

Fine-Grained Characterization of Faults Causing Long Latency Crashes in Programs

Guanpeng Li, Qining Lu and Karthik Pattabiraman

University of British Columbia, Canada

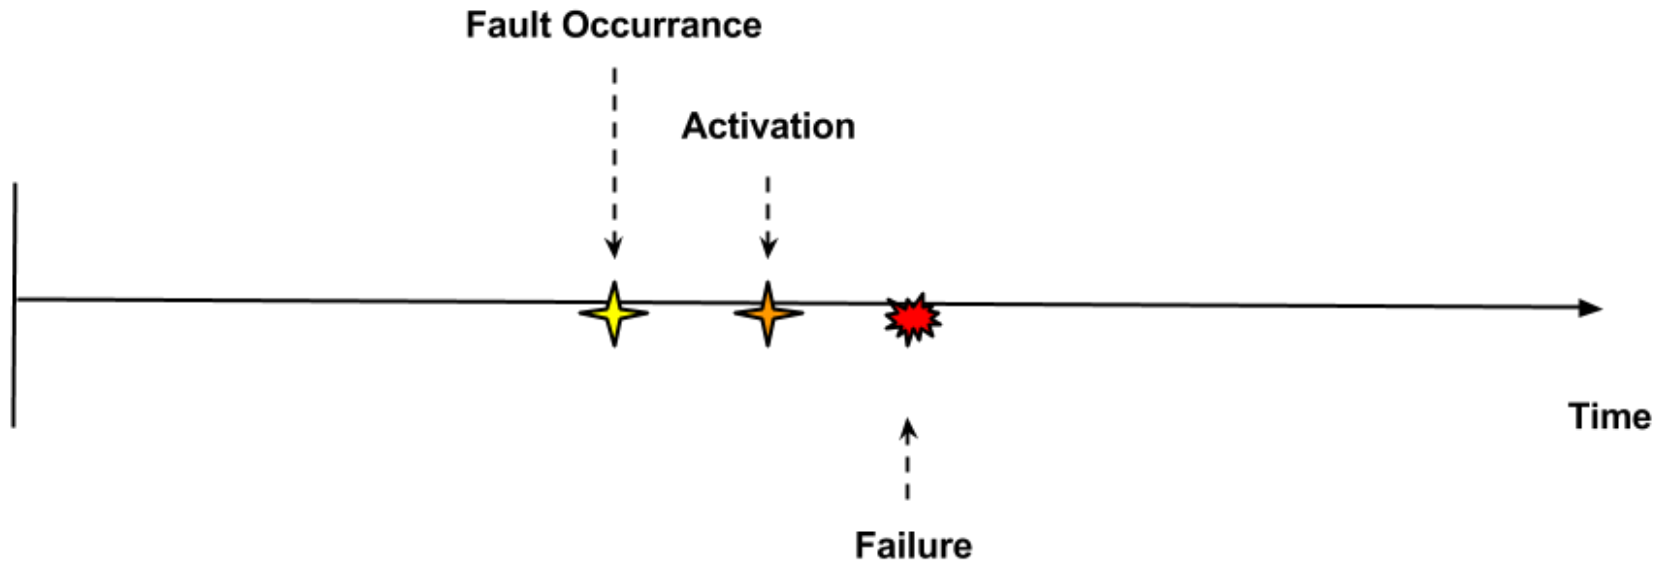


Soft Errors

