

CS 521: Topics in PL

Probabilistic &

Approximate

Computing

<http://misailo.web.engr.illinois.edu/courses/cs521>

Zoo:

SYSTEMS FOR ACCURACY-AWARE OPTIMIZATION

Original
Computation

Accuracy
Requirement

Accuracy-Aware Optimization

- **Find** an approximate program
- **Various** automatic or user-guided approaches

Optimized Computation +



Original
Computation

Typical
Inputs

Accuracy
Requirement

Testing-based Optimization

- **Transform** original computation
- **Validate** transformed computation

Optimized Computation +



For **all** inputs

Original
Computation

~~Typical
Inputs~~

Accuracy
Requirement

Analysis-driven Compiler

SAS '11, POPL '12, OOPSLA '13, OOPSLA '14, OOPSLA '19, CGO '20

- **Statically analyze** computation's accuracy
- **Transform** computation by solving a mathematical optimization problem

Optimized Computation +



Background: Compiler Autotuning

Search for program with maximum performance by reordering instructions, compiler parameters, and program configurations

- There are so many ways to tile an array (e.g., fit different cache sizes)
- Which optimizations to try `-O1`, `-O2`, `-O3`, remove some, add some?

Empirical process: explores the complexity of the system stack:

- Try new configuration
- If better than previous, save; and
- Search for more profitable configuration

Compiler Autotuning

Try new configuration: select one combination out of the space of all possible combinations

- Often too large to try them all
- The results will depend on the inputs you used

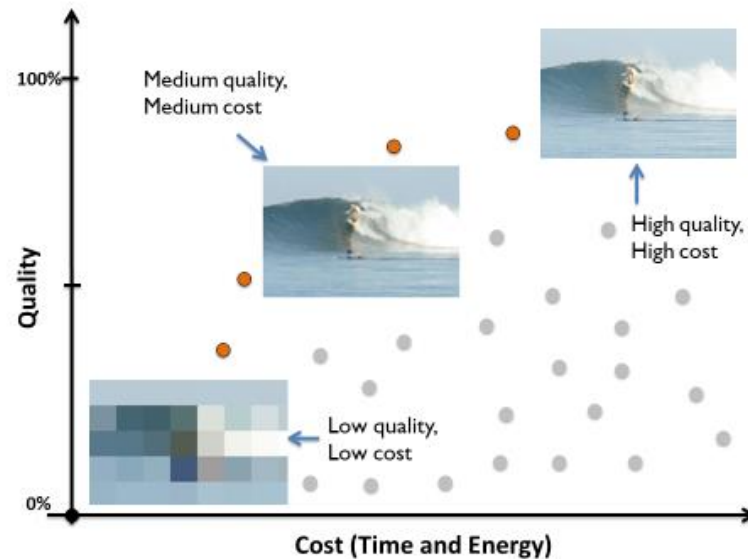
If better: (traditionally) compare performance or energy

- Uses fitness function which orders the configurations

Search for more: various heuristic algorithms, these days mainly based on machine learning and heuristic search (e.g., genetic programming in OpenTuner)

Compiler Autotuning

Accuracy opens up a new dimension for search



- Increases the number of options to try
- Includes (input-specific) accuracy metric in the fitness fun.
- Finds the configurations with best tradeoffs.

Multiobjective Optimization (Reminder)

Functions to optimize are called **objectives**

- Accuracy Loss – lower is better (or accuracy – higher is better)
- Speedup – higher is better (or normalized time – lower is better)
- Energy saving – higher is better (or consumption – lower is better)

They are the functions of program configuration – setting of knobs

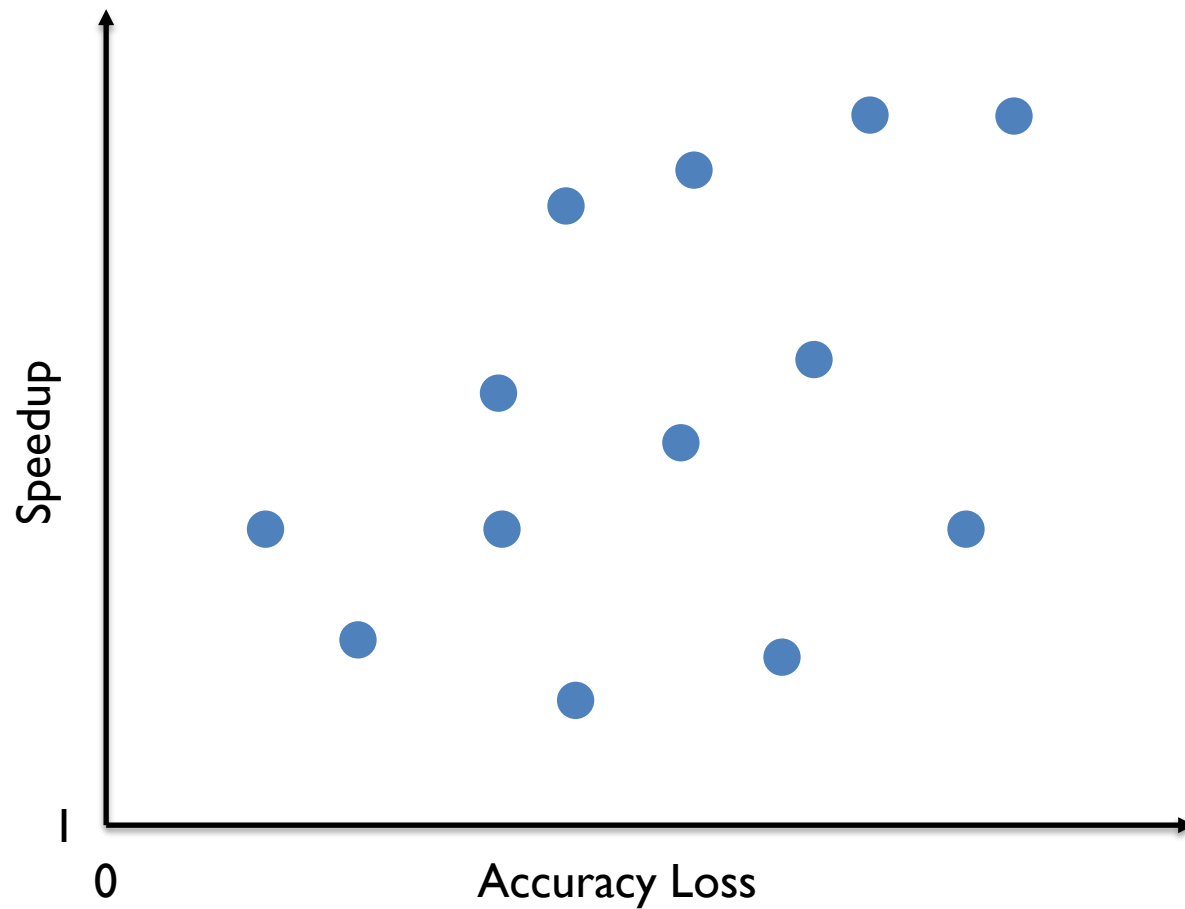
Two candidate program configurations X and Y :

- X **Pareto dominates** Y if X is as good as Y in all objectives, and is better in at least one objective

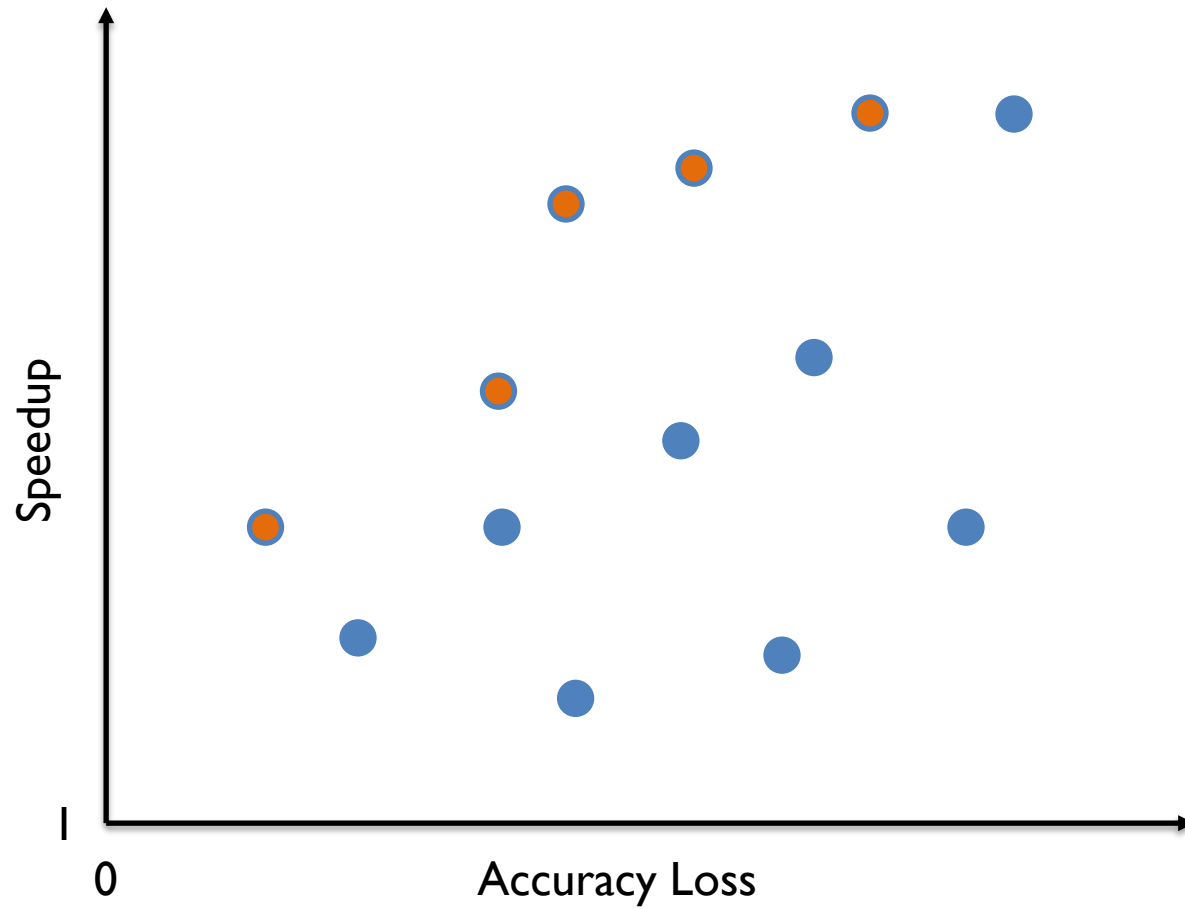
Pareto frontier: the set of points that are not dominated by other points

We will come back and formalize these notions later in the lecture!

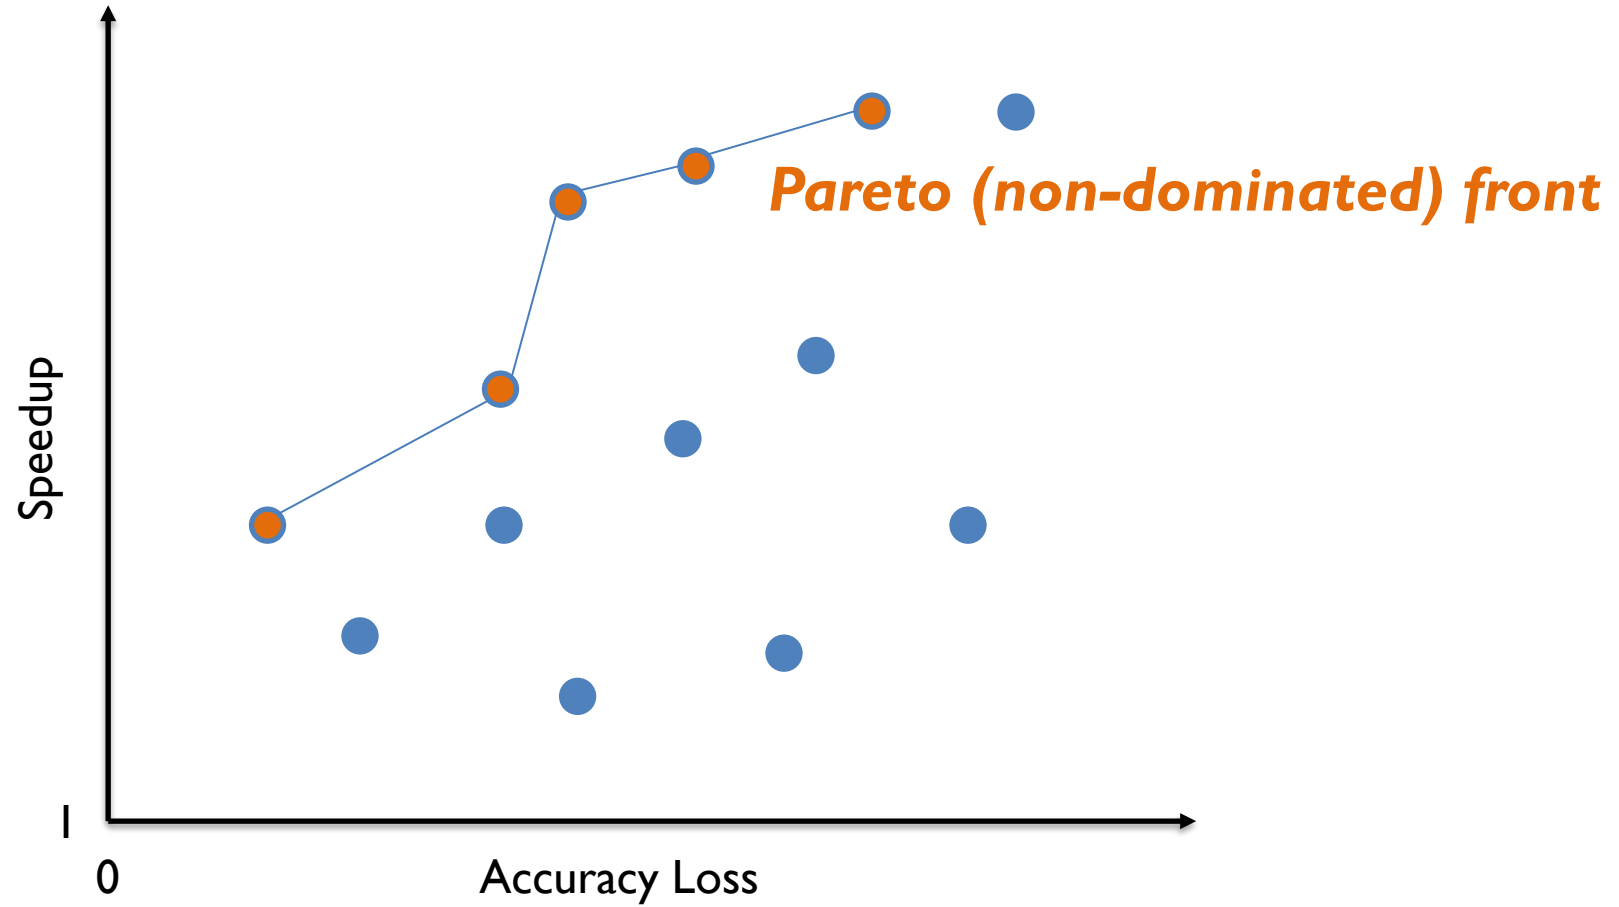
Example



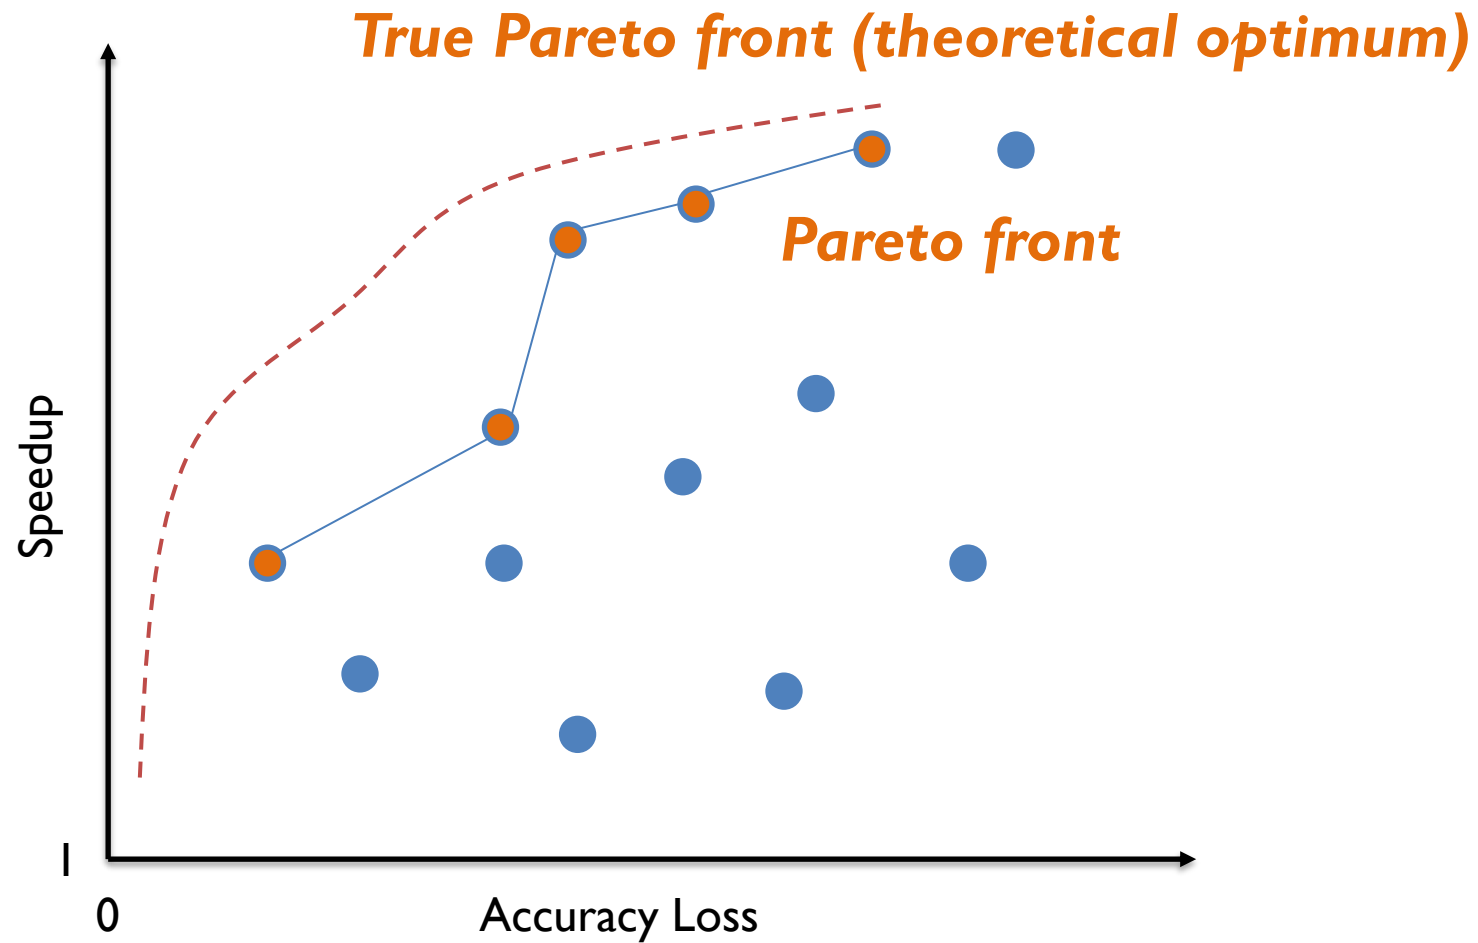
Example



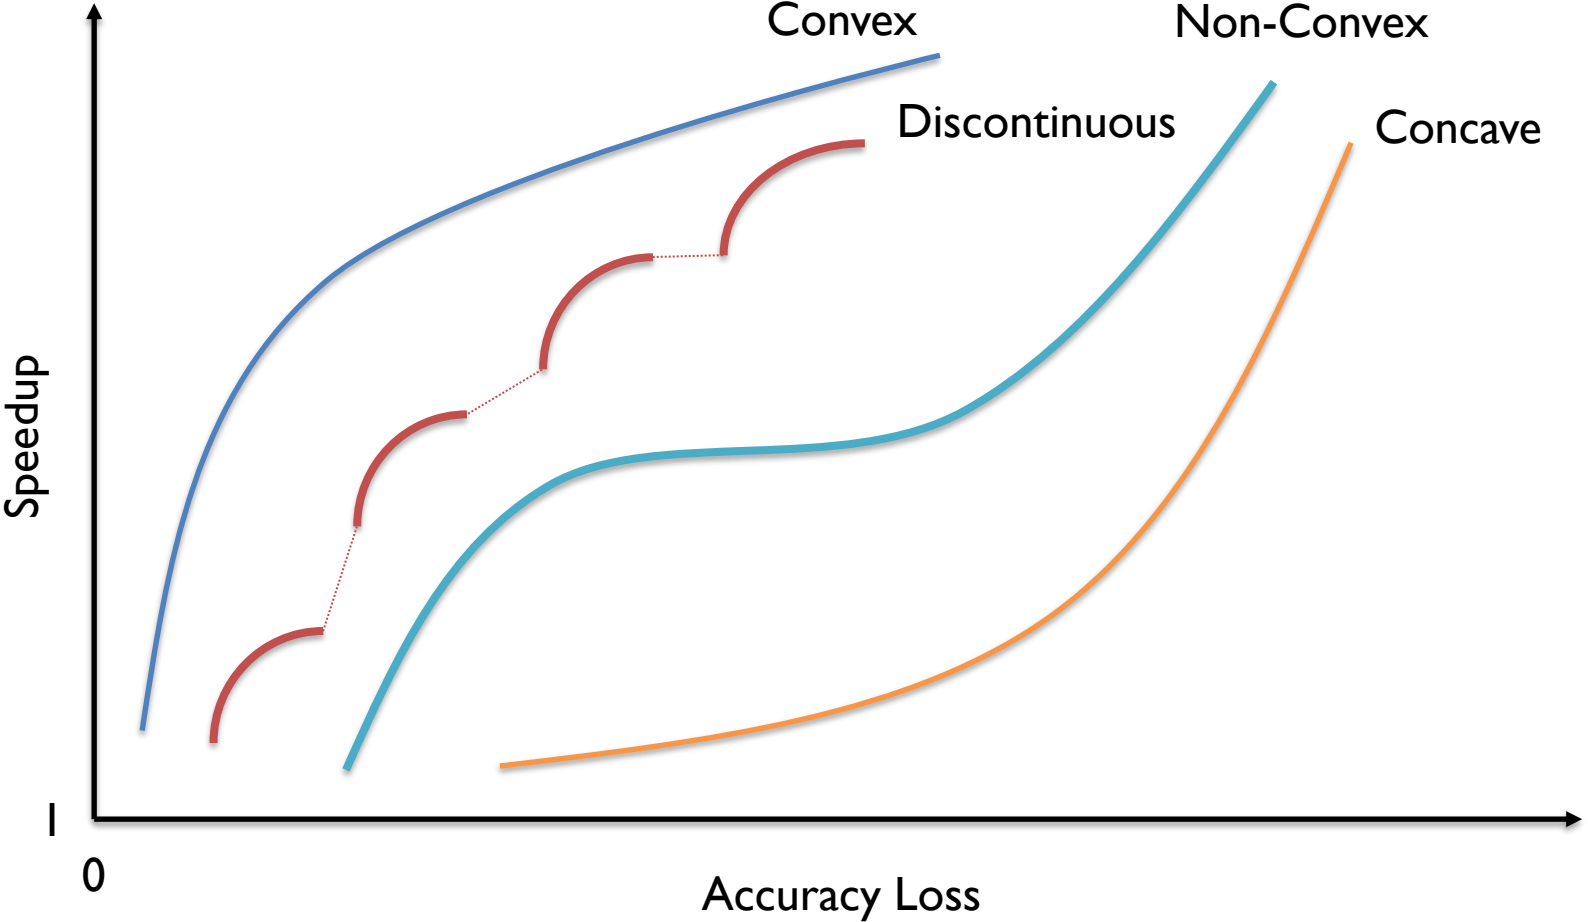
Example



Example



Pareto Fronts (aka Tradeoff curves)



A BIT OF FORMALISM

Based on Knowels, Thiele, Zitzler

A Tutorial on the Performance Assessment of Stochastic Multiobjective Optimizers (2006)

Optimization Problem

Optimization Problem is a Quadruple (X, Z, f, \preceq) :

- X : decision space, and $x \in X$ is a **decision vector**
- Z : objective space, and $z \in Z$ is a **objective vector**
- $f: X \rightarrow Z$ is a function that assigns to each decision vector x an objective vector $z = f(x)$
- We can think of it $z = (f_1, \dots, f_n) = f(x_1, \dots, x_m)$ while assuming $Z = R^n$
- \preceq is a binary relation over Z that defines a **partial order** of the objective space (it also induces a preorder on the decision space)

Weak Dominance

When $n = 1$ (single objective function):

- Optimization problem: $(X, \mathbb{R}, f, \preceq)$
- \preceq is our good old \leq on reals; there always exists a unique maximum

When $n > 1$ (multiple objective functions)

- Typically define \preceq as $z^{(a)} \preceq z^{(b)} \equiv \forall i \in \{1 \dots n\} z_i^{(a)} \leq z_i^{(b)}$
- Known as **weak Pareto dominance**: $z^{(b)}$ weakly dominates $z^{(a)}$

Optimization Problem

Goal:

Find solution x^* that is mapped to a maximal element $z^* = f(x^*)$ in the set $f(X) = \{z \in Z \mid \exists x \in X : z = f(x)\}$

Think: x is *program configuration*,

z is pair (accuracy, speedup), and

f computes (or records) accuracy and time of the execution.

- We can define the problem similar for searching minimal element (accuracy loss, run time)
- We can also make three dimensional tradeoff space accuracy, performance, energy, or even multidimensional

Our Optimization Problem

Select Program Configuration $X \in \text{Configs}$ to

maximize $(\text{Speedup}(X, i), \text{Accuracy}(X, i))$
forall $i \in \text{InputSet}$

But these are most often competing objectives.

Consider turning into weighted single optimization problem ($w_{1,2}$ express preference):

maximize $w_1 \times \text{Speedup}(X, i) + w_2 \times \text{Accuracy}(X, i)$
forall $i \in \text{InputSet}$

To maintain accuracy guarantees rephrase: for every accuracy loss threshold δ

maximize $\text{Speedup}(X, i)$
subject to $\text{AccuracyLoss}(X, i) \leq \delta$
forall $i \in \text{InputSet}$

Dominance

$z^{(a)} \preceq z^{(b)}$ for objective vectors of size n is defined as

$$\forall i \in \{1 \dots n\} \cdot z_i^{(a)} \leq_{\mathbb{R}} z_i^{(b)}$$

It is also called weak Pareto dominance

A strong Pareto dominance $z^{(a)} \prec z^{(b)}$ is defined as above, but cannot have any element being equal.

Read:

- $z^{(a)} \preceq z^{(b)}$ we say that $z^{(b)}$ **weakly dominates** $z^{(a)}$
- $z^{(a)} \prec z^{(b)}$ we say that $z^{(b)}$ **dominates** $z^{(a)}$

We can similarly define this relation for the cases when we want to maximize one but minimize another objective.

