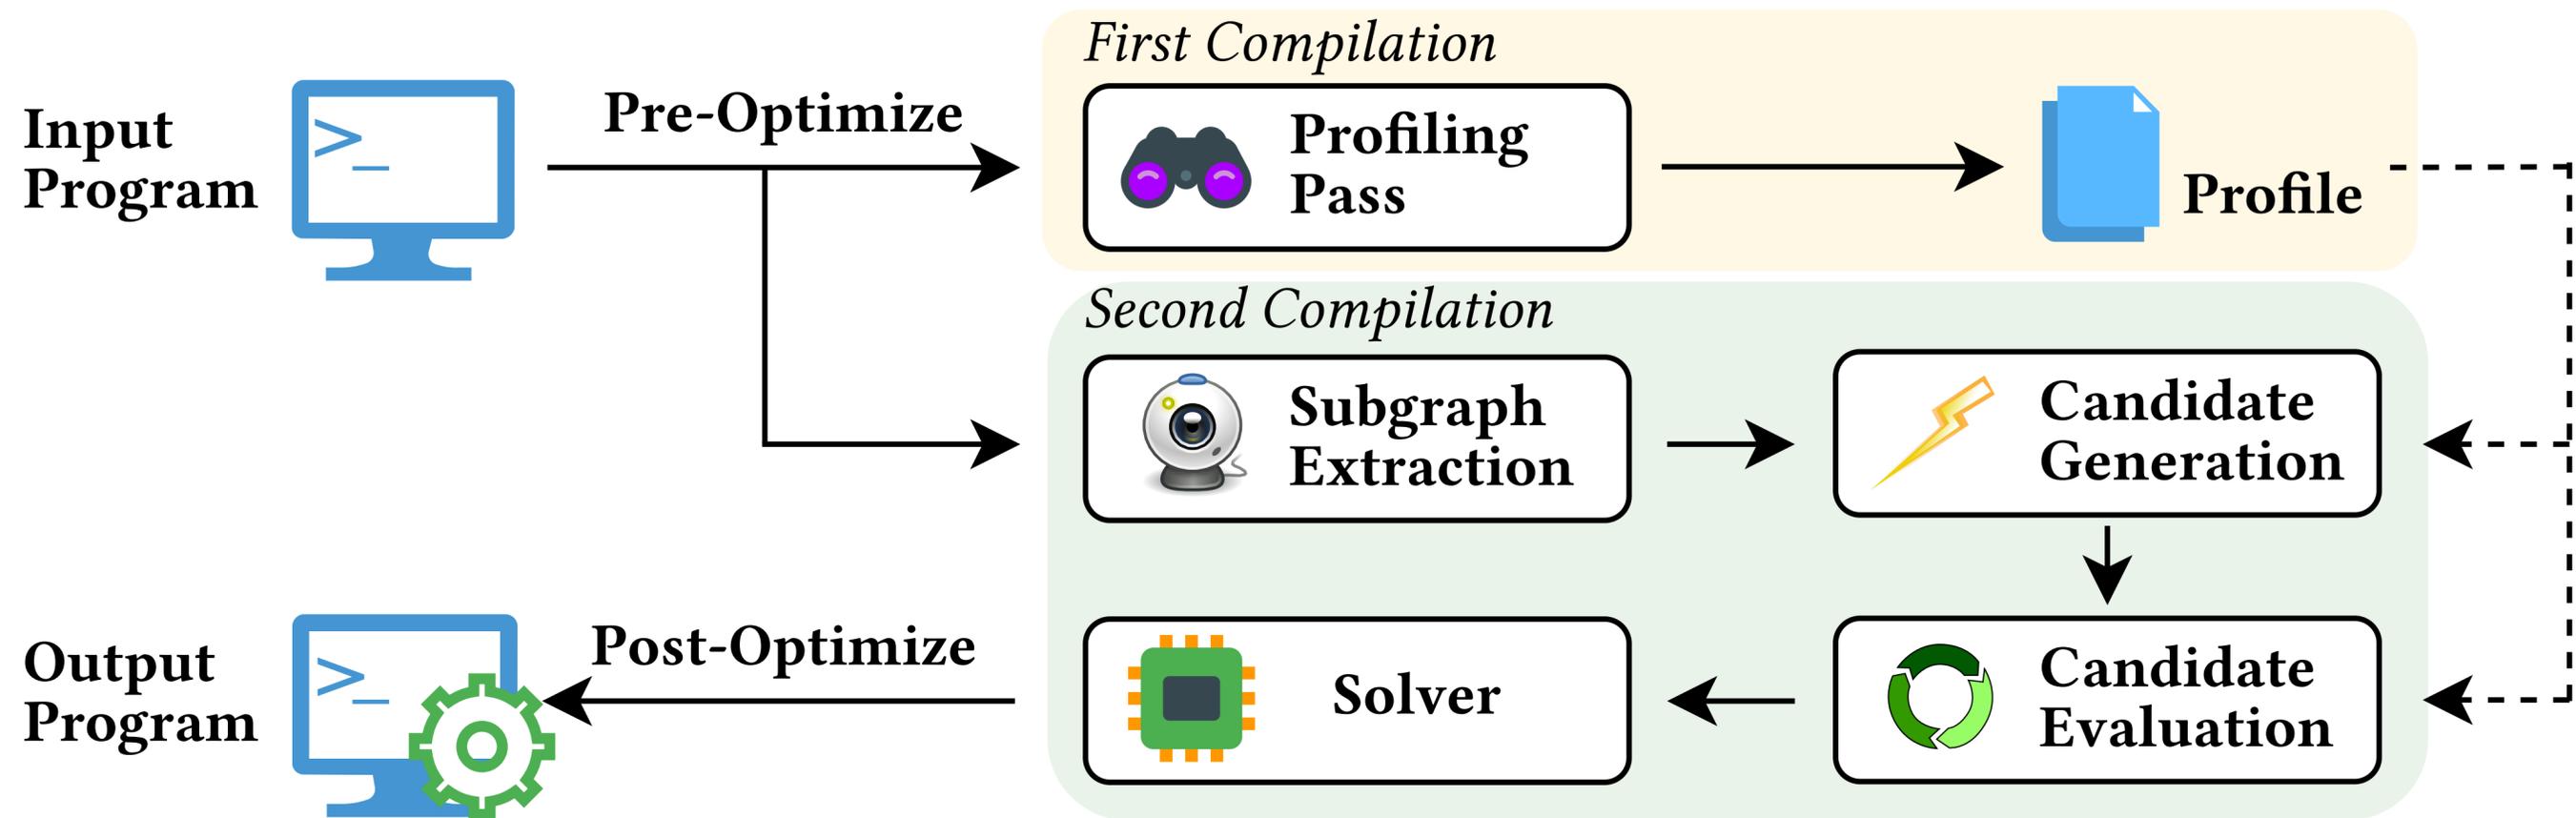
# Solver (Global Selection of Rewrites)

- Objective: **minimize predicted global error** given a set of cost-error pairs
- Similar to *knapsack*: dynamic programming builds a **cost-error frontier**



Can we **automate** numerical rewriting techniques in compilers?

How much performance are we **leaving on the table** due to *suboptimal* choices of precision and expressions?



