

Analyzing Floating-Point Programs

Debasmita Lohar

Formal Systems II: Applications, Summer 2025

In this Part

you will learn the fundamentals of analyzing floating-point programs*:

27.06.

- Challenges
- Basics of Floating-Point Arithmetic
- Dataflow analysis
 - Interval Arithmetic (IA)
 - Floating-Point IA
 - Analysis of roundoff errors with IA
 - Demo: Daisy

More on Dataflow Analysis

Today

- Pro's and con's of Intervals
- Affine Arithmetic (AA)
- Analysis of roundoff errors with AA
- Interval Subdivision
- Other Approaches and Recent Directions
- Demo: Daisy

Pro's and con's of Intervals

Pro's:

- Conceptually simple (important for correctness)
- Fast efficient machine-independent implementation
- Several successful applications

Con's:

- Can be imprecise: not relational, i.e. cannot track correlations between variables

$$x = [-1, 2]$$

$$x - x = [-1, 2] - [-1, 2]$$

$$= [-3, 3]$$

$$\neq [0, 0]$$



In a chained computation intervals can become too wide to be useful!

Affine Arithmetic (AA)

- Improves over intervals in that it tracks *linear* correlations
- Represents each range as a linear (affine) form:

$$\hat{x} := x_0 + \sum_{i=1}^n x_i \epsilon_i \quad \epsilon \in [-1, 1]$$

- x_0 is the *central value*
 - $x_i \epsilon_i$ are *noise terms*, tracking deviation from x_0
 - ϵ_i are symbolic variables, tracking correlations
- The range represented by an affine form:

$$[\hat{x}] = [x_0 - \sum_{i=1}^n |x_i|, \quad x_0 + \sum_{i=1}^n |x_i|]$$

- AA tracks linear correlations: $\hat{x} = x_0 + x_1 \epsilon_1$

$$\hat{x} - \hat{x} = x_0 + x_1 \epsilon_1 - (x_0 + x_1 \epsilon_1) = x_0 - x_0 + x_1 \epsilon_1 - x_1 \epsilon_1 = 0$$

Quiz



$$\hat{x} := x_0 + \sum_{i=1}^n x_i \epsilon_i \quad \epsilon \in [-1, 1]$$

What is the affine form corresponding to the following interval?

$$x = [2, 10]$$

a)



$$\hat{x} = 6 + 4\epsilon_1$$

b)

$$\hat{x} = 2 + 8\epsilon_1$$

c)

$$\hat{x} = 6 + 8\epsilon_1$$

AA Arithmetic Operations

- Addition, subtraction, unary minus are linear and can thus be computed exactly:

$$\hat{z} = \hat{x} + \hat{y} = x_0 + \sum_{i=1}^n x_i \epsilon_i + y_0 + \sum_{i=1}^n y_i \epsilon_i$$

$$z_0 = x_0 + y_0 \quad z_i \epsilon_i = (x_i + y_i) \epsilon_i$$

- Multiplication is nonlinear:

$$\hat{z} = \hat{x} \cdot \hat{y} = (x_0 + \sum_{i=1}^n x_i \epsilon_i) \cdot (y_0 + \sum_{i=1}^n y_i \epsilon_i)$$


$$= (x_0 \cdot y_0) + (x_0 \cdot \sum_{i=1}^n y_i \epsilon_i) + (y_0 \cdot \sum_{i=1}^n x_i \epsilon_i) + (\sum_{i=1}^n x_i \epsilon_i \cdot \sum_{i=1}^n y_i \epsilon_i)$$

$$= (x_0 \cdot y_0) + \sum_{i=1}^n (x_0 y_i + y_0 x_i) \epsilon_i + \sum_{1 \leq i, j \leq n} |x_i y_j| \epsilon_{n+1}$$

over-approximation!

AA Arithmetic Operations

- Unary operations (sqrt, inverse, log, exp, etc.):
 - ➔ compute linear approximation



Quiz



What is the result of the following computation?

$$\hat{x} = 3 + 2\epsilon_1 + 4\epsilon_2$$

$$\hat{y} = 1 + 1\epsilon_2 + 5\epsilon_3$$

$$\hat{z} = \hat{x} - \hat{y} =$$

a)

$$\hat{z} = 2 + 3\epsilon_2$$

c)

$$\hat{z} = 2 + 2\epsilon_1 + 3\epsilon_2$$



b)

$$\hat{z} = 2 + 2\epsilon_1 + 3\epsilon_2 - 5\epsilon_3$$

Roundoff Analysis with AA

For arithmetic operations: $eval(x \odot y)$

- real-valued range: $xrange \odot yrange$
- error: $propagatedError + newRoundoffError$

Recall the roundoff error analysis

Roundoff Analysis with AA

For arithmetic operations: $eval(x \odot y)$

- real-valued range: $xrange \odot yrange$
- error: $propagatedError + newRoundoffError$

To improve the precision, we have several options to replace IA with AA:

Option 1: AA for ranges and errors

- ➔ operations over AA use AA rules
- ➔ error propagation remains the same

Roundoff Analysis with AA

For arithmetic operations: $eval(x \odot y)$

- real-valued range: $xrange \odot yrange$
- error: $propagatedError + newRoundoffError$