Dominance

We just learned about **Pareto Dominance (and weak dominance)**

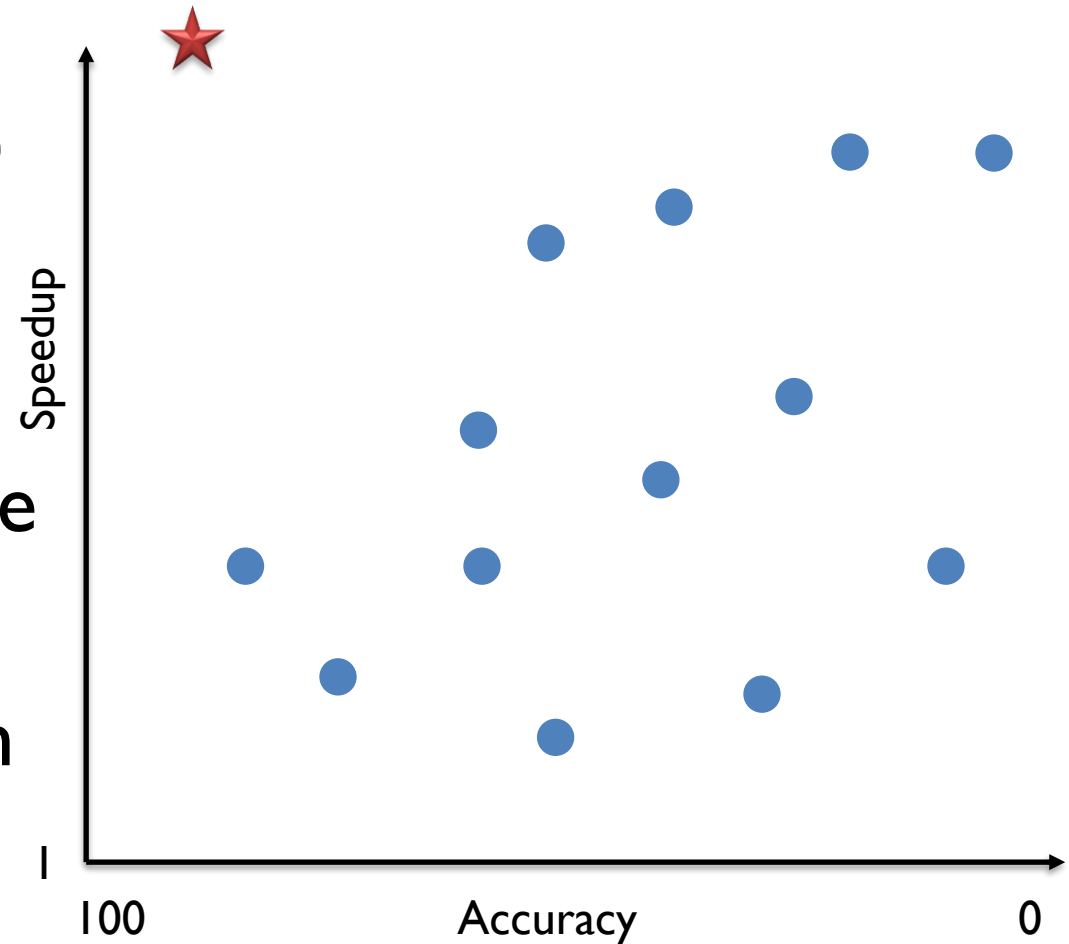
Incomparable points:

neither $z^{(a)} \preceq z^{(b)}$ nor $z^{(b)} \preceq z^{(a)}$

Indifferent: both points have the same value in all objectives

Strict domination: $z^{(a)}$ is better than $z^{(b)}$ in all objectives

Utopia point (★): result of individual objective optimizations



Pareto Set Approximations

In optimization we are interested in the entire Pareto-optimal set, not just individual solutions

- The set comprises the non-dominated objectives and decisions:*
 $A = \{(z, x) \mid \exists x \in X \exists z \in Z \text{ s.t. } z = f(x) \text{ and } z \text{ is not dominated}\}$
- We want to find mutually incomparable solutions
- Each such solution is a Pareto set approximation

We can extend the optimization problem: we want the best set of Pareto points (over other sets)

- Think: we want the best tradeoff curve across all that can be computed

*With a small abuse of notation, $z_B \in B$ refers to $(z_B, x_B) \in B$ for some x_B but the decision vector x_B is not necessary in this context; Alternatively, one could write $(z_B, _) \in B$. We treat the case $x_B \in B$ the same way.

```

SimpleCull( $\mathcal{C}$ )
 $\mathcal{C}_{min} := \emptyset$ ;
while  $\mathcal{C} \neq \emptyset$  do
     $c := \text{RemoveElement}(\mathcal{C})$ ; dominated := false;
    foreach  $d \in \mathcal{C}_{min}$  do
        if  $c \preceq d$  then  $\mathcal{C}_{min} := \mathcal{C}_{min} \setminus \{d\}$ ;
        else if  $d \preceq c$  then dominated := true; break;
    end
    if not dominated then  $\mathcal{C}_{min} := \mathcal{C}_{min} \cup \{c\}$ ;
end
return  $\mathcal{C}_{min}$ ;

```

Algorithm 2: Simple Cull minimisation algorithm

Comparing Pareto Sets Approximations

Let A and B the sets of Pareto-optimal points (e.g., produced by different search algorithms or multiple runs of a randomized algorithm)

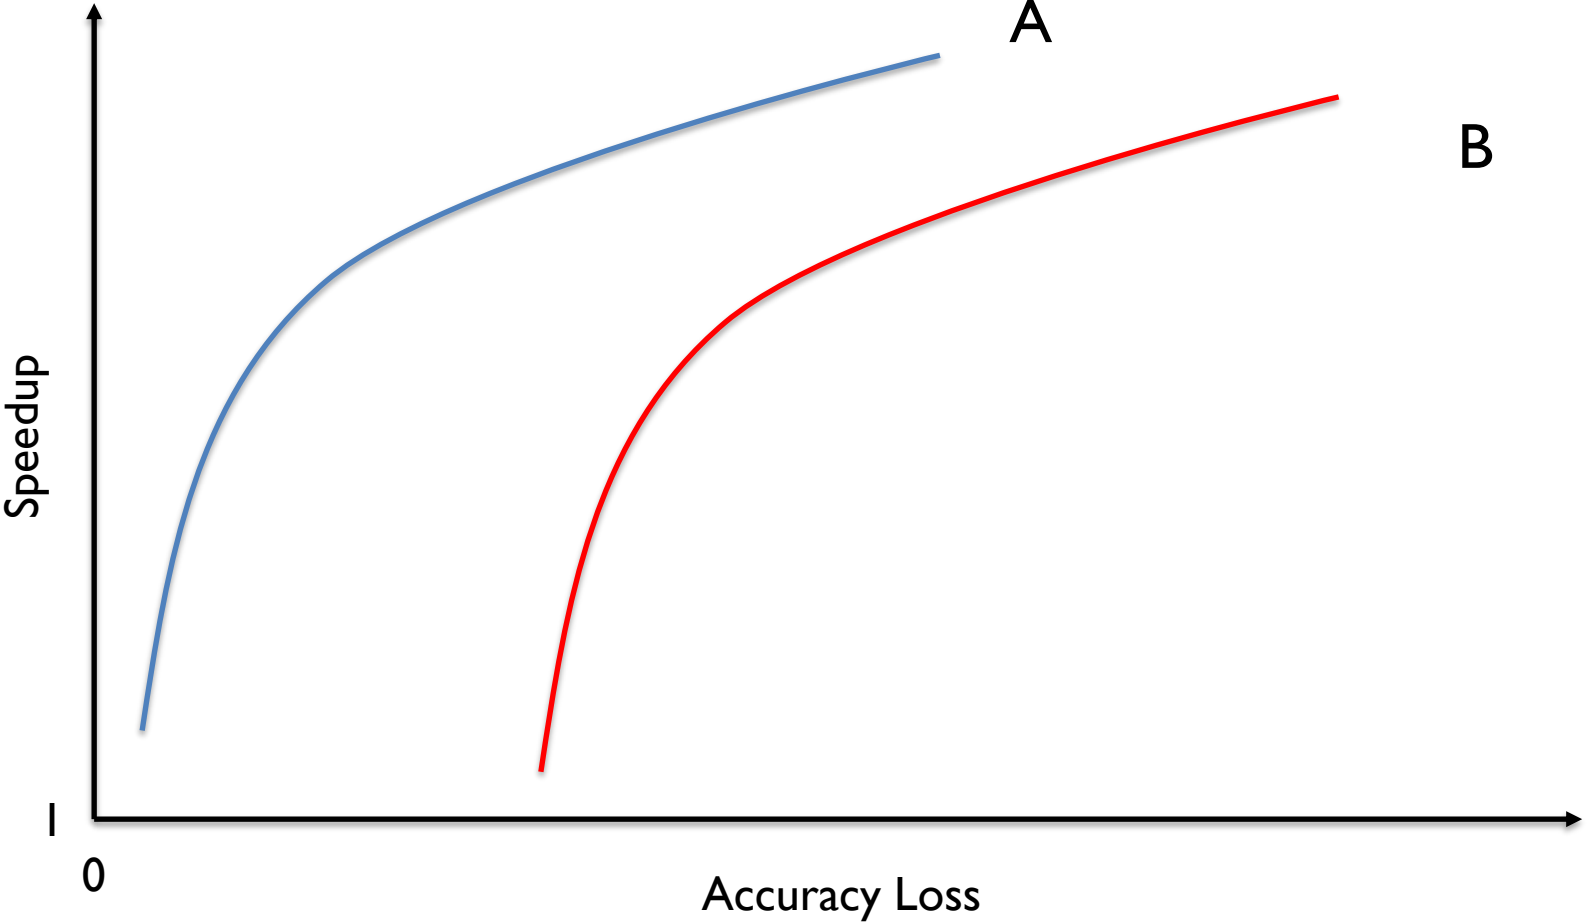
Is this enough? Typically no, we may need to define **quality indicators** to compare 'incomparable' sets

- There is no standard quality indicator, but needs to be selected based on context

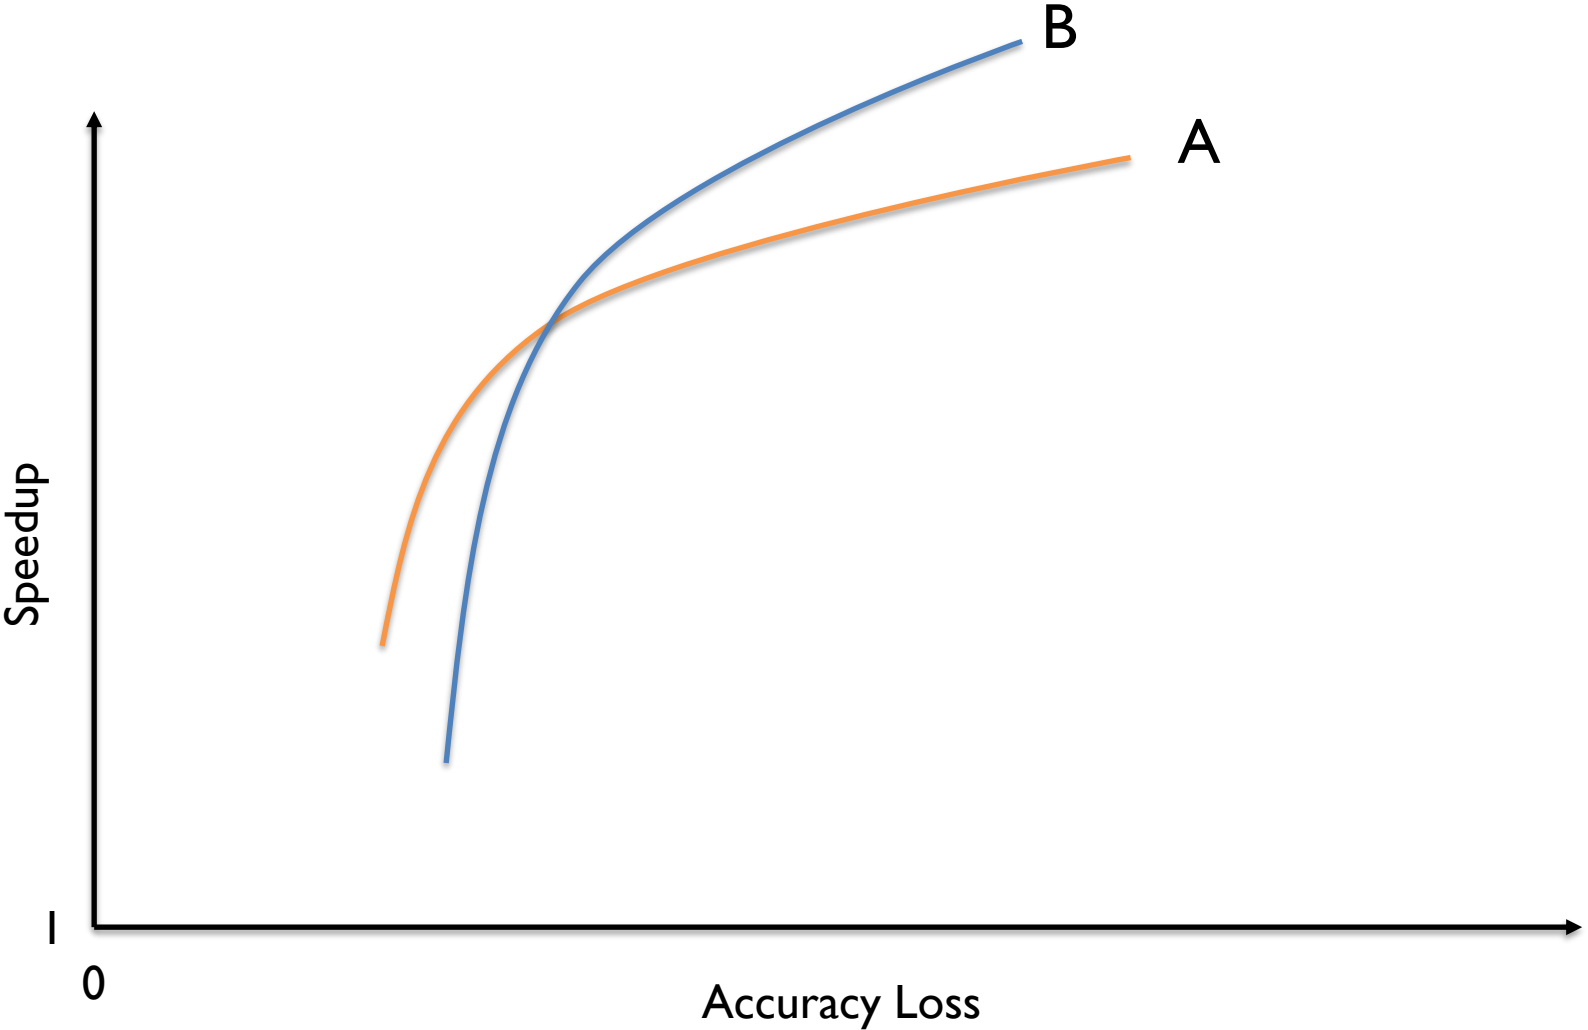
Hypervolume Indicator: intuitively, a volume (in our case area) of dominated solutions covered by the Pareto set.

- Need to select a reference point (or points). In our case, (speedup, accuracy) pairs (1.0, 100%) and (1.0, max-acceptable-accuracy) are intuitive choices
- Can order the Pareto sets $I(A) > I(B) \Rightarrow A \triangleright B$ (i.e., A is better than B)
- For randomized search algorithms, can compute and compare expected indicators i.e., $\mathbb{E} I(A) > \mathbb{E} I(B) \Rightarrow A \triangleright B$

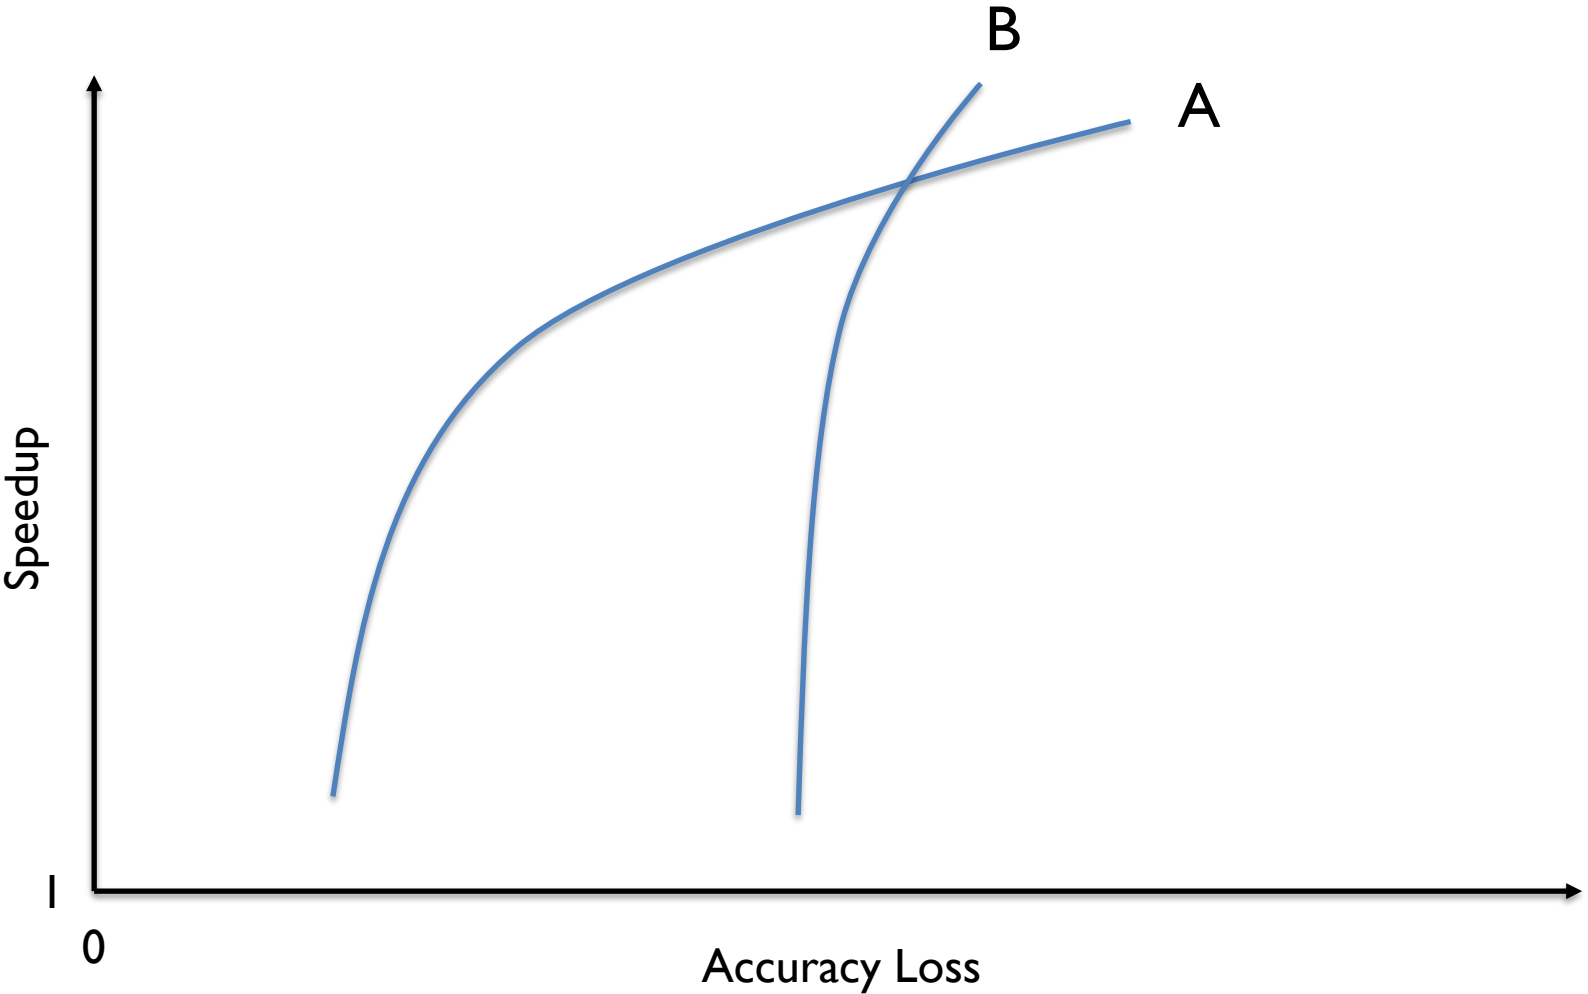
Pareto Fronts (aka Tradeoff curves)



Pareto Fronts (aka Tradeoff curves)



Pareto Fronts (aka Tradeoff curves)



Comparing Pareto Set Approximations

Let A and B the sets of Pareto-optimal points (e.g., produced by different search algorithms or multiple runs of a randomized algorithm)

We can define the relations for the sets:

- **A dominates B** ($B \prec A$) iff every $z_B \in B$ is dominated by some $z_A \in A$
- **Weak domination** ($B \preceq A$) is defined similarly
- **A and B are indifferent:** A weakly dominates B and B weakly dominates A
- **A is better than B** ($B \triangleleft A$): every $z_B \in B$ is weakly dominated by at least one $z_A \in A$ and A and B are not indifferent
- **A and B are incomparable:** neither set weakly dominates the other
- Is this enough? Typically no, we may need to define quality indicators to compare 'incomparable' sets.

Note on Our Optimization

Since we execute the programs, the input distributions will impact the approximation sets

Alternatively, if we combine with static analysis, some of the tradeoffs will end up being conservatively set

The search algorithms (e.g., auto-tuners) will impact what solutions we find – especially if they are randomized

The distribution between the ‘training’ and ‘test’ inputs may change, impacting accuracy and performance

LET'S START WITH PRACTICE THEN

Petabricks

Language for algorithmic choice (expresses options to tune) and an autotuner (using genetic search)

Precursor to OpenTuner (popular autotuner: <http://opentuner.org>)

Hand-coded algorithmic compositions are commonplace. A typical example of such a composition can be found in the C++ Standard Template Library (STL)¹ routine `std::sort`, which uses merge sort until the list is smaller than 15 elements and then switches to insertion sort. Our tests have shown that higher cutoffs (around 60-150) perform much better on current architectures. However, because the optimal cutoff is dependent on architecture, cost of the comparison routine, element size, and parallelism, no single hard-coded value will suffice.

Petabricks

Language for algorithmic choice (expresses options to tune) and an autotuner (using genetic search)

Precursor to OpenTuner (popular autotuner)

Classes of algorithms that can benefit from approximation:

- Polyalgorithms
- NP-Complete Algorithms
- Iterative Algorithms
- Signal Processing

Petabricks Autotuner

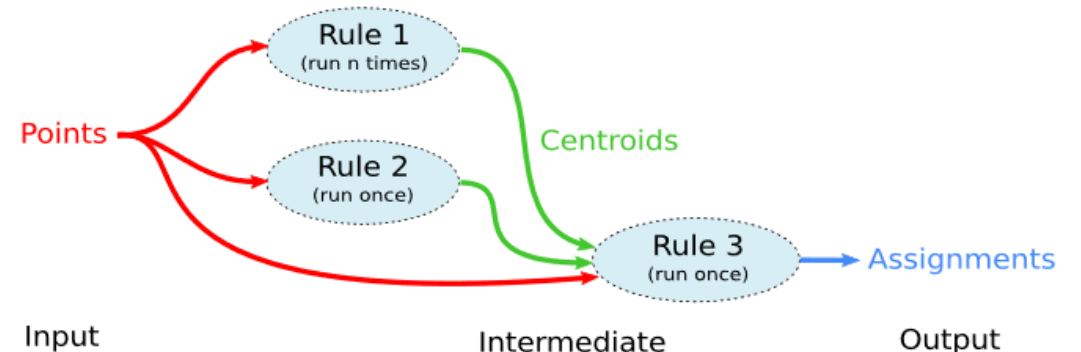
Language and Compiler Support
for Auto-Tuning Variable-Accuracy
Algorithms (CGO 2011)

```
transform kmeans
accuracy_metric kmeansaccuracy
accuracy_variable k
from Points[n,2] // Array of points (each column
                // stores x and y coordinates)
through Centroids[k,2]
to Assignments[n]
{
  ... (Rules 1 and 2 same as in Figure 1) ...
  // Rule 3:
  // The kmeans iterative algorithm
  to(Assignments a) from(Points p, Centroids c) {
    for_enough {
      int change;
      AssignClusters(a, change, p, c, a);
      if (change==0) return; // Reached fixed point
      NewClusterLocations(c, p, a);
    }
  }
}
```

The rules contained in the body of the transform define the various pathways to construct the Assignments data from the initial Points data.

```
// Rule 1:
// One possible initial condition: Random
// set of points
to(Centroids.column(i) c) from(Points p) {
  c=p.column(rand(0,n))
}

// Rule 2:
// Another initial condition: Centerplus initial
// centers (kmeans++)
to(Centroids c) from(Points p) {
  CenterPlus(c, p);
}
```



```
transform kmeansaccuracy
from Assignments[n], Points[n,2]
to Accuracy
{
  Accuracy from(Assignments a, Points p){
    return sqrt(2*n/SumClusterDistanceSquared(a,p));
  }
}
```

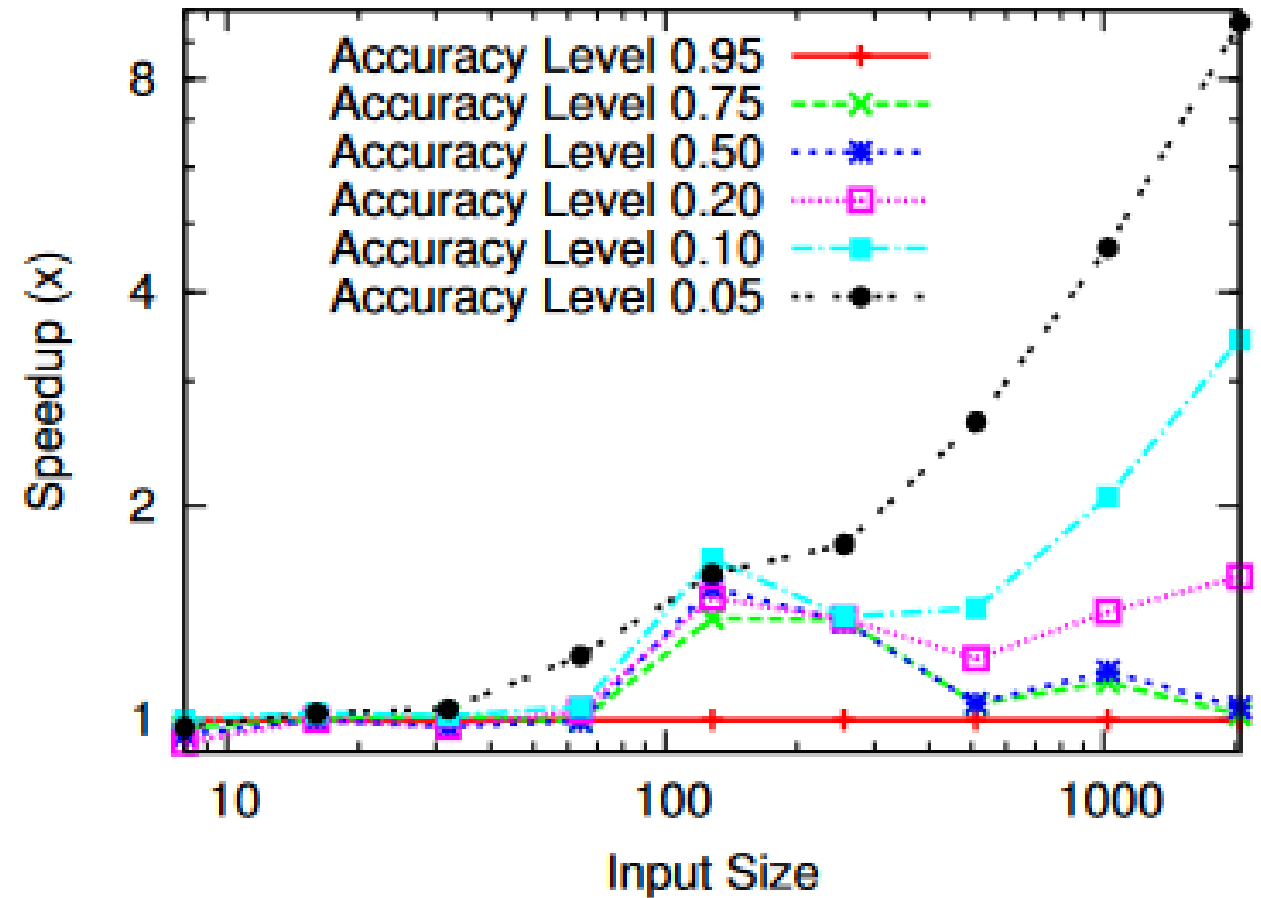
Petabricks Autotuner

Language and Compiler Support
for Auto-Tuning Variable-Accuracy
Algorithms (CGO 2011)

```
transform kmeans
accuracy_metric kmeansaccuracy
accuracy_variable k
from Points[n,2] // Array of points (each column
                // stores x and y coordinates)
through Centroids[k,2]
to Assignments[n]
{
... (Rules 1 and 2 same as in Figure 1) ...

// Rule 3:
// The kmeans iterative algorithm
to(Assignments a) from(Points p, Centroids c) {
  for_enough {
    int change;
    AssignClusters(a, change, p, c, a);
    if (change==0) return; // Reached fixed point
    NewClusterLocations(c, p, a);
  }
}
}

transform kmeansaccuracy
from Assignments[n], Points[n,2]
to Accuracy
{
  Accuracy from(Assignments a, Points p){
    return sqrt(2*n/SumClusterDistanceSquared(a,p));
  }
}
```



Next Step

What if a language **does not** expose approximation choices?

Let a compiler find and expose some by modifying the program!