# Evaluation

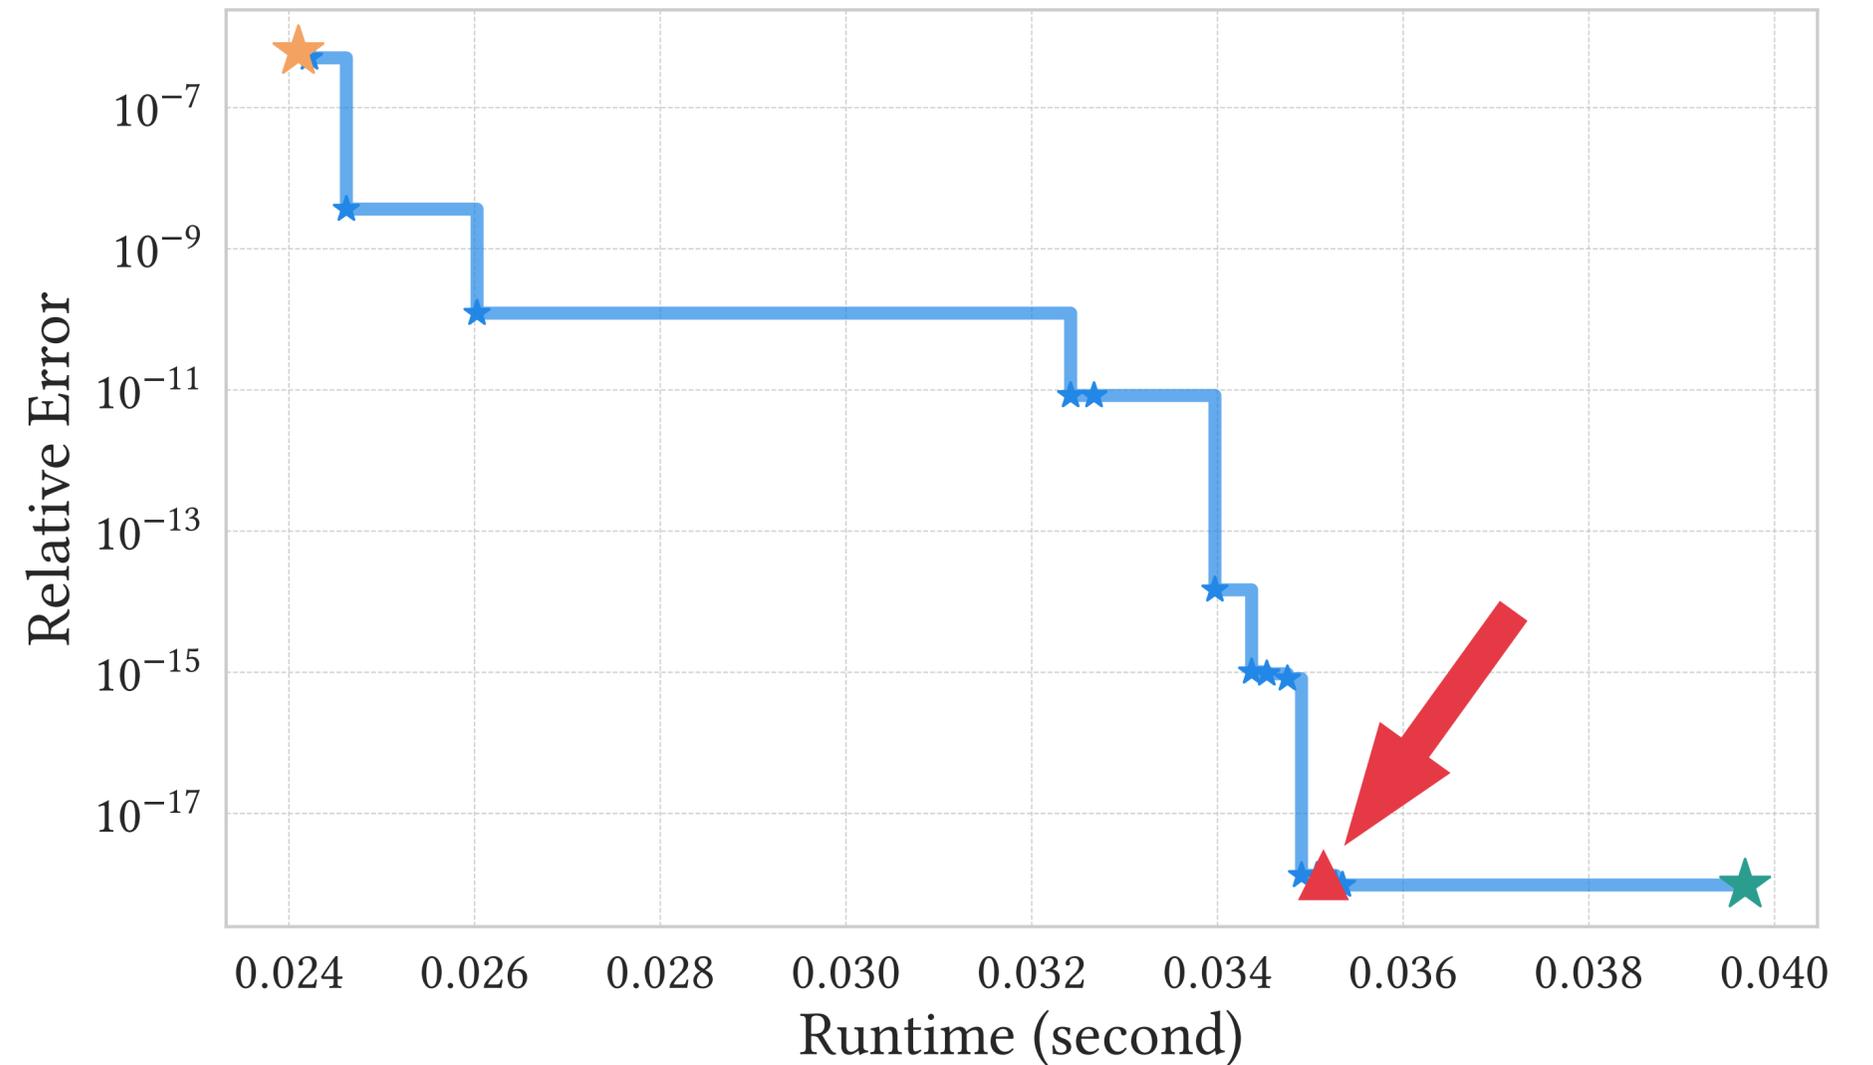
---

- Structured Critical Value Hypothesis: ***small surrogate profiles*** reveal numerical structure
- Coverage (FPBench): Accuracy improvements on **58%** of benchmarks;  
Maximum speedup of **1.82x** (error < 1e-6)
- Quaternion Differentiator: **1.46x** speedup (error < 1e-6)
- LULESH: **Bitwise-identical to MPFR-512** in FP64; no substantial slowdown



