

Statistical Debugging for Real-World Performance Problems

Linhai Song¹ and Shan Lu²

¹University of Wisconsin-Madison

²University of Chicago

What are Performance Problems?

- Definition of Performance Problems (PPs):
 - Implementation mistakes causing inefficiency
- An example

rows=0 causing
no cache allocated

MySQL Bug DB

```
void ha_partition::start_bulk_insert(int rows) {  
    .....  
- if (!rows) //check whether rows is 0  
- DEBUG_VOID_RETURN;  
- rows= rows/m_tot_parts + 1;  
+ rows= rows ? rows/m_tot_parts + 1 : 0;  
    ..... // fast path using caches  
}
```

MySQL Bug 26527

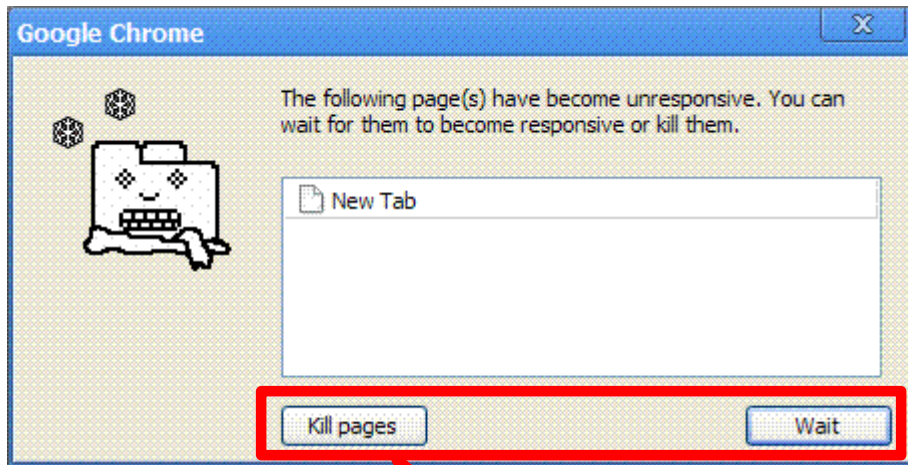
The screenshot shows the MySQL Bug DB interface. The bug title is "MySQL Bug 26527". The bug was submitted on 21 Feb 2007 14:49 and modified on 29 Oct 2007 18:41. The reporter is Guillaume Lefranc, and the assigned person is Alexey Botchkov. The bug is categorized as "Server: Partition" with a severity of "S5 (Performance)". The description states: "Inserting data with LOAD DATA INFILE is painfully slow with partitioned table and sometimes crawl to a stop. I haven't tried with SQL dumps to see if the problem repeats." The "How to repeat" section includes a SQL script to create a table with 12 partitions. A callout box points to the script with the text "20 X Slower". At the bottom, a callout box points to the performance comparison text: "Inserting 64GB of data takes more than 1 day with this setup. Repeat without partitioning : Insert took lh30 avg."

Fighting PPs is Important

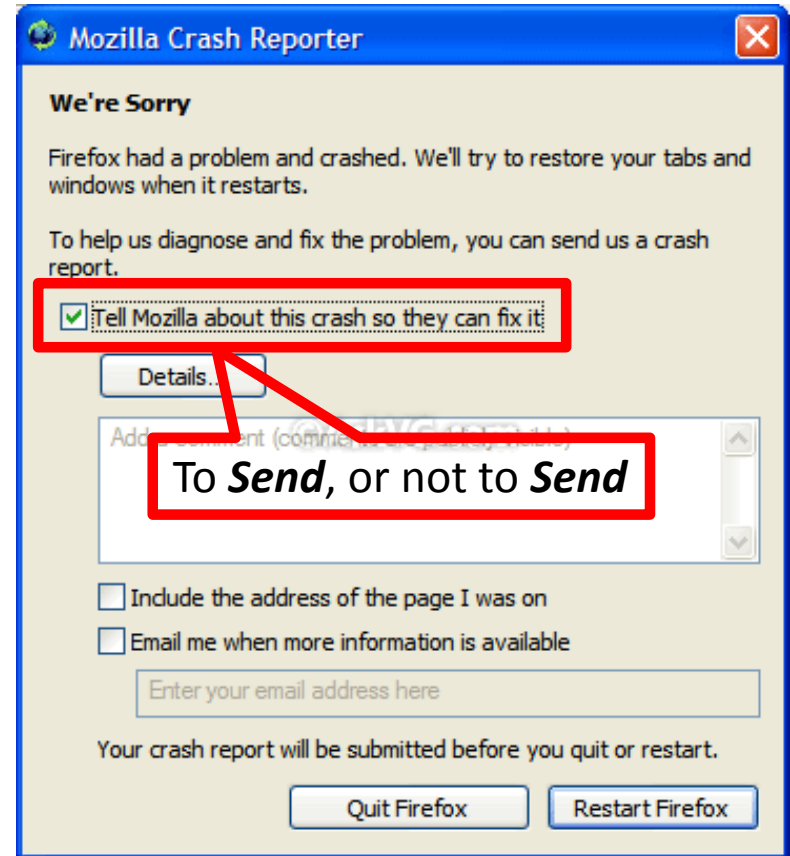
- PPs widely exists in production run software
 - 5 to 60 Mozilla PPs are fixed each month
- PPs are getting more important
 - Hardware is not getting faster (per-core)
 - Software is getting more complex
 - Energy saving is getting more urgent



Performance Diagnosis

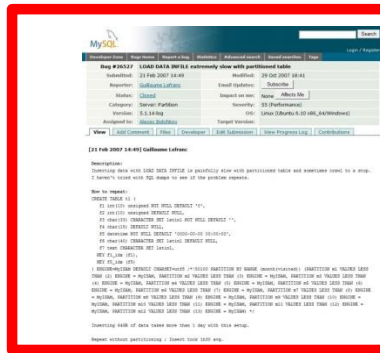


To *Kill*, or to *Wait*

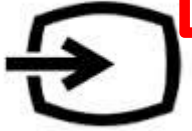


Performance Diagnosis

- Identifying the causes of performance failures
 - In-house and on-line diagnosis




INPUT



In-house



On-line



Performance Diagnosis is Challenging

- The state of the art is *preliminary*
 - Profilers
 - Only tools mentioned in bug reports we studied
 - Output time-consuming functions, not root causes
- More effective tools are necessary

```
void ha_partition::start_bulk_insert(int rows) {  
    .....  
-   if (!rows)  
-       DEBUG_VOID_RETURN;  
-   rows= rows/m_tot_parts + 1;  
+   rows= rows ? rows/m_tot_parts + 1 : 0;  
    ..... // fast path using caches  
}
```