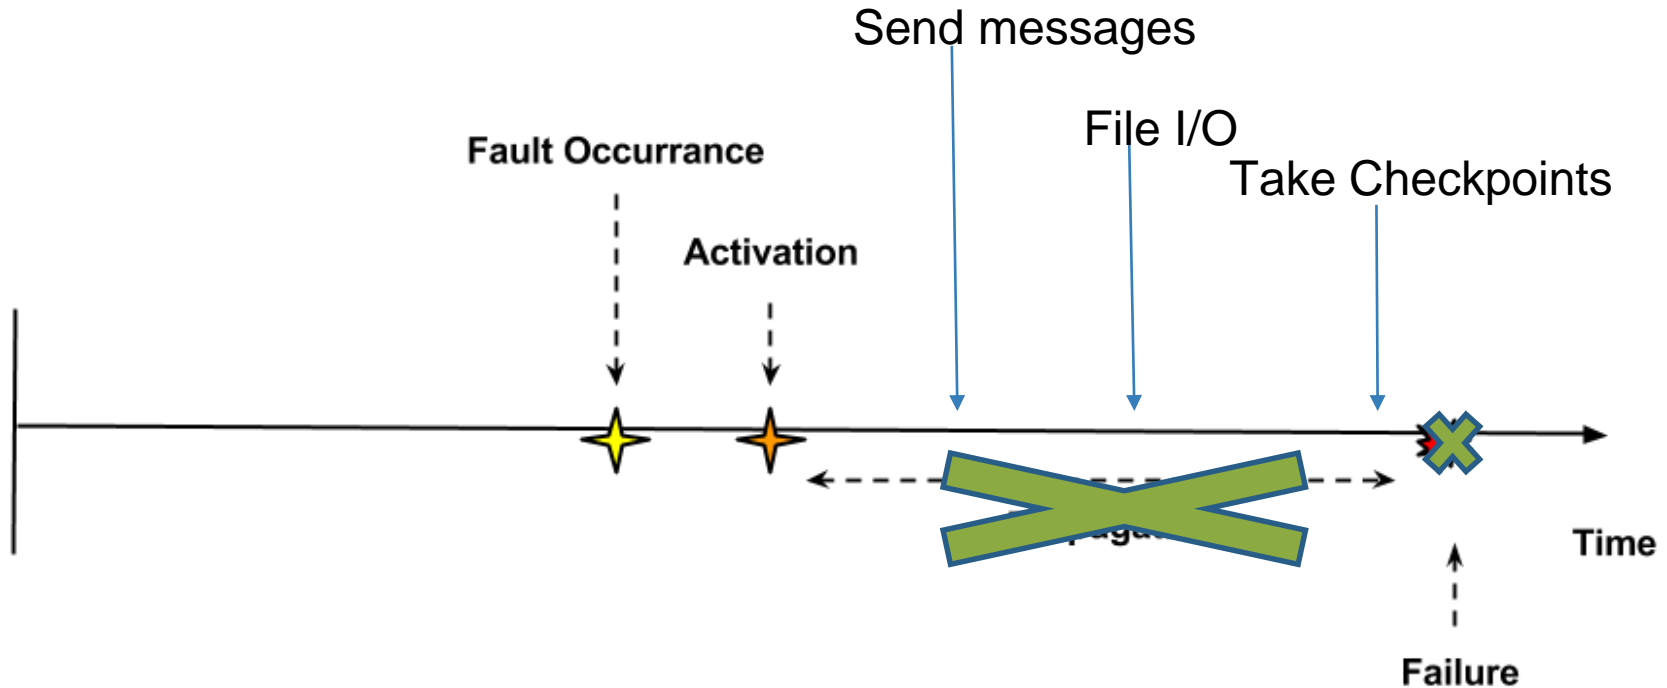
Soft error rate will increase by nearly an order of magnitude as chip feature sizes decrease from 22nm to 14nm.

[Feng et. al., ASPLOS'10]

Fail-stop Assumption



But, in reality ...