# Evaluation: Quaternion Differentiator

- Given the **original** program, Poseidon produces tradeoffs

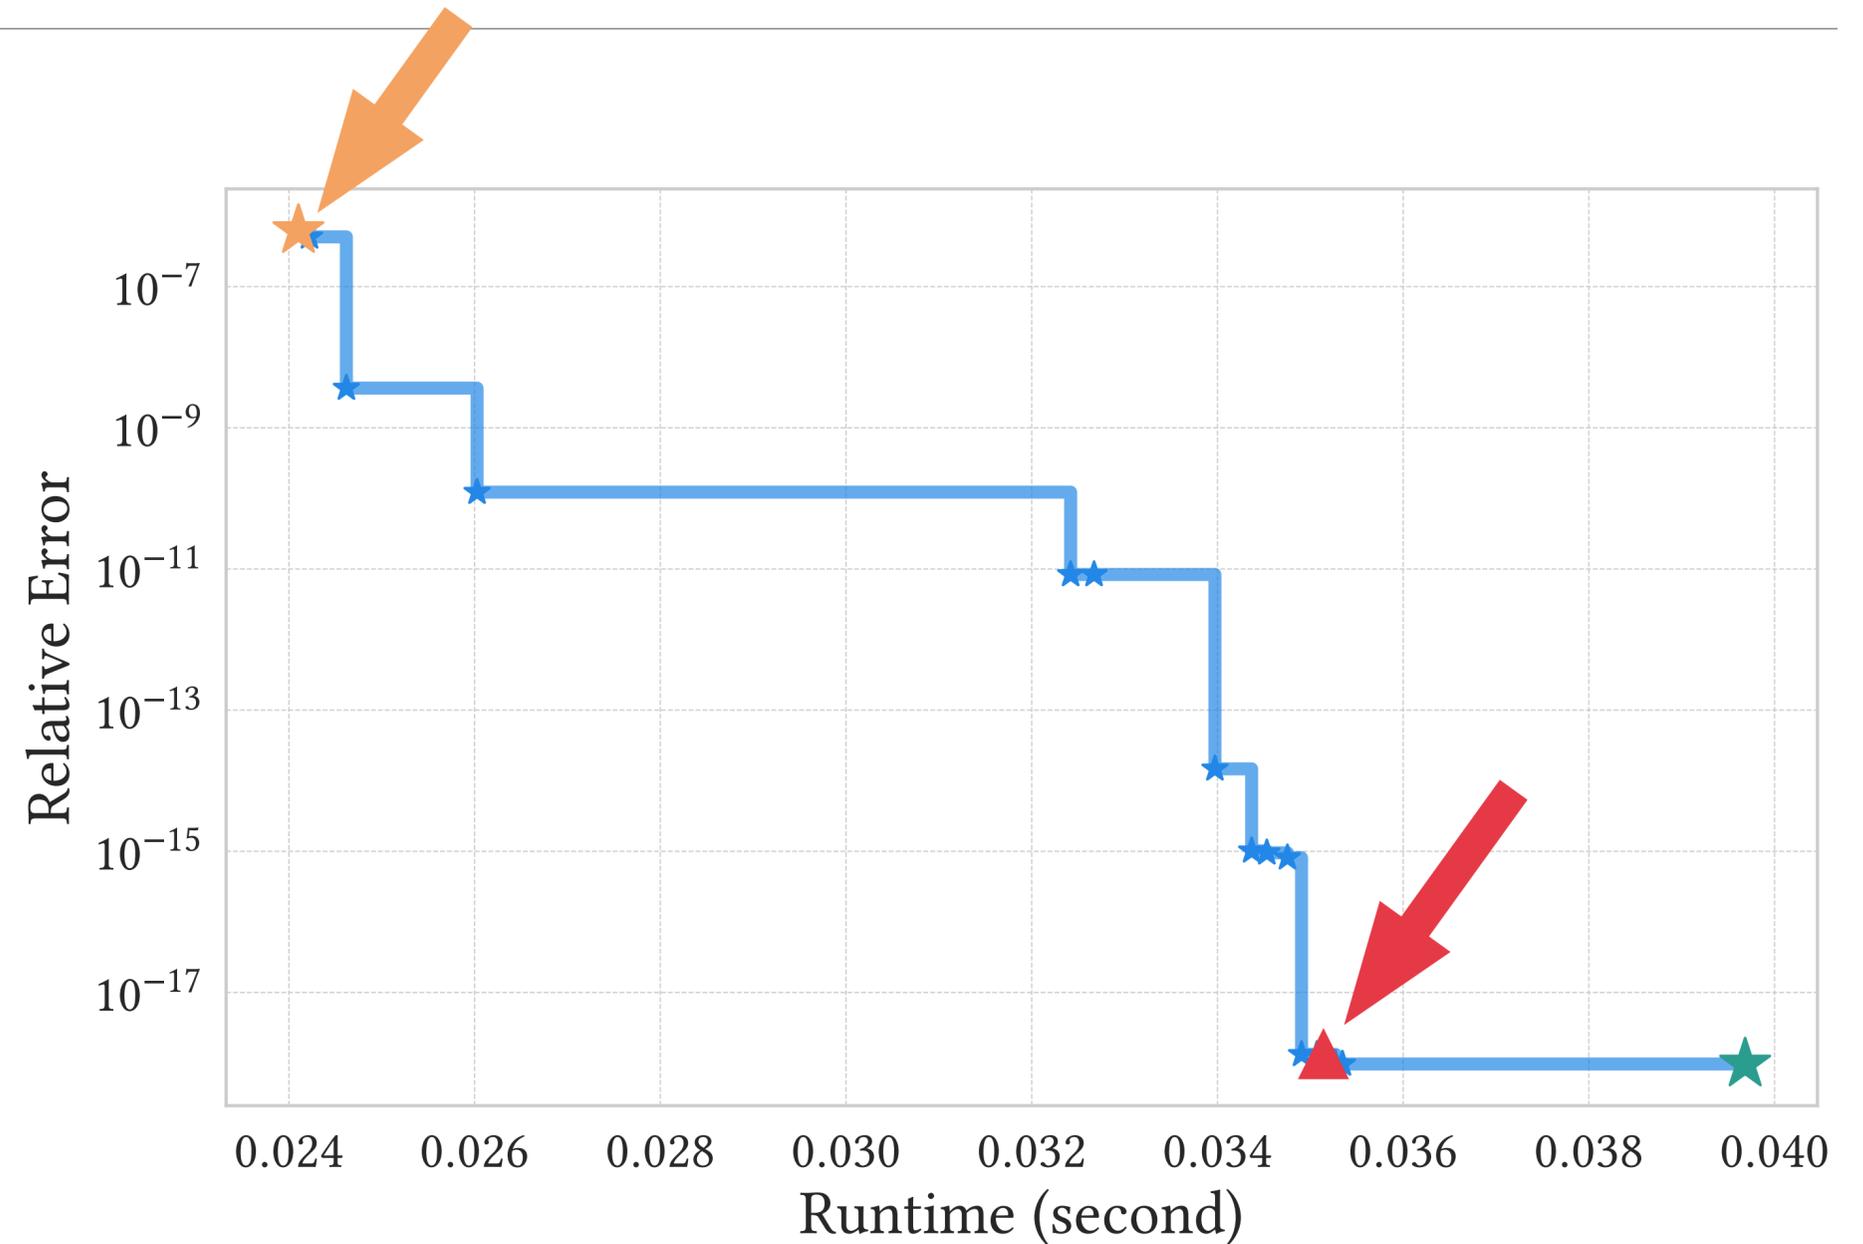


*Lower is more accurate; left is faster*



# Evaluation: Quaternion Differentiator

- Given the **original** program, Poseidon produces tradeoffs
- Trade accuracy for performance
  - **1.46x** with error of  $6e-7$