Original
Program

Typical
Inputs

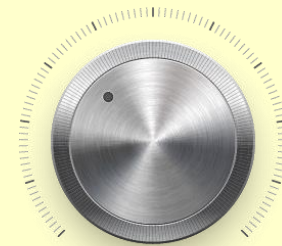
Accuracy
Specification

SpeedPress

- **Transforms** programs with perforation
- **Validates** new programs using testing

Quality of Service Profiling (ICSE 2010)
Managing Performance vs. Accuracy Trade-offs With Loop Perforation (FSE 2011)

Optimized Program +



x264 Video Encoder Example

**Typical
Inputs**



**Accuracy
Specification**

- **Quality Metric:**
e.g. PSNR and bit rate
- **Quality Loss:**
e.g. relative difference < 10%

Phases of Approximate Compiler:

Find perforatable loops

- **Identify Opportunity:** Run **performance** profiler
Identify time consuming loops
- **Sensitivity Testing:** Perforate **one loop** at a time
Filter out loops that do not satisfy accuracy requirement
- **Search for Optimal Knobs:** Perforate **multiple loops**
Find combinations of loops that maximize performance
Return a tradeoff curve of best solutions found

Validate Perforated Loops

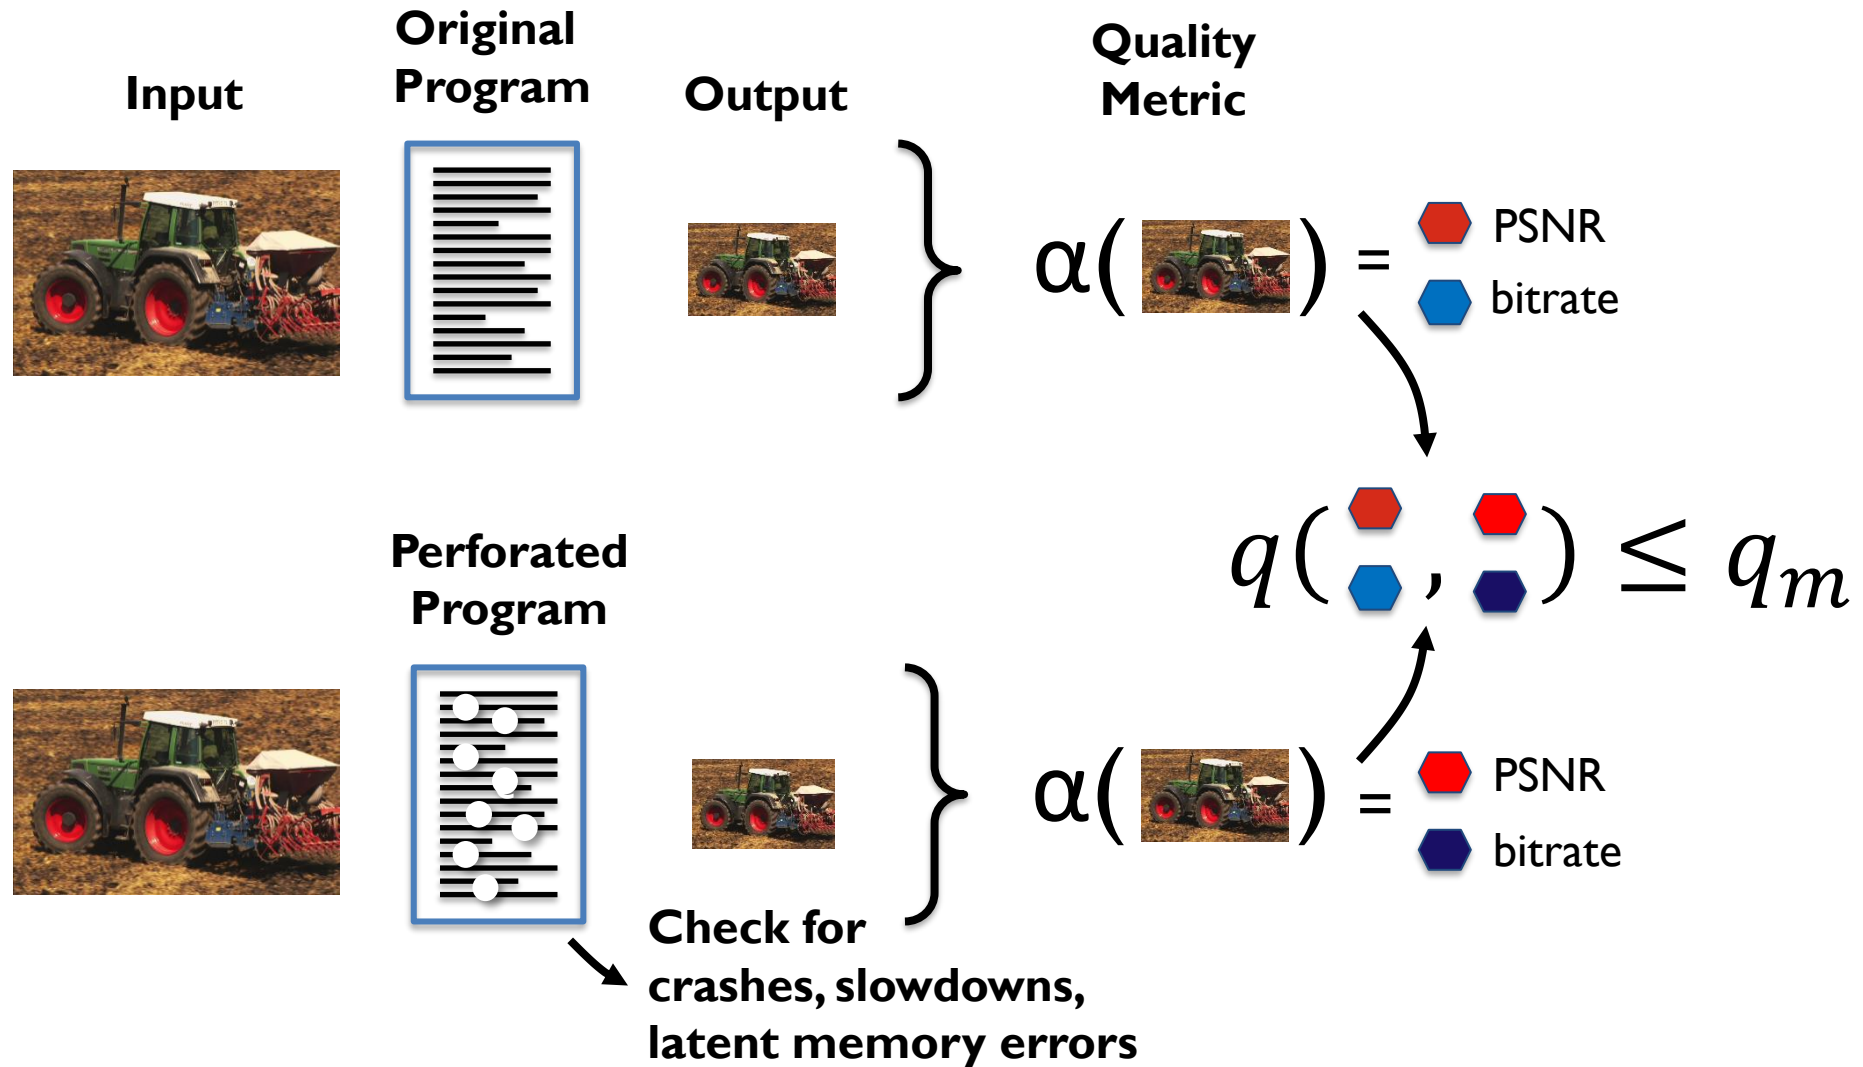
Filter out loops that do not satisfy requirement

Criticality (Sensitivity) Testing: Ensure that the program with perforated loop **does not:**

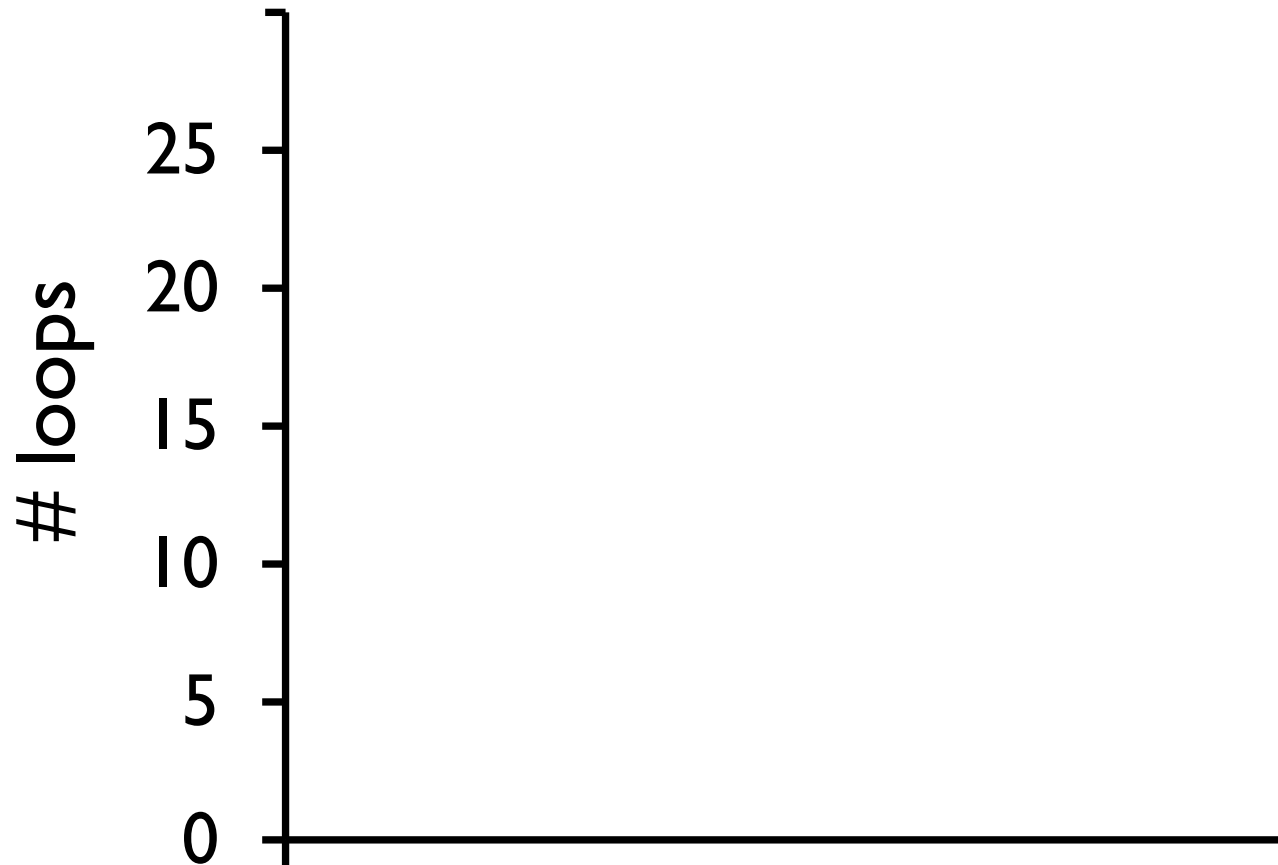
- Crash or return error
- Runs slower than original (or not terminates)
- Causes other errors identified by dynamic analysis (e.g., latent memory errors)
- Produces unacceptable result (e.g., NaN, inf...)
- Produces inaccurate result (according to accuracy metric)

Criticality (Sensitivity) Testing:

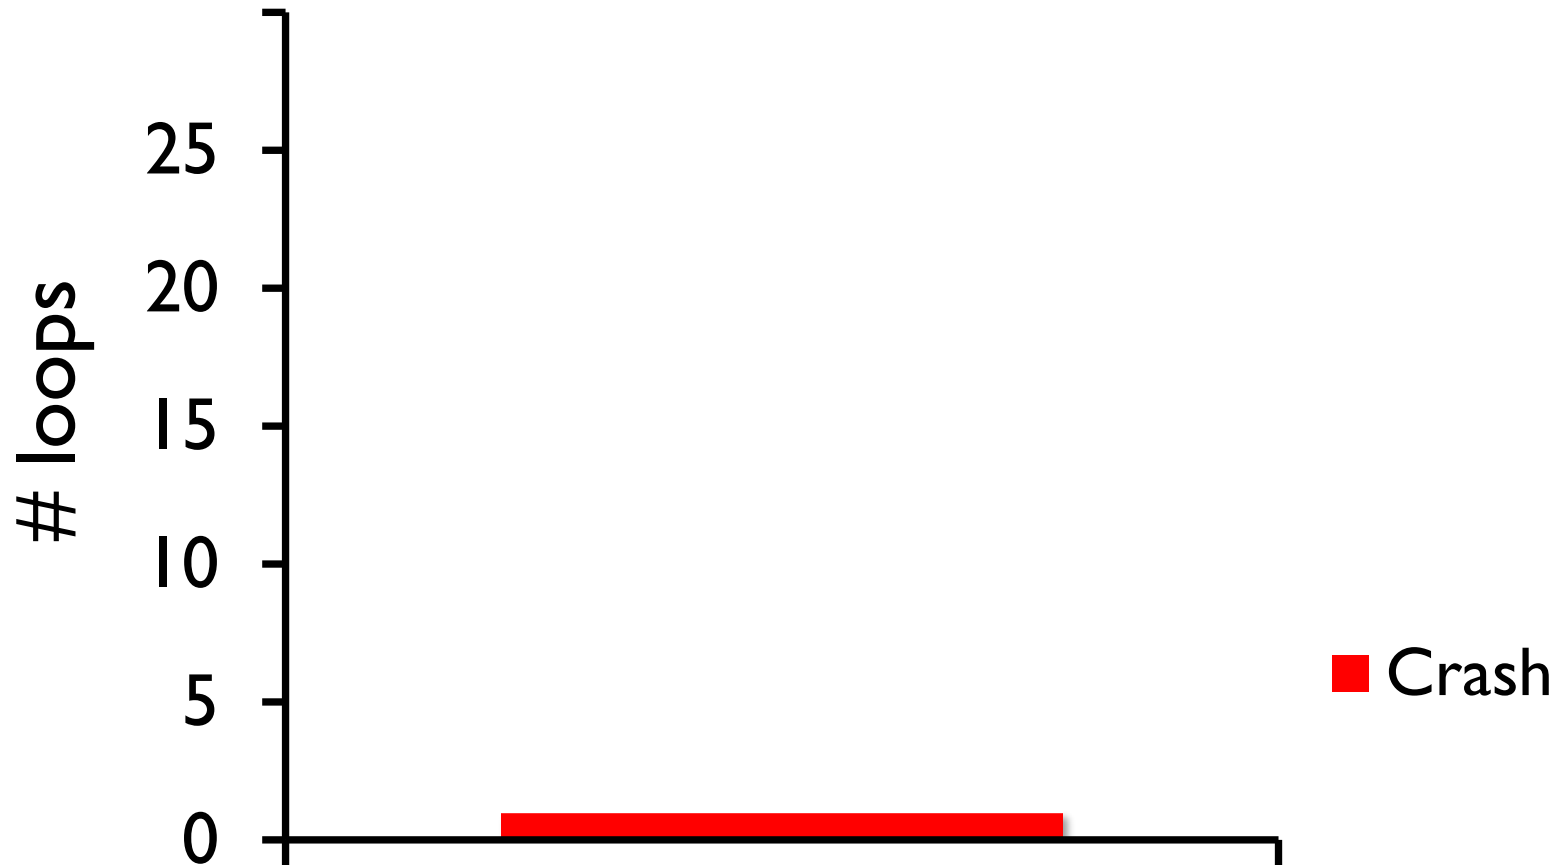
Filter out loops that do not satisfy requirement



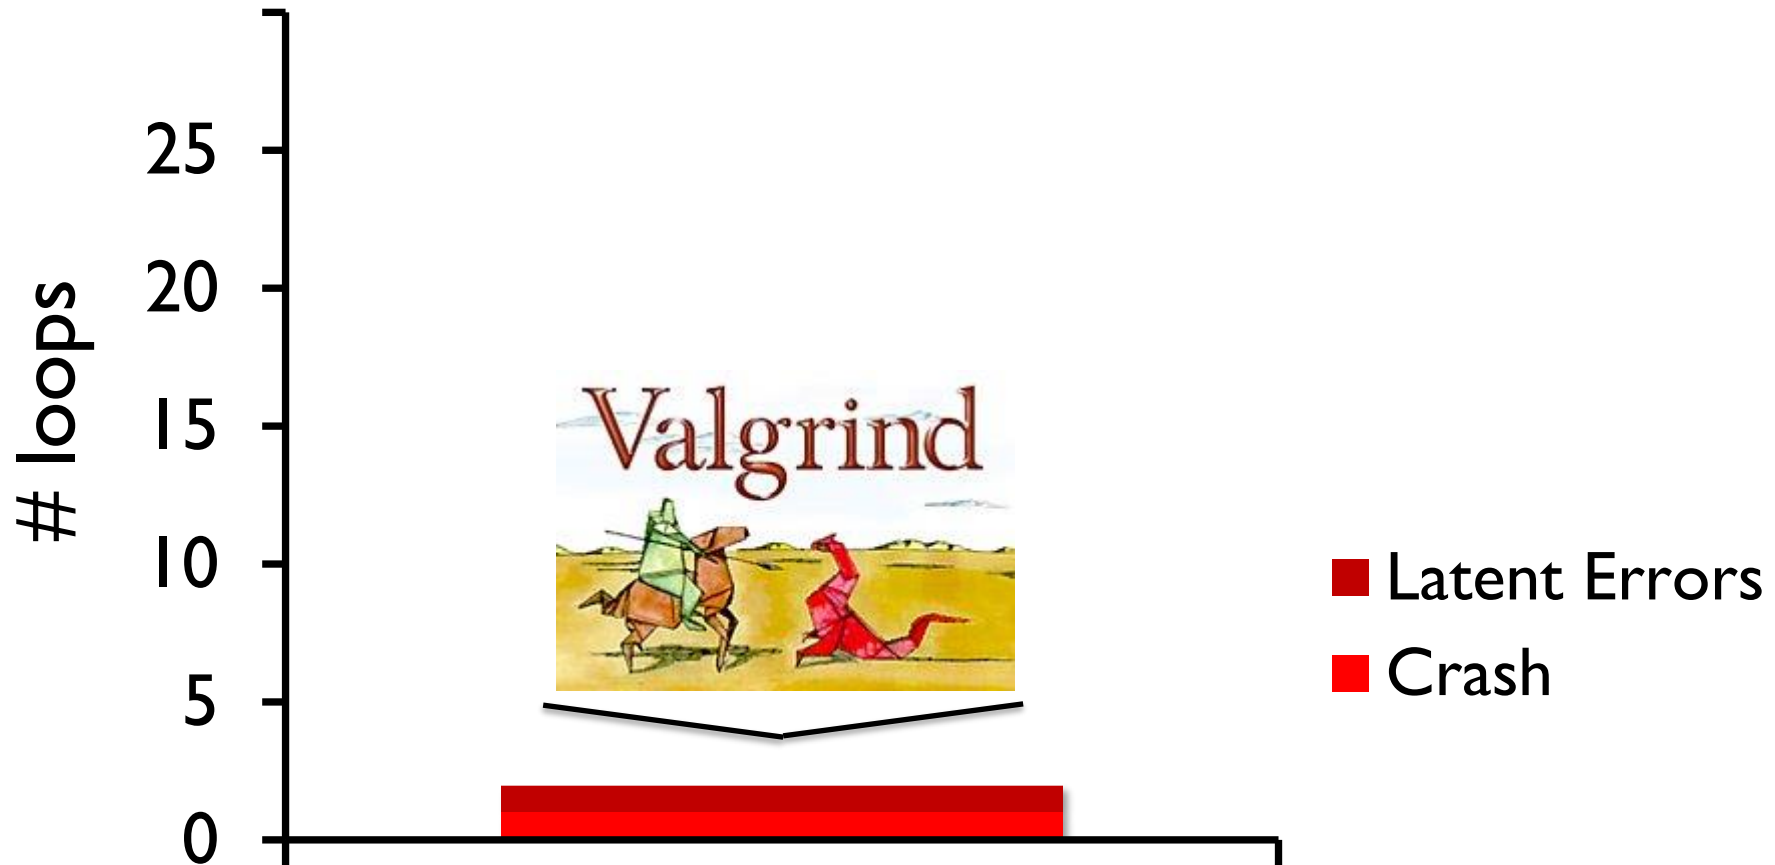
Perforating Individual Loops in **x264** (**Quality Loss < 0.1**)



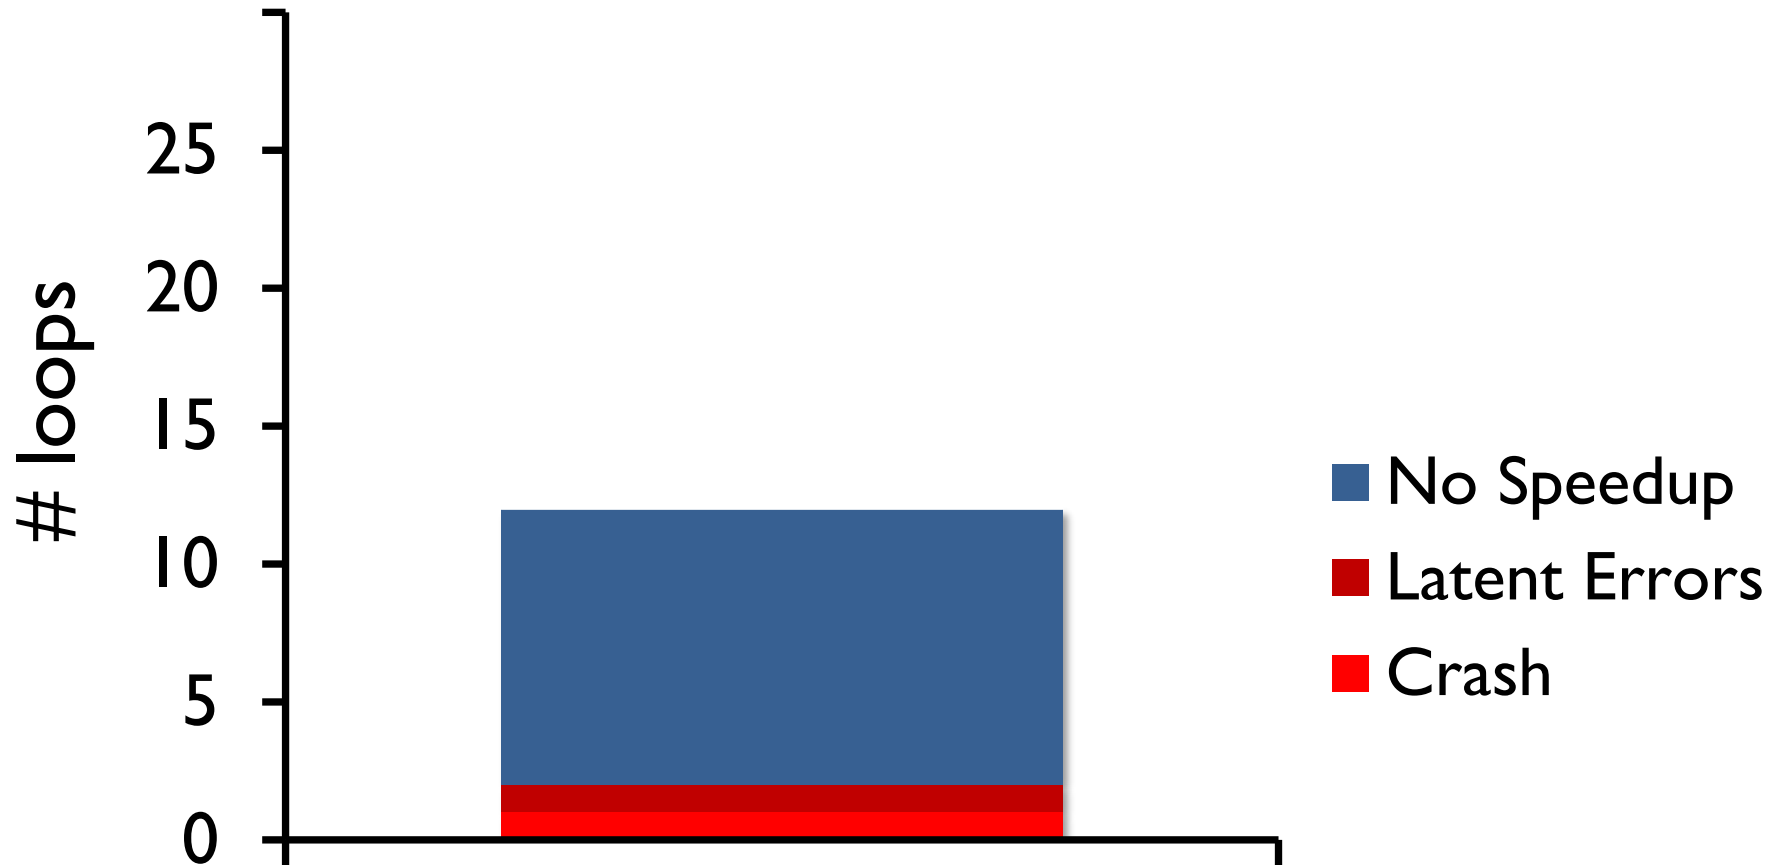
Perforating Individual Loops in **x264** (**Quality Loss < 0.1**)



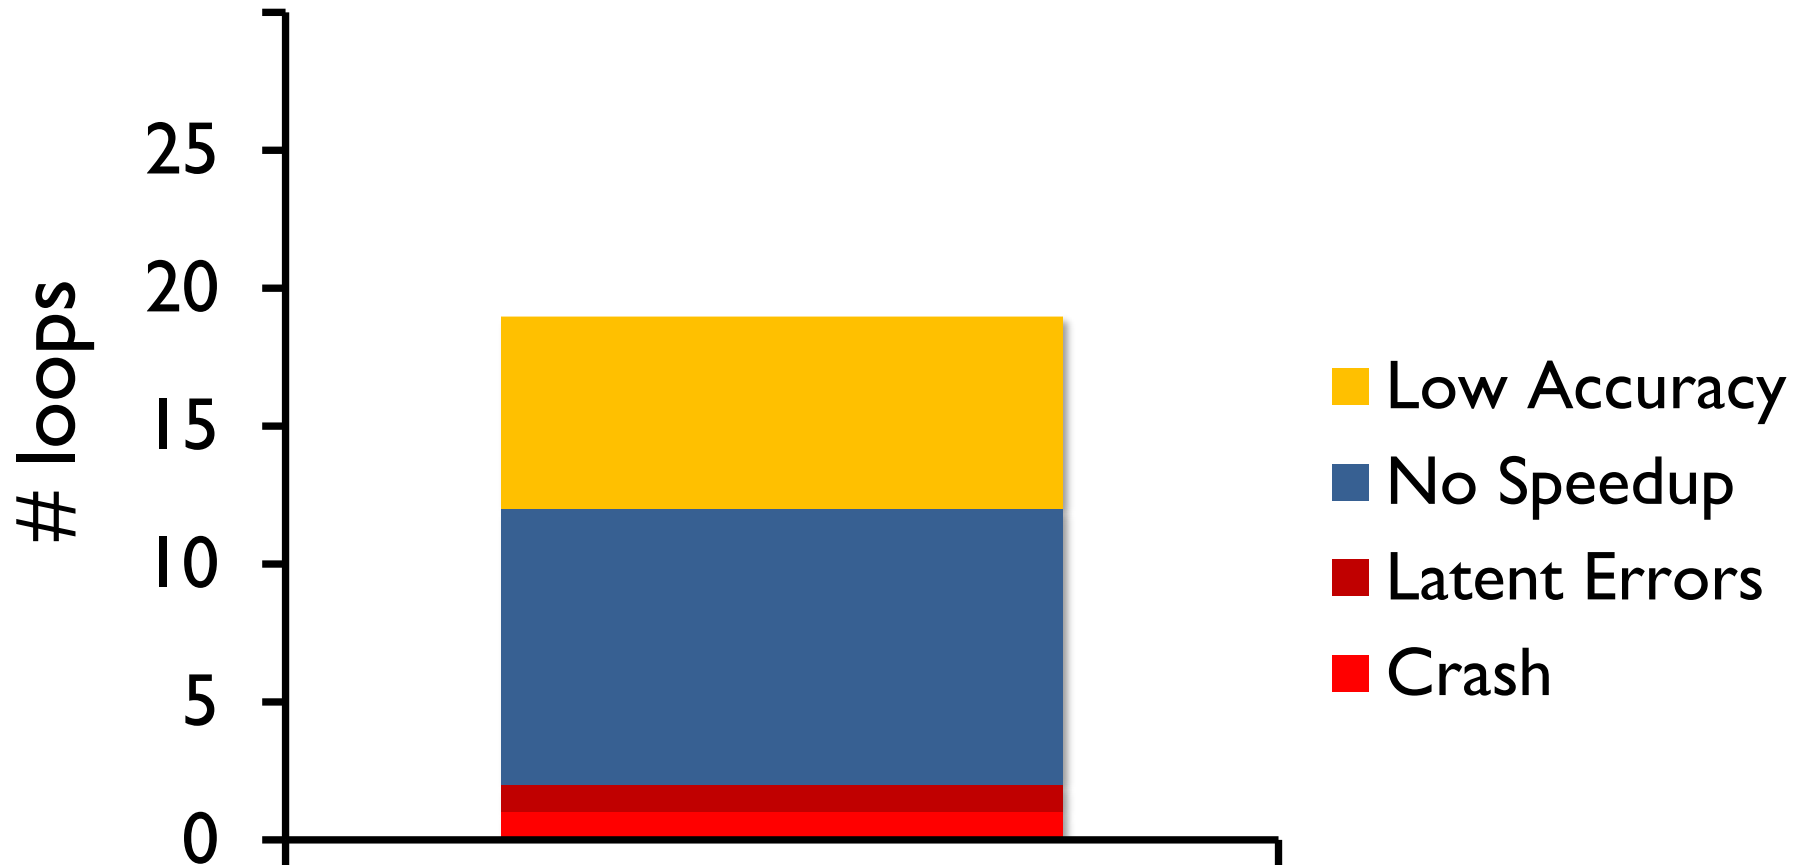
Perforating Individual Loops in **x264** (**Quality Loss < 0.1**)