To improve the precision, we have several options to replace IA with AA:

Option 1: AA for ranges and errors

- ➔ operations over AA use AA rules
- ➔ error propagation remains the same

Option 2: AA for ranges and IA for errors

- ➔ operations over AA use AA rules
- ➔ error propagation remains nearly the same
 - ➔ need to cast AA into an interval

Roundoff Analysis with AA

For arithmetic operations: $eval(x \odot y)$

- real-valued range: $xrange \odot yrange$
- error: $propagatedError + newRoundoffError$

To improve the precision, we have several options to replace IA with AA:

Option 1: AA for ranges and errors

- ➔ operations over AA use AA rules
- ➔ error propagation remains the same

Option 2: AA for ranges and IA for errors

- ➔ operations over AA use AA rules
- ➔ error propagation remains nearly the same
 - ➔ need to cast AA into an interval

Option 3: IA for ranges and AA for errors

- ➔ operations over AA use AA rules
- ➔ error propagation remains nearly the same
 - ➔ need to cast an interval into an AA

Different Kinds of Approximations

Affine arithmetic is not universally better than interval arithmetic.

They commit different kinds of (over-)approximations:

- ➔ interval arithmetic loses correlations ($x - x \neq 0$)
 - ➔ affine arithmetic over-approximates nonlinear operations
-
- AA over-approximation can be larger than due to interval's loss of correlations
 - AA gives best possible result for linear programs
 - AA over-approximation is smaller when input intervals are smaller
 - ✓ Daisy uses AA by default to track errors, and intervals to track ranges

Interval Subdivision

is an additional technique to reduce over-approximations

→ standard technique in numerical analysis

Idea:

- split input intervals into smaller subintervals
- form cartesian product to get all possible subdomains
- run analysis on each subdomain
- worst-case error is the overall maximum

Interval Subdivision

Example:

$$x = [-10, 10], y = [0, 5]$$

Splitting each input domain into 4:

$$x_1 = [-10, -5], x_2 = [-5, 0], x_3 = [0, 5], x_4 = [5, 10]$$

$$y_1 = [0, 1.25], y_2 = [1.25, 2.5], y_3 = [2.5, 3.75],$$

$$y_4 = [3.75, 5]$$

Run analysis on each subdomain:

$$(x_1, y_1), (x_1, y_2), (x_1, y_3), (x_1, y_4), (x_2, y_1), \dots$$

Alternative Roundoff Error Analysis

Recall the absolute error:

$$\max_{x \in I} |f(x) - \tilde{f}(\tilde{x})|$$

This is fundamentally an **optimization problem**

- ➔ suggest an alternative static analysis based on global optimization

Problem:

- ➔ \tilde{f} is highly discontinuous and complex
- ➔ formulation as-is cannot be handled by optimization tools

Idea:

- ➔ approximate \tilde{f} using floating-point abstraction ($x \circ_{fp} y = (x \circ y) \cdot (1 + \delta)$)
- ➔ simplify formula further using first-order Taylor approximation

Recent Research Directions

Probabilistic Error Analysis

The worst-case analysis may be too pessimistic!

✓ for applications that may tolerate large infrequent errors.

Probabilistic Analysis of Errors:

- considers probability distribution of x
- propagates the distribution through the program using probabilistic AA

$$\hat{x} := x_0 + \sum_{i=1}^n x_i \epsilon_i \quad \epsilon \in [-1, 1]$$

noise symbol propagates discretized distributions


$$d_x := \langle [a_1, b_1], w_1 \rangle, \dots, \langle [a_n, b_n], w_n \rangle$$

set of <interval, weight>

- combined with interval subdivisions

Rewriting Optimization

If the accuracy is still not acceptable:

- ➔ increase precision and check again
 - ➔ rewrite equations, check accuracy
- 

Examples:

$$\begin{array}{lcl} a + (b + c) & \xrightarrow{\text{mathematical identities}} & (a + b) + c \\ a * (b + c) & \xrightarrow{\hspace{1.5cm}} & (a * b) + (a * c) \end{array}$$

Finding an optimal order is computationally infeasible!

- ➔ Daisy uses a genetic (heuristic) search to find better expressions
 - tournament selection of an expression based on fitness
 - randomly mutate
 - evaluate fitness of the new ones