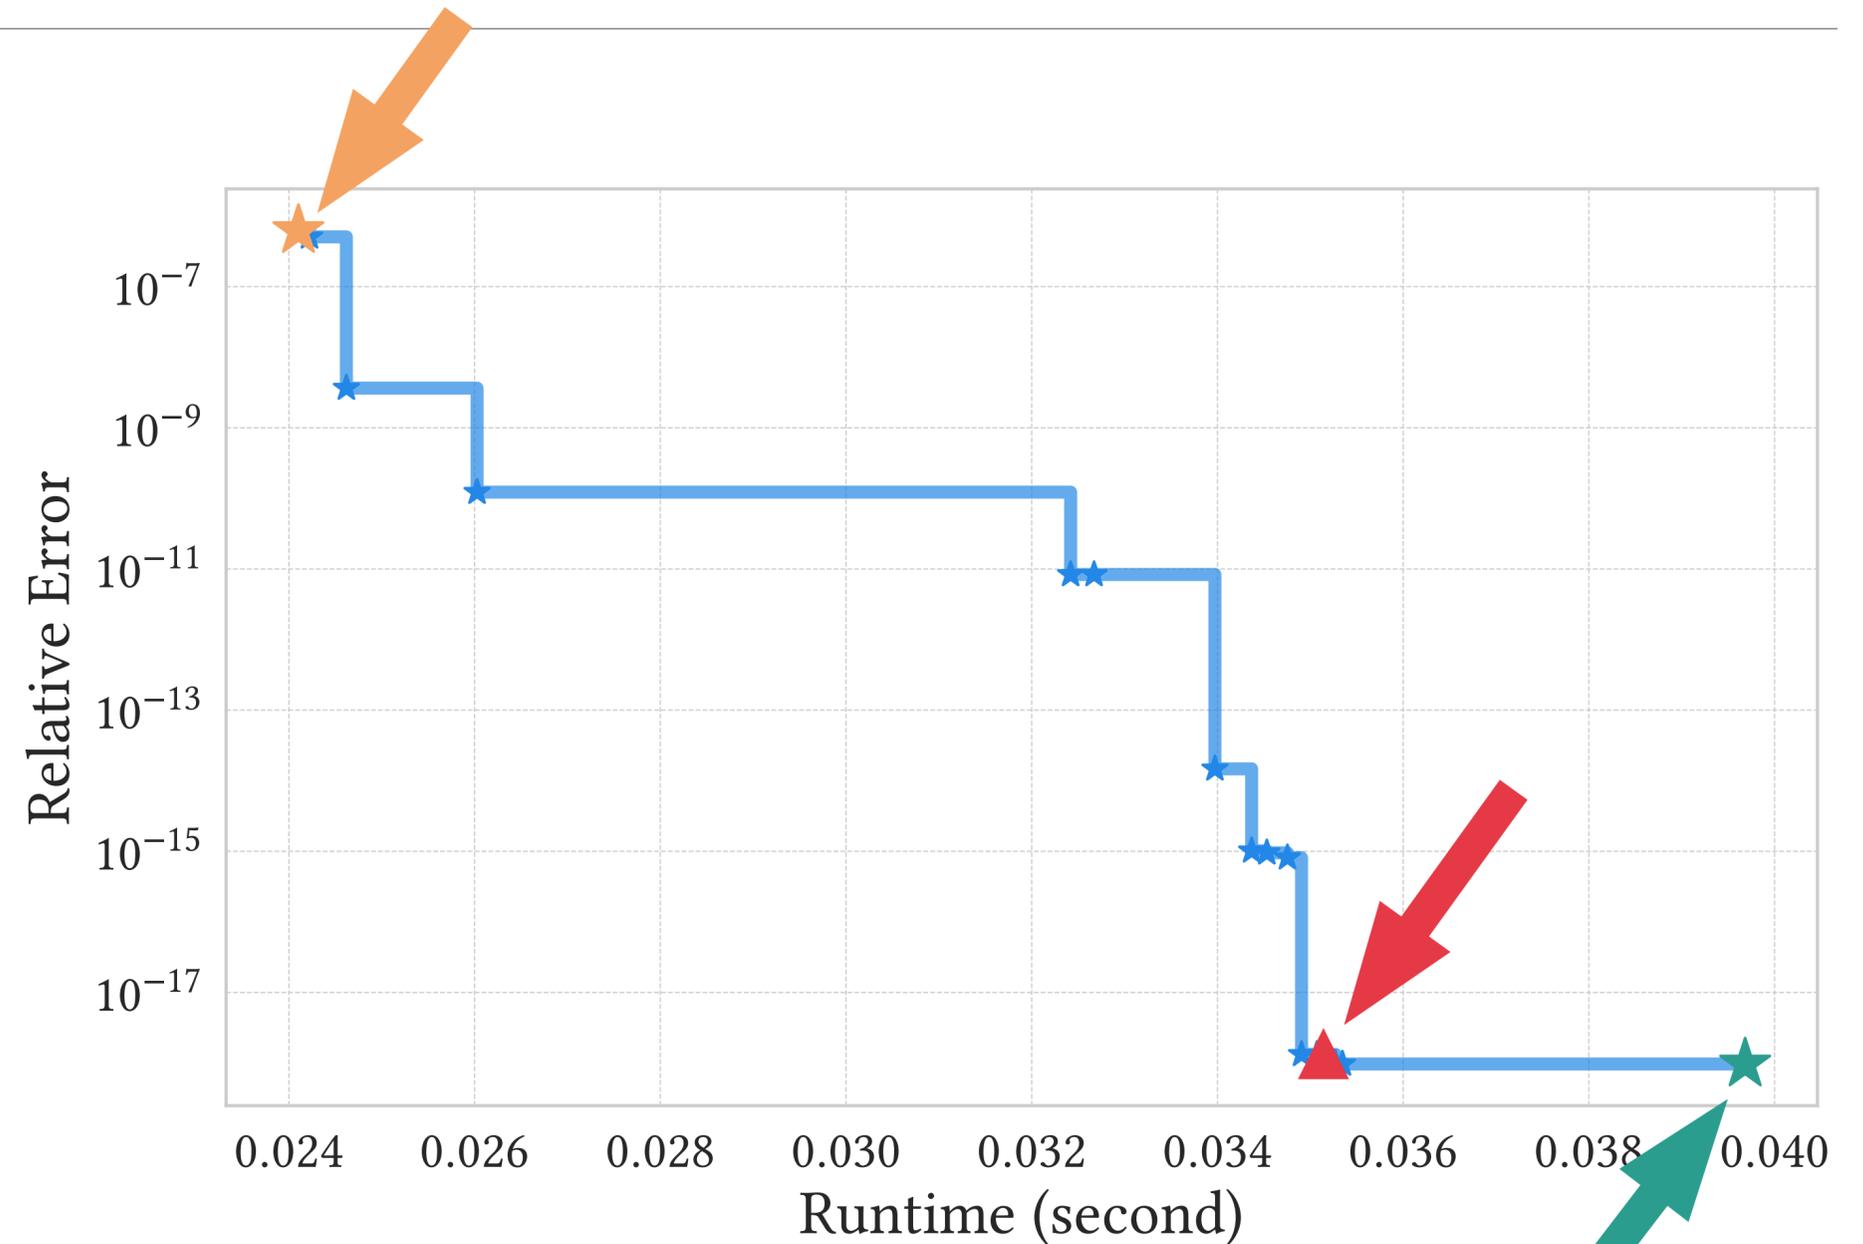


*Lower is more accurate; left is faster*



# Evaluation: Quaternion Differentiator

- Given the **original** program, Poseidon produces tradeoffs
- Trade accuracy for performance
  - **1.46x** with error of  $6e-7$
- Trade performance for accuracy
  - Error reduced by **27%** with **1.13x** more compute time