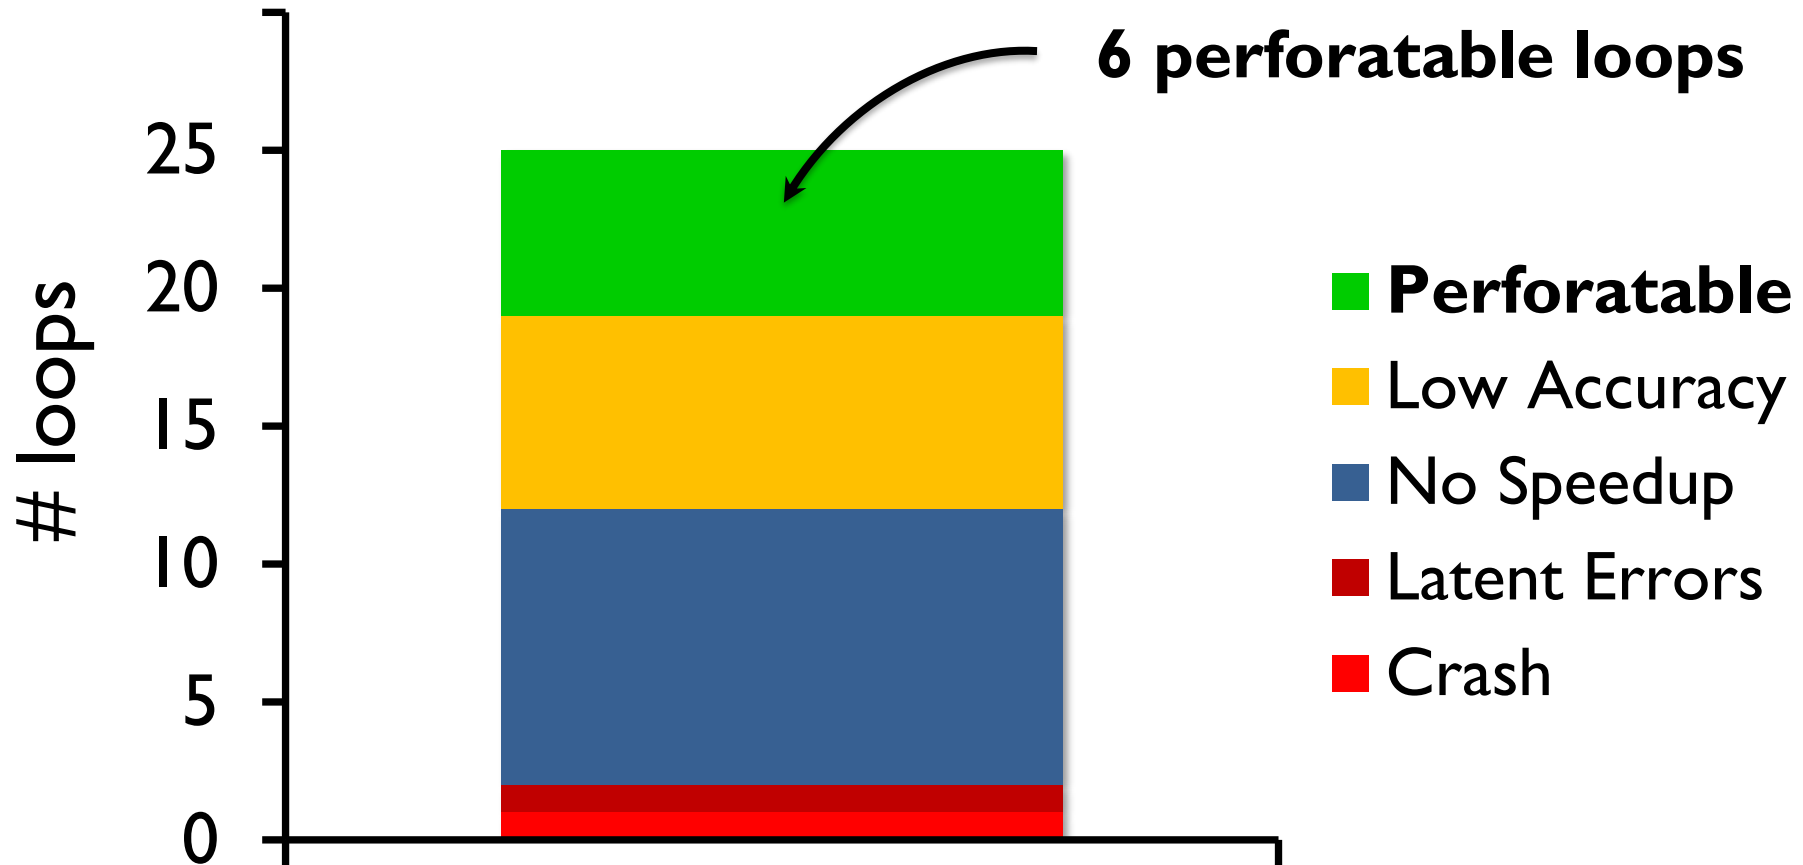
Perforating Individual Loops in **x264** (**Quality Loss < 0.1**)



Perforating Individual Loops in **x264** (**Quality Loss < 0.1**)



Perforating Individual Loops in **x264** (**Quality Loss < 0.1**)



Status

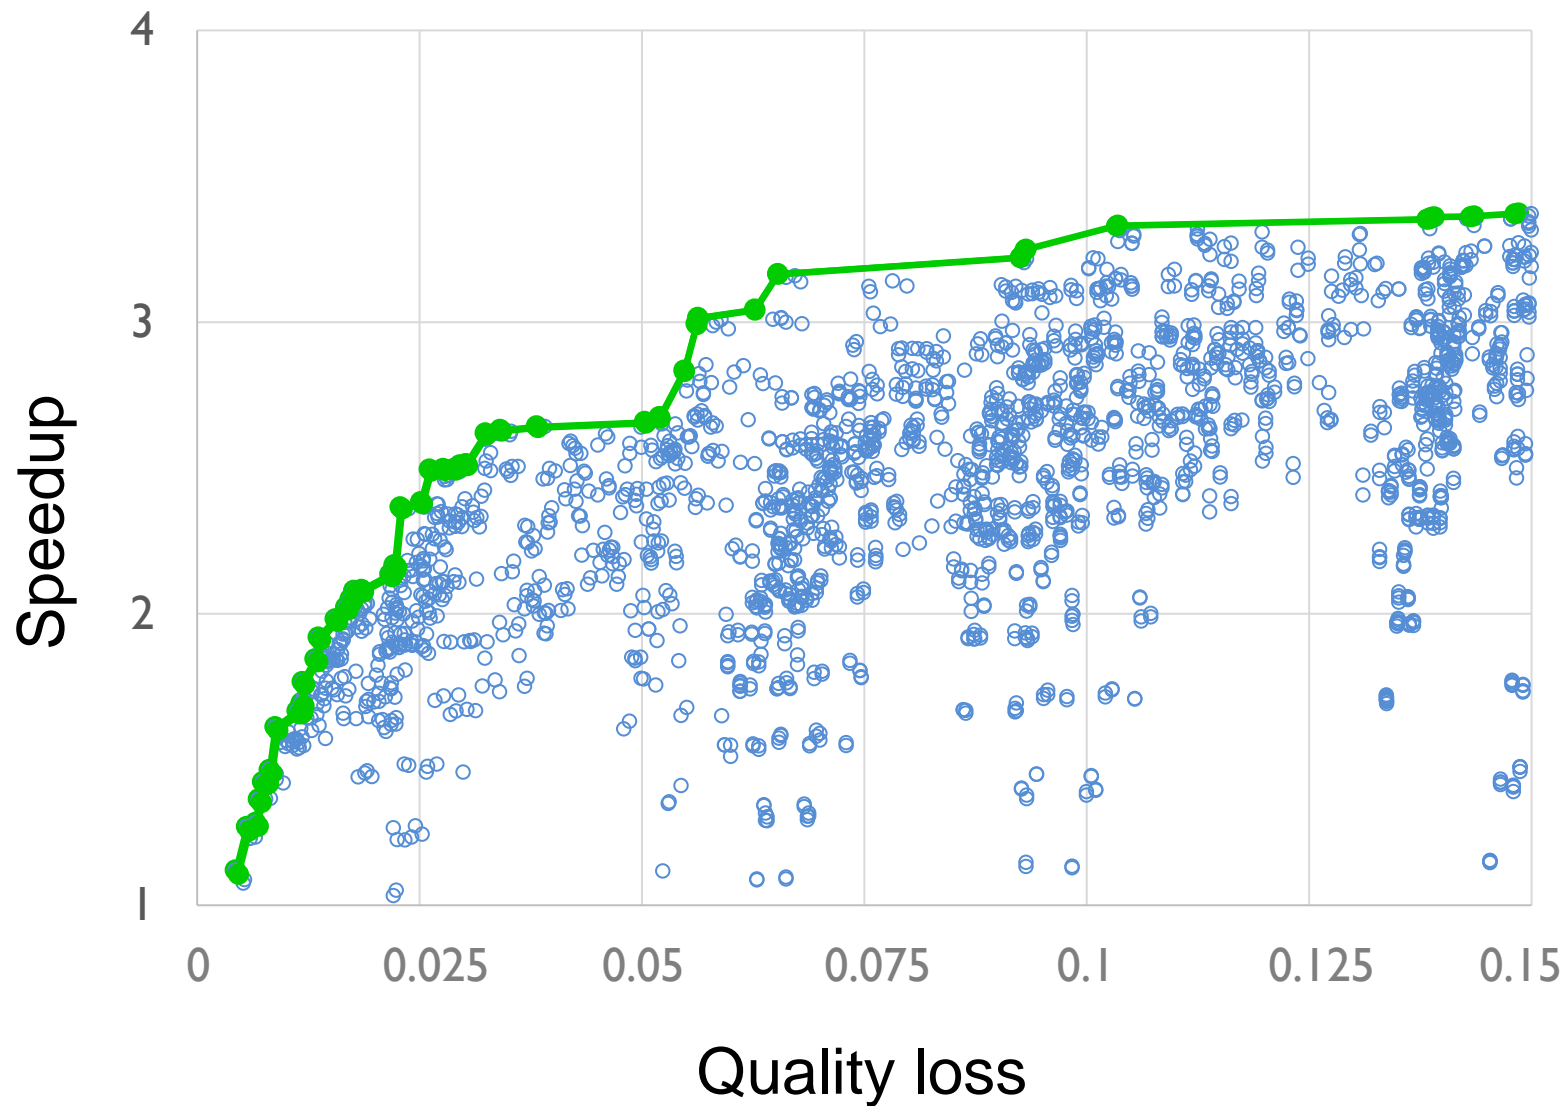
We found approximate computations and exposed individual knobs

Next, let us combine the knob values to utilize the approximation
“budget”

Search Strategies and Algorithms

- Greedy
 - Exhaustive
 - Combined
 - Hill-climbing
 - Simulated annealing
 - Genetic algorithm
 - Reinforcement learning
 - ...
- We had the comfort to do a bounded-exhaustive evaluation to explore the tradeoff space

Navigate Tradeoff Space



Applications

From PARSEC Suite

x264	video encoder
bodytrack	human motion tracking
swaptions	financial analysis
ferret	image search
canneal	electronic circuit placement
streamcluster	point clustering
blackscholes	financial analysis

Inputs

Augmented or Replaced Existing Sets

x264	from Internet
bodytrack	augmented
swaptions	randomly generated
ferret	provided inputs
canneal	augmented (autogenerated)
streamcluster	from Internet
blackscholes	provided inputs

Metrics

Application Specific

x264	PSNR + Size
bodytrack	weighted relative difference
swaptions	relative difference
ferret	recall
canneal	relative difference
streamcluster	clustering metric
blackscholes	relative difference

Loop Perforation

(Quality Loss < 10%)

x264	3.2x
bodytrack	6.9x
swaptions	5.0x
ferret	1.1x
canneal	1.2x
streamcluster	1.2x

Loop Perforation

(Quality Loss < 10%)

x264	3.2x	motion estimation
bodytrack	6.9x	particle filtering
swaptions	5.0x	MC simulation
ferret	1.1x	image similarity
canneal	1.2x	simulated annealing
streamcluster	1.2x	cluster center search

Loop Perforation

(Quality Loss < 10%)

x264
bodytrack
swaptions
ferret
canneal
streamcluster

Tasks of most perforated loops:

- Distance metrics
- Search-space enumeration
- Iterative improvement
- Redundant executions

Main Observations

- **Approximate Kernel Computations**
(have specific structure + functionality)
- **Accuracy vs Performance Knob**
(tune how aggressively to approximate kernel)
- **Magnitude and Frequency of Errors**
(kernels rarely exhibit large output deviations)

**Approximate
Program Analysis =**

Accuracy + Safety

Accuracy and Guarantees

Logic-Based (*worst-case*)

“for all inputs...”

Probabilistic (*worst-case or average-case*)

“for all inputs, with probability at least p ...”

“for inputs distributed as...”

Statistical (*average-case*)

“for inputs distributed as... with confidence c ”

“for tested inputs... with confidence c ”

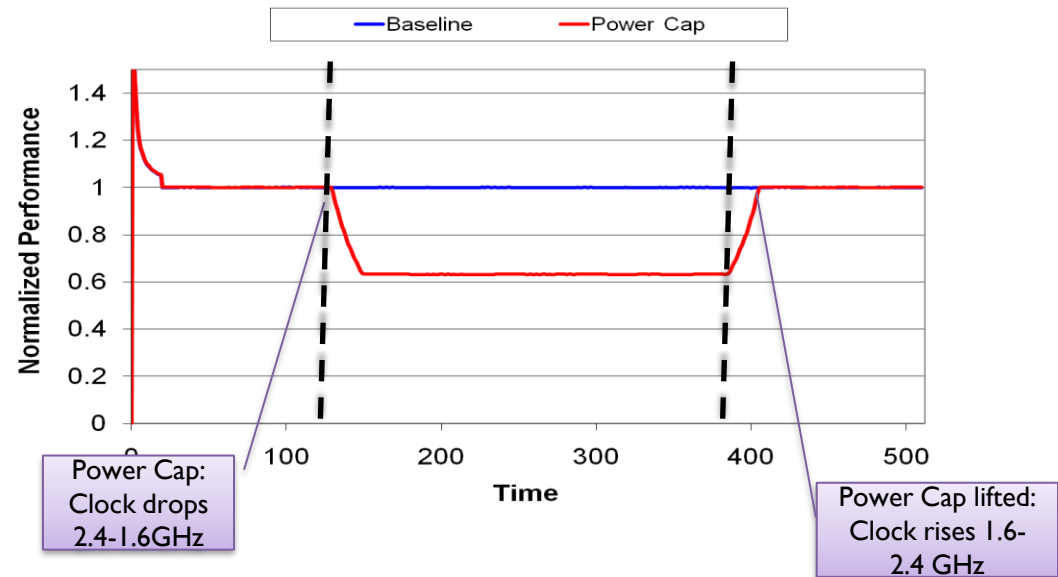
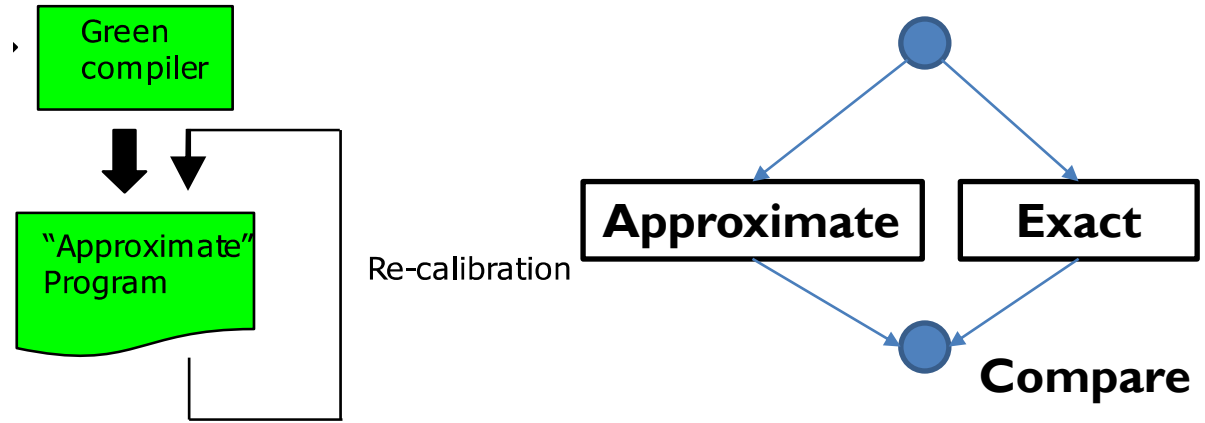
Empirical (*typical-case*)

“for typical inputs...”

Goals of Runtime Adaptation

Accuracy (Green)

Time or Energy
(Loop perforation)



Green : Framework for Controlled Approximations (PLDI'10) *

End-to-end framework for controlled application on approximations

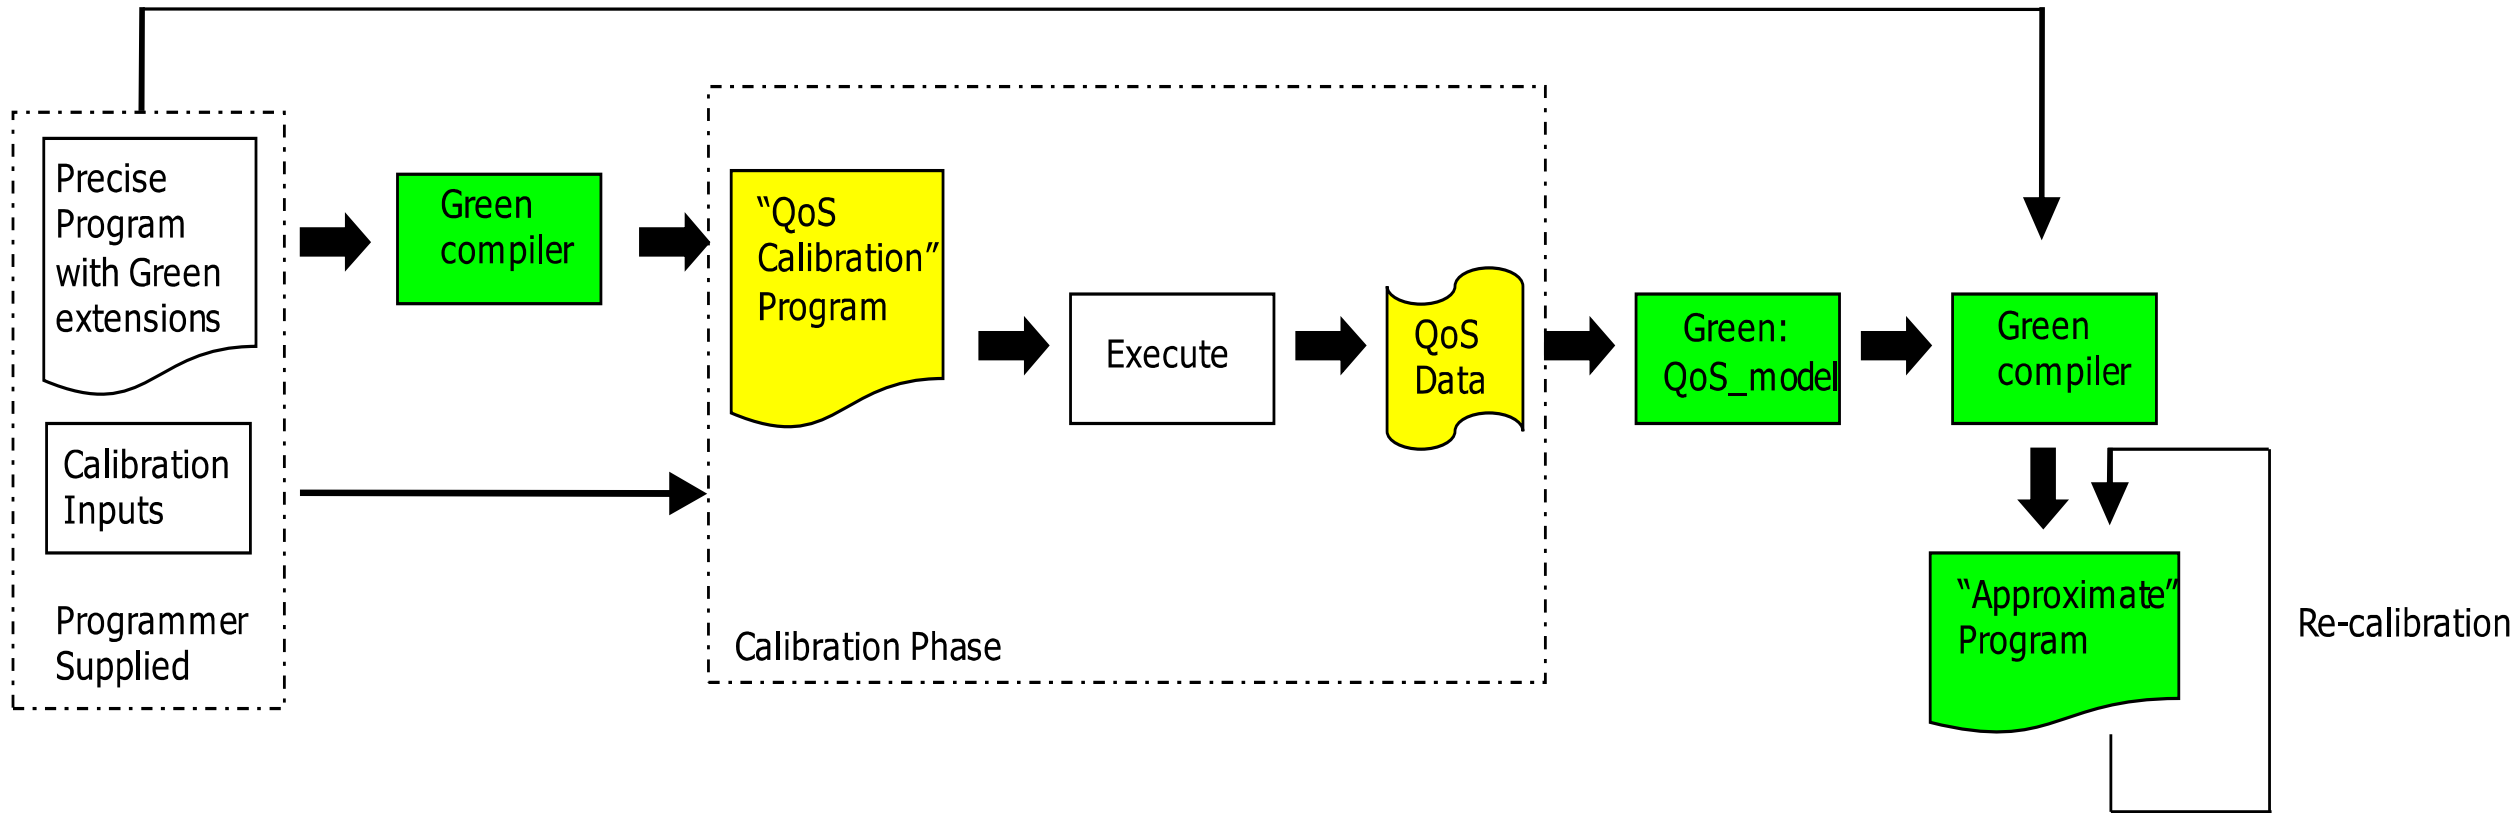
- Loop and function approximations

Relatively easy for programmers to use

Hooks for expert programmers and custom policies

Online mechanism to reactively adapt approximation policy to meet QoS

Green Framework



Recalibration

Concern:
Overhead for
running non-
approximate

Address: Run
infrequently,
restructure the
code

Original code:

```
#approx_loop (*QoS_Compute, Calibrate_QoS,  
             QoS_SLA, Sample_QoS, static)  
for(i=0; i<N; i++) {  
    pi_est += factor/(2*i+1);  
    factor /= -3.0;  
}
```

Calibration code:

```
for(i=0; i<N; i++) {  
    loop_body;  
    if ((i%Calibrate_QoS)==0) {  
        QoS_Compute(0, i, ...);  
    }  
}  
QoS_loss = QoS_Compute(1, i, ...);  
store(i, QoS_loss);
```

Approximation code:

```
count++;  
recalib=false;  
if (count%Sample_QoS==0) {  
    recalib=true;  
}  
  
for (i=0; i<N; i++) {  
    loop_body;  
    if (QoS_Lp_Approx(i, QoS_SLA, true)) {  
        if (!recalib) {  
            // Terminate the loop early  
            break;  
        } else {  
            // For recalibration, log the QoS value  
            // and do not terminate the loop early  
            if(!stored_approx_QoS) {  
                QoS_Compute(0, i, ...);  
                stored_approx_QoS = 1;  
            }  
        }  
    }  
  
    if(recalib) {  
        QoS_loss=QoS_Compute(1, i, ...);  
        QoS_ReCalibrate(QoS_loss, QoS_SLA);  
    }  
}
```

Default QoS_Lp_Approx:

```
QoS_Lp_Approx(loop_count, QoS_SLA, static) {  
    if (loop_count<M)  
        return false;  
    else {  
        if (static)  
            return true;  
        else {  
            // adaptive approximation  
        }  
    }  
}}
```

Default QoS_ReCalibrate:

```
QoS_ReCalibrate(QoS_loss, QoS_SLA) {  
    if (QoS_loss>QoS_SLA) {  
        // low QoS case  
        increase_accuracy();  
    } else if (QoS_loss<0.9*QoS_SLA) {  
        // high QoS case  
        decrease_accuracy();  
    } else {  
        // do nothing  
    }  
}
```

Figure 3. An end-to-end example of applying loop approximation to the Pi estimation program.