Not in profiling results

MySQL Bug 26527

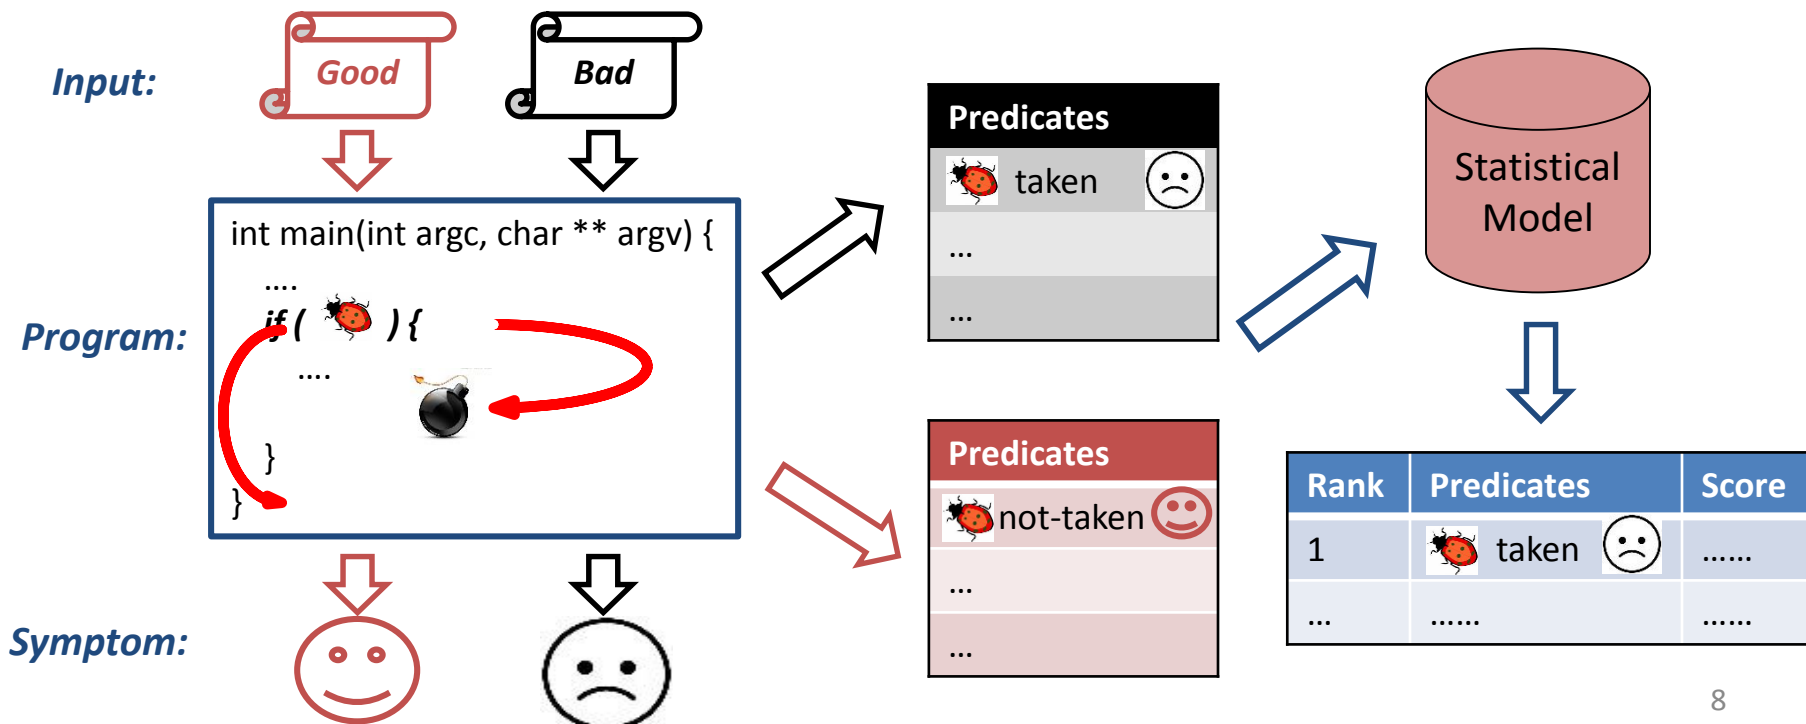
How Can We Do Better?

Can we learn from functional bug diagnosis?



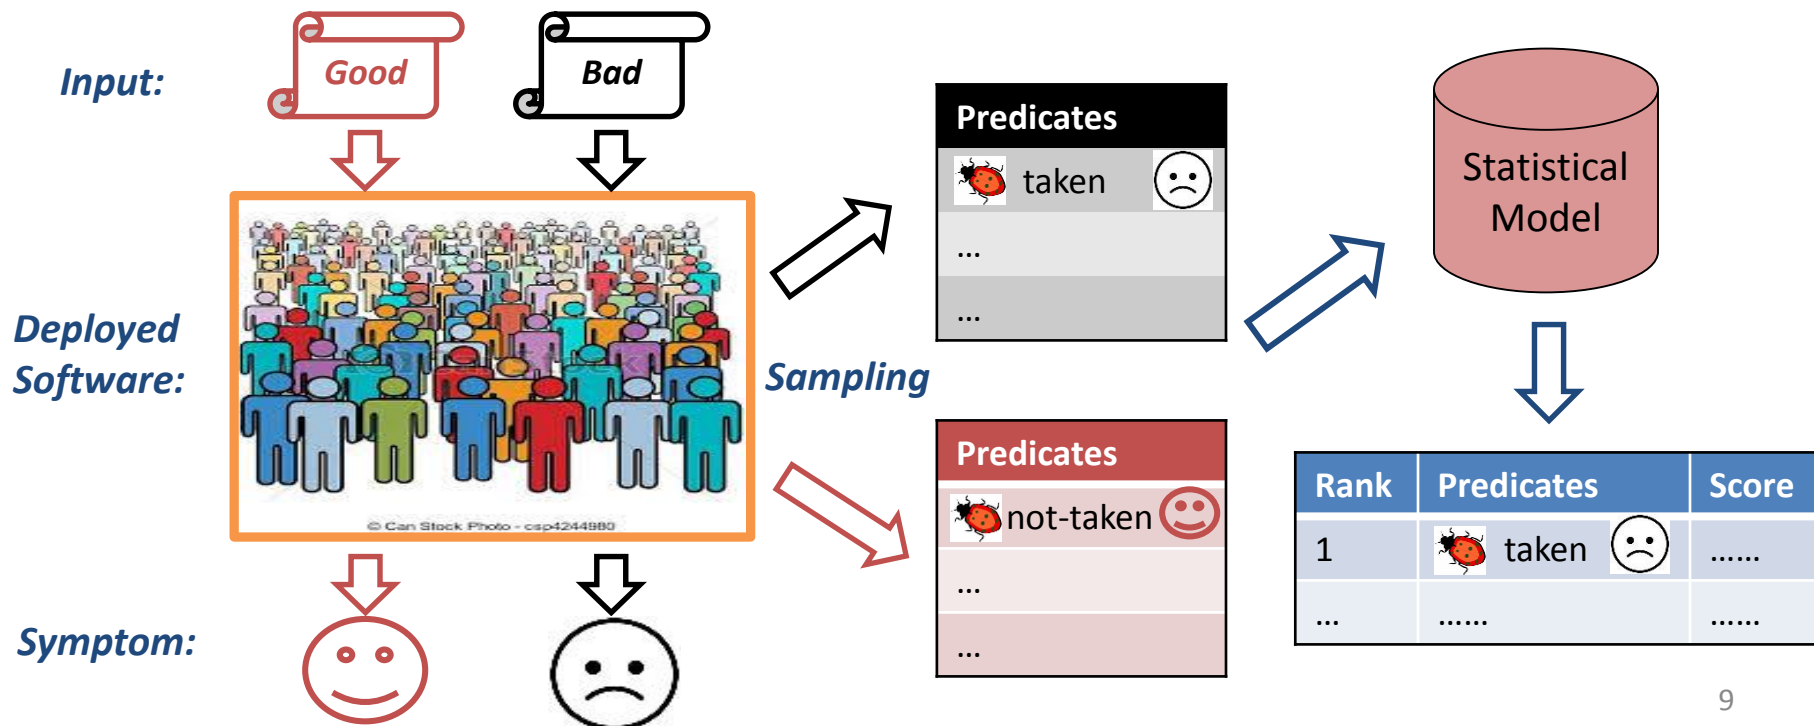
How to Diagnose Functional Bugs

- The state of the art is *mature*
- Statistical Debugging(SD)
 - Among the most effective



How to Diagnose Functional Bugs

- The state of the art is *mature*
- Statistical Debugging(SD)
 - Among the most effective