Traditional Solutions

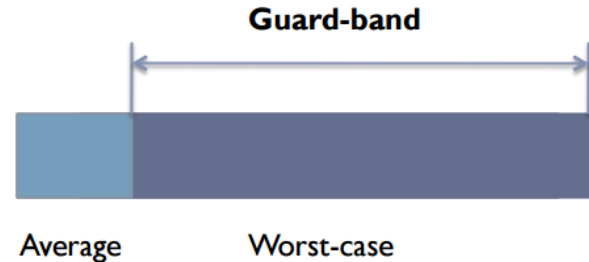
▶ Duplication

Hardware duplication (DMR) can result in 2X slowdown and/or energy consumption

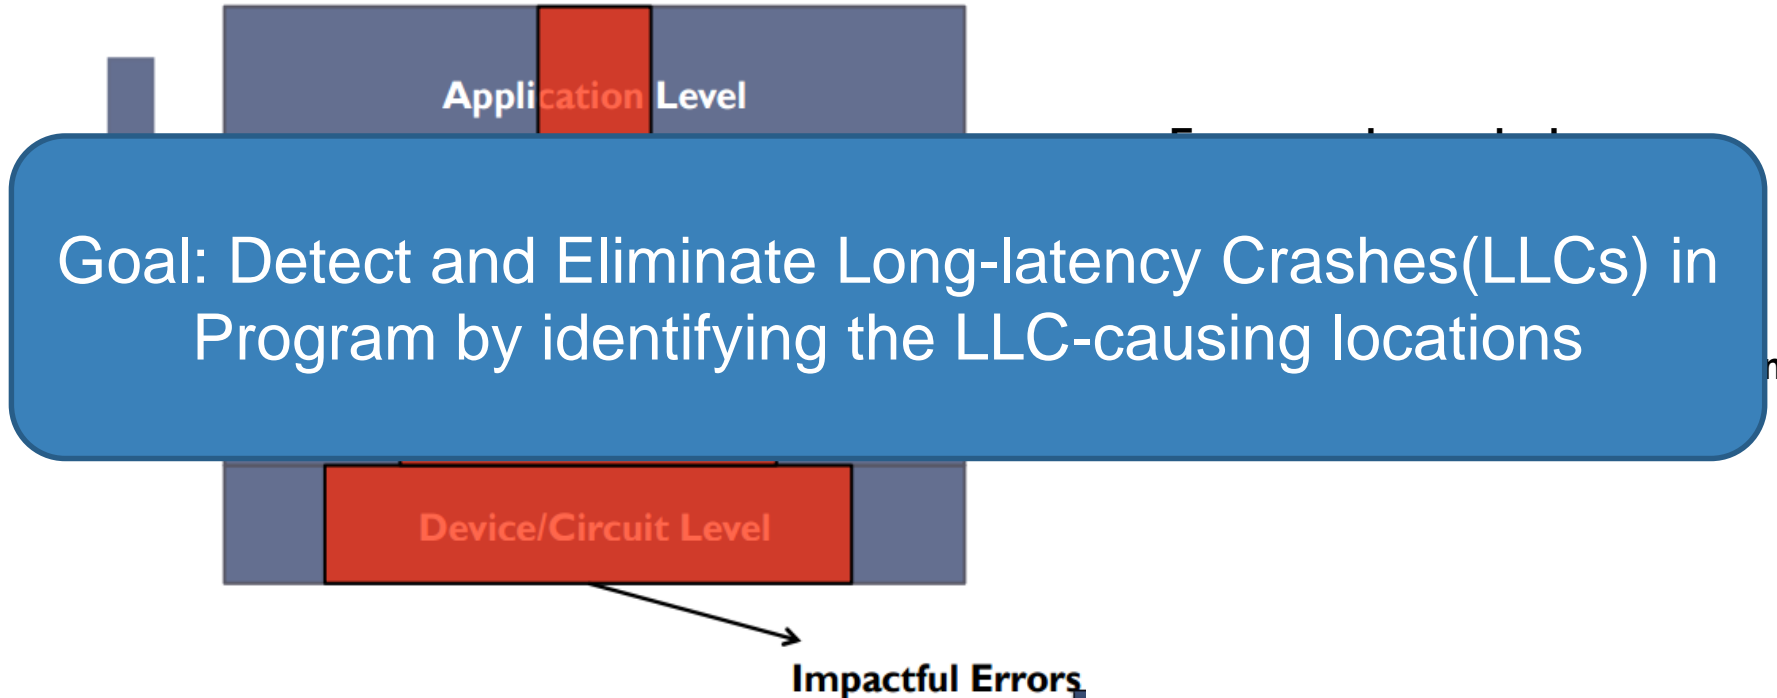


▶ Guard-banding

Guard-banding wastes power and performance as gap between average and worst-case widens due to variations



Why Software ?