Recalibration

```
QoS ReCalibrate(QoS_loss, QoS_SLA) {
    // n_m: number of monitored queries
    // n_l: number of low QoS queries in monitored queries
    if (n_m==0) {
        // Set Sample QoS to 1 to trigger QoS ReCalibrate
        // for the next 100 consecutive queries
        Saved_Sample_QoS=Sample QoS;
        Sample QoS=1;
    }
    n_m++;
    if (QoS_loss !=0)
        n_l++;
    if (n_m==100) {
        QoS_loss=n_l/n_m;
        if(QoS_loss>QoS SLA) {
            // low QoS case
            increase_accuracy();
        } else if (QoS_loss < 0.9*QoS SLA) {
            // high QoS case
            decrease_accuracy();
        } else {
            ; // no change
        }
        Sample QoS=Saved_Sample_QoS;
    }
}
```

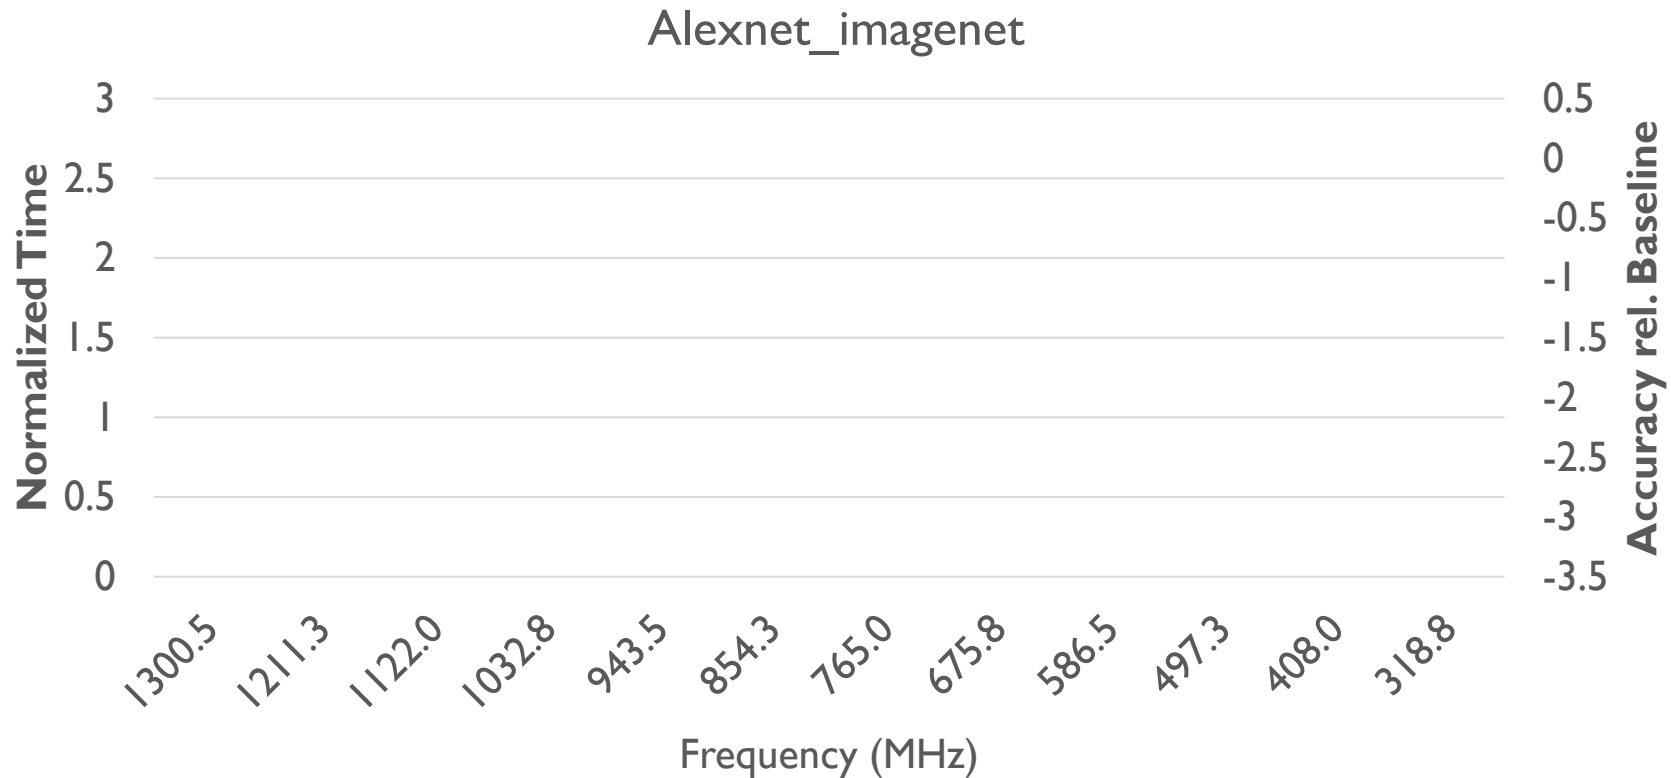
Figure 9. Customized QoS_ReCalibrate for Bing Search.

Runtime Adaptation for Accuracy

Key concerns:

- **Reexecuting infrequently to reduce the overhead**
checking every result is expensive, rely on spatial and temporal locality
- **The computation needs to be amenable for re-execution:**
think no side effects or crashes due to approximation

Runtime Approximation for Time/Energy

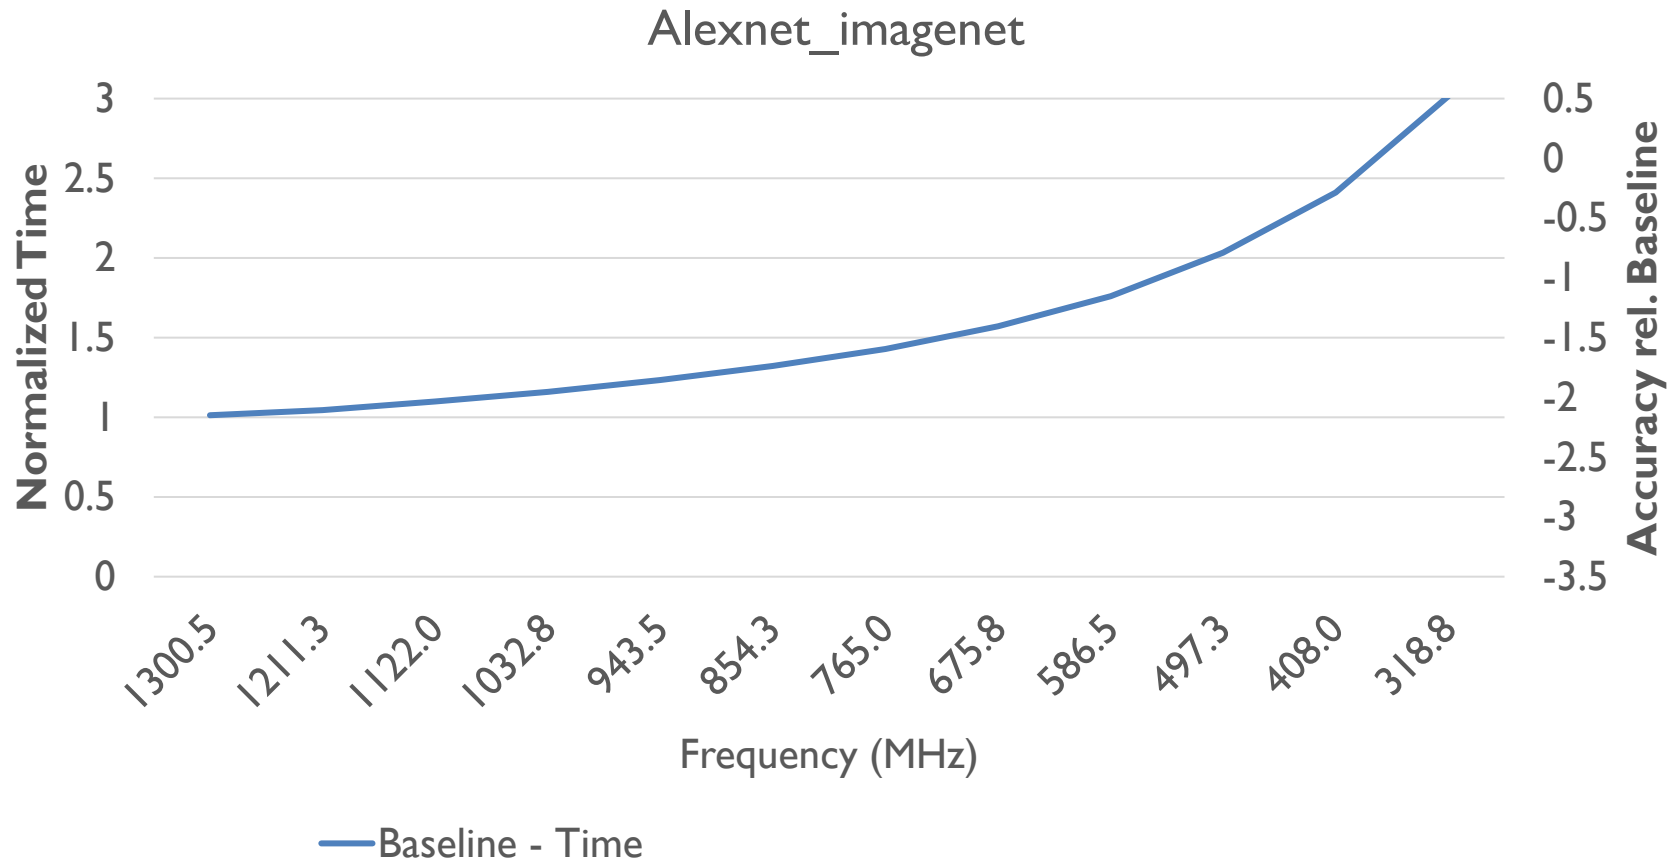


At regular time intervals we gradually reduce the frequency of the SoC



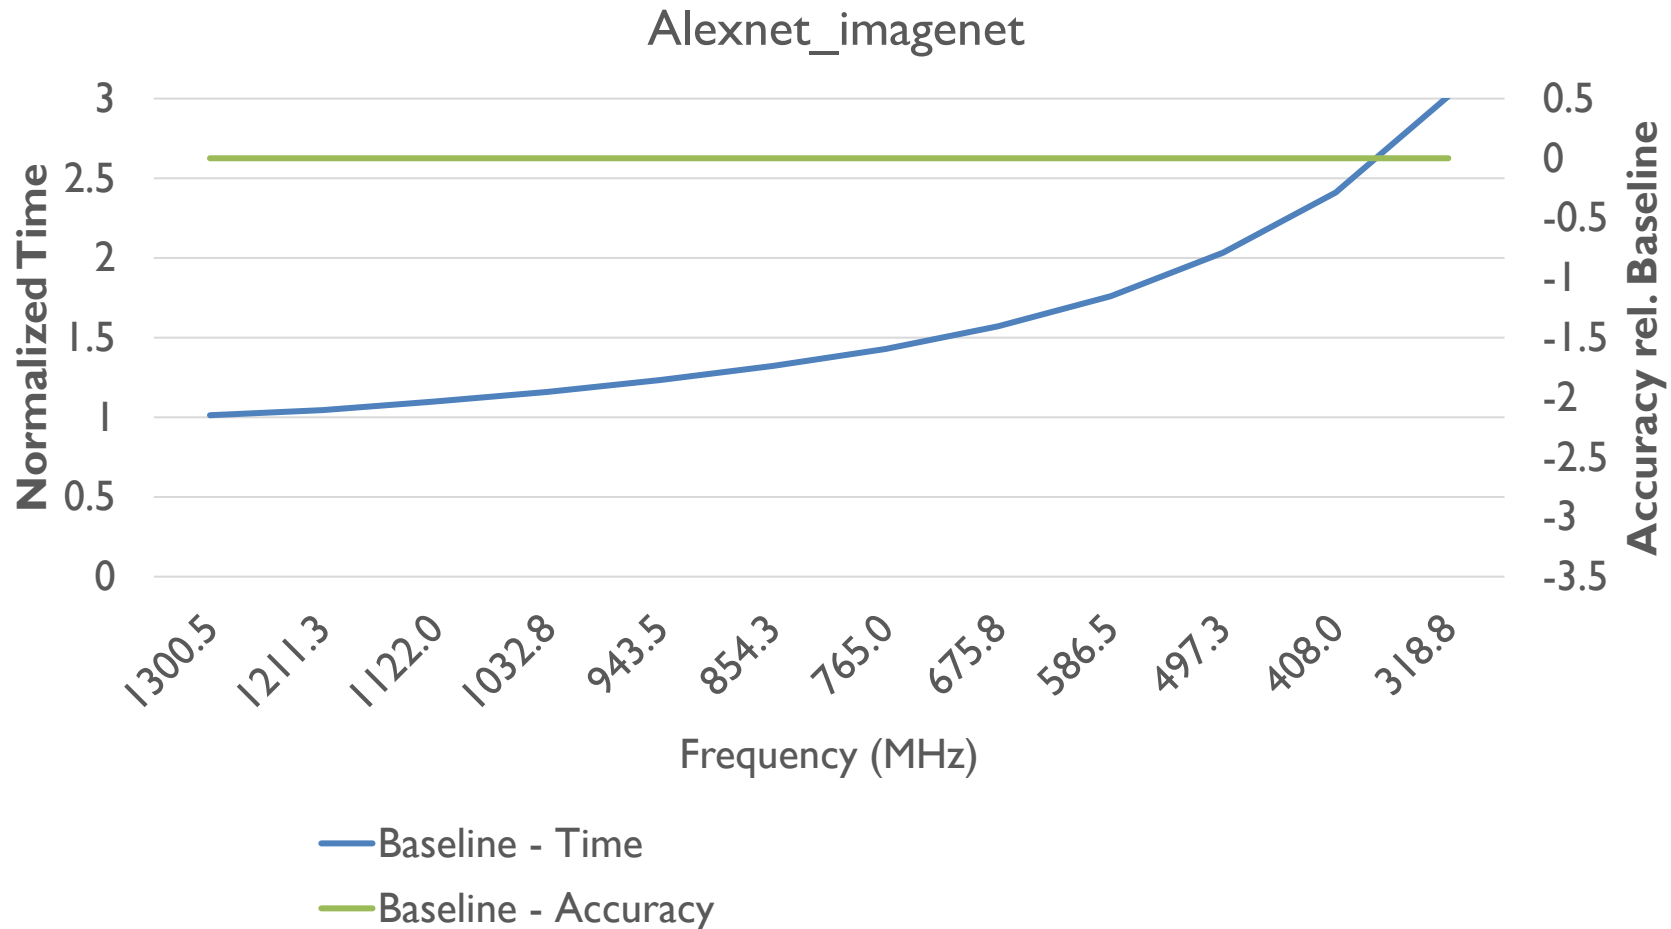
When you notice a disruption, read the value from the tradeoff curve that would negate the disruption

Runtime Approximation for Time/Energy



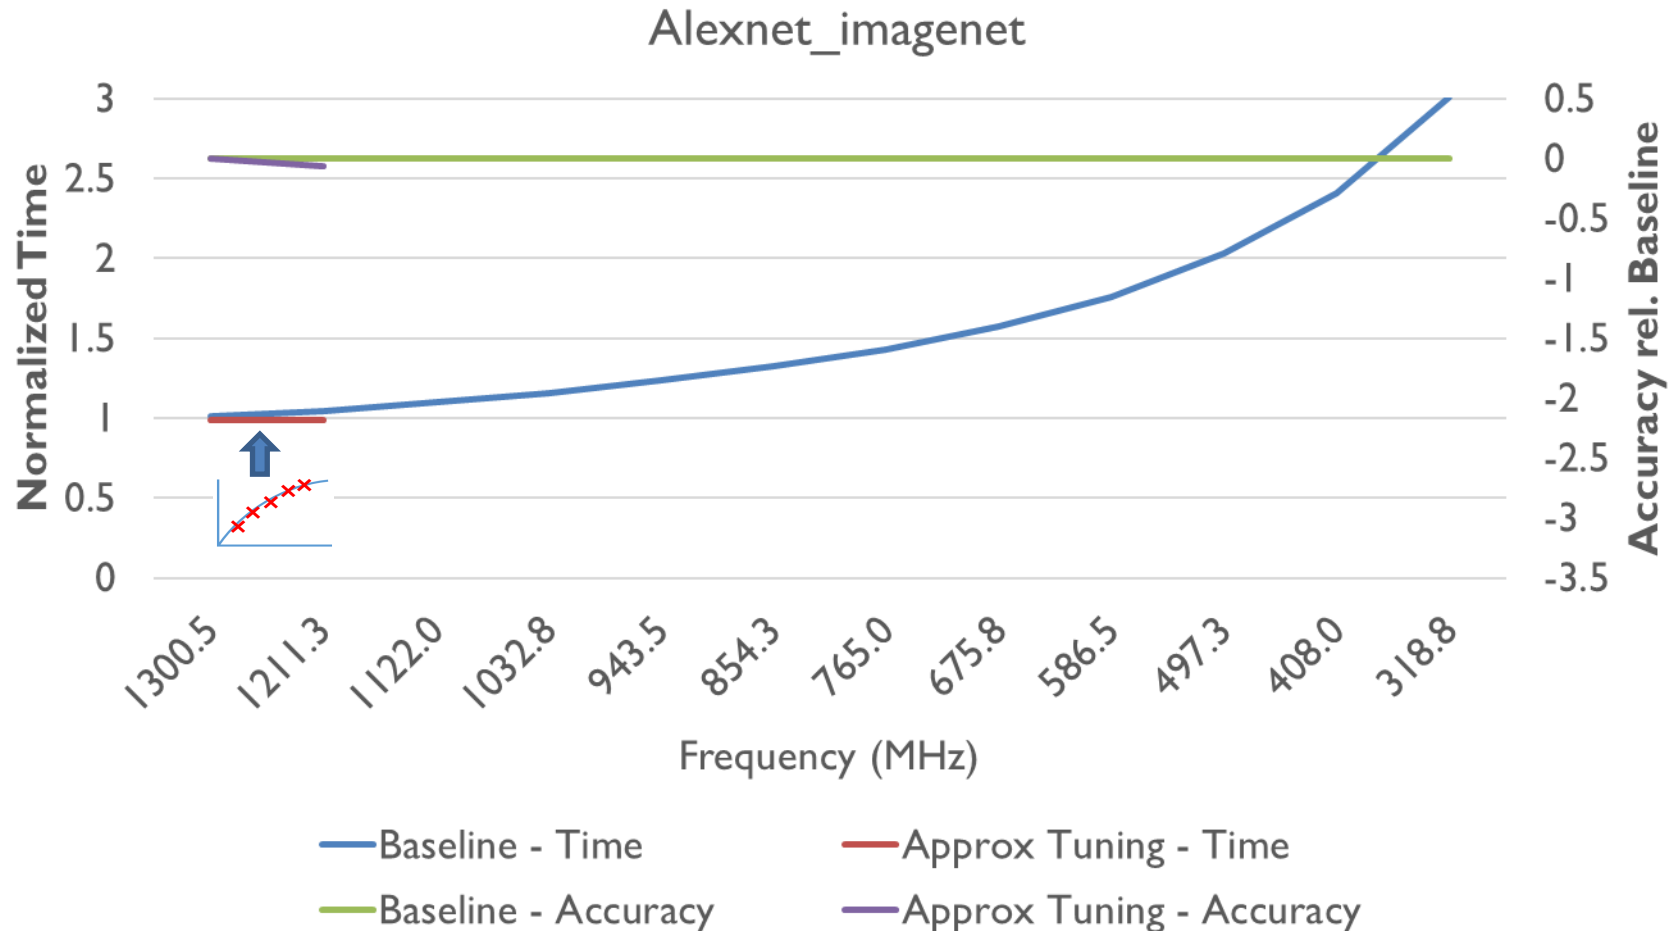
Runtime tuning helps maintain responsiveness in face of frequency changes

Runtime Approximation for Time/Energy



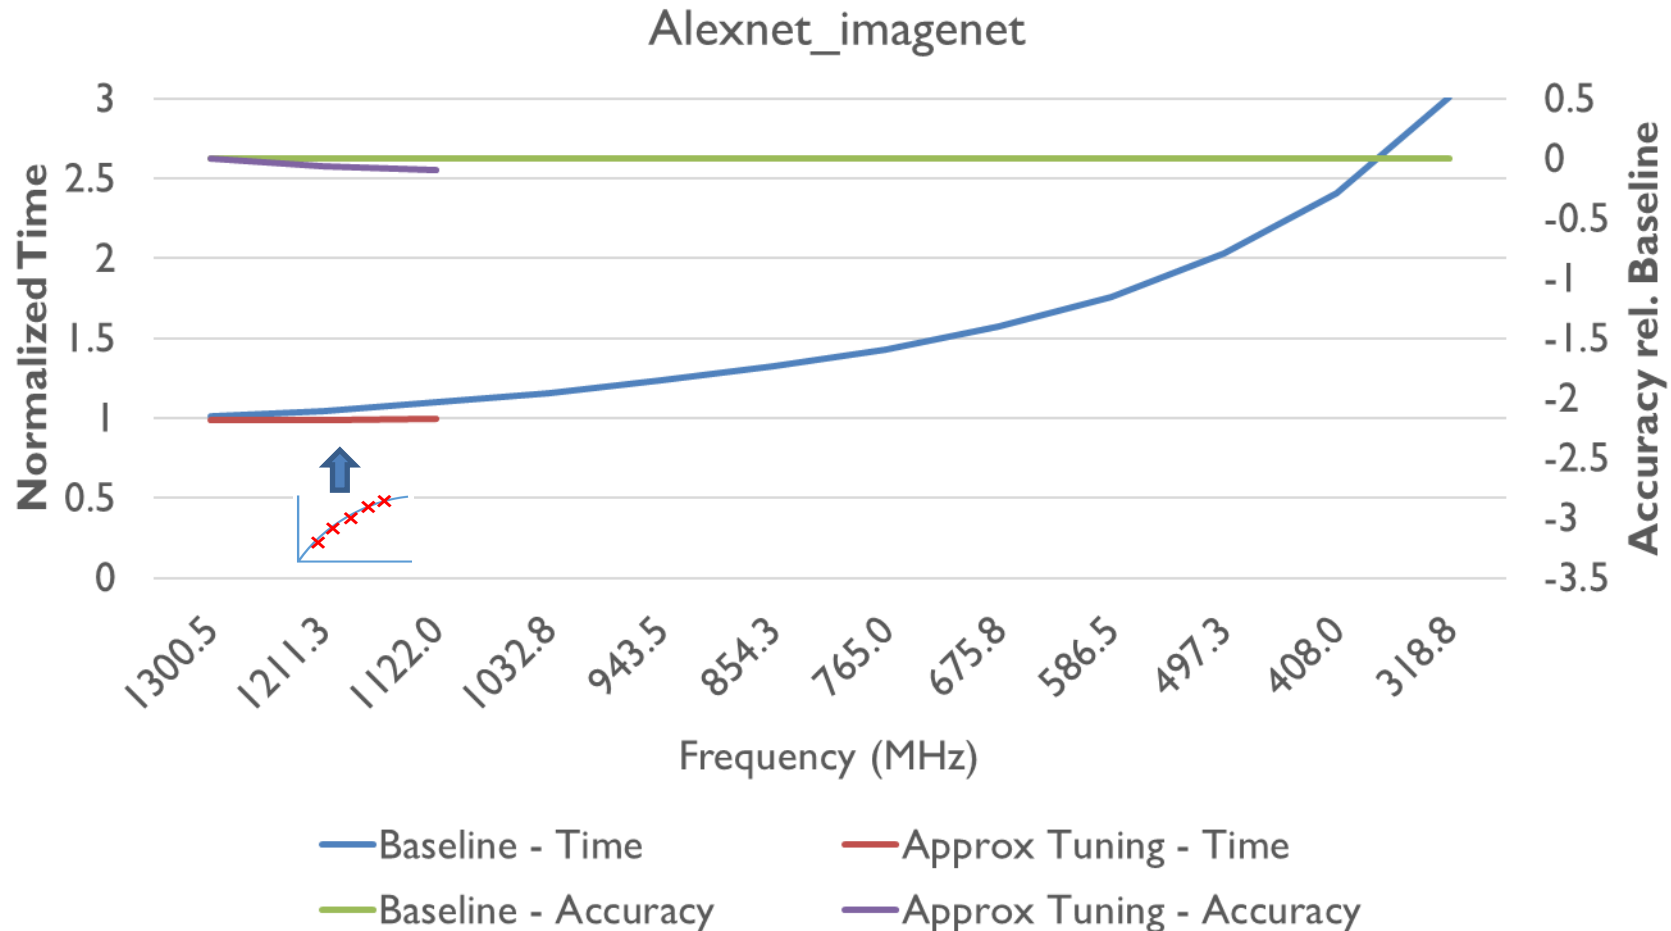
Runtime tuning helps maintain responsiveness in face of frequency changes

Runtime Approximation for Time/Energy



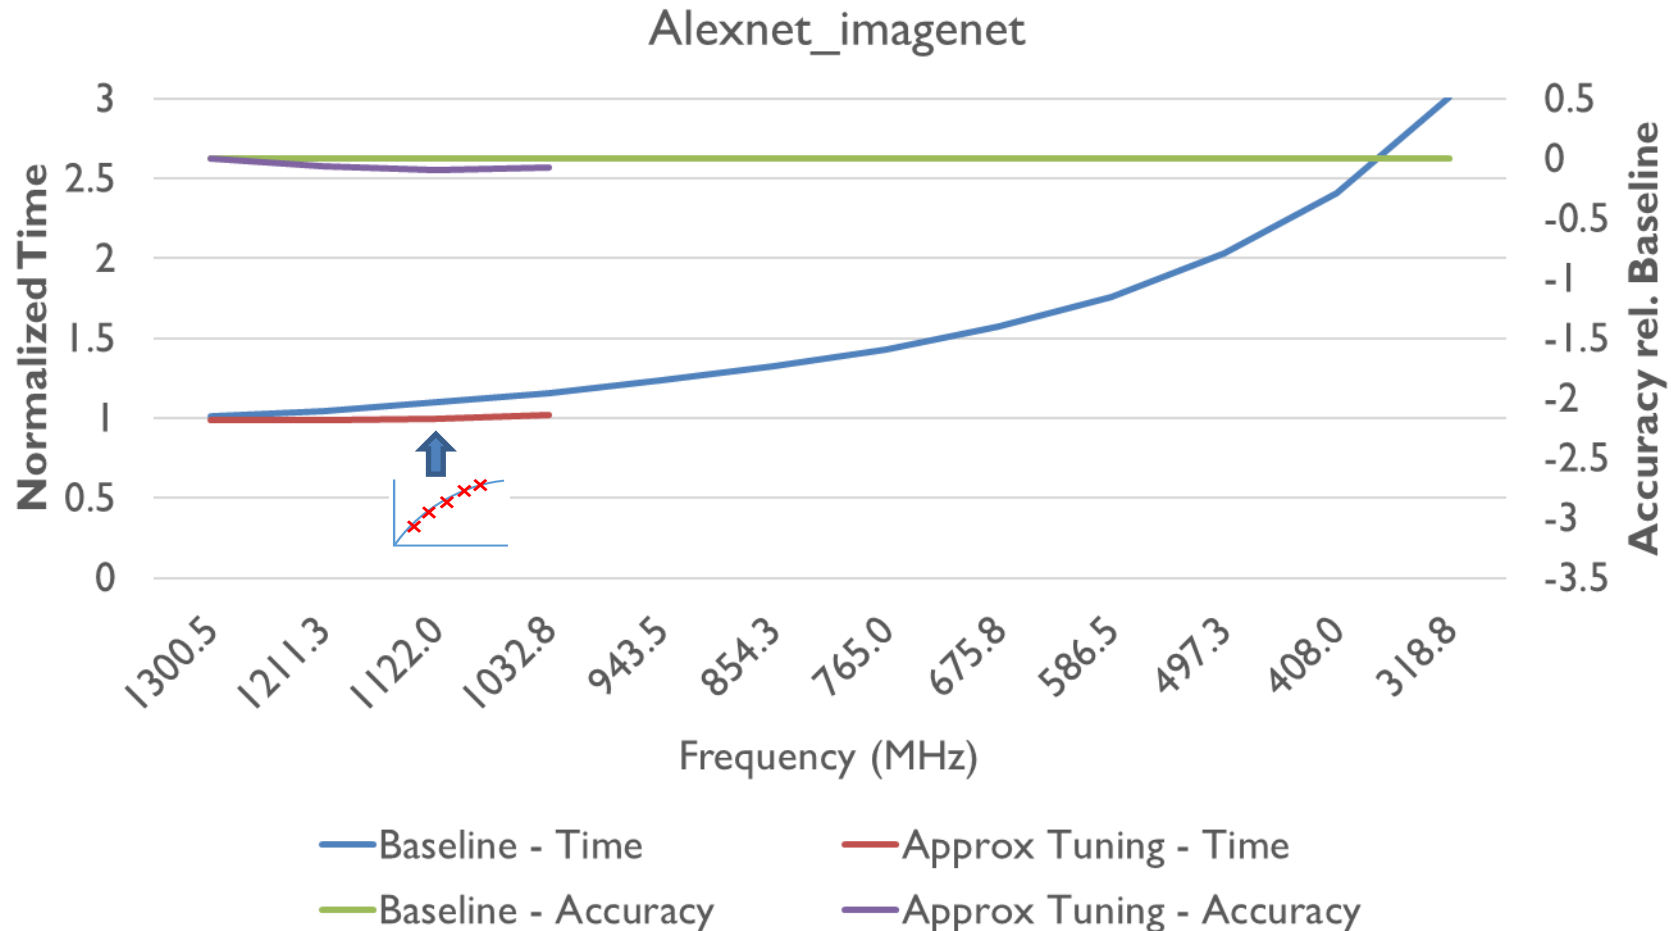
Runtime tuning helps maintain responsiveness in face of frequency changes