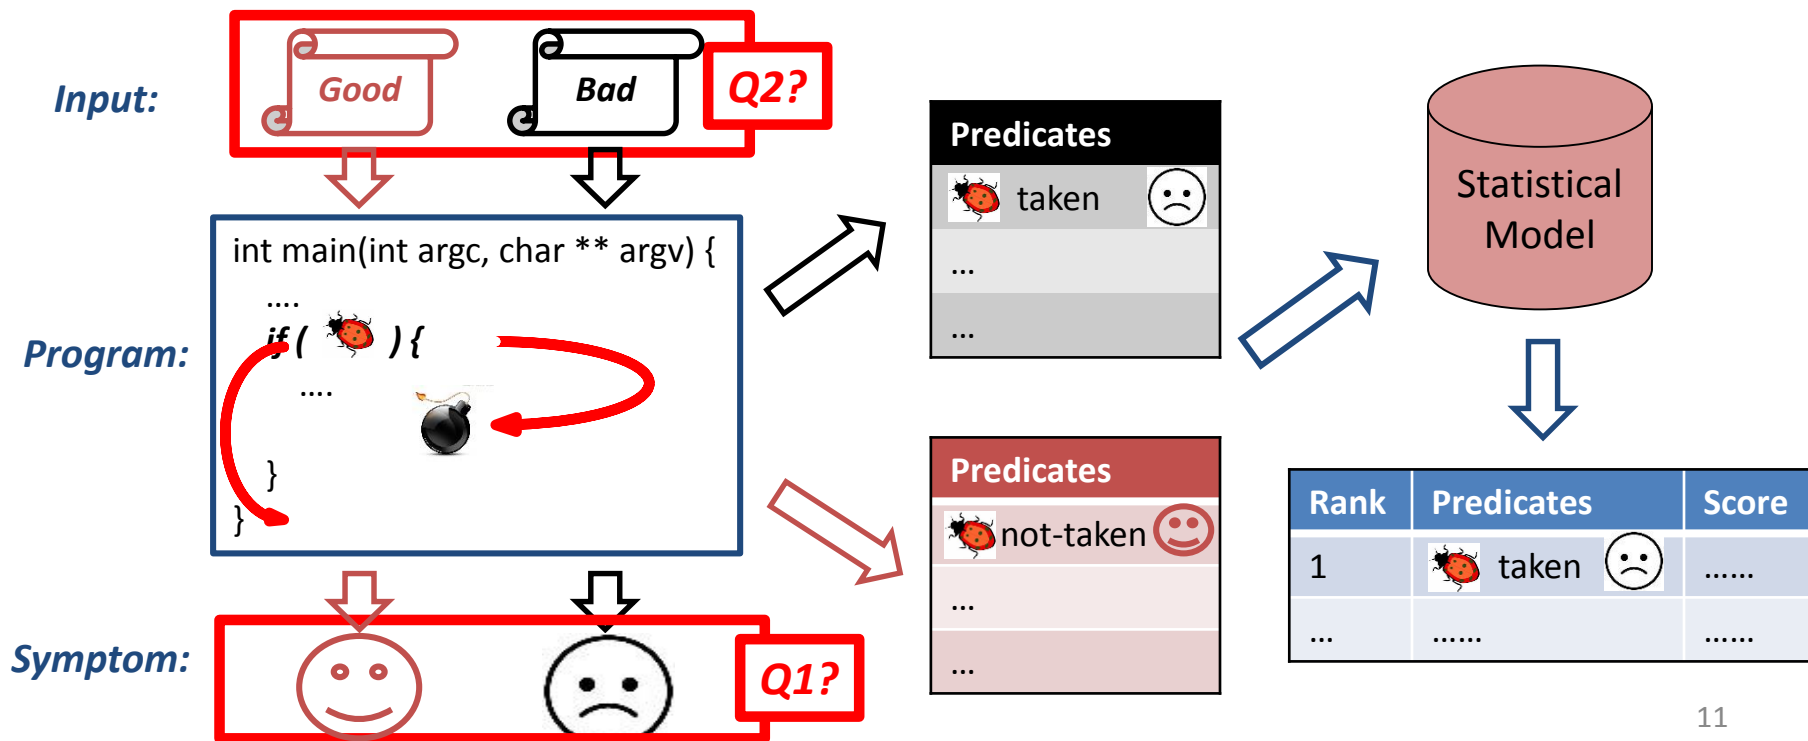
Apply SD to Performance Diagnosis?

Since statistical debugging is effective for functional diagnosis, maybe it will also be effective for performance diagnosis.



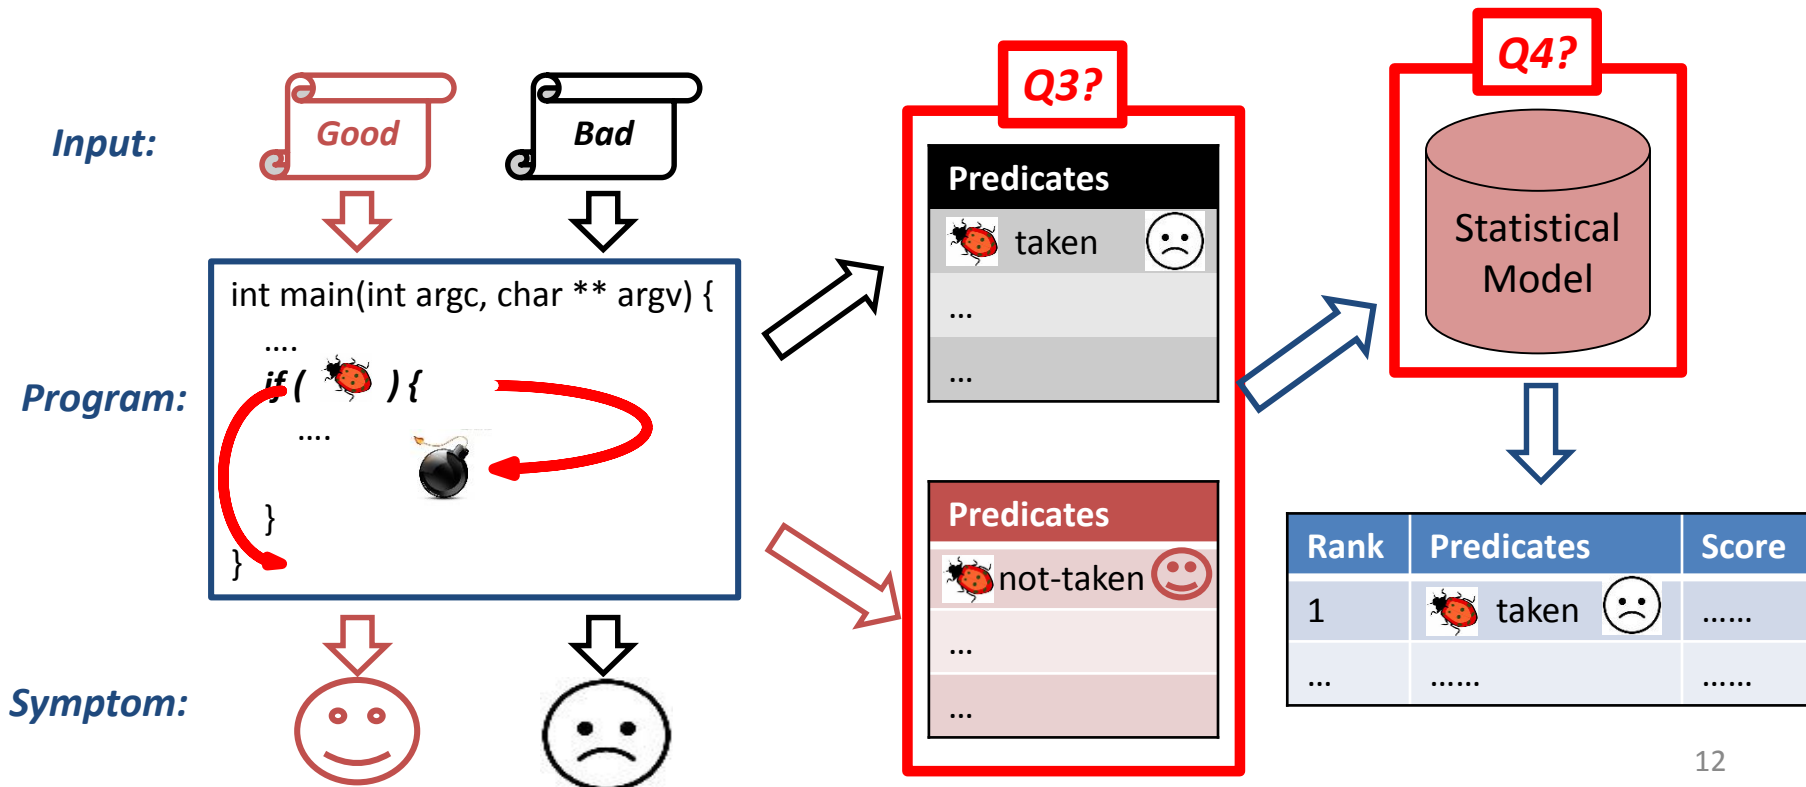
Is it Feasible?

- Q1: How to tell success runs from failure runs?
- Q2: How to obtain good and bad inputs?