Challenges



Statistical Fault Injection

- Good for resiliency characterization
- Takes long time to find LLCs



What we do



Code patterns leading to LLC fall into very few dominant patterns



Static analysis to identify the patterns



Selective sampling to filter out false-positives



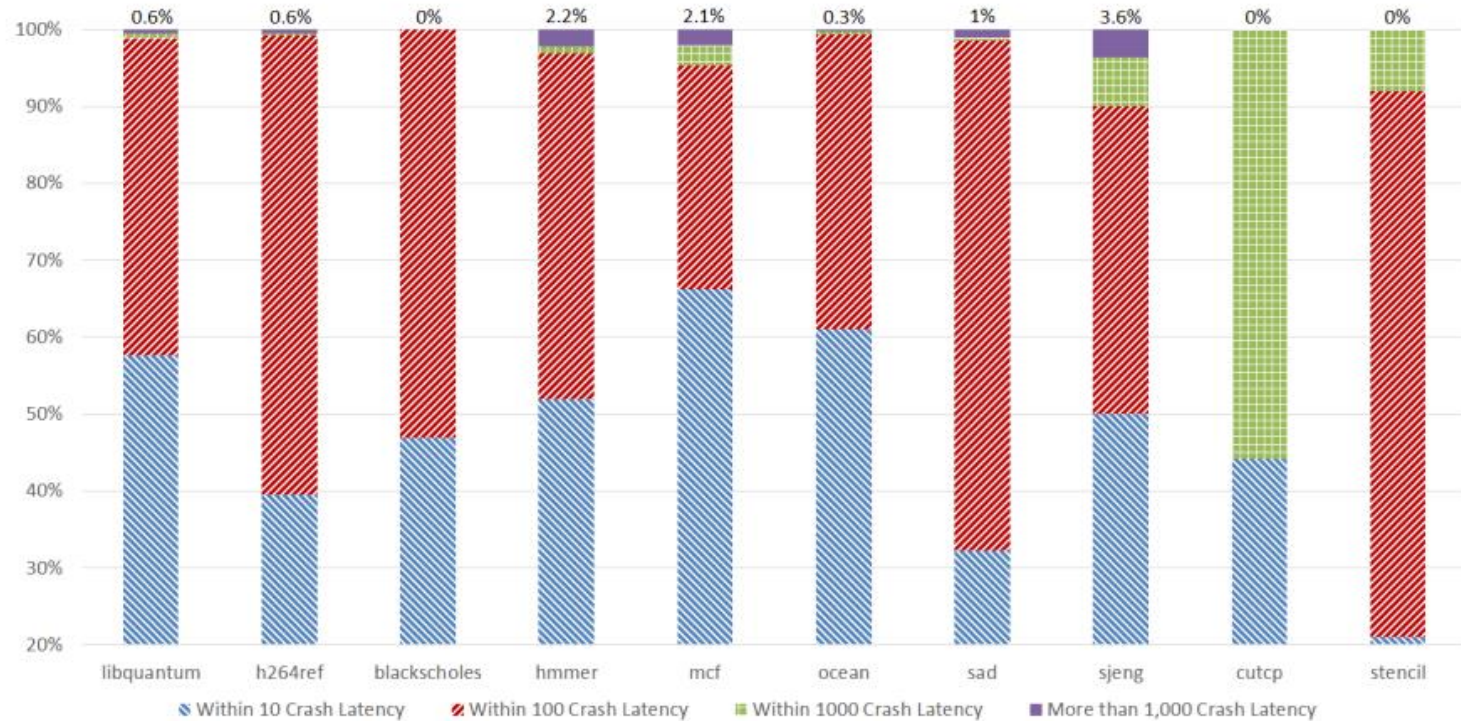
Initial Fault Injection Study

- Choose 5 from 10 benchmark applications
- 1,000 random fault injections per application
- 1 fault injection per run – single bit flip

Fault Model

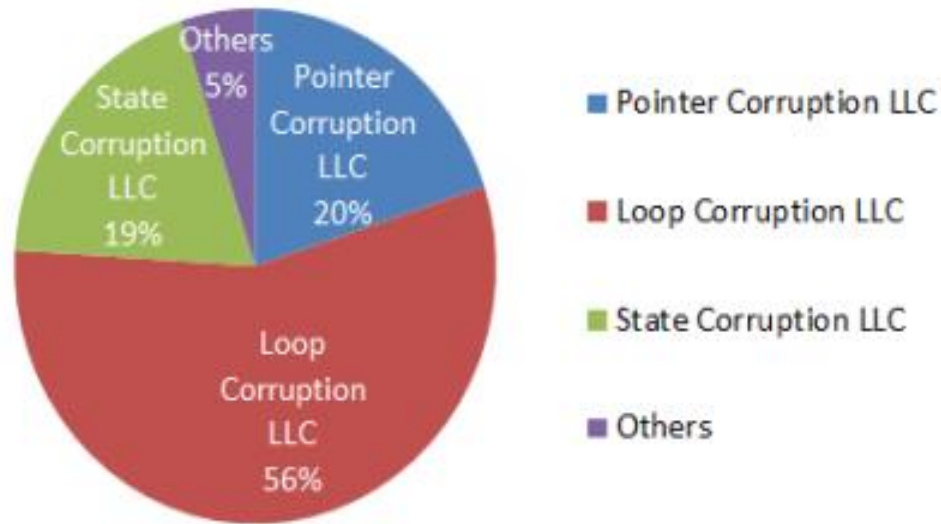
- Faults occur in computational components or load/store units in CPU
- Assume memory and cache - ECC protected
- **LLVM Fault Injector (LLFI)** [Wei – DSN'14]