Runtime Approximation for Time/Energy



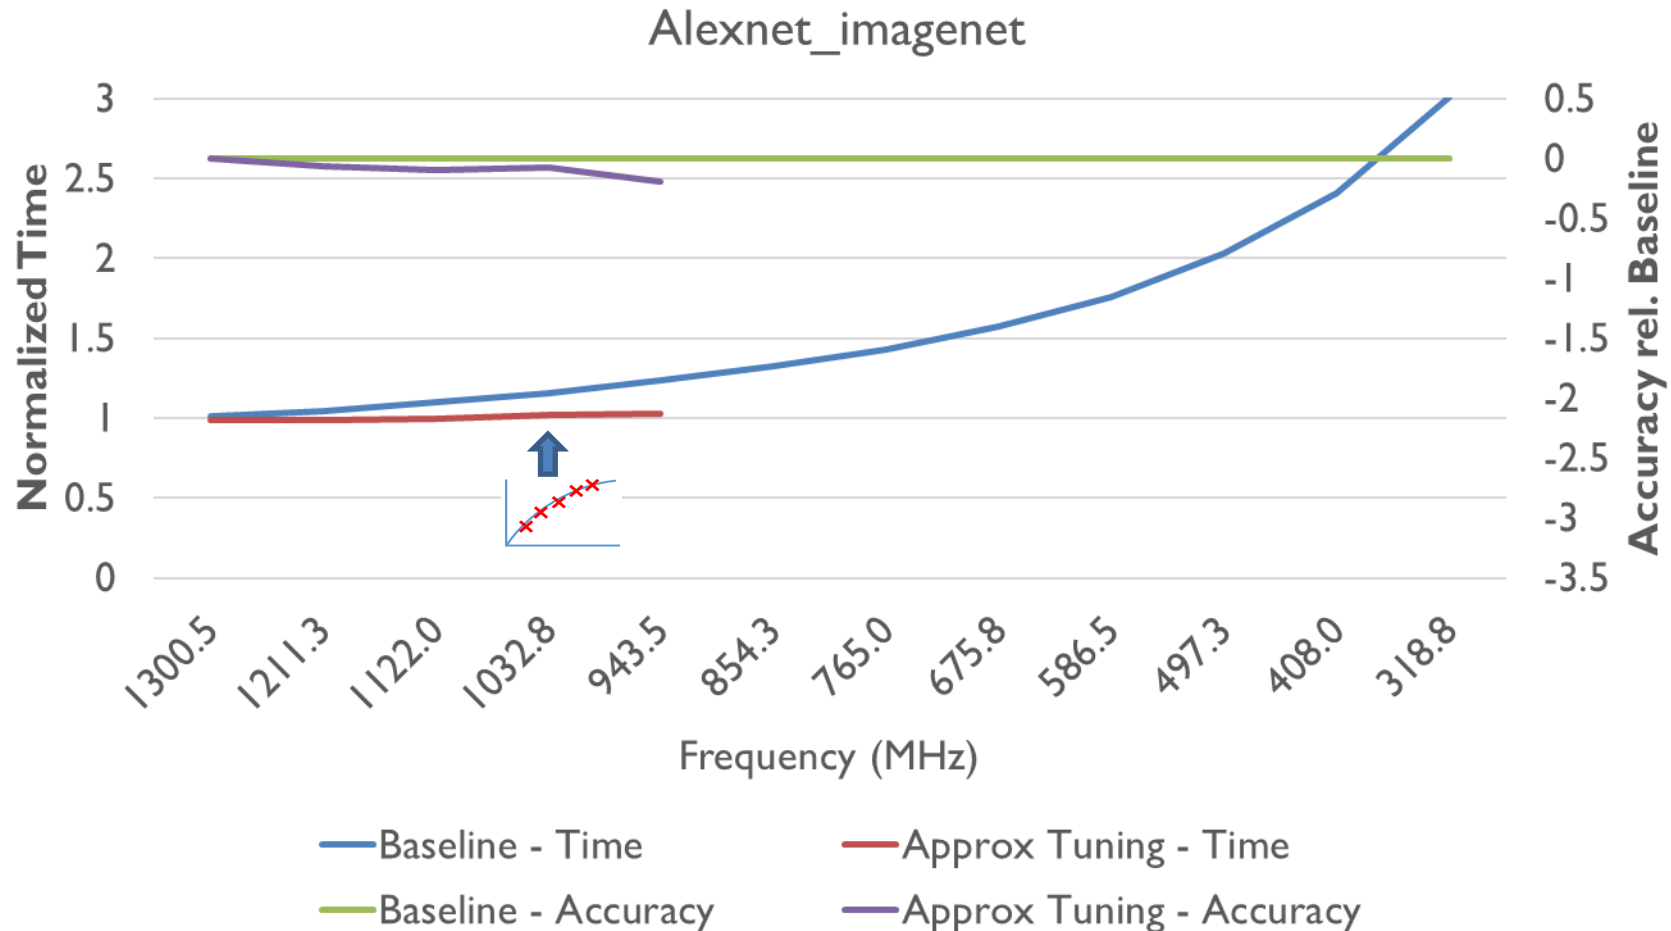
Runtime tuning helps maintain responsiveness in face of frequency changes

Runtime Approximation for Time/Energy



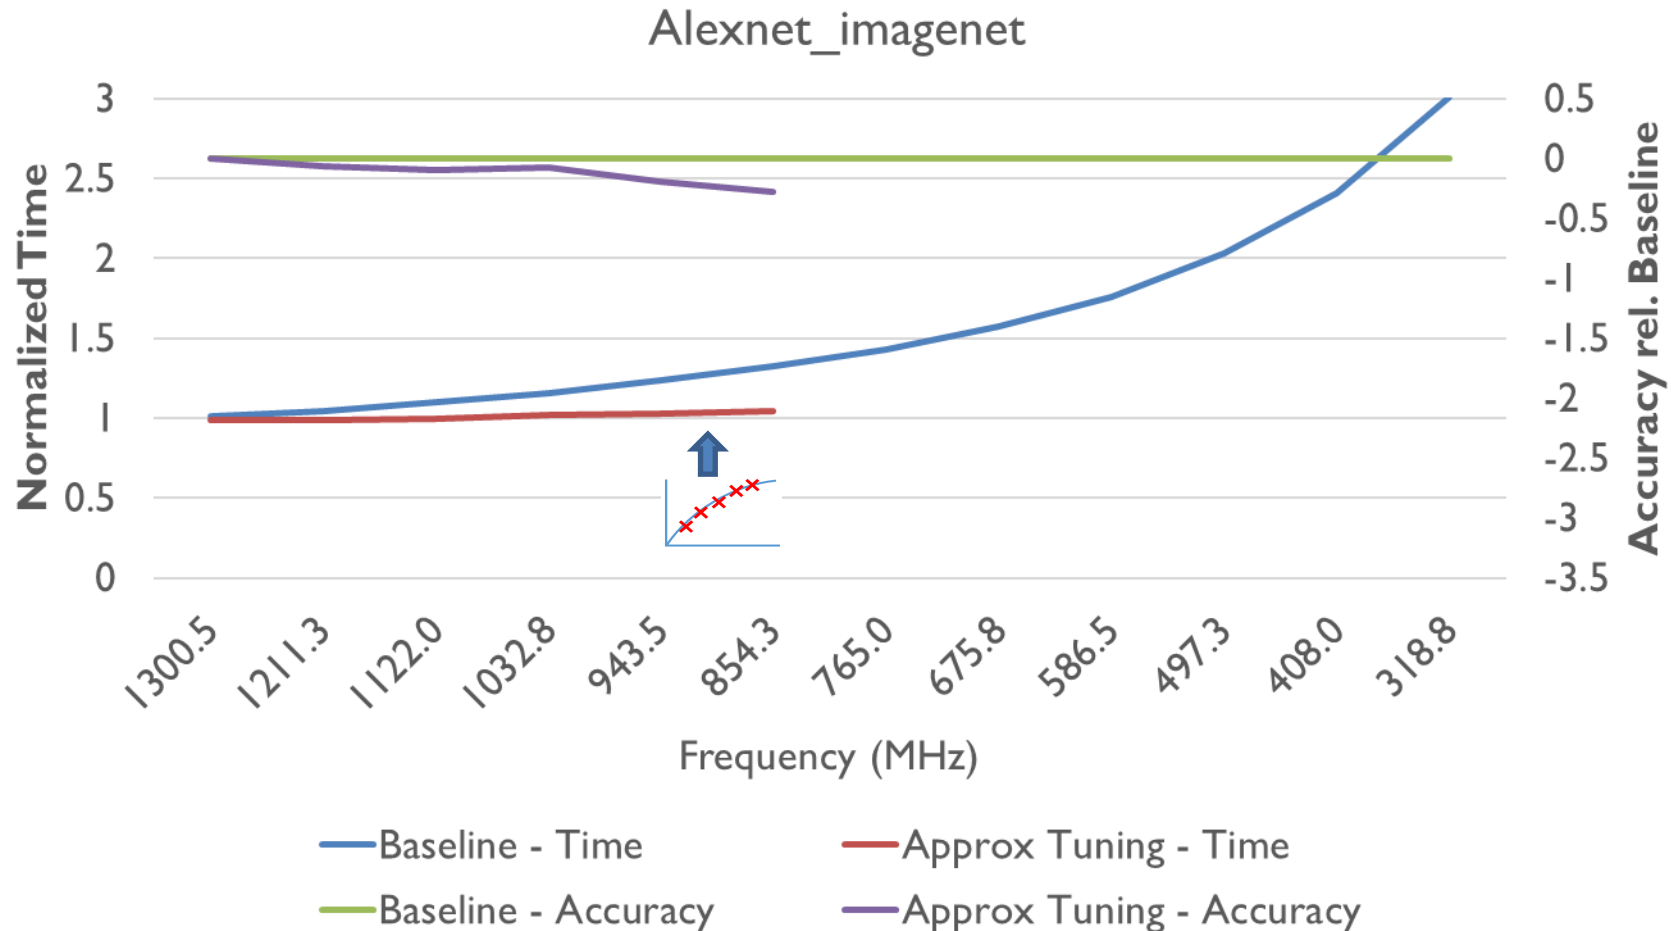
Runtime tuning helps maintain responsiveness in face of frequency changes

Runtime Approximation for Time/Energy



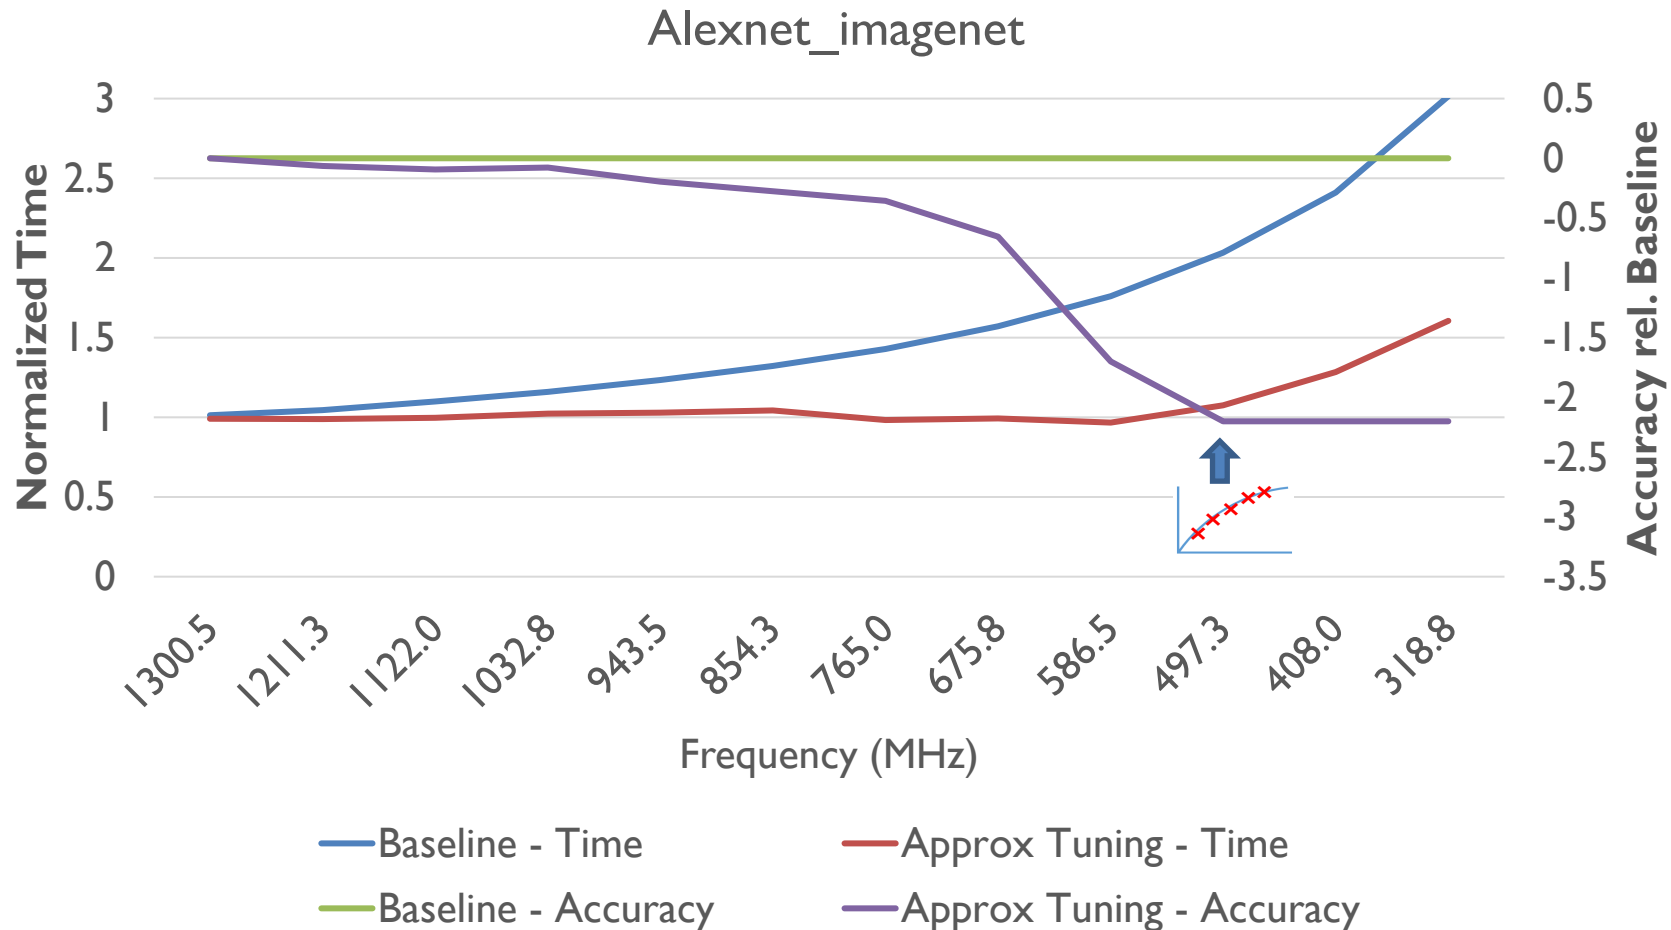
Runtime tuning helps maintain responsiveness in face of frequency changes

Runtime Approximation for Time/Energy



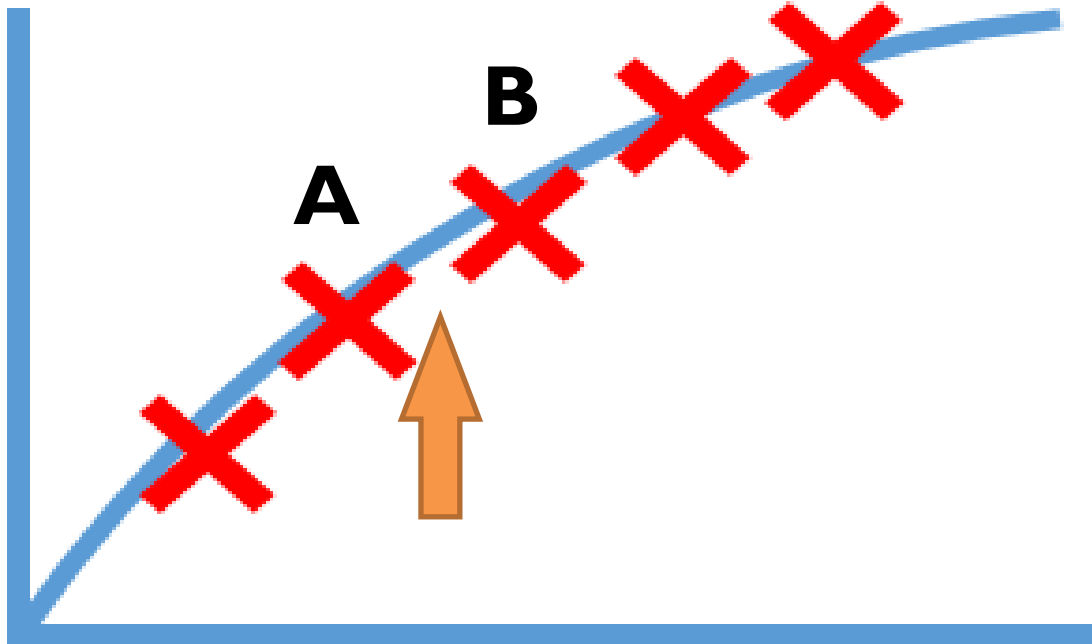
Runtime tuning helps maintain responsiveness in face of frequency changes

Runtime Approximation for Time/Energy



Runtime tuning helps maintain responsiveness in face of frequency changes

What if you don't have the exact point?



Solution 1: Select more conservative, suffer some performance drop (point A)

Solution 2: Select more aggressive, lose some more accuracy and make program even faster (point B)

Solution 3: We can use randomization

- Choose point A with probability p and
- Choose point B with probability $1-p$

Why would this work over a long sequence of runs?

For **all** inputs

Original
Computation

~~Typical
Inputs~~

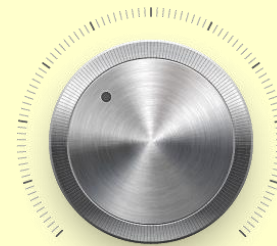
Accuracy
Requirement

Analysis-driven Compiler

SAS '11, POPL '12, OOPSLA '13, OOPSLA '14, OOPSLA '19, CGO '20

- **Statically analyze** computation's accuracy
- **Transform** computation by solving a mathematical optimization problem

Optimized Computation +



Recall: Our Optimization Problem

Select Program Configuration $X \in \text{Configs}$ to

maximize $(\text{Speedup}(X, i), \text{Accuracy}(X, i))$

forall $i \in \text{InputSet}$

But these are most often competing objectives.

Consider turning into weighted single optimization problem ($w_{1,2}$ express preference):

maximize $w_1 \times \text{Speedup}(X, i) + w_2 \times \text{Accuracy}(X, i)$

forall $i \in \text{InputSet}$

To maintain accuracy guarantees rephrase: for every accuracy loss threshold δ

maximize $\text{Speedup}(X, i)$

subject to $\text{AccuracyLoss}(X, i) \leq \delta$

forall $i \in \text{InputSet}$

Let's Extend It a Bit:

maximize $Speedup(X, i)$

subject to $AccuracyLoss(X, i) \leq \delta$

$SafetyPredicate(X, i) = True$

forall $i \in InputSet$

Let's Extend It a Bit:

maximize $Speedup(X, i)$

subject to $AccuracyLoss(X, i) \leq \delta$

$SafetyPredicate(X) = True$

forall $i \in InputSet$

Approximate Program Safety:

Information-flow Type Systems

Relational Logic Reasoning

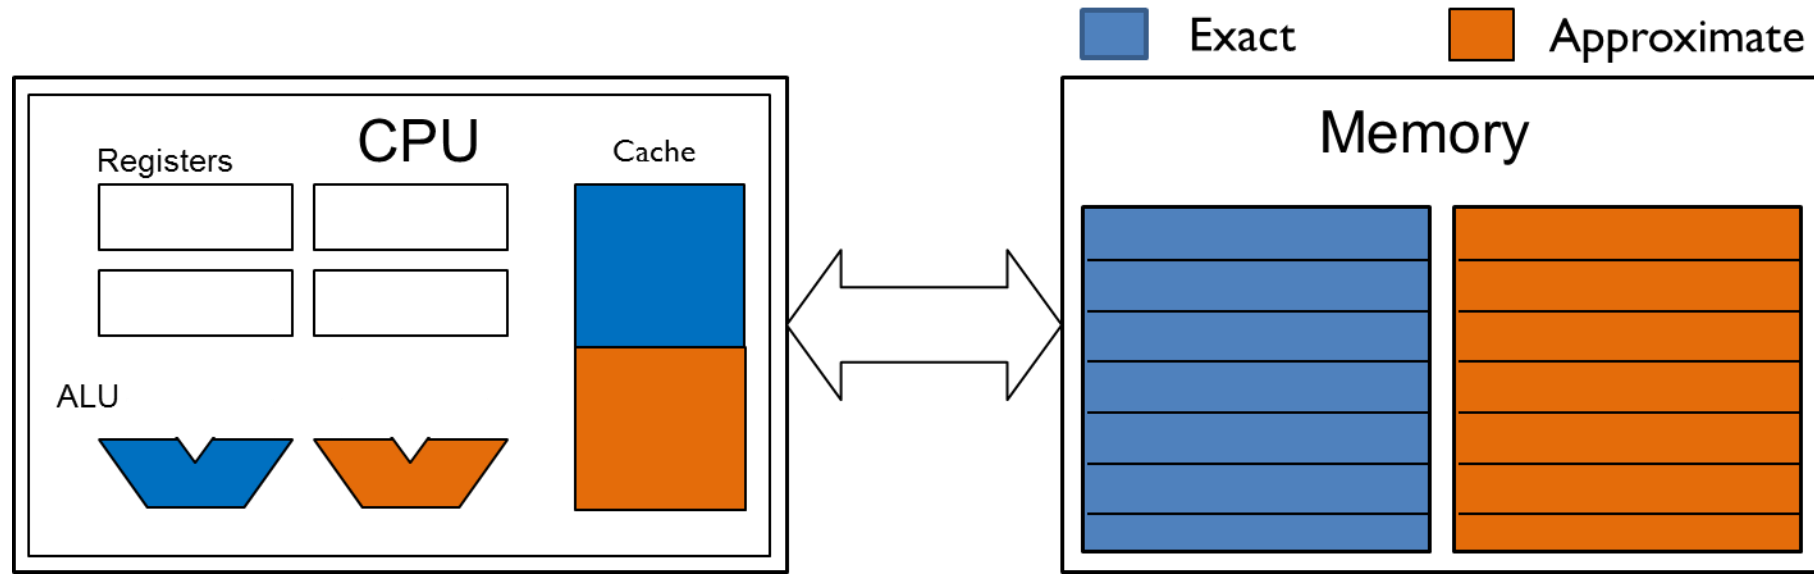
EnerJ Type System

Idea:

Isolate code and data that **must be precise**
from those that **can be approximated**

Sampson, Dietl, Fortuna, Gnanapragasam, Ceze, Grossman
EnerJ: Approximate Data Types for Safe and General Low-Power Computation
(PLDI 2011)

Approximate Hardware Model from EnerJ



```
reliability spec {  
  operator (+.) = 1 - 10^-7;  
  operator (-.) = 1 - 10^-7;  
  operator (*.) = 1 - 10^-7;  
  operator (<.) = 1 - 10^-7;  
  memory rel {rd = 1, wr = 1};  
  memory urel {rd = 1 - 10^-7, wr = 1};  
}
```

Recall – hardware approximations:

- Soft errors
- Timing errors
- Voltage variations
- Aging, Refresh rates, ...

All can be modeled as wrong bits
(permanent or transient)

EnerJ Type System

Idea:

Isolate code and data that **must be precise**
from those that **can be approximated**

Variable annotations (extends Java annotation system)

@Approx int a = approximate_code();

int p;

p = a; <----- not ok

EnerJ Example

```
@Approximable class FloatSet {  
    @Context float[] nums = ...; // can be approximate or exact  
    float mean() {  
        float total = 0.0f;  
        for (int i = 0; i < nums.length; ++i) total += nums[i];  
        return total / nums.length; }  
}
```

```
@Approx float mean APPROX() {  
    @Approx float total = 0.0f;  
    for (int i = 0; i < nums.length; i += 2) total += nums[i];  
    return 2 * total / nums.length; }  
}
```

Use:

```
@approx FloatSet afs
```

```
@precise FloatSet fs
```

```
afs.mean() // call approx
```

```
fs.mean() // call precise
```

Formulate the Rules

Types: @approx, @precise,

@context (the precision of the type depends on the precision of the enclosing object),

@top (either of approx. or precise), @lost (information about type lost)

Subtyping. Subtyping is defined using an ordering of the precision qualifiers and subclassing.

The following rules define the ordering of precision qualifiers:

$q <:_q q'$ ordering of precision qualifiers

$$\frac{q \neq \text{top}}{q <:_q \text{lost}}$$

$$\frac{}{q <:_q \text{top}}$$

$$\frac{}{q <:_q q}$$

Recall that `top` qualifies the common supertype of `precise` and `approx` types. Every qualifier other than `top` is below `lost`; every qualifier is below `top`; and the relation is reflexive. Note that the `precise` and `approx` qualifiers are not related.

EnerJ Type System

Idea:

Isolate code and data that **must be precise** from those that **can be approximated**

```
@Approx FloatSet afs;
```

```
@Precise FloatSet fs;
```

```
@Approx int a= fs.mean(); // call approx
```

```
@Precise int p = fs.mean(); // call precise
```

```
print (a);
```

```
print (p);
```

EnerJ Type System

Idea:

Isolate code and data that **must be precise** from those that **can be approximated**

```
@Approx int a = approximate_code();
```

```
int p;
```

```
if (a > 3) { p = 1; } else { p = 2; }
```

Control flow dependency (implicit flow)



Typing Rule

$$\frac{{}^s\Gamma \vdash e_0 : \text{precise } P \quad {}^s\Gamma \vdash e_1 : T \quad {}^s\Gamma \vdash e_2 : T}{{}^s\Gamma \vdash \text{if}(e_0) \{e_1\} \text{ else } \{e_2\} : T}$$

Approximate data cannot change the control flow

EnerJ Type System

Idea:

Isolate code and data that **must be precise**
from those that **can be approximated**

```
@Approx int a = approximate_code();
```

```
int p;
```

```
p = endorse(a); <----- ok
```

Like “(cast_type) a” in Java

EnerJ Type System

Consequence:

Then the approximate parts may be optimized automatically, but the developer needs to **ensure the endorsed values are valid.**

```
@Approx int a = approximate_code();
```

```
int p;
```

```
p = endorse(a); <----- ok
```

```
if ( isValid(p) ) { ... } else { errorHandle(a) }
```

EnerJ Type System

Consequence:

Then the approximate parts may be optimized automatically, but the developer needs to **ensure the endorsed values are valid.**

```
@Approx int a = approximate_code();  
int p;  
try {  
    p = endorse(a);    <----- ok  
    check ( isValid(p) ) { ... } // do the computation  
    recover { errorHandle(a) }
```

EnerJ Type System

Motivation:

Security information flow type systems – prevent the program from leaking information about **private** variables into **public** variables.

See: [https://en.wikipedia.org/wiki/Information_flow_\(information_theory\)#Information_flow_control](https://en.wikipedia.org/wiki/Information_flow_(information_theory)#Information_flow_control)

Noninterference [Goguen and Meseguer 1982]:

“one group of users, using a certain set of commands is **noninterfering** with another group of users if the first group does with those commands can no effect on what the second group of users can see.”

General Verification?

For safety properties or worst-case accuracy properties

Two (not so simple steps):

1. Verify the original program wrt the property
2. Verify the approximate program wrt the property

But: how often can you do the step 1?

Relative Safety

If the original program **satisfies all assertions**,
then the relaxed program satisfies all assertions

Relative Safety vs. Just Safety

Established **through any means:**
verification, testing, code review



If the original program **satisfies all assertions**,
then the relaxed program satisfies all assertions

**Any inconsistent behavior must be
in the original program!**

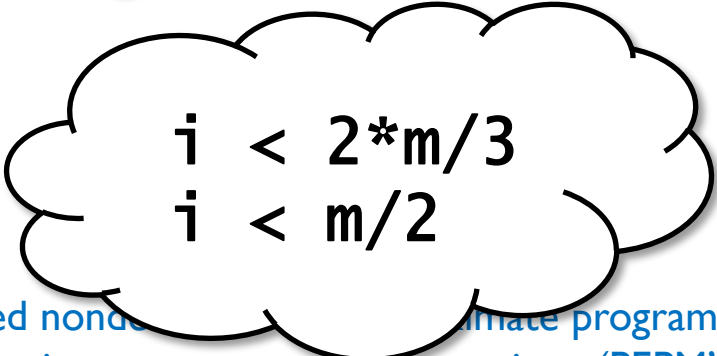
How does it work?