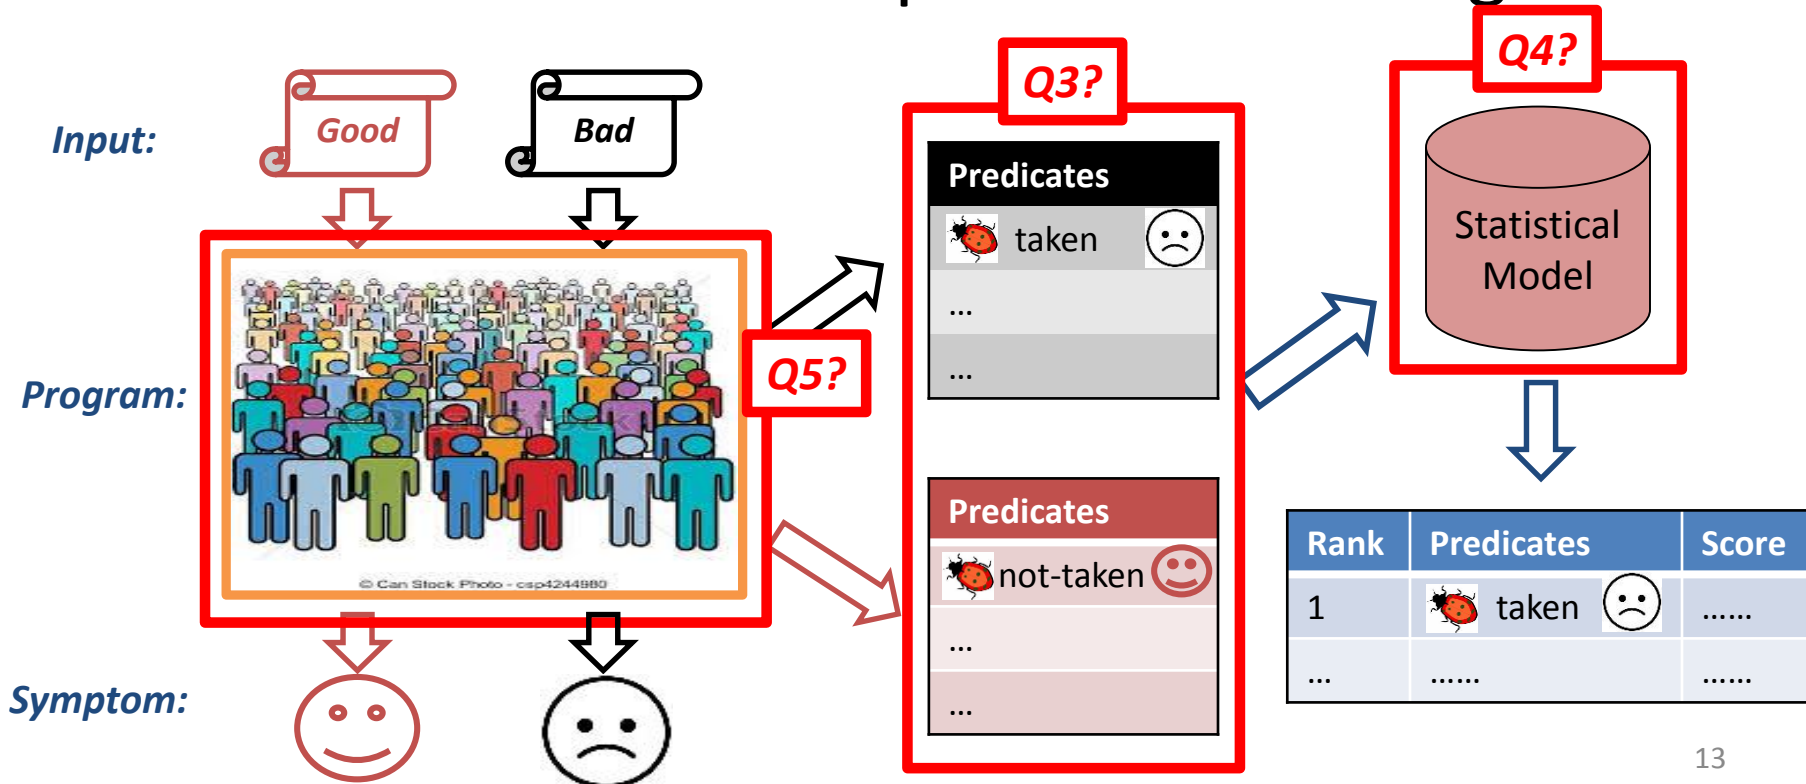
How to Apply?

- Q3: What predicates to collect?
- Q4: What statistical model to use?



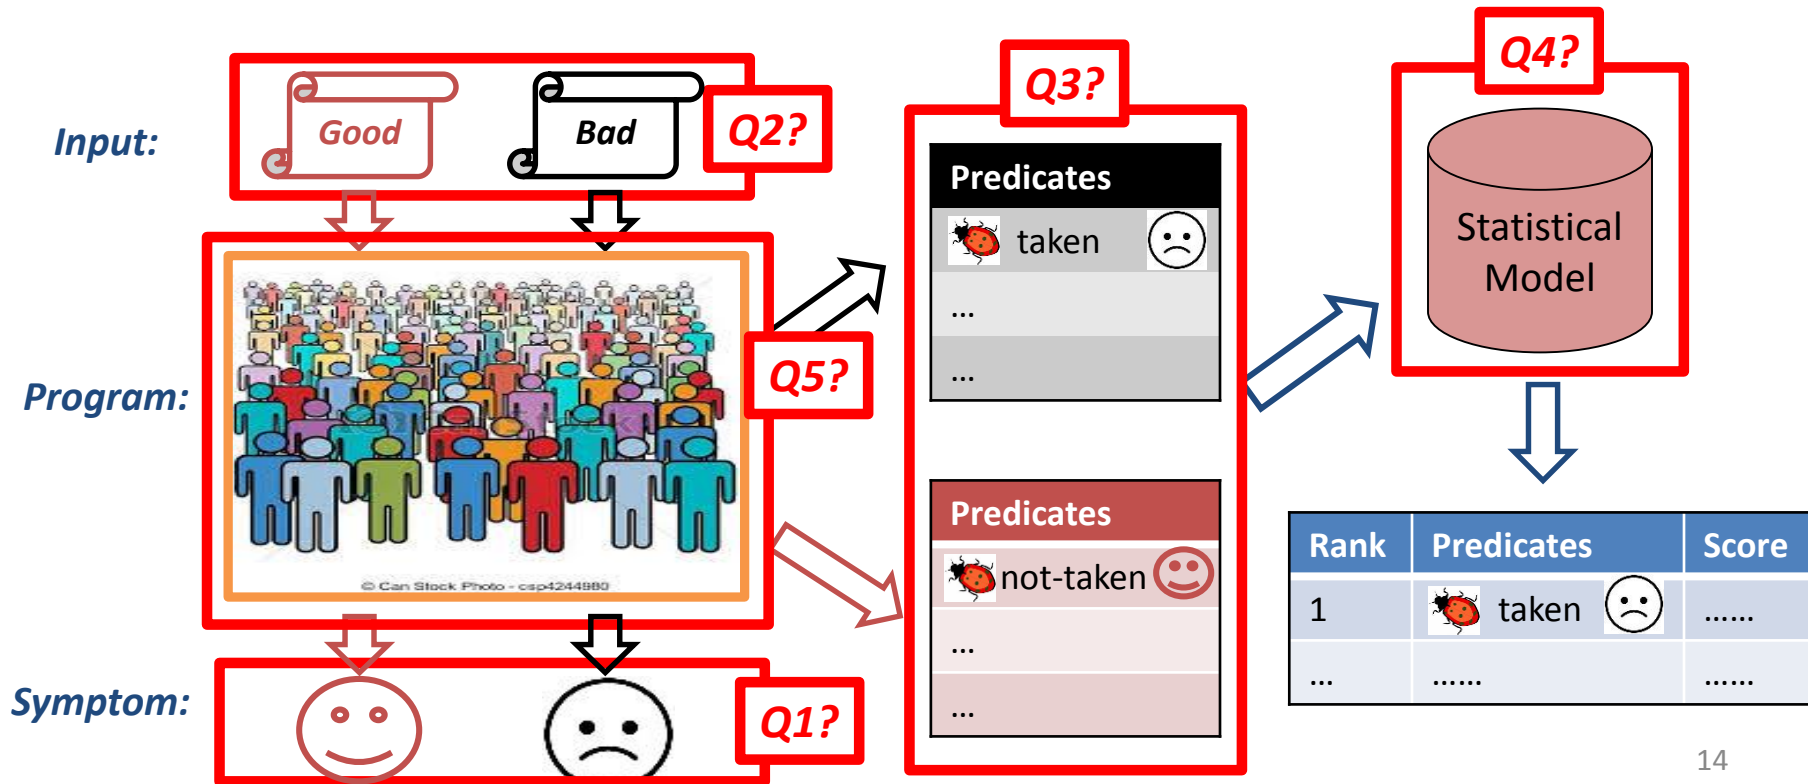
How to Apply?

- Q3: What predicates to collect?
- Q4: What statistical model to use?
- Q5: How to do on-line performance diagnosis?




Contributions (I)

We answer all these questions by studying 65 user-reported performance bugs.