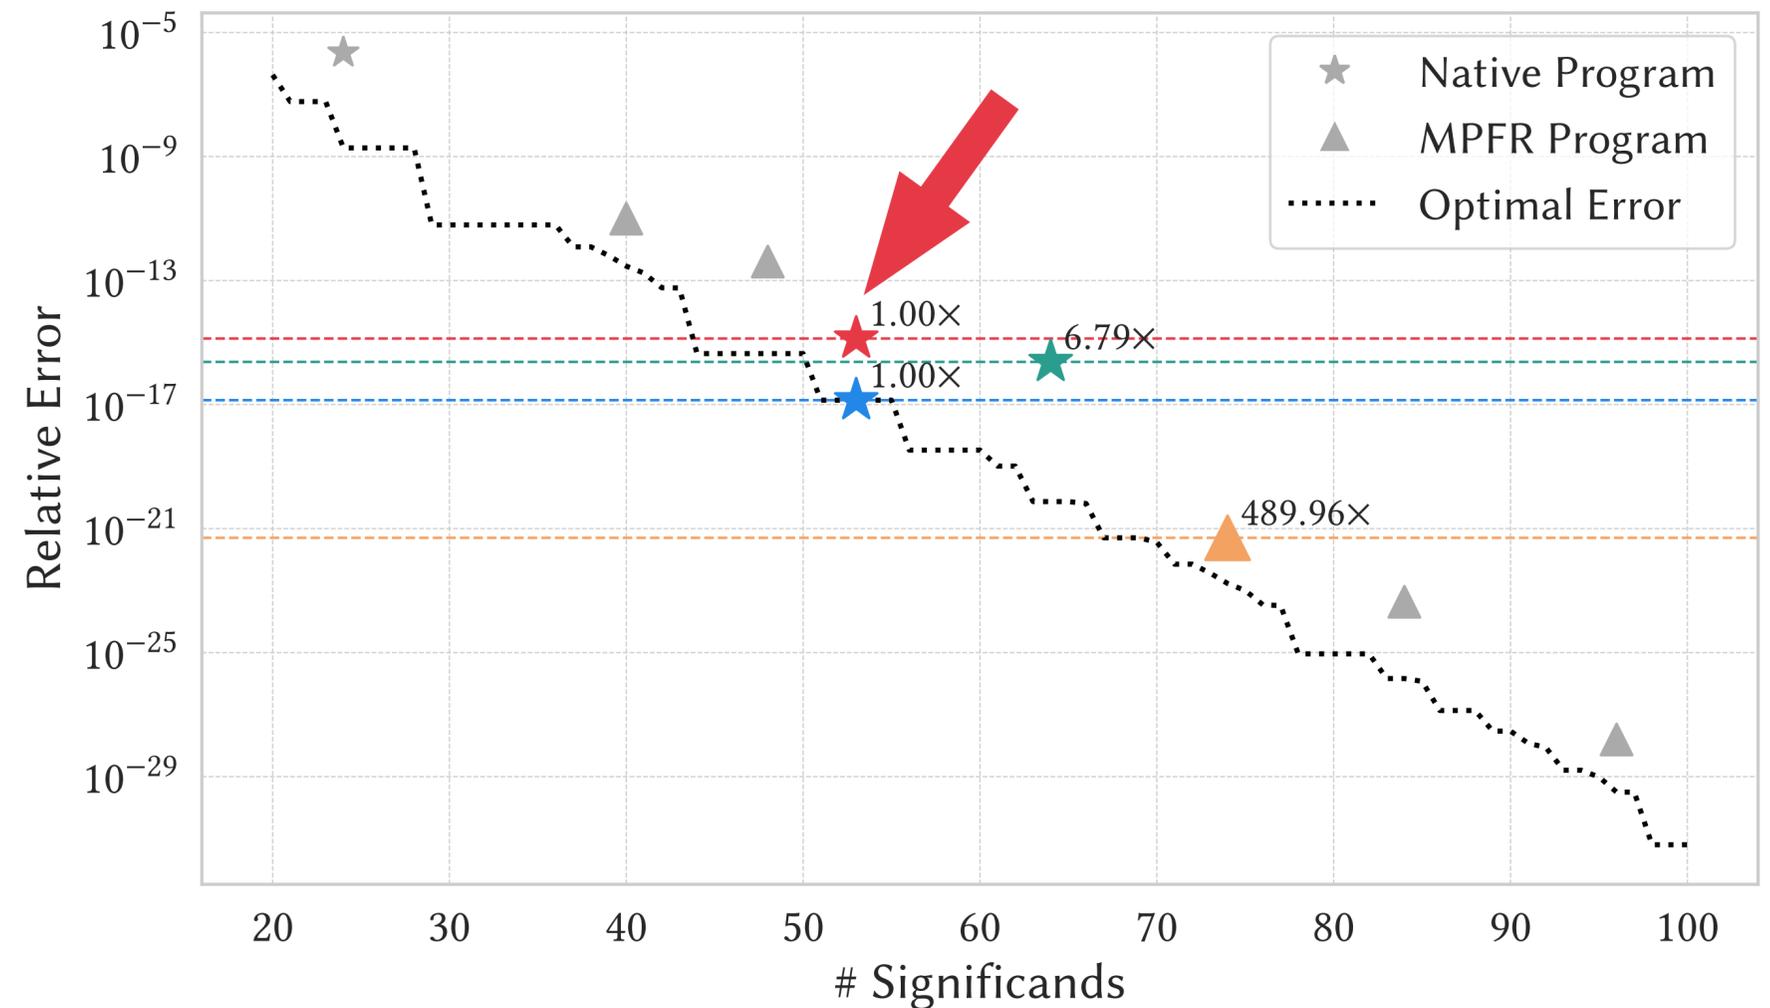


*Lower is more accurate; left is faster*



# Evaluation: LULESH

- Large DOE proxy application that simulates Lagrangian hydrodynamics
- 5600+ LOC with over 200 loops
- Original (FP64): **12 ULPs\*** off