Propagation Latency (dynamic insns)



Propagation latency is application-specific

Patterns Leading to LLCs



What we do



Code patterns leading to LLC fall into very few dominant patterns

CrashFinder Static

Static analysis to identify the patterns

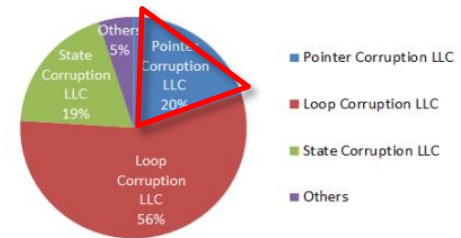
Selective sampling to filter out false-positives



Pointer Corruption LLC: Example

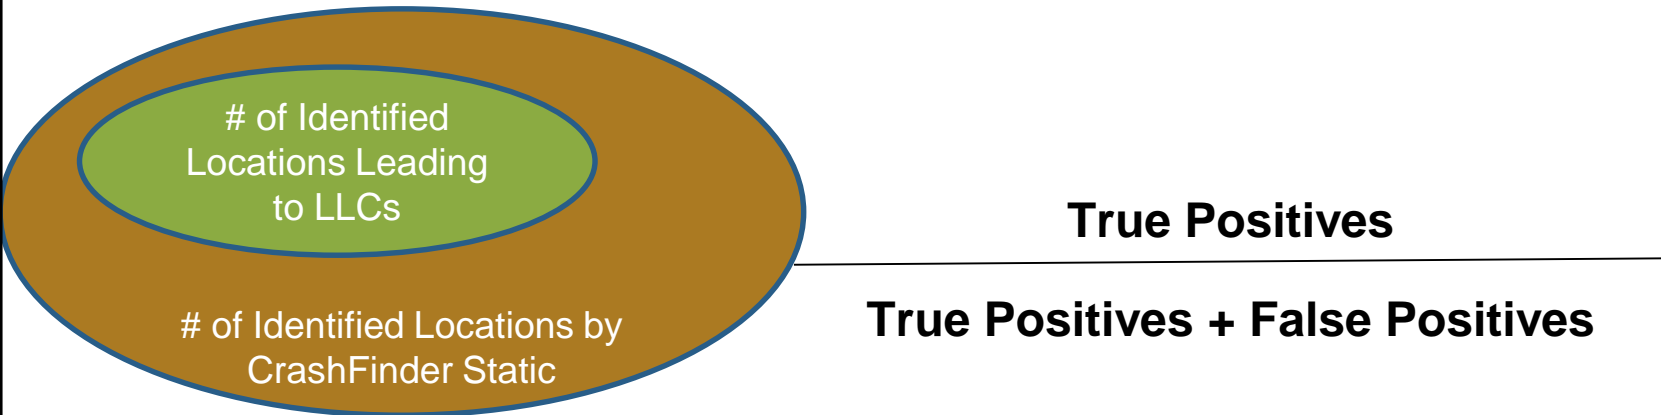
```
1 static unsigned int state[N+1];
2 static unsigned int *next;
3 ...
4 unsigned int reloadMT(void)
5 {
6     ...
7     register unsigned int *p0 = state;
8     next = state+1;
9     ...
10    *p0++ = *pM++ ^ ... ;
11    ...
12 }
13 ...
14 unsigned int randomMT(void)
15 {
16     unsigned int y;
17     ...
18     y = *next++;
19     ...
20 }
21 ...
```

[From sjeng program]