Mixed-Precision Tuning

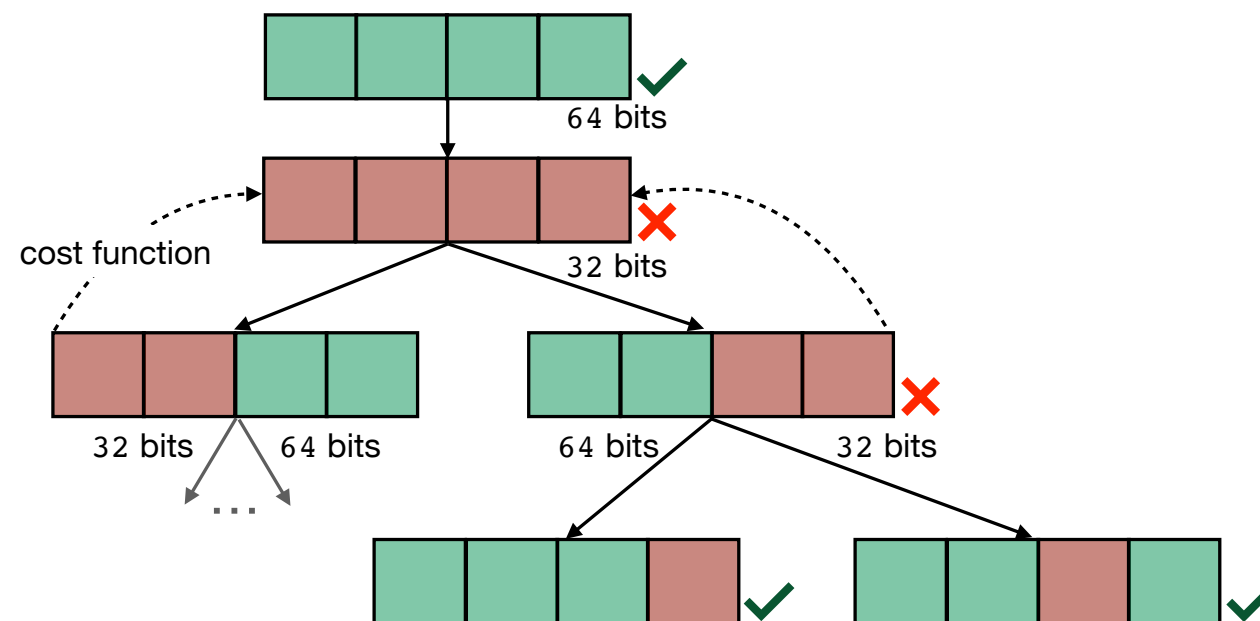
Assigning one precision to all variables may be suboptimal:

➔ Optimize precision to increase performance

minimize: precision and static cost

subject to: $\epsilon_n \leq \epsilon_{target}$

➔ Daisy uses a delta debugging search to find the optimal precision



Limitations

Static analyses compute sound worst-case and probabilistic roundoff errors, and support rewriting and mixed-precision tuning.

But, they have limitations in:

- scalability for large programs
 - (any) scalable analysis + static analysis*
- efficiency for high-precision analysis
- fine-grained optimizations
 - tailor to application contexts**

*A Two-Phase Approach for Conditional Floating-Point Verification, D. Lohar, C. Jeangoudoux, J. Sobel, E. Darulova, M. Christakis, TACAS'21

**Sound Mixed Fixed-Point Quantization of Neural Networks", D. Lohar, C. Jeangoudoux, A. Volkova, E. Darulova, EMSOFT'23

**Of Good Demons and Bad Angels: Guaranteeing Safe Control under Finite Precision, S. Teuber, D. Lohar, B. Beckert, FMCAD'25

Demo ~35min



Installation:

- `sudo docker pull dlohar/daisy`
- `sudo docker run -it dlohar/daisy`

`float-analysis-class/Test.scala:`

```
import daisy.lang._
import Real._
object Test {
  def test(x: Real, y: Real): Real = {
    require(1 <= x && x <= 3 && -5 <= y && y <= 4)

    val z = x * y + x + y + 2*x*x
    z
  } ensuring (res => res +/- 1e-5)
}
```

Demo



~35min



Running Daisy:

- `sbt script`
- `./daisy float-analysis-class/Test.scala`
 - `--rangeMethod=affine --errorMethod=affine`
 - `--precision=Float32`

Demo



~35min



Running Daisy:

- `sbt script`
- `./daisy float-analysis-class/Test.scala`
 `--rangeMethod=affine --errorMethod=affine`
 `--precision=Float32`
- `./daisy float-analysis-class/Test.scala`
 `--rangeMethod=affine --errorMethod=affine`
 `--precision=Float32 --subdiv --divLimit=10`
 - for subdivision Z3 is needed which is currently not installed!

Demo



~35min



Running Daisy:

- `sbt script`
- `./daisy float-analysis-class/Test.scala`
`--rangeMethod=affine --errorMethod=affine`
`--precision=Float32`
- `./daisy float-analysis-class/Test.scala`
`--rangeMethod=affine --errorMethod=affine`
`--precision=Float32 --subdiv --divLimit=10`
- `./daisy float-analysis-class/Test.scala`
`--precision=Float32 --rewrite`
- `./daisy float-analysis-class/Test.scala`
`--precision=Float32 --mixed-tuning --codegen`

Demo



~35min



Exercise 1

- Find the lowest uniform precision that guarantees the error bound.
- Which range and error analysis method works better?
- Can you optimize more to improve accuracy/efficiency?
- Generate an optimized code (scala, apfixed)

```
import daisy.lang._
import Real._
object Test {
  def test(x: Real, y: Real): Real = {
    require(1 <= x && x <= 3 && -5 <= y && y <= 4)

    val z = x * y + x + y + 2*x*x
    z
  } ensuring (res => res +/- 1.4e-14)
}
```