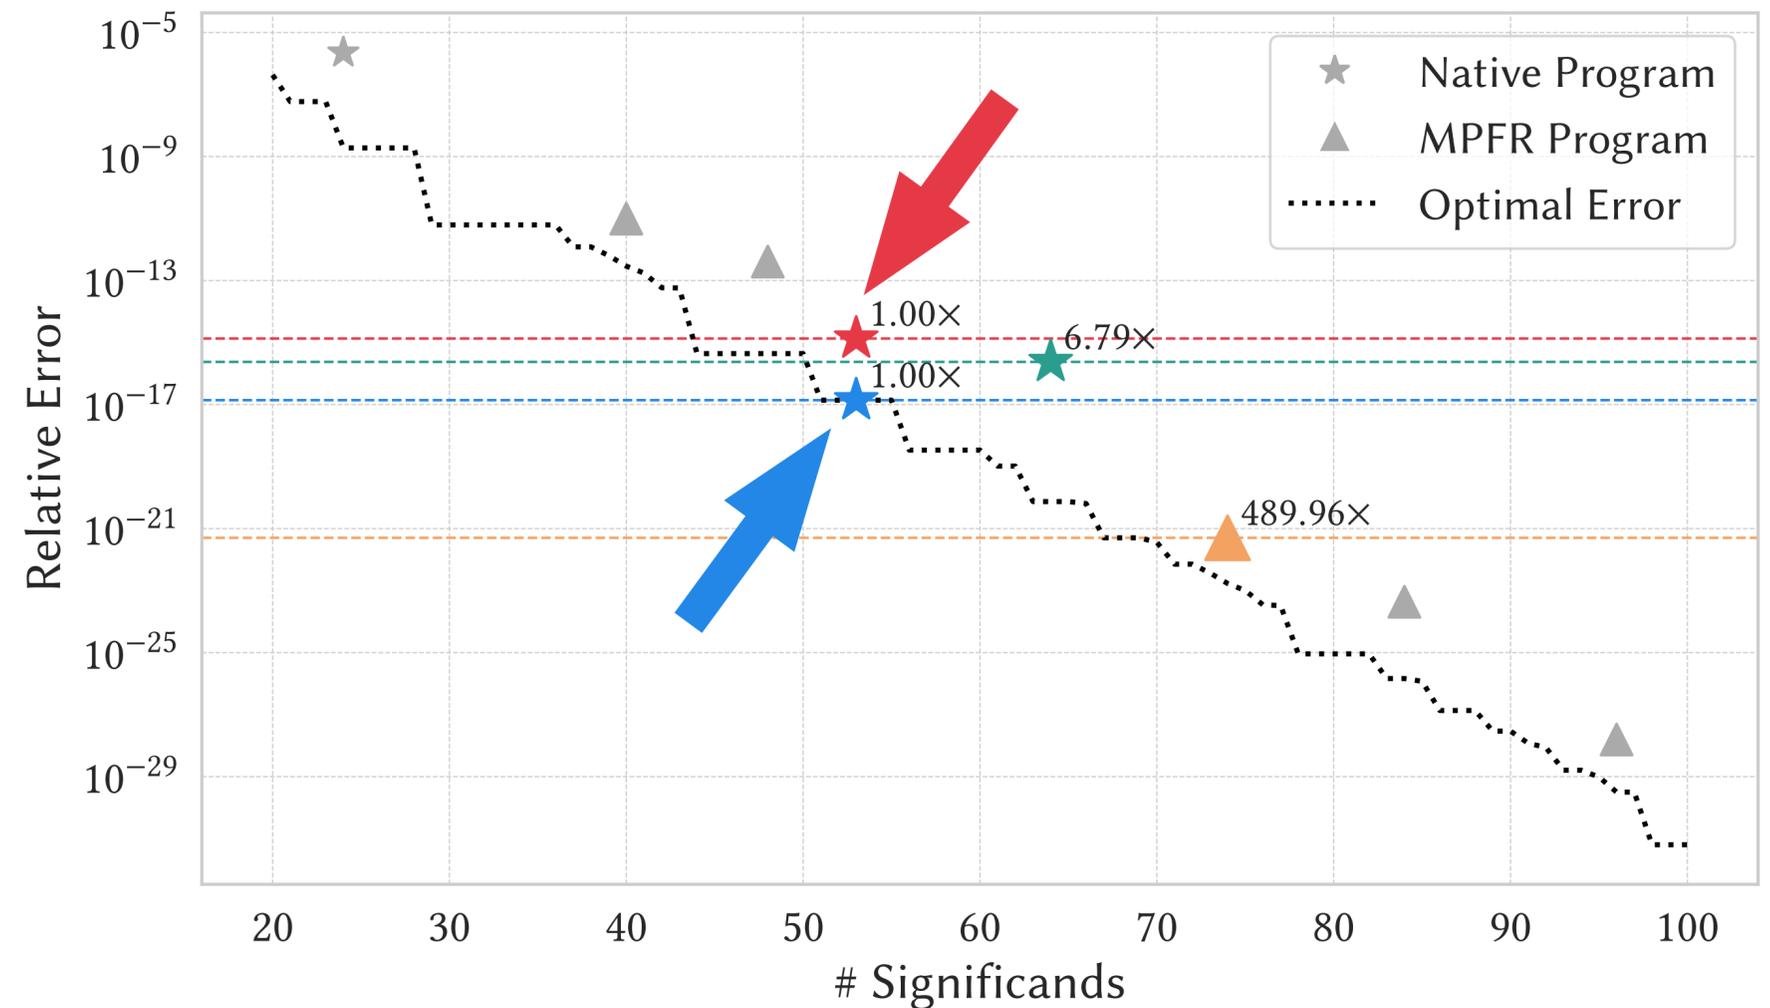


\* Ground truth: 64-bit rounding of MPFR-512 computation.



# Evaluation: LULESH

- Large DOE proxy application that simulates Lagrangian hydrodynamics
- 5600+ LOC with over 200 loops
- Original (FP64): **12 ULPs\*** off
- Poseidon (FP64): **Optimal\***  
**FP64 error**; little slowdown