Contributions (II)

- Is it feasible to apply SD to PPs?
 - Easy to tell failure runs based on users' reports (Q1)
 - Inputs are provided during reporting (Q2)
- How to apply SD to PPs?
 - In-house diagnosis
 - 3 predicates (Q3)
 - 2 statistical models (Q4)
 - On-line diagnosis (Q5)
 - Same diagnosis capability with <10% overhead
 - Not sacrifice diagnosis latency (**Unique**)



Outline

- **Overview**
- Is it feasible to apply SD for PPs?
- How to conduct SD for PPs?
 - In-house diagnosis scenario
 - On-line diagnosis scenario
- **Conclusion**

Methodology

- Application and Bug Source

App.	# Bugs ^[1]	# Bug User Perceived
<i>Apache</i>	25	16
<i>Chrome</i>	10	5
<i>GCC</i>	11	9
<i>Mozilla</i>	36	19
<i>MySQL</i>	28	17

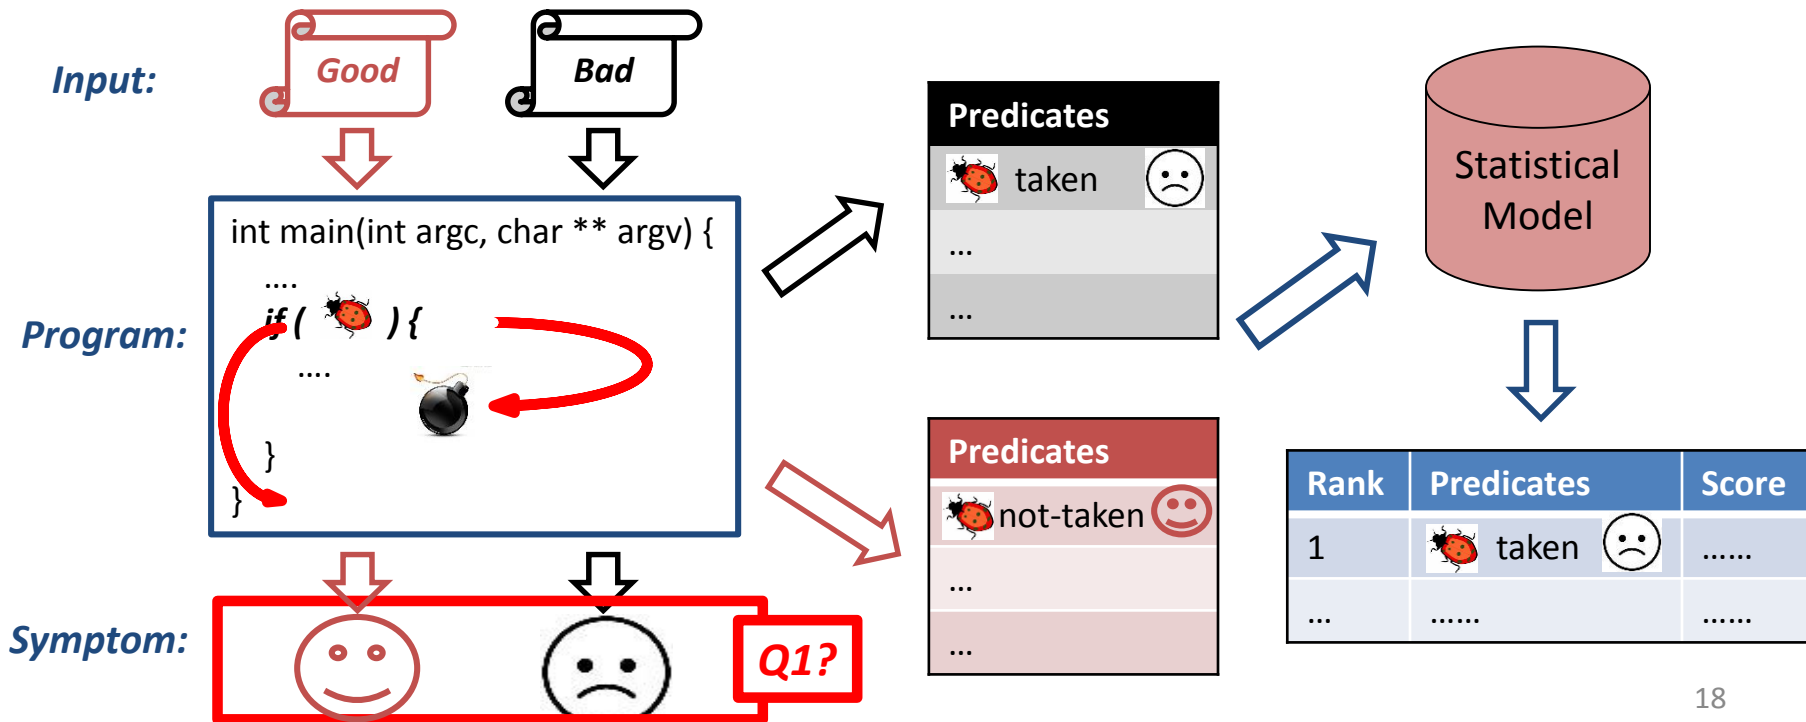
Total: 110

65

[1] Guoliang Jin, Linhai Song, Xiaoming Shi, Joel Scherpelz, and Shan Lu. Understanding and Detecting Real-World Performance Bugs. In PLDI'2012.

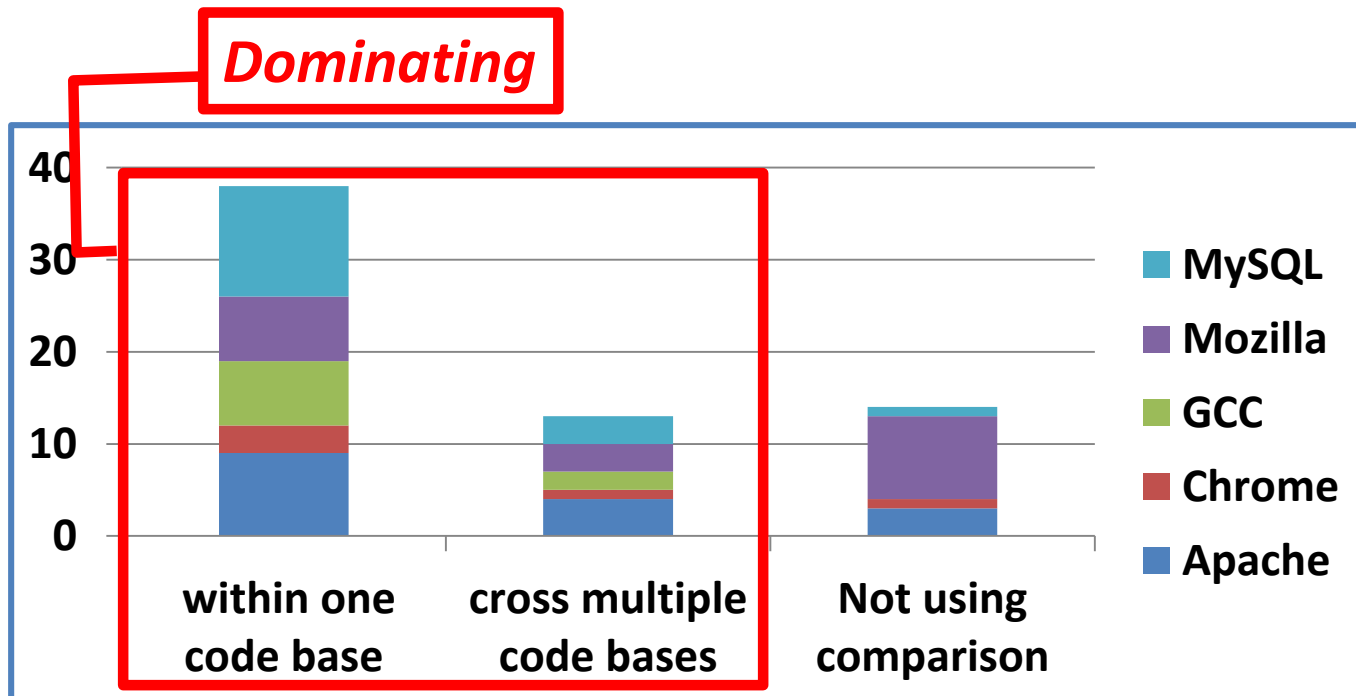
Is It Feasible? (Part I)

- **Q1: How to tell success runs from failure runs?**
 - A large workload? Or inefficient implementation?




Q1: How to Identify Failure Runs?