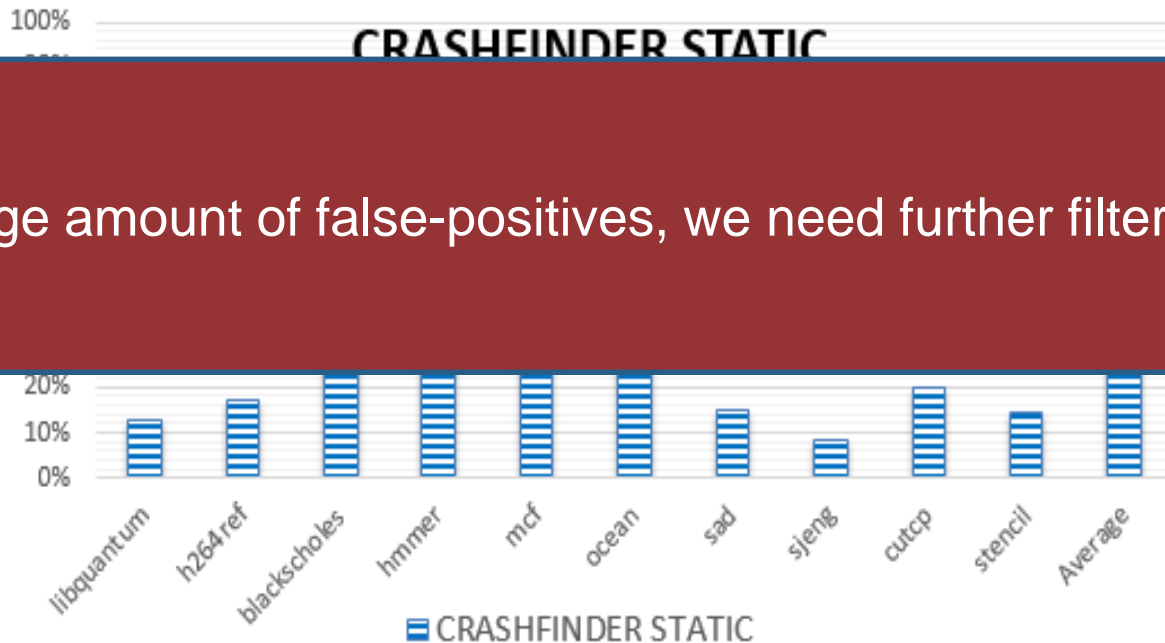


Precision

- 200 random fault injections on each static location identified by the technique
- 10 applications from 4 benchmark suites



$$\text{Precision} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Positives}}$$



What we do



Code patterns leading to LLC fall into very few dominant patterns



CrashFinder Static

Static analysis to identify the patterns



Selective sampling to filter out false-positives

2 Heuristics:
H1 & H2



H1: Instruction Sampling

Similar behavior in similar control flow

More efficient to sample by unique function call sequence

```
1 void multig(long my_id) {
2   ...
3   while ((!flag1) && (!flag2)) {
4     ...
5     relax();
6     copy_red();
7     relax();
8     copy_black();
9     ...
10  }
11 }
12
13 void relax() {
14   ...
15   for (...) {
16     ...
17     * t1a = (double *) t2a[i];
18     ...
19   }
20 }
```

```
sample candidate 1 → relax()
                   copy_red()
sample candidate 2 → relax()
                   copy_black()
sample candidate 3 → relax()
                   copy_red()
sample candidate 4 → relax()
                   copy_black()
                   ...
sample candidate N → relax()
                   copy_red()
sample candidate (N+1) → relax()
                       copy_black()
```