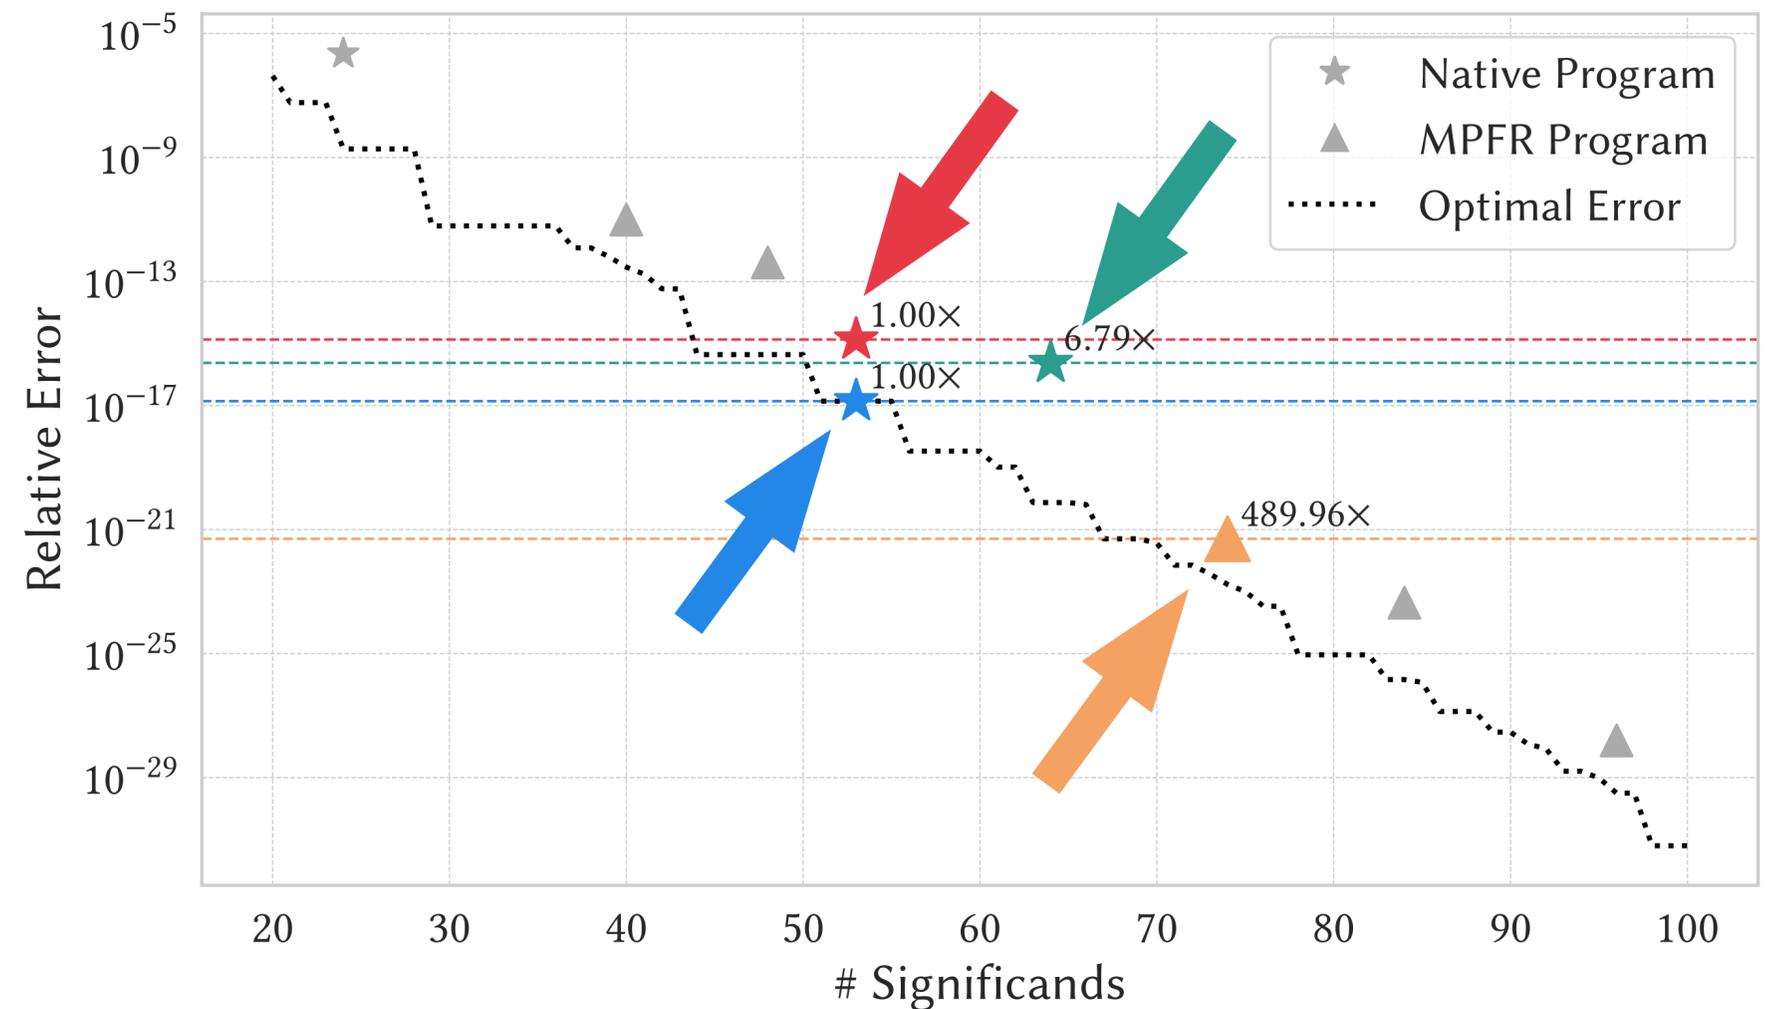


\* Ground truth: 64-bit rounding of MPFR-512 computation.



# Evaluation: LULESH

- Large DOE proxy application that simulates Lagrangian hydrodynamics
- 5600+ LOC with over 200 loops
- Original (FP64): **12 ULPs\*** off
- Poseidon (FP64): **Optimal\*** **FP64 error**; little slowdown
- FP80: **6.79x** slower; still **2 ULPs\*** off
- MPFR-74: **~500x** slower



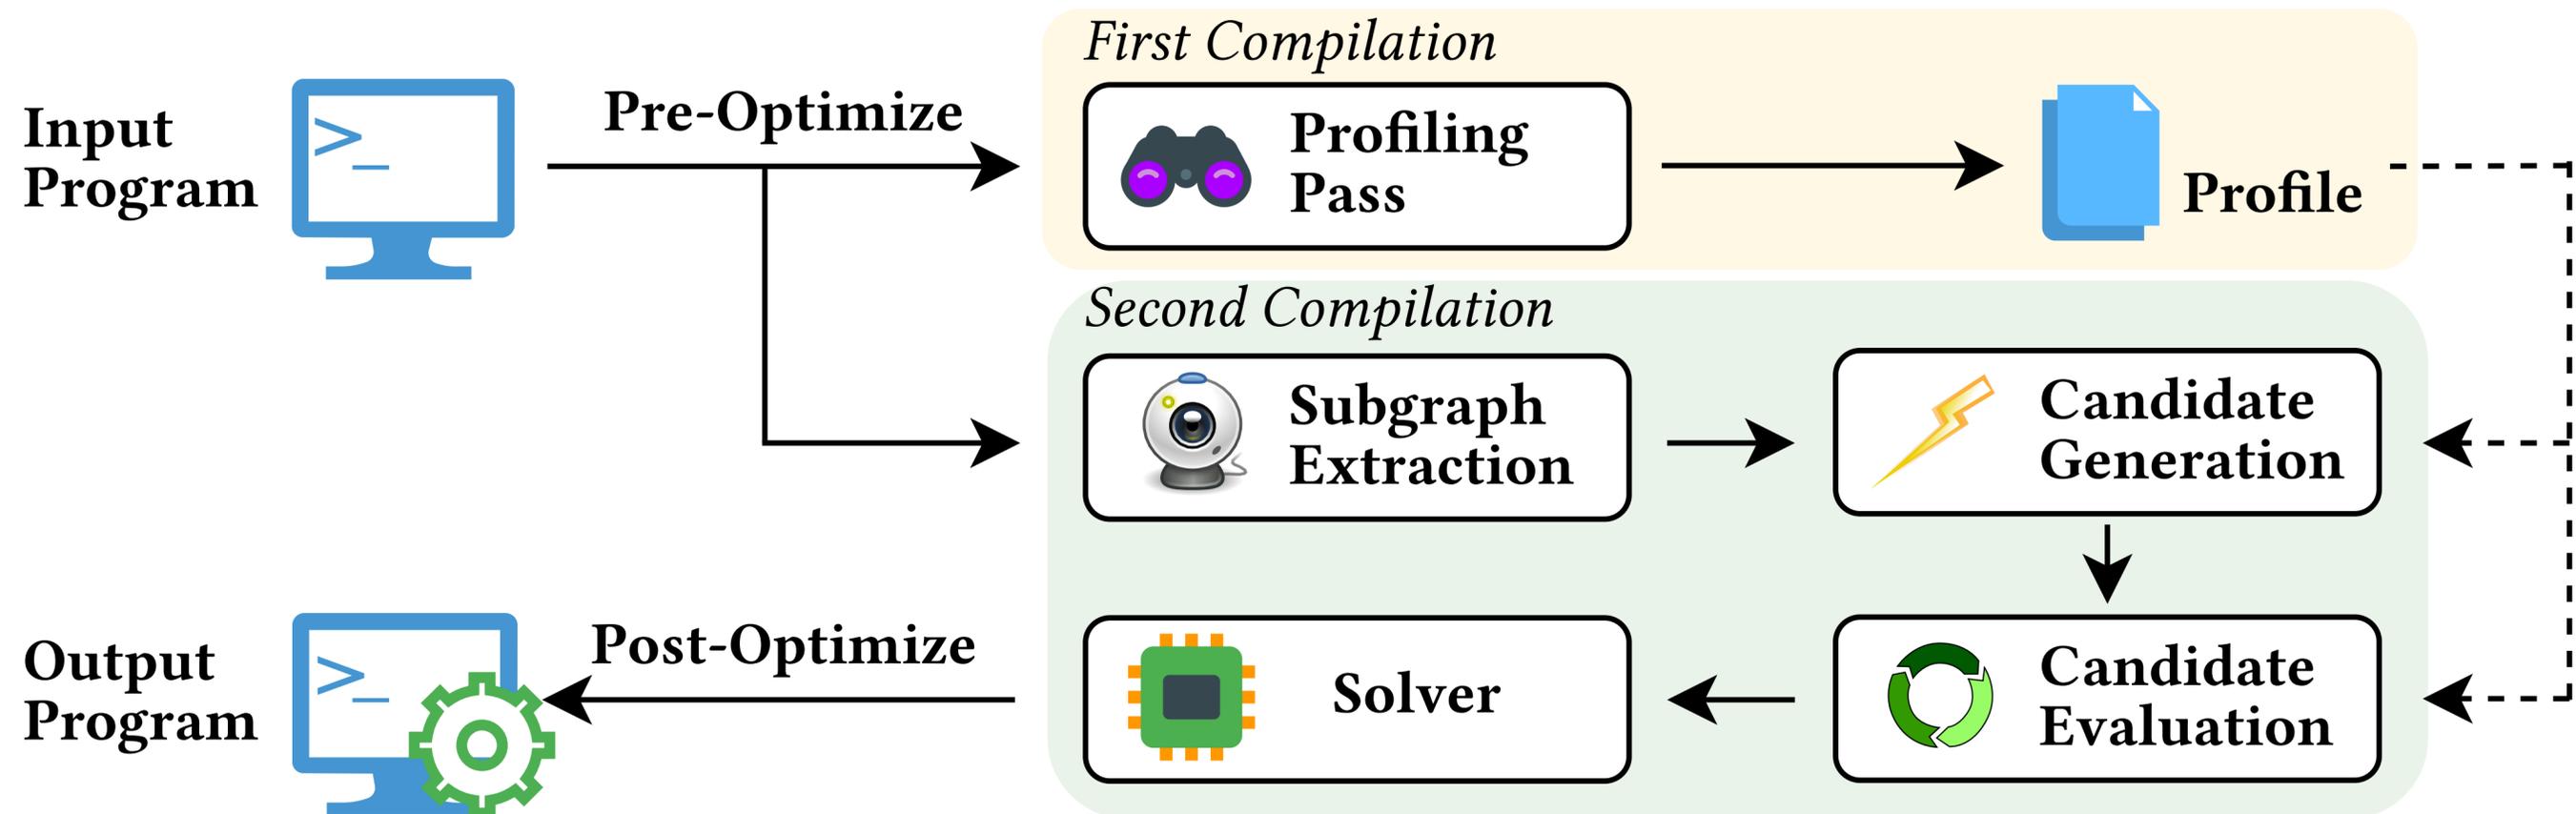
\* Ground truth: 64-bit rounding of MPFR-512 computation.