- The majority is observed through comparison



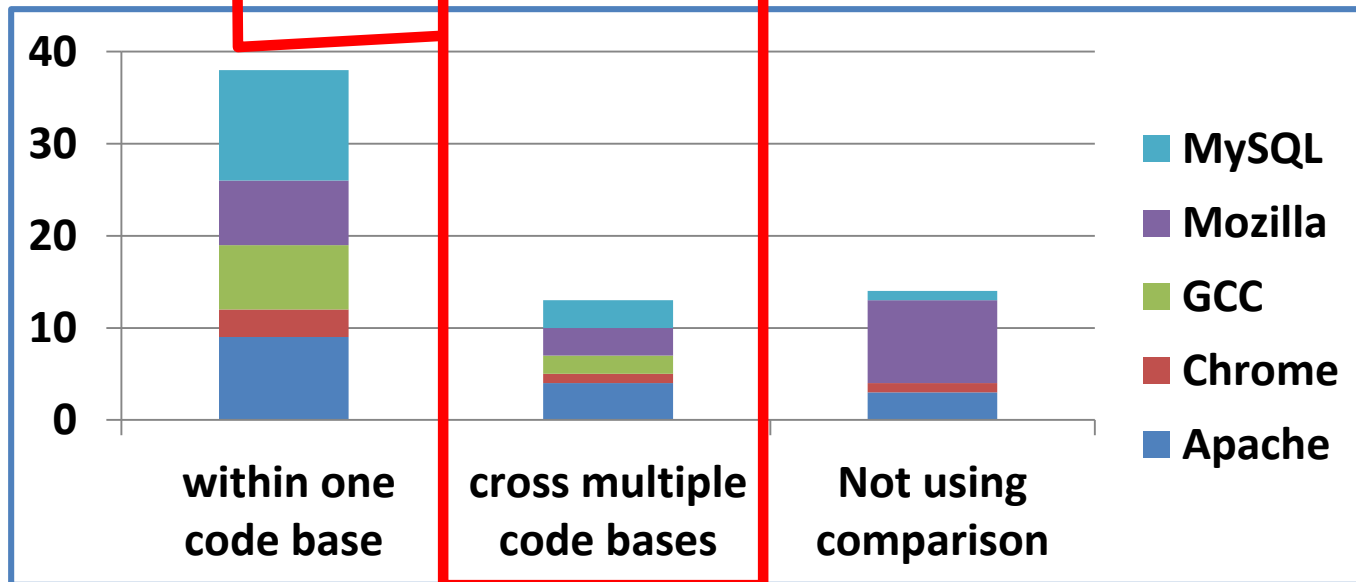
Comparison within One Code Base

- the same input with different *configuration*
- inputs with different *sizes*
- inputs with slightly different *functionality*



Comparison across Multiple Code Bases

- same applications' different *versions*
- different *applications*

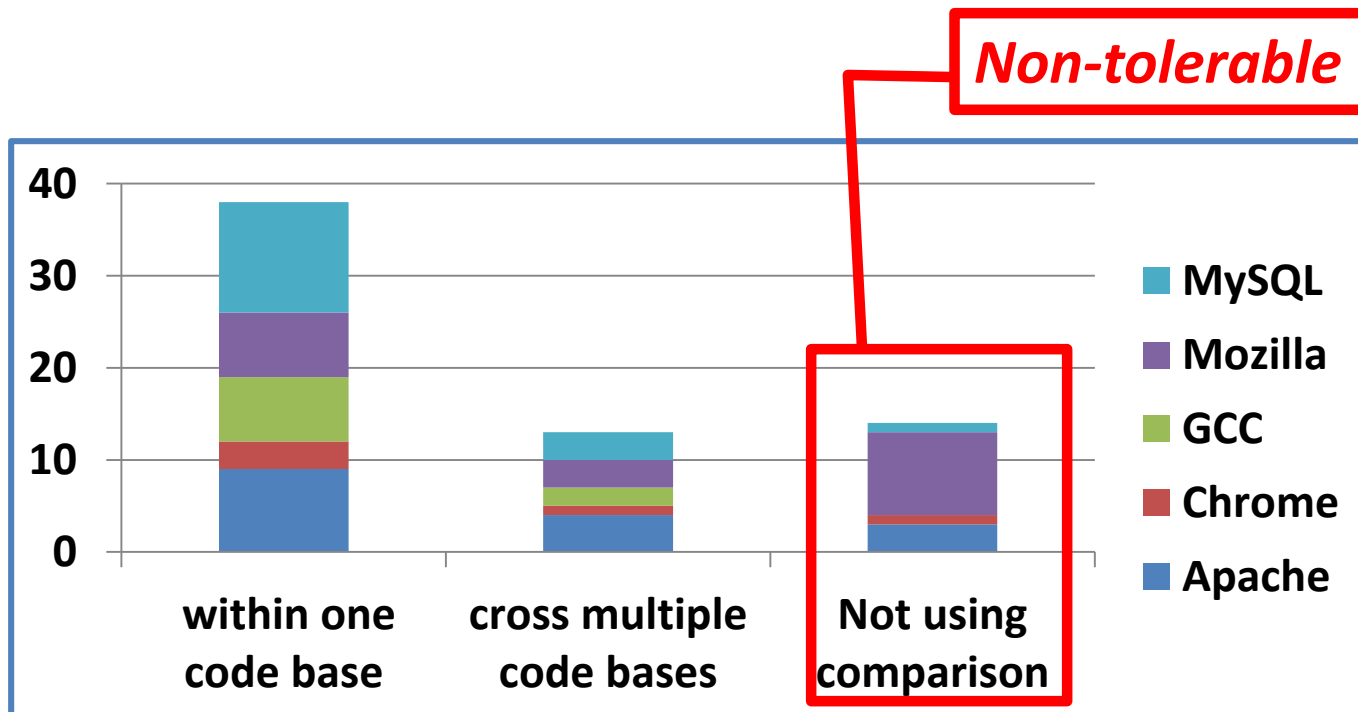


Not Using Comparison

Mozilla#299742: “it frozen the GUI to crawl”

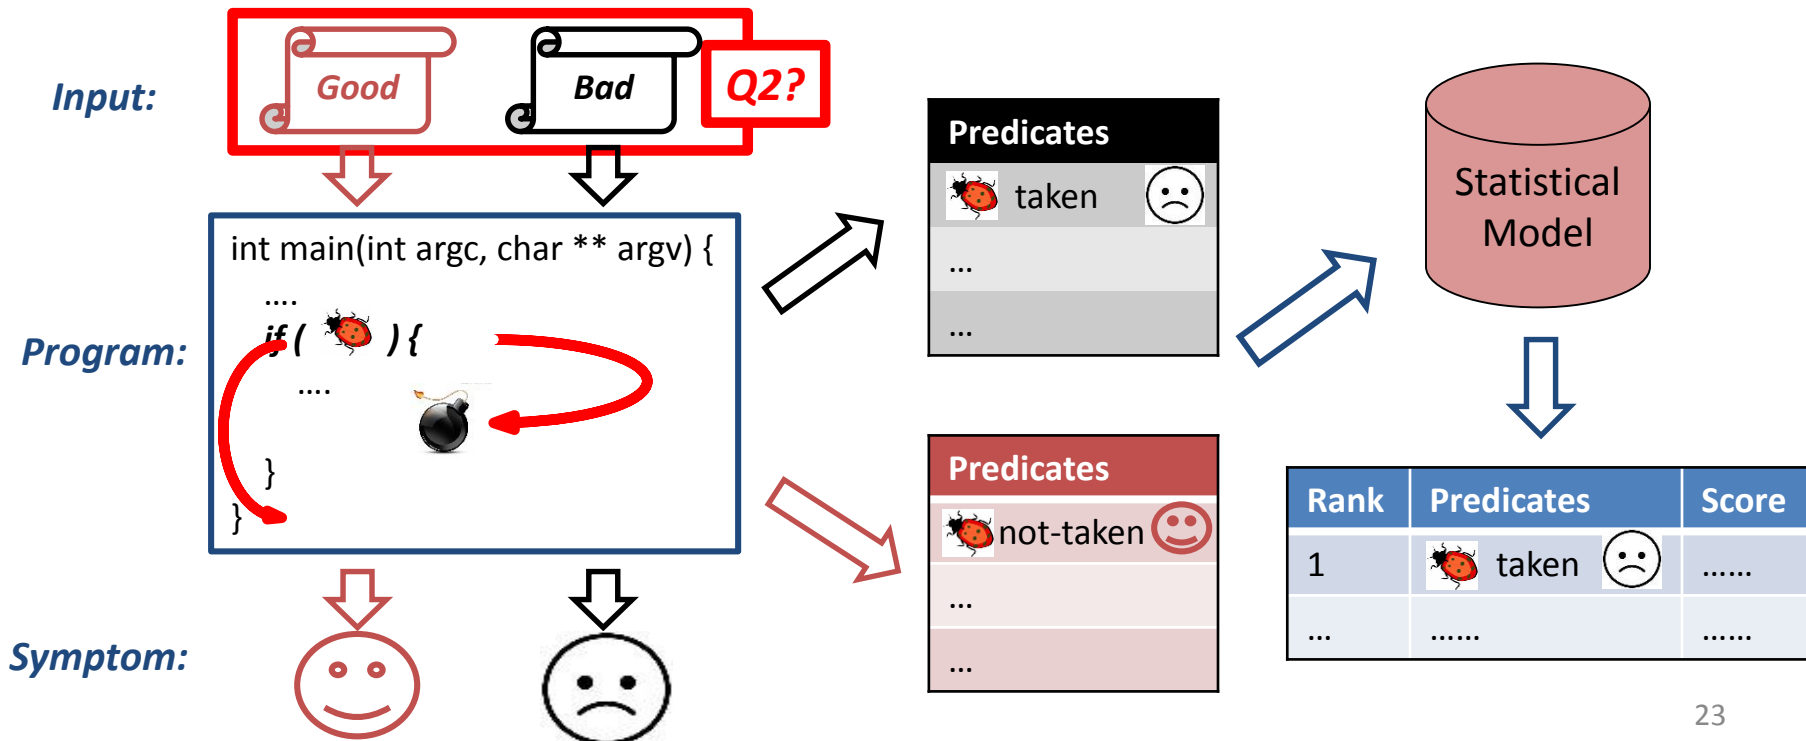
Mozi ***Easy to tell failures from successes!*** with the page”