H2: Bit Sampling



What we do



Code patterns leading to LLC fall into very few dominant patterns



CrashFinder Static

Static analysis to identify the patterns



Selective sampling to filter out false-positives

2 Heuristics:
H1 & H2

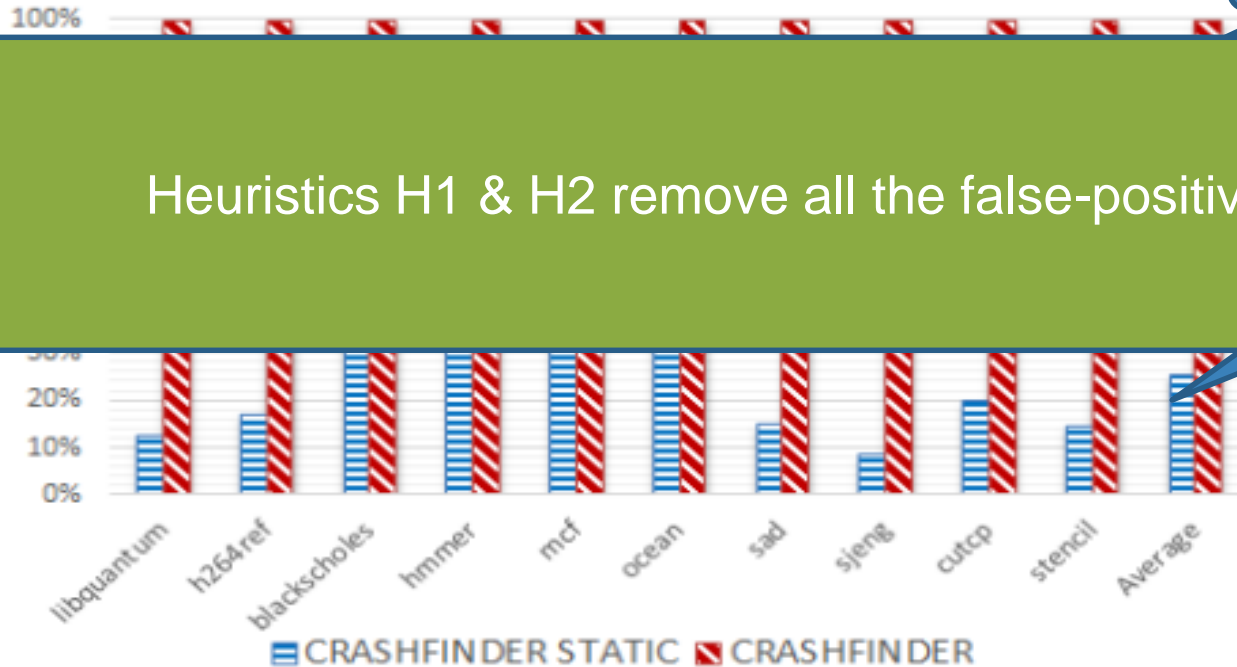
CrashFinder

- LLVM compiler pass
- No annotation required - Fully automated
- Supports C and C++ programs

$$\text{Precision} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Positives}}$$

CrashFinder
100%

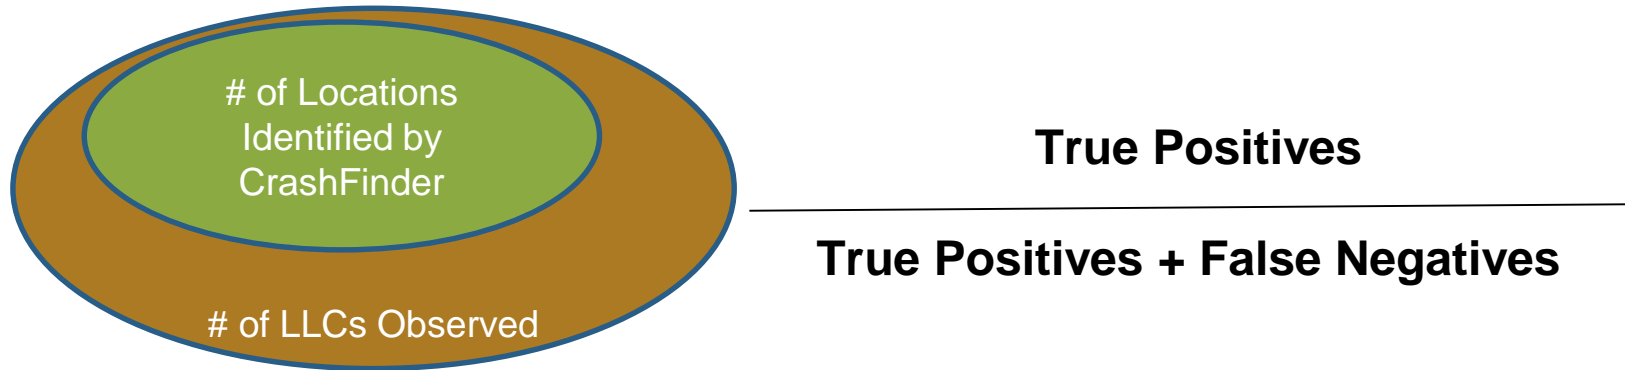
Heuristics H1 & H2 remove all the false-positives !



Recall

Experiment

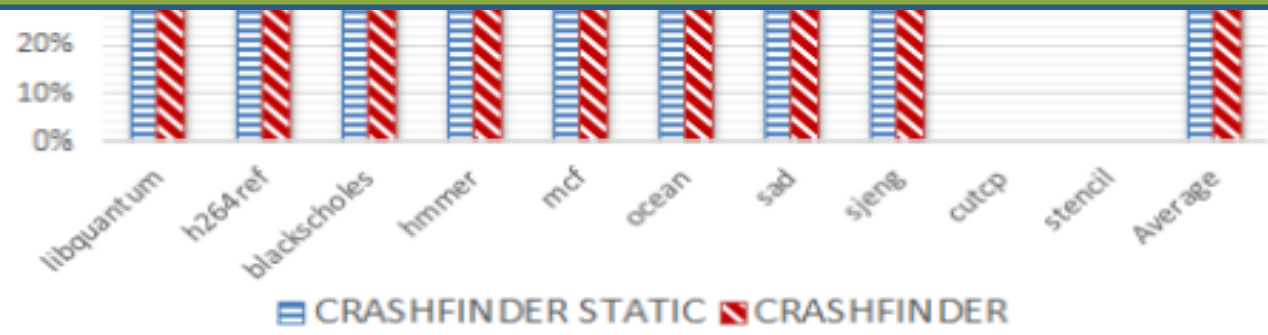
- 3,000 random fault injections on each application
- Total of 10 benchmarks