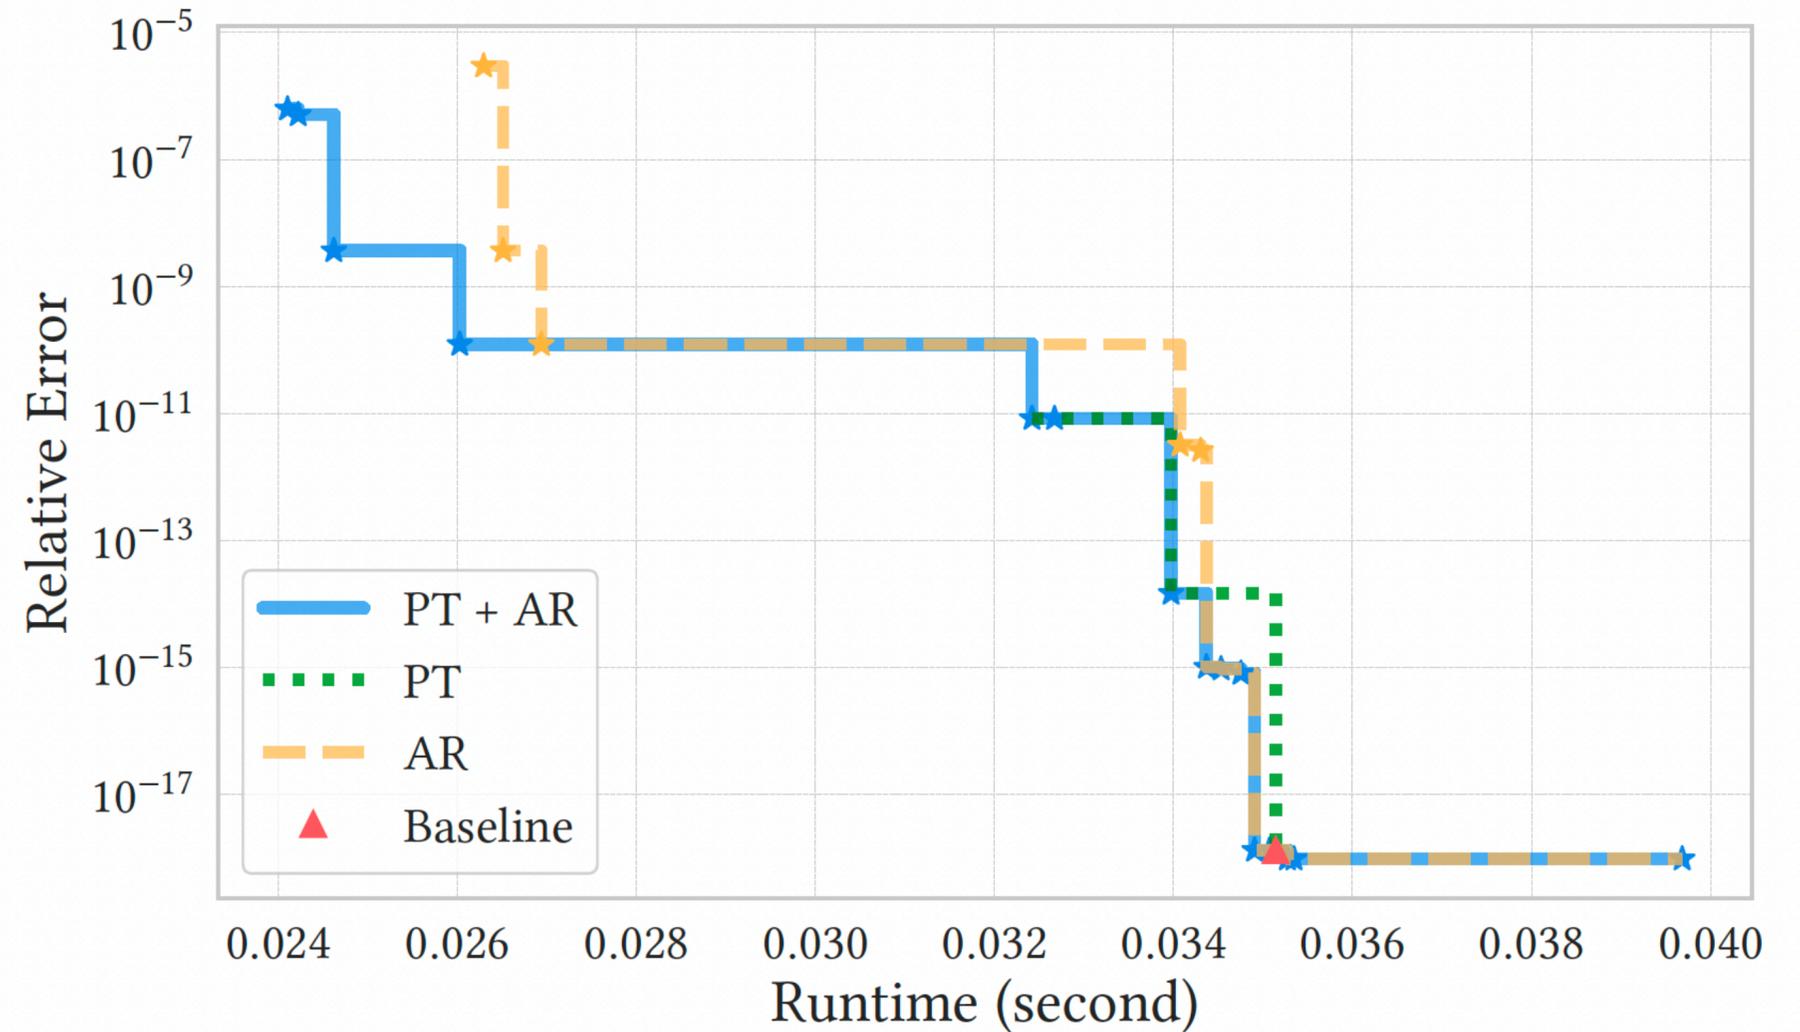
# Poseidon (Open Source on GitHub: [PRONTOLab/Poseidon](https://github.com/PRONTOLab/Poseidon))

- The first framework that automates numerical rewrites within a production compiler
- **Profiling + Compiler Optimization** provides context, scope, and scale
- Outsized performance benefit without sacrificing the accuracy, and *vice versa*



# Evaluation: Quaternion Differentiator

- PT + AR: Precision Tuning + Algebraic Rewrites
- PT: Precision Tuning only
- AR: Algebraic Rewrites only
- **PT + AR has the best frontier!**



Lower-left is better