MySQL46461: “causing the test suit to fail due to timeout”



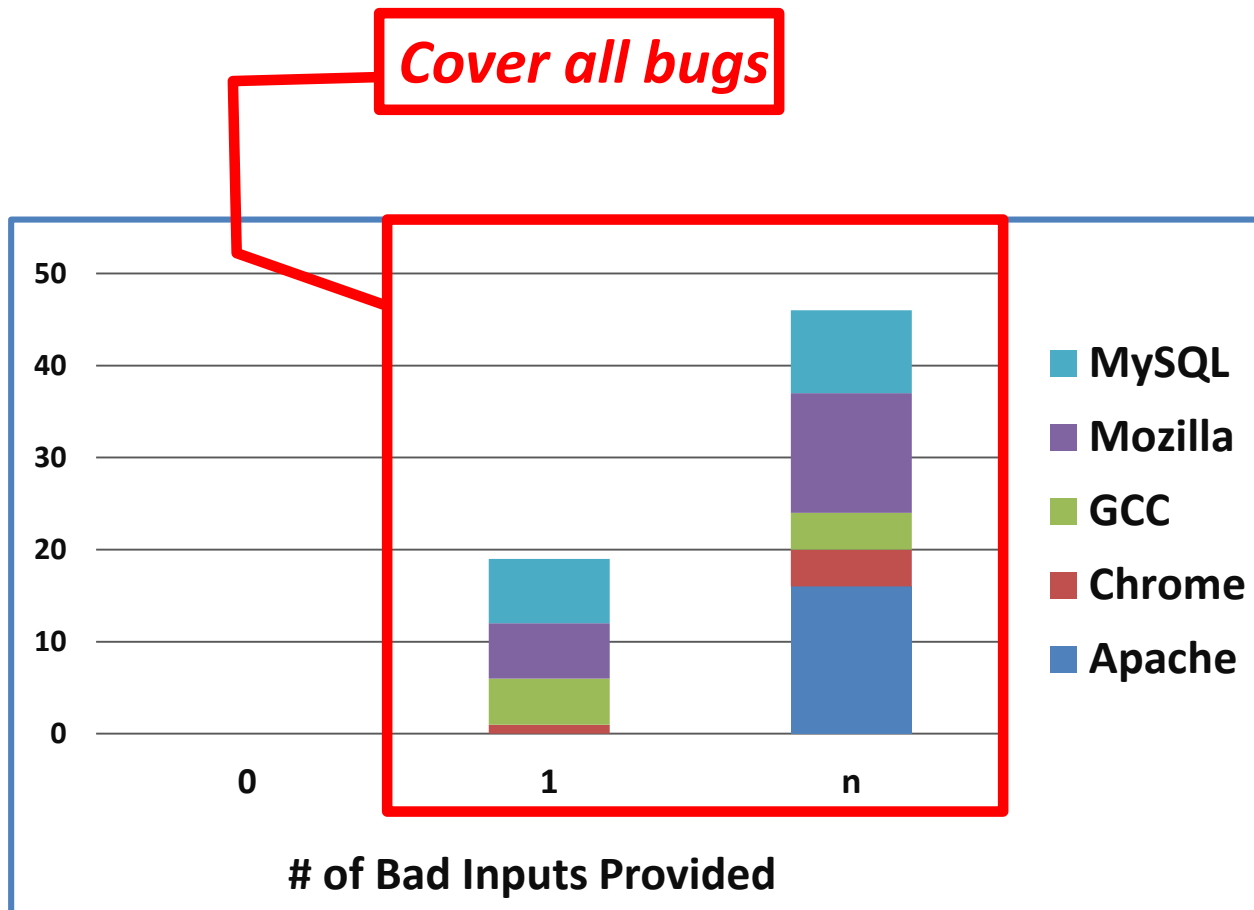
Is It Feasible? (Part II)

- Q1: How to tell success runs from failure runs?
- **Q2: How to obtain good and bad inputs?**



How to Obtain Bad Inputs?

- Bad inputs are provided in all bug reports




How to Obtain Good Inputs? (I)

- 60% contain good inputs



How to Obtain Good Inputs? (II)