1. Define operational and axiomatic semantics of a product program (it has the behaviors of the original and approximated programs)
2. Specify the approximation as a perturbation of the variables
3. Capture when the flow of control of the approximate program diverges from the original program
4. Capture when the flow of control of the approximate program converges back with the original program
5. The predicate to verify is on the variables that are from both programs

Relational Safety Verification

```
for ((i=0; i < m; i++) {  
    sum = sum + x[i]  
}
```

```
avg = sum / m
```



$i < 2*m/3$
 $i < m/2$

Relational Safety Verification

```
relax (m) st ( $0 < m \leq \text{old}(m)$ )
```

```
for (i=0; i < m; i++) {  
    sum = sum + x[i]  
}
```


```
avg = sum / m
```

Relational Safety Verification

```
relax (m) st (0 < m <= old(m))
```

```
for (i=0; i < m; i++) {  
    sum = sum + x[i]  
}
```

```
avg = sum / m
```



Transformed execution accesses only (a subset of) memory locations that the original execution would have accessed

Relational Safety Verification

```
relax (m) st (0 < m <= old(m))
```

```
for (i=0; i < m; i++) {  
    sum = sum + x[i]  
}
```

```
avg = sum / m
```

The difference between the variable in the original and approximate runs is at most δ

$$|sum\langle o \rangle - sum\langle r \rangle| \leq \delta$$

Control-Flow Divergence

```
// when b is possibly approximate
if (b) {
    x = y - d
} else {
    x = y + d
}
```

```
relate ( |x<o> - x<r>| <= ?? )
```

Need to consider all four options: (1) original then – relaxed then; (2) original then – relaxed else (3) original else – relaxed then (3) original else – relaxed else

Control-Flow Divergence

```
// when b is possibly approximate
if (b) {
    x = y - d
} else {
    x = y + d
}
```

$x_{<o>} = y - d$	$x_{<o>} = y - d$
$x_{<r>} = y - d$	$x_{<r>} = y + d$
-----	-----
$x_{<o>} = y + d$	$x_{<o>} = y + d$
$x_{<r>} = y + d$	$x_{<r>} = y - d$

```
relate ( |x<o> - x<r>| <= ?? )
```

Need to consider all four options: (1) original then – relaxed then; (2) original then – relaxed else (3) original else – relaxed then (3) original else – relaxed else

Control-Flow Divergence

```
// when b is possibly approximate
if (b) {
    x = y - d
} else {
    x = y + d
}
```

$x_{<o>} = y - d$	$x_{<o>} = y - d$
$x_{<r>} = y - d$	$x_{<r>} = y + d$
-----	-----
$x_{<o>} = y + d$	$x_{<o>} = y + d$
$x_{<r>} = y + d$	$x_{<r>} = y - d$

```
relate ( |x<o> - x<r>| <= 2*d )
```

Need to consider all four options: (1) original then – relaxed then; (2) original then – relaxed else (3) original else – relaxed then (3) original else – relaxed else

Recall: Accuracy and Guarantees

Logic-Based (*worst-case*)

“for all inputs...”

Probabilistic (*worst-case or average-case*)

“for all inputs, with probability at least p ...”

“for inputs distributed as...”

Statistical (*average-case*)

“for inputs distributed as... with confidence c ”

“for tested inputs... with confidence c ”

Empirical (*typical-case*)

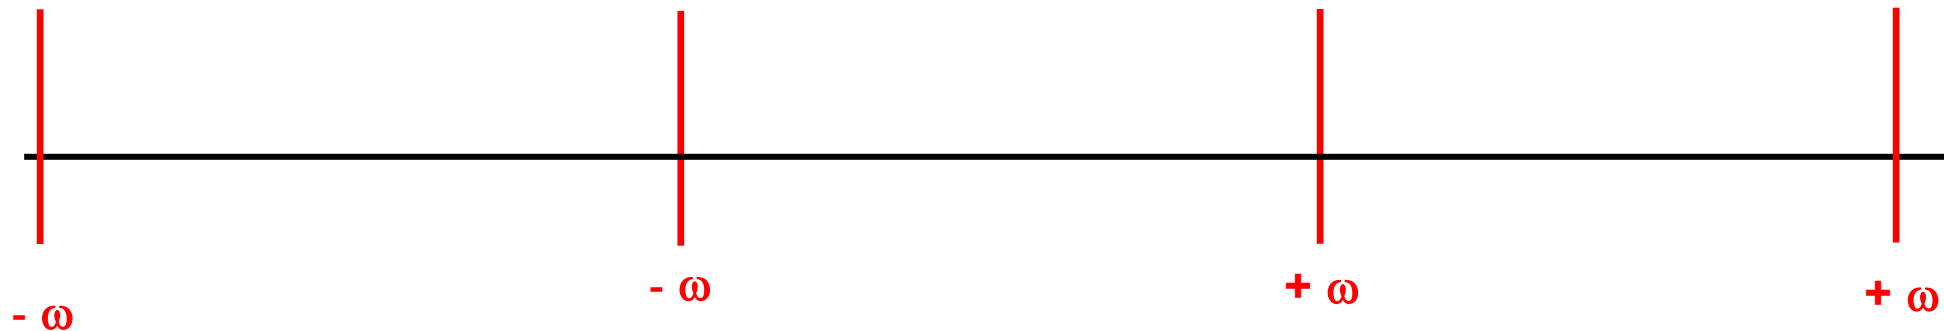
“for typical inputs...”

Computation difference

$$D = \text{res}_{\text{original}} - \text{res}_{\text{transformed}}$$

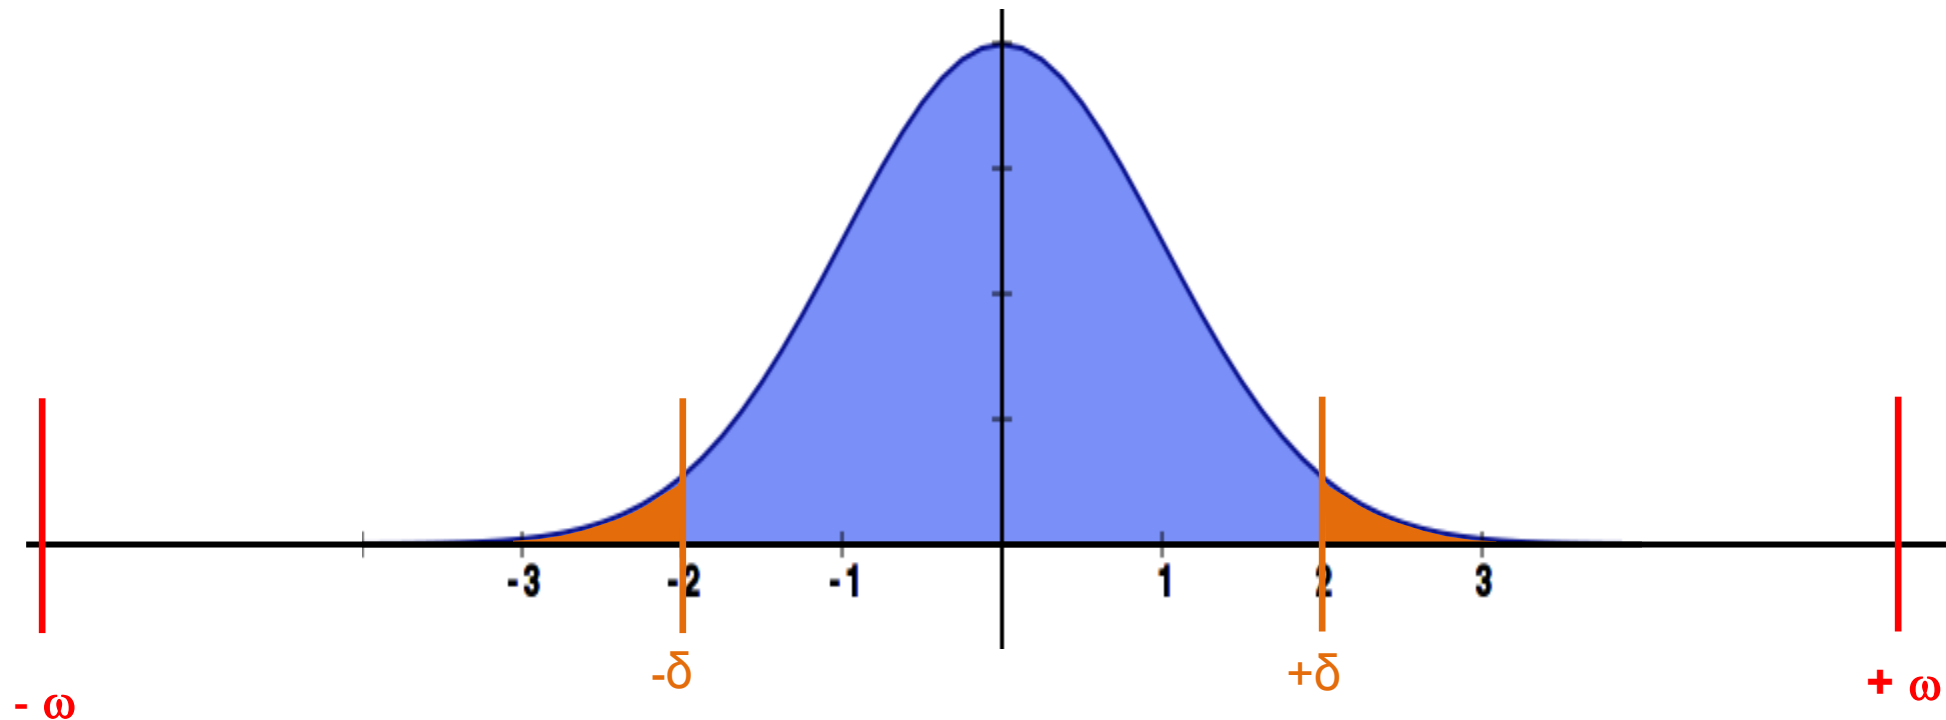
Worst-Case Analysis

$$D = \text{res}_{\text{original}} - \text{res}_{\text{transformed}}$$



Alternative: Probabilistic Analysis

$$D = \text{res}_{\text{original}} - \text{res}_{\text{transformed}}$$



Analysis-Based Optimizations

Accuracy Specification

Reliability Function computes result correctly
with probability > 0.99

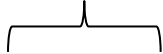
Absolute Error Absolute error of function's result < 2.0

**Reliability and
Absolute Error** Absolute error of function's result < 2.0
with probability > 0.99

```
int { $\Delta f \leq 2 ; 0.99 * R(\Delta x = 0, \Delta y = 0)$ } f(int x, int y);
```

Reliability Specification

Reliability
degradation


`int < 0.99 * R($\Delta x = 0$, $\Delta y = 0$, $\Delta src = 0$) >`
`interpolation(int x, int y, int src[][]);`

The function computes result correctly
with probability at least **0.99**

Reliability Specification

Reliability degradation Parameter Reliability

`int < 0.99 * R($\Delta x = 0, \Delta y = 0, \Delta src = 0$) >`
`interpolation(int x, int y, int src[][]);`

Probability that the parameters have correct values before function starts executing
(facilitates function composition)

Reliability Specification

Reliability degradation Parameter Reliability

`int < 0.99 * R($\Delta x = 0, \Delta y = 0, \Delta src = 0$) >`
`interpolation(int x, int y, int src[][]);`

- **Reliability factor:** $R(\Delta v_1 \leq d_1, \dots, \Delta v_n \leq d_n)$

$$\Delta v \equiv |v_{exact} - v_{approx}|$$

Numerical bound

Function Optimization Problem

Find Function Configuration q :

max EnergySavings (q)

s. t. Reliability (q) \geq ReliabilityBound

AbsoluteError (q) \leq ErrorBound

Analysis of the Function

Specifications

Image Scaling

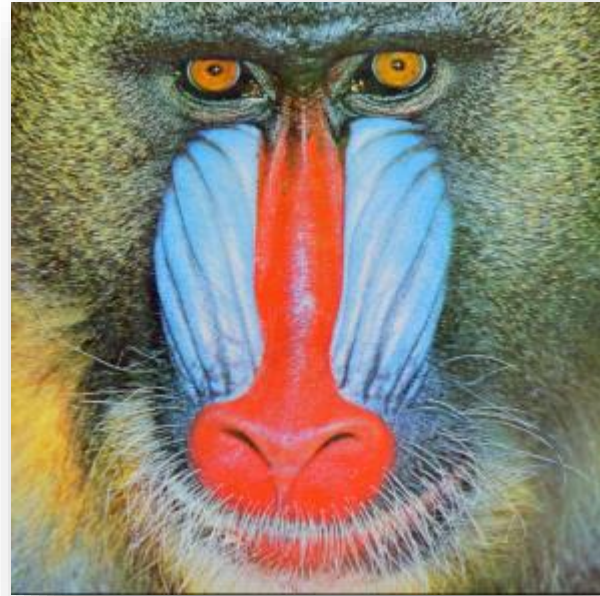
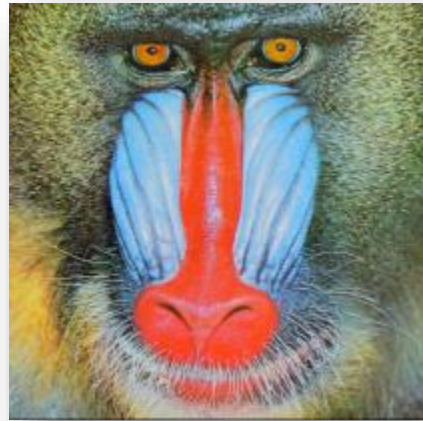
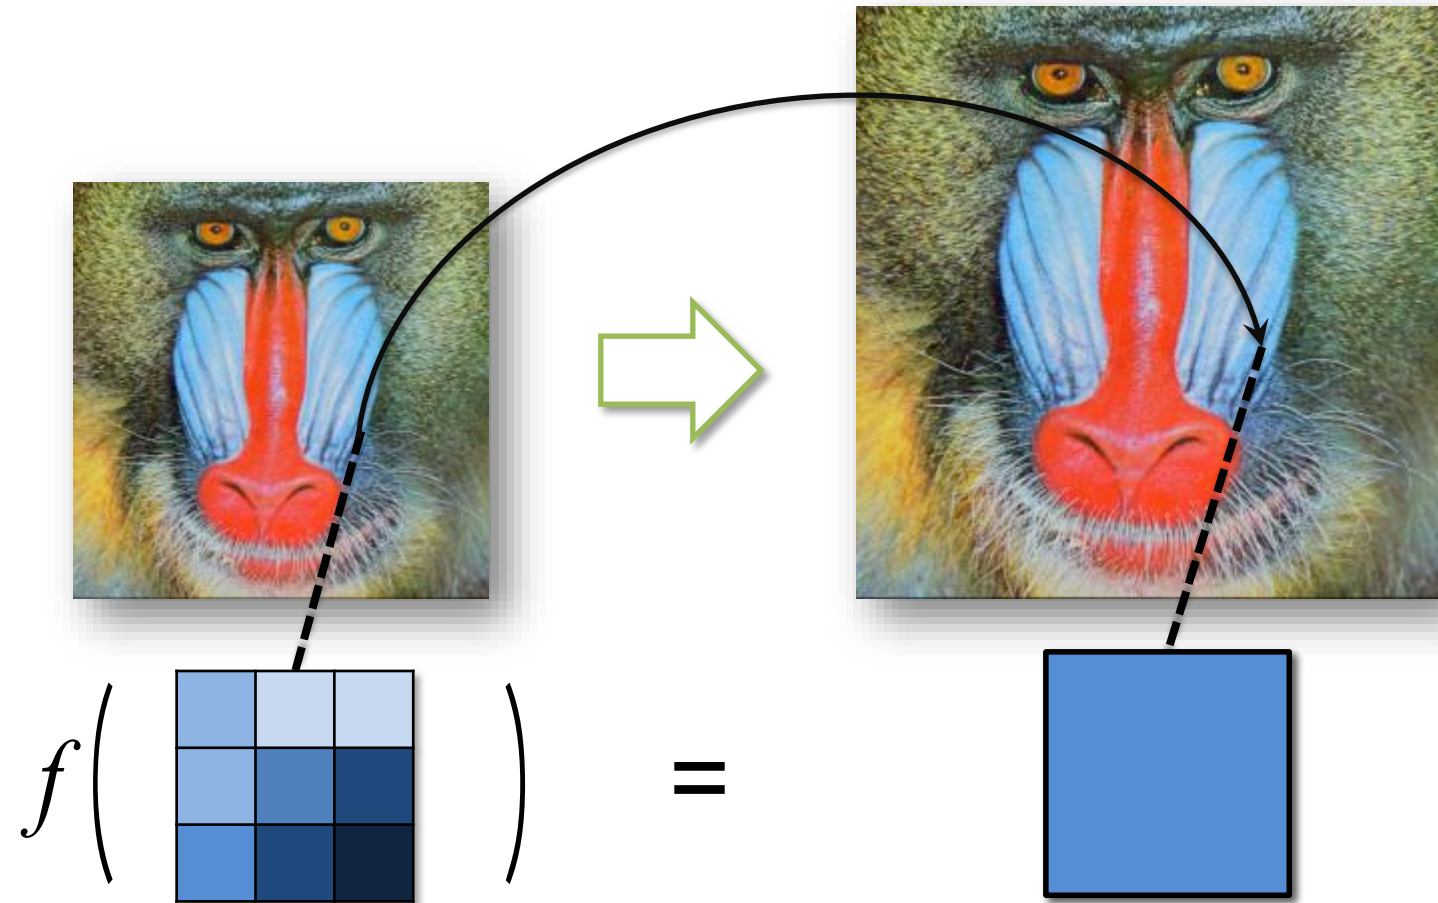


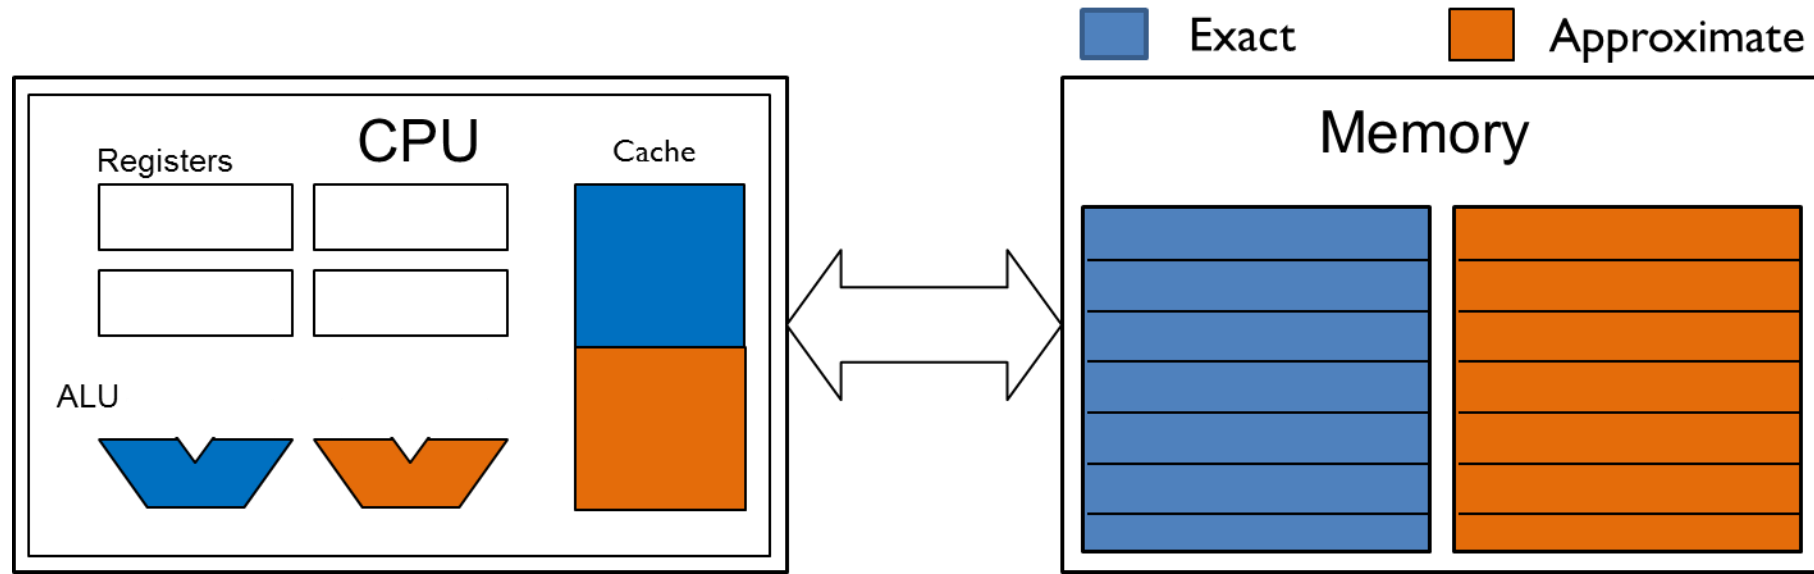
Image Scaling: Interpolation Function



Interpolation Function

```
int interpolation(int dst_x, int dst_y, int src[][])  
{  
    int x = src_location_x(dst_x, src),  
        y = src_location_y(dst_y, src);  
  
    int up    = src[y - 1][x],  
        down  = src[y + 1][x],  
        left  = src[y][x - 1],  
        right = src[y][x + 1];  
  
    int val = up + down + left + right;  
  
    return 0.25 * val;  
}
```

Approximate Hardware Model from EnerJ



```
reliability spec {  
  operator (+.) = 1 - 10^-7;  
  operator (-.) = 1 - 10^-7;  
  operator (*.) = 1 - 10^-7;  
  operator (<.) = 1 - 10^-7;  
  memory rel {rd = 1, wr = 1};  
  memory urel {rd = 1 - 10^-7, wr = 1};  
}
```

Recall – hardware approximations:

- Soft errors
- Timing errors
- Voltage variations
- Aging, Refresh rates, ...

All can be modeled as wrong bits
(permanent or transient)

Run Function on Approximate Hardware

```
int interpolation(int dst_x, int dst_y, int src[][])  
{  
    int x = src_location_x(dst_x, src),  
        y = src_location_y(dst_y, src);  
  
    int up    = src[y - 1][x],  
        down  = src[y + 1][x],  
        left  = src[y][x - 1],  
        right = src[y][x + 1];  
  
    int val = up + down + left + right;  
  
    return 0.25 * val;  
}
```

Run Function on Approximate Hardware

```
int interpolation(int@ dst_x, int@ dst_y, int@ src[][])  
{  
    int@ x = src_location_x(dst_x, src),  
        y = src_location_y(dst_y, src);  
  
    int@ up    = src[y -. 1][x],  
        down  = src[y +. 1][x],  
        left  = src[y][x -. 1],  
        right = src[y][x +. 1];  
  
    int@ val = up +. down +. left +. right;  
  
    return 0.25 *. val;  
}
```

Function Configuration

Binary vector $\mathbf{q} = (q_1, q_2, \dots, q_n)$

Variable Declarations:

- q_i - if 1, variable is stored in approximate memory
if 0, variable is stored in exact memory

Arithmetic Operations:

- q_i - if 1, the operation is approximate,
if 0, the operation is exact

Function Configuration

```
int interpolation(int dst_x, int dst_y, int src[][])  
{  
    int x = src_location_x(dst_x, src);  
    int y = src_location_y(dst_y, src);  
  
    int up    = src[y - 1][x];  
    int down  = src[y + 1][x];  
    int left  = src[y][x - 1];  
    int right = src[y][x + 1];  
  
    int val = up + down + left + right;  
  
    return 0.25 * val;  
}
```

Function Configuration

```
int interpolation(intqdstx dst_x, intqdsty dst_y, intqsrc src[][])  
{  
    intqx x = src_location_x(dst_x, src);  
    intqy y = src_location_y(dst_y, src);  
  
    intqup up = src[y - 1][x];  
    intqdown down = src[y + 1][x];  
    intqleft left = src[y][x - 1];  
    intqright right = src[y][x + 1];  
  
    intqval val = up + down + left + right;  
    return 0.25 * val;  
}
```

Each assignment of vector q denotes
a different approximate function

```
int interpolation(int $q_{dstx}$  dst_x, int $q_{dsty}$  dst_y, int $q_{src}$  src[][])  
{  
    int $q_x$  x = src_location_x(dst_x, src);  
    int $q_y$  y = src_location_y(dst_y, src);  
  
    int $q_{up}$  up = src[y -  $q_7$  1][x];  
    int $q_{down}$  down = src[y +  $q_6$  1][x];  
    int $q_{left}$  left = src[y][x -  $q_5$  1];  
    int $q_{right}$  right = src[y][x +  $q_4$  1];  
  
    int $q_{val}$  val = up +  $q_1$  down +  $q_2$  left +  $q_3$  right;  
  
    return 0.25 *  $q_0$  val;  
}
```

Reliability Analysis

Motivation

- Efficiently represent reliability of **all approximate** function versions
- Construct constraints to separate **those** function versions **that satisfy specification**

Reliability Analysis

Approximate hardware specification:

- Reliability of arithmetic operations: $r_{op} \in (0, 1]$
- Reliability of memory reads and writes: $r_{rd}, r_{wr} \in (0, 1]$

```
operator (*) = 0.9999;  
memory approx {rd = 0.99998, wr = 0.99999};
```

Analysis:

- Sound **static analysis**, operates backward
- Constructs **symbolic expressions** that characterize reliability of kernel's traces

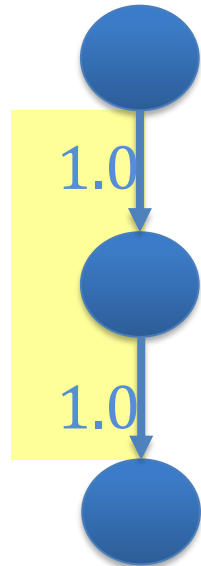
Reliability Analysis

Statement

`return val * 0.25;`

Exact Statement

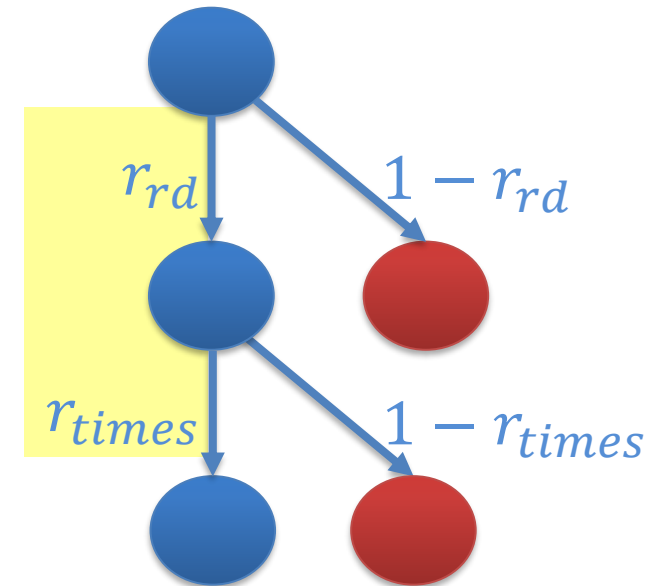
*val and * exact*



Approximate Statement

*val and * approximate*

Read val
Multiply
Return result



Reliability Analysis

Statement `return val * 0.25;`

Exact Statement

Approximate Statements

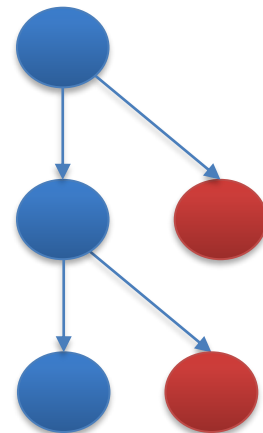
*val and *
exact*



Statement
reliability

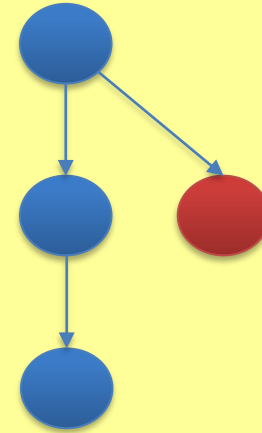
1.0

*val and *
approximate*



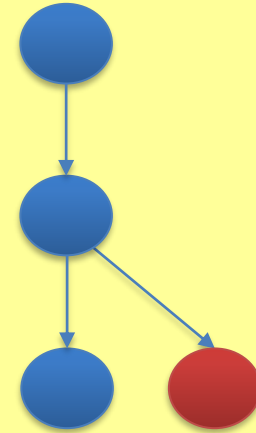
$r_{rd} \cdot r_{times}$

*val
approximate*



r_{rd}

**
approximate*

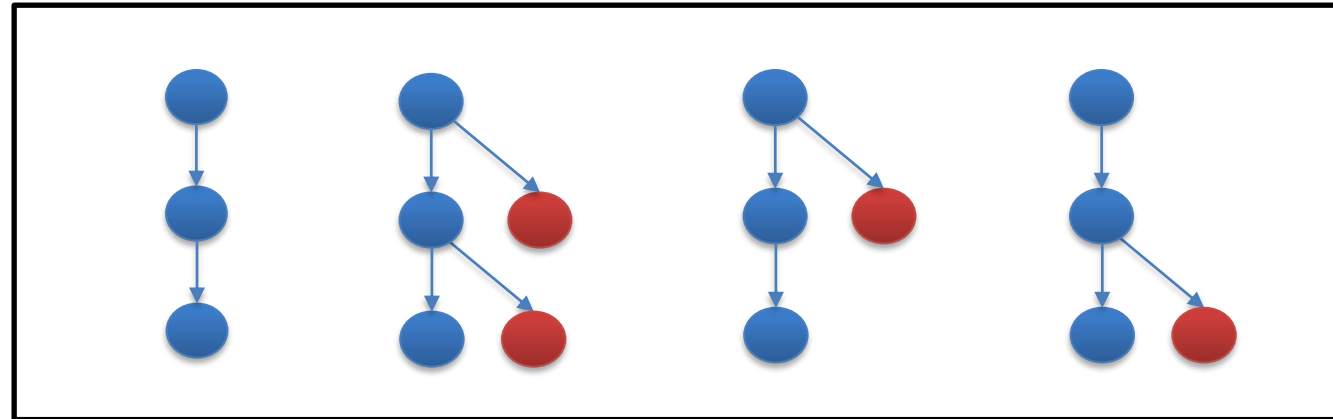


r_{times}

Reliability Analysis

Statement

```
return val * 0.25;
```



Encode approximation choice:

- Variable declaration: `int q_{val} val;`
- Multiplication: `val * q_* 0.25;`

Reliability Analysis

Statement

```
return val * 0.25;
```

**Reliability
Expression**

$$(r_{rd})^{q_{val}} \cdot (r_{times})^{q_*} \cdot R(\Delta val = 0)$$

Encode approximation choice:

- Variable declaration: `int q_{val} val;`
- Multiplication: `val * q_* 0.25;`

Reliability Analysis

Statement

```
return val * 0.25;
```

**Reliability
Expression**

$$(r_{rd})^{q_{val}} \cdot (r_{times})^{q_*} \cdot R(\Delta val = 0)$$

Reliability of reading `val` from
either exact or approximate memory:

$$(r_{rd})^0 = 1.0$$

$$(r_{rd})^1 = r_{rd}$$

Reliability Analysis

Statement

```
return val * 0.25;
```

**Reliability
Expression**

$$(r_{rd})^{q_{val}} \cdot (r_{times})^{q_*} \cdot R(\Delta val = 0)$$

Reliability of either exact or
approximate multiplication

Reliability Analysis

Statement

```
return val * 0.25;
```

**Reliability
Expression**

$$(r_{rd})^{q_{val}} \cdot (r_{times})^{q_*} \cdot \mathbf{R}(\Delta\mathbf{val} = \mathbf{0})$$

Probability that previous statements
computed **val** correctly



Interpolation Function

```
int interpolation(intqdstx dst_x, intqdsty dst_y, intqsrc src[][])  
{
```

$$(r_{rd})^{q_{val}} \cdot (r_{times})^{q_*} \cdot \mathbf{R}(\Delta val = 0)$$

```
return val *q* 0.25;  
}
```

Interpolation Function

```
int interpolation(intqdstx dst_x, intqdsty dst_y, intqsrc src[][])  
{
```

$$\begin{aligned} & (r_{rd})^{q_{val}} \cdot (r_{times})^{q_*} \cdot (r_{plus})^{q_1+q_2+q_3} \cdot (r_{rd})^{q_{up}+q_{down}+q_{left}+q_{right}} \\ & \cdot \mathbf{R(\Delta_{up} = 0, \Delta_{down} = 0, \Delta_{left} = 0, \Delta_{right} = 0)} \end{aligned}$$

```
intqval val = up + q1 down + q2 left + q3 right;
```

$$(r_{rd})^{q_{val}} \cdot (r_{times})^{q_*} \cdot \mathbf{R(\Delta_{val} = 0)}$$

```
return val * q* 0.25;
```

```
}
```

Reliability Expression

Function's Reliability Expression:

$$\begin{array}{l} \Delta \text{dst_x} = 0, \\ \Delta \text{dst_y} = 0, \\ \Delta \text{src} = 0 \end{array}$$

$$r_1^{q_1} \cdot r_2^{q_2} \cdot \dots \cdot r_n^{q_n} \cdot R \left(P_{param} \right)$$

**Probability operations
executed reliably
(for all approximate versions
of the function)**

**Probability parameters
have correct values
at function start**

Reliability Constraint

Relates developer's specification and analysis result:

$$r_{spec} \cdot R(P_{spec}) \leq r_1^{q_1} \cdot r_2^{q_2} \cdot \dots \cdot r_n^{q_n} \cdot R(P_{param})$$

Reliability Constraint

$$r_{spec} \leq r_1^{q_1} \cdot r_2^{q_2} \cdot \dots \cdot r_n^{q_n}$$

and

$$\underbrace{R(P_{spec}) \leq R(P_{param})}$$

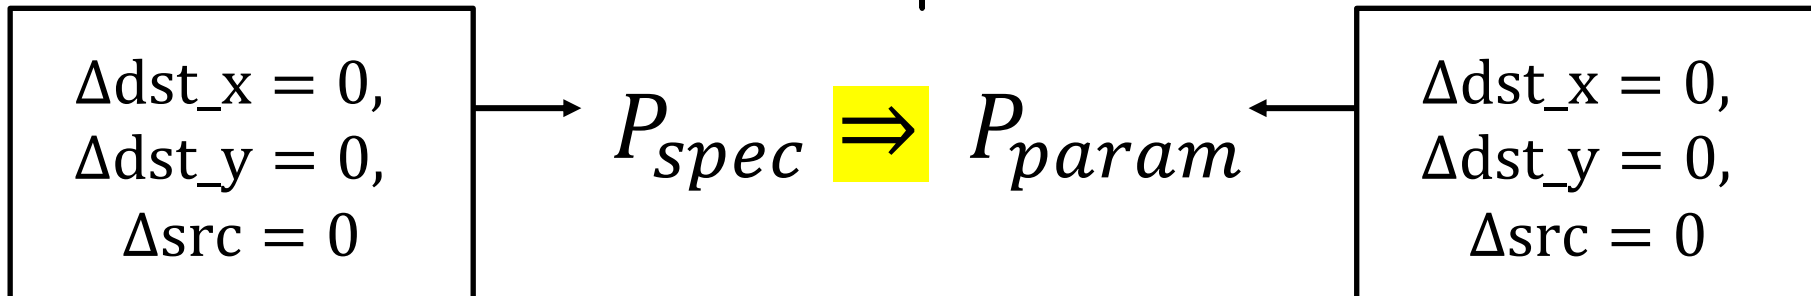
Can Immediately Solve

Reliability Constraint

$$r_{spec} \leq r_1^{q_1} \cdot r_2^{q_2} \cdot \dots \cdot r_n^{q_n}$$

and

$$R(P_{spec}) \leq R(P_{param})$$



Reliability Constraint

$$r_{spec} \leq r_1^{q_1} \cdot r_2^{q_2} \cdot \dots \cdot r_n^{q_n}$$

Denotes approximate function versions that satisfy the developer's specification

Reliability Constraint

for the optimization problem

$$\log(r_{spec}) \leq q_1 \cdot \log(r_1) + q_2 \cdot \log(r_2) + \dots + q_n \cdot \log(r_n)$$

**Denotes approximate function versions that
satisfy the developer's specification**

Reliability and Control Flow

Conditionals

Constraints for each program path

Analysis removes redundant constraints

(most constraints can be removed - OOPSLA '13)

Bounded Loops

Statically known loop bound

Analysis unrolls loop

Optimization Granularity

Optimize blocks of code instead of individual instructions

Function Optimization Problem

Find Function Configuration q : ✓

max EnergySavings (q)

Reliability (q) \geq ReliabilityBound ✓

AbsoluteError (q) \leq ErrorBound

Absolute Error Analysis

Reduced-precision floating-point instructions:

- Almost always incorrect, but error is bounded
- Hardware specification: number of significant mantissa bits

Analysis:

- Bounds worst-case numerical deviation
- Embeds accuracy predicate in reliability factor:

$$r_1^{q_1} \cdot \dots \cdot r_n^{q_n} \cdot R(\Delta x = 0, \Delta y = 0)$$

Absolute Error Analysis

Reduced-precision floating-point instructions:

- Almost always incorrect, but error is bounded
- Hardware specification: number of significant mantissa bits

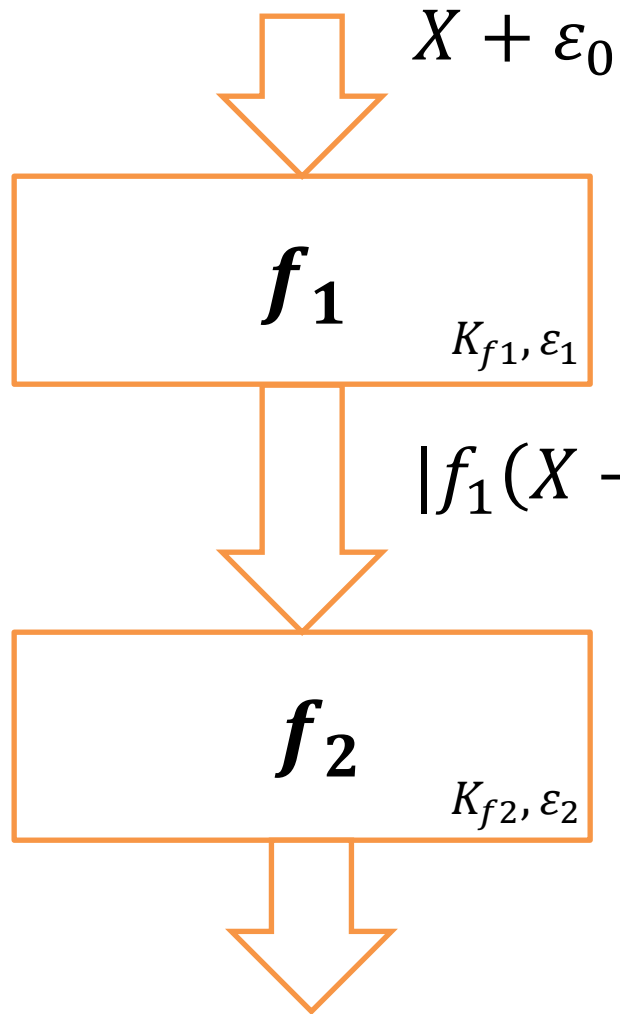
Analysis:

- Bounds worst-case numerical deviation
- Embeds accuracy predicate in reliability factor:

$$r_1^{q_1} \cdot \dots \cdot r_n^{q_n} \cdot \underbrace{R(\Delta x + 2 \cdot \Delta y + q_1 \cdot \xi_{x,y} < d)}_{\text{Linear function of } q_1, \dots, q_n}$$

Linear function of q_1, \dots, q_n

Recall Error Propagation



$$|f_1(X + \varepsilon_0) + \varepsilon_1 - f_1(X)| \leq K_{f_1} \cdot \varepsilon_0 + \varepsilon_1$$

$$K_f = \max_{x \in \text{Inputs}} \left| \frac{df}{dx} \right|$$

$$|f_2(f_1(X + \varepsilon_0) + \varepsilon_1) + \varepsilon_2 - f_2(f_1(X))|$$

$$\leq K_{f_2} \cdot K_{f_1} \cdot \varepsilon_0 + K_{f_2} \cdot \varepsilon_1 + \varepsilon_2$$

Error Propagation for Some Common Functions

$$K_f = \max_{x \in Inputs} \left| \frac{df}{dx} \right|$$

$$K_{fi} = \max_{x \in Inputs} \left| \frac{\partial f(x_1 \dots x_n)}{\partial x_i} \right|$$

$f(x_1, x_2)$

Err

$x \cdot const$

$\Delta x \cdot const$

$x + y$

$\Delta x + \Delta y$

$x \cdot y$

$\Delta x \cdot \max(|y + \Delta y|) + \Delta y \cdot \max(|x + \Delta x|)$

Interpolation Function

```
int interpolation(intqdstx dst_x, intqdsty dst_y, intqsrc src [][])  
{
```

$$\left(r_{rd}\right)^{q_{val}} \cdot \left(r_{times}\right)^{q_*} \cdot \mathbf{R}\left(0.25 \cdot \Delta val + q_* \cdot e_* \leq E\right)$$

```
    return val *q* 0.25;  
}
```

Interpolation Function

```
int interpolation(intqdstx dst_x, intqdsty dst_y, intqsrc src[][])
{
```

$$\left(r_{rd} \right)^{q_{val}} \cdot \left(r_{times} \right)^{q_*} \cdot \left(r_{plus} \right)^{q_1+q_2+q_3} \cdot \left(r_{rd} \right)^{q_{up}+q_{down}+q_{left}+q_{right}} \cdot R \left(\begin{array}{l} 0.25 \cdot (\Delta_{up} + \Delta_{down} + \Delta_{left} + \Delta_{right}) + \\ 0.25 \cdot (q_1 \cdot e_{+1} + q_2 \cdot e_{+2} + q_3 \cdot e_{+3}) + q_* \cdot e_* \leq E \end{array} \right)$$

```
intqval val = up + q1 down + q2 left + q3 right;
```

$$\left(r_{rd} \right)^{q_{val}} \cdot \left(r_{times} \right)^{q_*} \cdot R(0.25 \cdot \Delta_{val} + q_* \cdot e_* \leq E)$$

```
return val * q* 0.25;
}
```

Function Optimization Problem

Find Function Configuration q : ✓

max EnergySavings (q)

Reliability (q) \geq ReliabilityBound ✓

AbsoluteError (q) \leq ErrorBound ✓

Energy Savings Analysis

Profile information:

- Collects traces from running representative inputs

Analysis:

- Estimates savings for instructions and variables from traces

instruction

$$q_\ell \cdot Count_\ell \cdot Saving_{ALU}$$

variable

$$q_m \cdot Size_m \cdot Saving_{MEM}$$

Energy Savings Analysis

Profile information:

- Collects traces from running representative inputs

Analysis:

- Estimates savings for instructions and variables from traces

$$c_{ALU} \sum_{\ell \in Instr} \text{instruction } q_{\ell} \cdot Count_{\ell} \cdot Saving_{ALU} + c_{MEM} \sum_{m \in Var} \text{variable } q_m \cdot Size_m \cdot Saving_{MEM}$$

Approximate hardware specification:

- Relative savings for operations and memories
- Percentage of system energy that ALU and memory consume

Function Optimization Problem

Find Function Configuration q : ✓

max EnergySavings (q) ✓

Reliability (q) \geq ReliabilityBound ✓

AbsoluteError (q) \leq ErrorBound ✓

Reduces to Integer Programming

Find Function Configuration q : ✓

max EnergySavings (q) ✓

Reliability (q) \geq ReliabilityBound ✓

AbsoluteError (q) \leq ErrorBound ✓

Solve using off-the-shelf solvers (we use Gurobi)

Evaluation

Benchmarks With Approximated Functions:

Scale	image scaling
DCT	discrete cosine transform
IDCT	inverse discrete cosine transform
Blackscholes	financial option price calculation
SOR	successive over-relaxation kernel

Approximate Hardware Specifications:

- 5 specifications of ALU, caches, and memories from the literature [\[Ener\] – PLDI '11\]](#)

Complexity of Optimization Problem

Benchmark	Function LOC	Search Space Size	Reliability Constraints
Scale	88	2^{74}	4
DCT	62	2^{35}	1
IDCT	93	2^{53}	1
Blackscholes	143	2^{80}	2
SOR	23	2^{10}	1

Solver finds optimal solutions in less than a second

Energy/Accuracy Tradeoffs

Optimizer computes estimated system savings

Maximum estimated savings for hardware specifications:

Benchmark	Reliability Degradation	System-Level Energy Savings
Scale	0.995	19.4%
DCT	0.99992	8.7%
IDCT	0.992	13.4%
Blackscholes	0.999	9.8%
SOR	0.995	19.8%

Pros:

- Can explore the space induced by much finer grained transformations (e.g., numerical precision)
- The results are valid for all inputs within range
- New analyses were developed in the meantime

Cons:

- Static analysis is *much* more conservative than testing
- The set of supported programs is limited

Analysis: Middle Road

What if we know the distribution of the inputs?

CASE I: Sum Computation

- Original sum computation

```
s = 0;
```

```
for (i = 0; i < n; i++) s = s + f(i);
```

- Perforated, extrapolated sum computation

```
s = 0;
```

```
for (i = 0; i < n; i += 2) s = s + f(i);
```

```
s = s * 2;
```

Step 1: Represent Result Difference

- Original sum computation

```
s = 0;  
for (i = 0; i < n; i++) s = s + f(i);
```

- Perforated, extrapolated sum computation

```
s = 0;  
for (i = 0; i < n; i += 2) s = s + f(i);  
s = s * 2;
```

- Perforation noise: $\mathbf{D} = \mathbf{s}_{\text{original}} - \mathbf{s}_{\text{perforated}}$

Step 2: Probabilistic Modeling

- Original sum computation

s = 0;

for (i = 0; i < n; i++) s = s + f(i);

- Perforated, extrapolated sum computation

s = 0;

for (i = 0; i < n; i += 2) s = s + f(i);

s = s * 2;

- Perforation noise: **D = s_{original} - s_{perforated}**

Step 2: Probabilistic Modeling

- Original sum computation

s = 0;

for (i = 0; i < n; i++) s = s + X_i;

- Perforated, extrapolated sum computation

s = 0;

for (i = 0; i < n; i += 2) s = s + X_i;

s = s * 2;



Specify assumptions

- Perforation noise: **D = s_{original} - s_{perforated}**

Analysis: Input/Output Relation

Perforation noise:

$$D = S_{\text{original}} - S_{\text{perforated}}$$

Analysis: Input/Output Relation

Perforation noise:

$$D = S_{\text{original}} - S_{\text{perforated}}$$

$$= X_0 + X_1 + X_2 + X_3 + X_4 + X_5 + X_6 + X_7 + \dots$$

$$- 2 \cdot (X_0 + X_2 + X_4 + X_6 + \dots)$$

Analysis: Input/Output Relation

Perforation noise*:

$$D = S_{\text{original}} - S_{\text{perforated}}$$

$$= X_0 + X_1 + X_2 + X_3 + X_4 + X_5 + X_6 + X_7 + \dots$$

$$- X_0 - X_0 - X_2 - X_2 - X_4 - X_4 - X_6 - X_6 - \dots$$

* Assuming for simplicity that the number of elements is even

Analysis: Input/Output Relation

Perforation noise*:

$$D = S_{\text{original}} - S_{\text{perforated}}$$

$$= X_0 + X_1 + X_2 + X_3 + X_4 + X_5 + X_6 + X_7 + \dots$$

$$- X_0 - X_0 - X_2 - X_2 - X_4 - X_4 - X_6 - X_6 - \dots$$

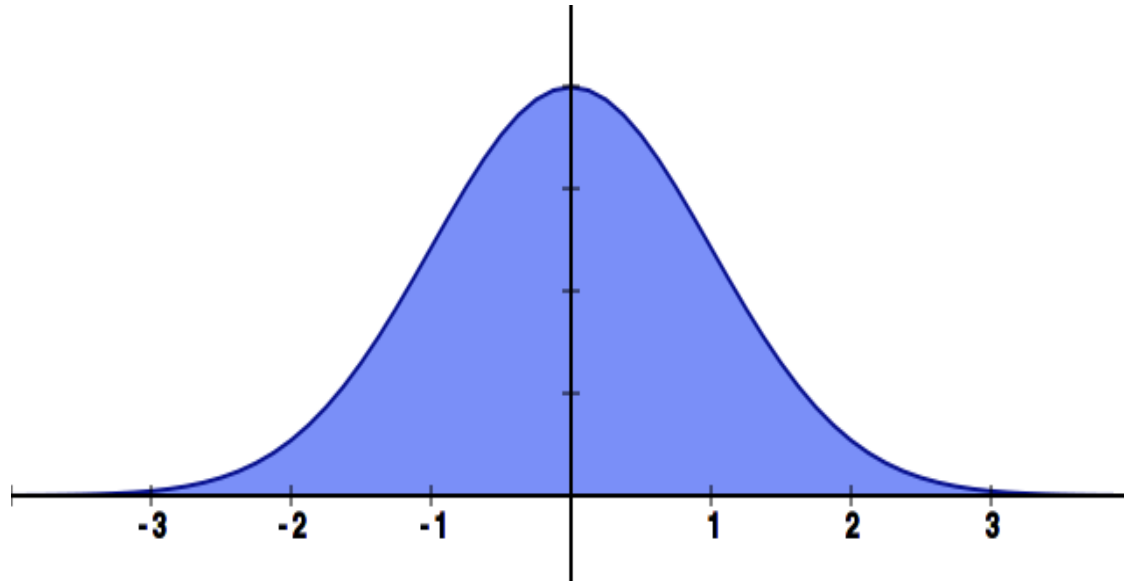
$$= \sum_{0 \leq i < \frac{n}{2}} (X_{2i+1} - X_{2i})$$

* Assuming for simplicity that the number of elements is even

Analysis Results

Perforation noise:

$$D = \varphi(X_0, X_2, \dots, X_{n-1})$$



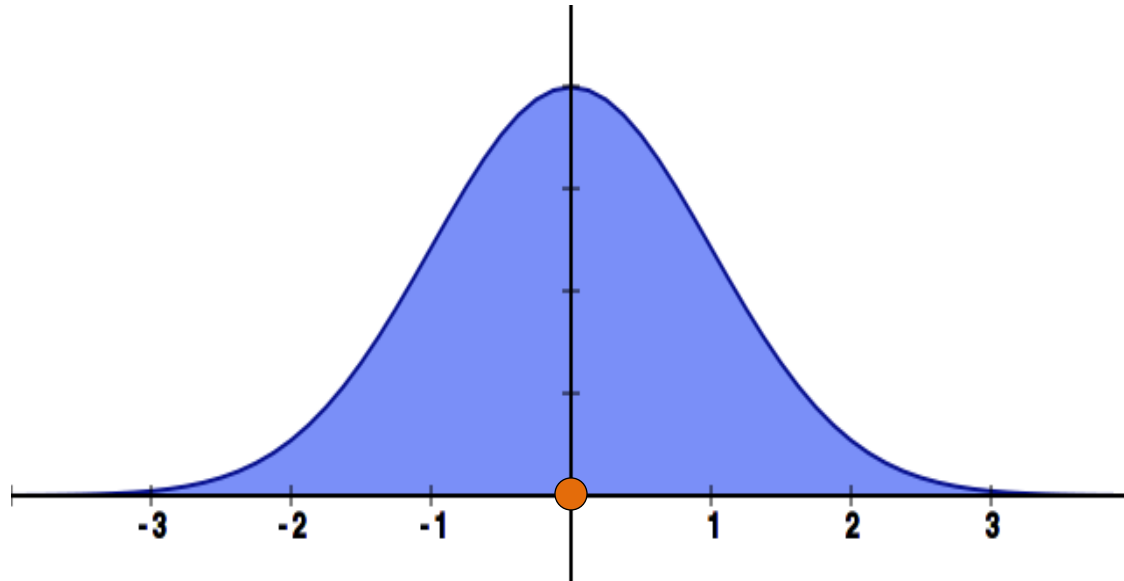
Analysis Results

Perforation noise:

$$D = \varphi(X_0, X_2, \dots, X_{n-1})$$

Location: Mean

$$E(D) = \mu$$



Analysis Results

Perforation noise:

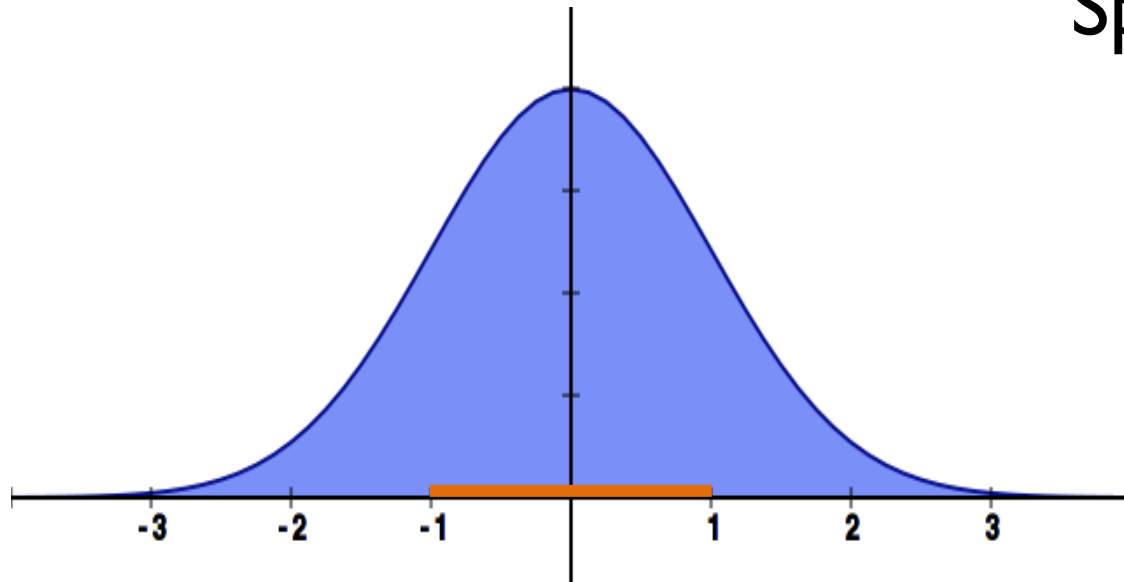
$$D = \varphi(X_0, X_2, \dots, X_{n-1})$$

Location: Mean

$$E(D) = \mu$$

Spread: Variance

$$\text{Var}(D) = \sigma^2$$



Analysis Results

Perforation noise:

$$D = \varphi(X_0, X_2, \dots, X_{n-1})$$

Location: Mean

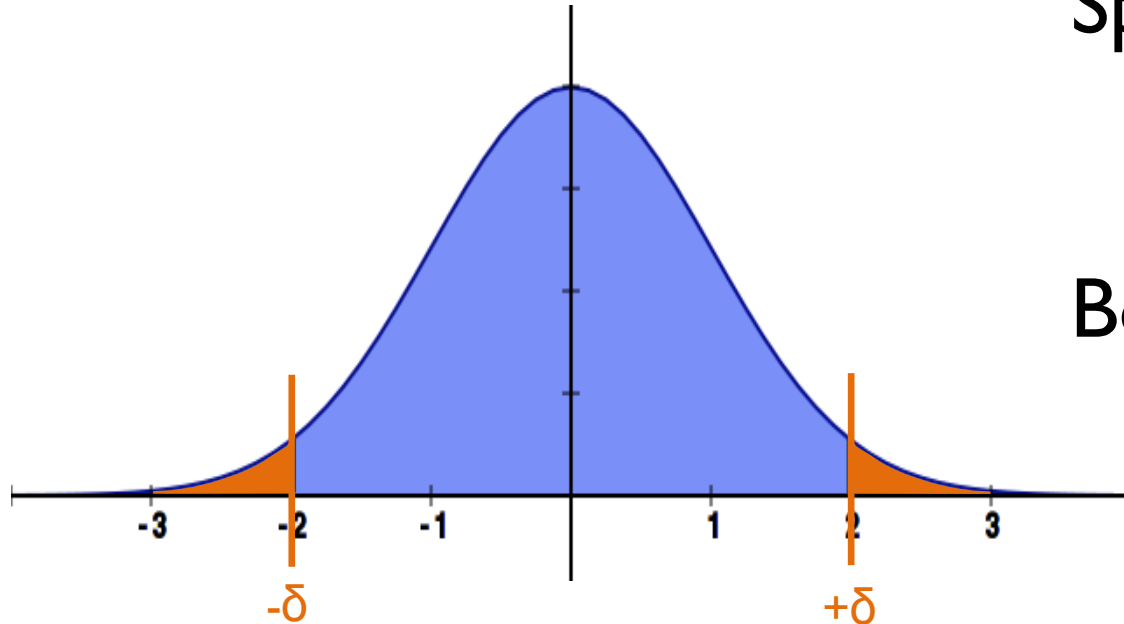
$$E(D) = \mu$$

Spread: Variance

$$\text{Var}(D) = \sigma^2$$

Bound: Distribution tail

$$\Pr[|D| > \delta] < \varepsilon$$



Next Time

Probabilistic programming:
Democratizing probabilistic inference