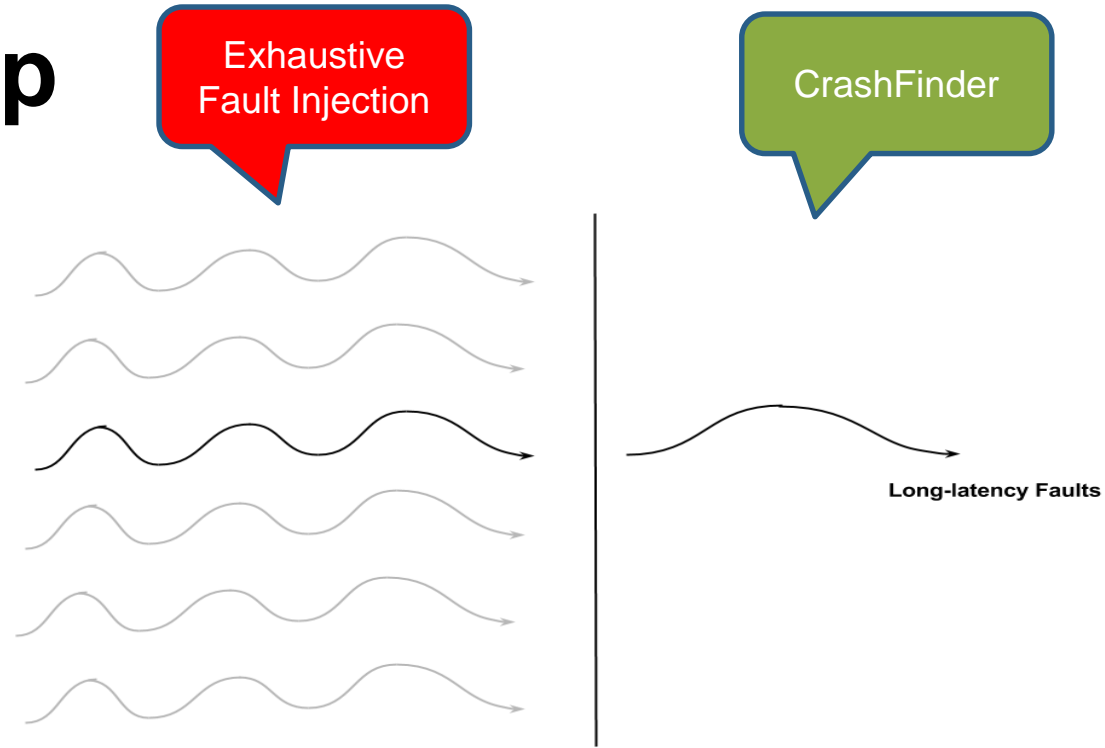
$$\text{Recall} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}}$$



- Able to identify most of the LLCs
- ~2% loss in recall between CrashFinder and CrashFinder Static



Speed Up



$$\text{Speed Up} = \frac{\text{Time taken by CrashFinder}}{\text{Time taken by exhaustive fault injection}}$$

Speed Up: Orders of Magnitude



Implications: Costs and Benefits

Performance Overhead [under submission]

- ~5% by selective duplication of LLC causing locations' backward slices

Availability [under submission]

- Avoids ~96% of checkpoint corruptions
- About 8 times reduction in unavailability
(unavailability = 1 - availability)

Related Work

- Long-latency faults have been observed, but no one has identified patterns leading to them [Chandra 2002] [Gu 2003] [Yim 2009]
- Relyzer [Hari 2011], SDCTune [Lu 2014] reduces fault injection space for SDCs
- Non-trivial to extend for LLCs which are much rarer

Summary

- Long-latency crashes (LLCs) fall into 3 dominant code patterns, which can be identified thro' static analysis
- Heuristics used in CrashFinder works well with ~90% recall and 100% precision (i.e., no false positives)
- Speedup of more than 9 orders of magnitude compared to exhaustive fault injection (current state of the art)

gpli@ece.ubc.ca

<https://github.com/DependableSystemsLab/Crashfinder>