Easy to design



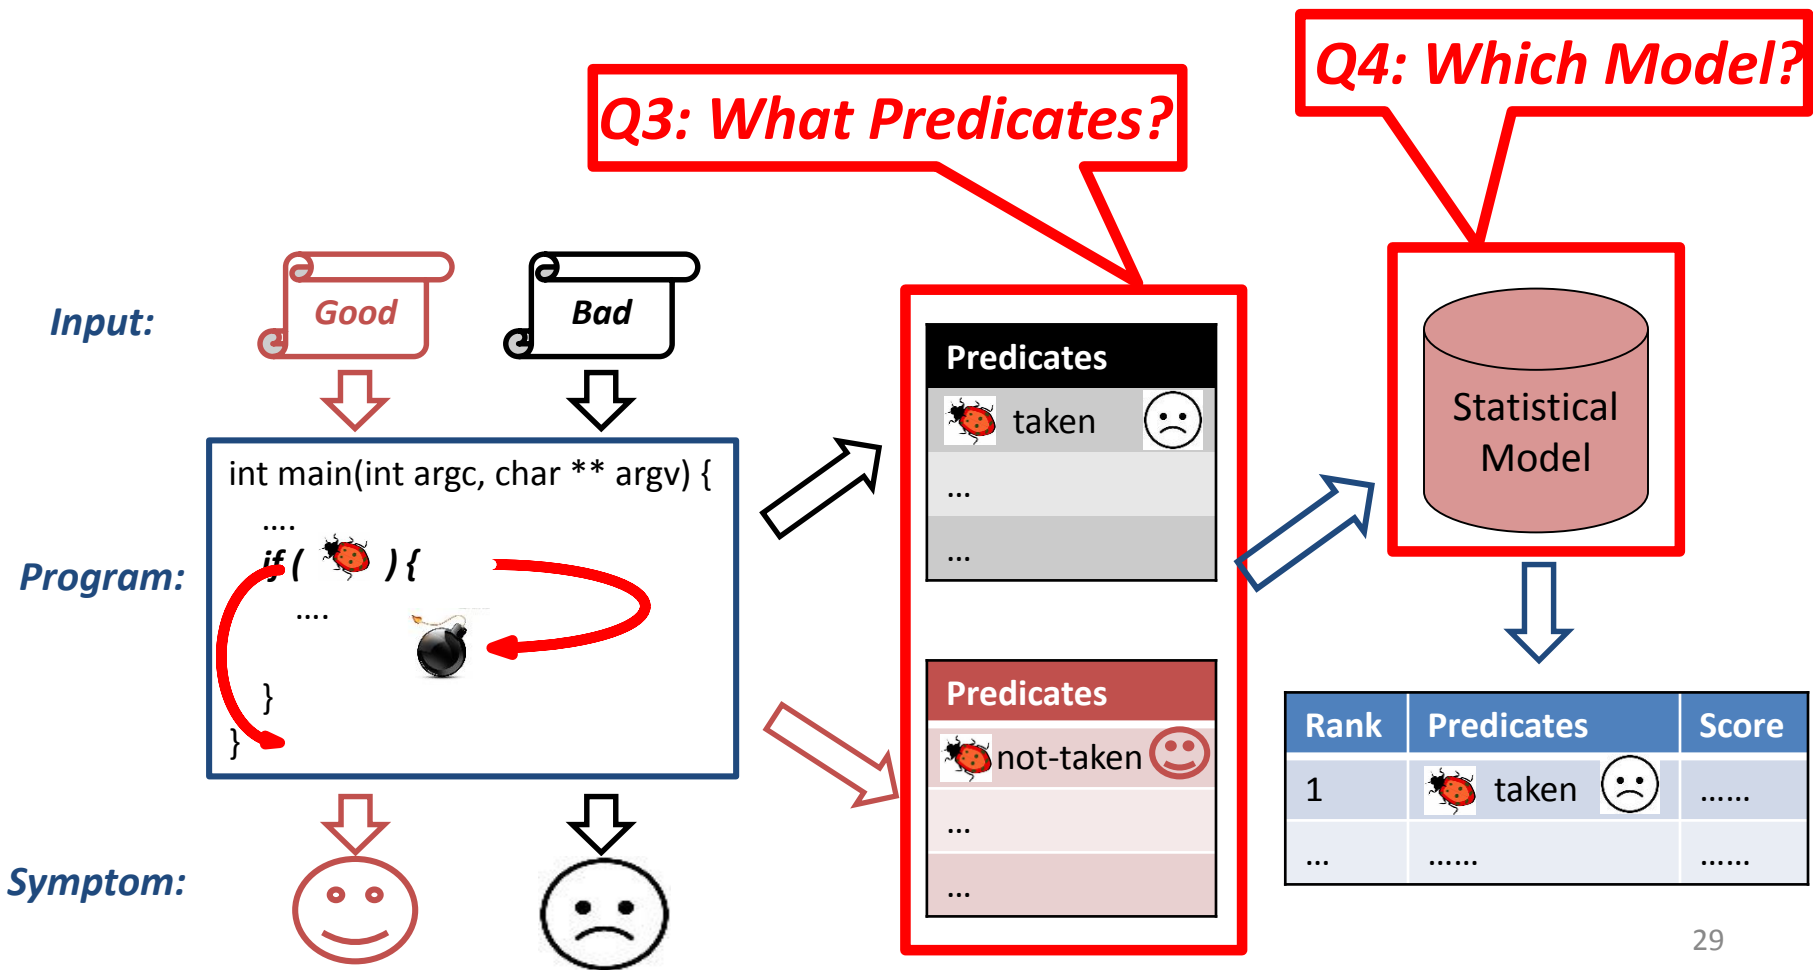
Other Findings and Implications

- Compared with functional bugs
 - More PPs observed through comparison
 - More PPs reported with good inputs
- Implications for SD
 - Easy to tell success runs from failure runs
 - Similar good inputs are provided
 - SD is a nature fit for PPs

Outline

- Overview
- Is it feasible to apply SD for PPs?
- **How to conduct SD for PPs?**
 - In-house diagnosis scenario
 - On-line diagnosis scenario
- Conclusion

In-house Statistical Debugging




Design

- Study widely used predicates and models

```
int x, y;  
...  
x = ...;
```

```
n=fopen(...);
```

```
if (p) ...  
else ...
```



Experimental Methodology

- Experimental setting
 - 10 success runs vs. 10 failure runs
 - 20 unique inputs
- Techniques under comparison
 - CBI[2, 3] for C programs
 - Pin for C++ programs
 - Compared with profiling results from OProfile

[2] Ben Liblit, Alex Aiken, Alice X Zhen, and Michael I Jordan. Bug Isolation via Remote Program Sampling. In PLDI'2003.

[3] Ben Liblit, Mayur Naik, Alice X Zheng, Alex Aiken, and Michael I Jordan. Scalable Statistical Bug Isolation. In PLDI'2005.

Experimental Results

BugID	Basic Model			ΔLDA			Profiler
	Branch	Return	S-pair	Branch	Return	S-pair	
Mozilla258793	√1	-	/	-	-	/	-
Mozilla299742	√1	-	/	-	-	/	-
Mozilla347306	-	-	-	√1	√1	√1	√1
Mozilla416628	-	-	-	√1	-	√1	√1
MySQL15811	-	-	/	√1	√1	/	√1
MySQL26527	√1	-	/	-	-	/	-
MySQL27287	-	-	/	√1	-	/	√1
MySQL40337	√1	-	/	-	-	/	-
MySQL42649	√1	-	/	-	-	/	-
MySQL44723	√1	-	/	-	-	/	-
Apache3278	√1	√1	√1	-	-	-	-
Apache34464	-	-	-	√3	√1	-	√5
...

Experimental Results

BugID	Basic Model			ΔLDA			Profiler
	Branch	Return	S-pair	Branch	Return	S-pair	
Mozilla258793	√1	-	/	-	-	/	-
Mozilla299742	√1	-	-	-	-	/	-
Mozilla347306	-	-	-	√1	√1	√1	√1
Mozilla416628	-	-	-	√1	-	√1	√1
MySQL15811	-	-	/	√1	√1	/	√1
MySQL26527	√1	-	/	-	-	/	-
MySQL27287	-	-	/	√1	-	/	√1
MySQL40337	√1	-	/	-	-	/	-
MySQL42649	√1	-	/	-	-	/	-
MySQL44723	√1	-	/	-	-	/	-
Apache3278	√1	√1	√1	-	-	-	-
Apache34464	-	-	-	√3	√1	-	√5
...

Most Useful

Experimental Results

	Basic Model			Δ LDA			Profiler
BugID	Branch	Return	S-pair	Branch	Return	S-pair	
Mozilla258793	√1	-	/	-	-	/	-
Mozilla299742	√1	-	/	-	-	/	-
Mozilla347306	-	-	-	√1	√1	√1	√1
Mozilla416628	-	-	-	√1	-	√1	√1
MySQL15811	-	-	/	√1	√1	/	√1
MySQL26527	√1	-	/	-	-	/	-
MySQL27287	-	-	/	√1	-	/	√1
MySQL40337	√1	-	/	-	-	/	-
MySQL42649	√1	-	/	-	-	/	-
MySQL44723	√1	-	/	-	-	/	-
Apache3278	√1	√1	√1	-	-	-	-
Apache34464	-	-	-	√3	√1	-	√5
...

Experimental Results

BugID	Basic Model			ΔLDA			Profiler
	Branch	Return	S-pair	Branch	Return	S-pair	
Mozilla258793	√1	-	/	-	-	/	-
Mozilla299742	√1	-	/	-	-	/	-
Mozilla347306	-	-	-	√1	√1	√1	√1
Mozilla416628	-	-	-	√1	-	√1	√1
MySQL15811	-	-	/	√1	√1	/	√1
MySQL26527	√1	-	/	-	-	/	-
MySQL27287	-	-	/	√1	-	/	√1
MySQL40337	√1	-	/	-	-	/	-
MySQL42649	√1	-	/	-	-	/	-
MySQL44723	√1	-	/	-	-	/	-
Apache3278	√1	√1	√1	-	-	-	-
Apache34464	-	-	-	√3	√1	-	√5
...

Experimental Results

BugID	Basic Model			ΔLDA			Profiler
	Branch	Return	S-pair	Branch	Return	S-pair	
Mozilla258793	√1	-	/	-	-	/	-
Mozilla299742	√1	-	/	-	-	/	-
Mozilla347306	-	-	-	√1	√1	√1	√1
Mozilla416628	-	-	-	√1	-	√1	√1
MySQL15811	-	-	/	√1	√1	/	√1
MySQL26527	√1	-	/	-	-	/	-
MySQL27287	-	-	/	√1	-	/	√1
MySQL40337	√1	-	/	-	-	/	-
MySQL42649	√1	-	/	-	-	/	-
MySQL44723	√1	-	/	-	-	/	-
Apache3278	√1	√1	√1	-	-	-	-
Apache34464	-	-	-	√3	√1	-	√5
...

Experimental Results

	Basic Model			ΔLDA			Profiler
BugID	Branch	Return	S-pair	Branch	Return	S-pair	
Mozilla258793	v1	-	/	-	-	/	-
Mozilla299742	v1	-	/	-	-	/	-
Mozilla347306	-	-	-	v1	v1	v1	v1
Mozilla416628	-	-	-	v1	-	v1	v1
MySQL15811	-	-	/	v1	v1	/	v1
MySQL26527	v1	-	/	-	-	/	-
MySQL27287	-	-	/	v1	-	/	v1
MySQL40337	v1	-	/	-	-	/	-
MySQL42649	v1	-	/	-	-	/	-
MySQL44723	v1	-	/	-	-	/	-
Apache3278	v1	v1	v1	-	-	-	-
Apache34464	-	-	-	v3	v1	-	v5
...

Experimental Results

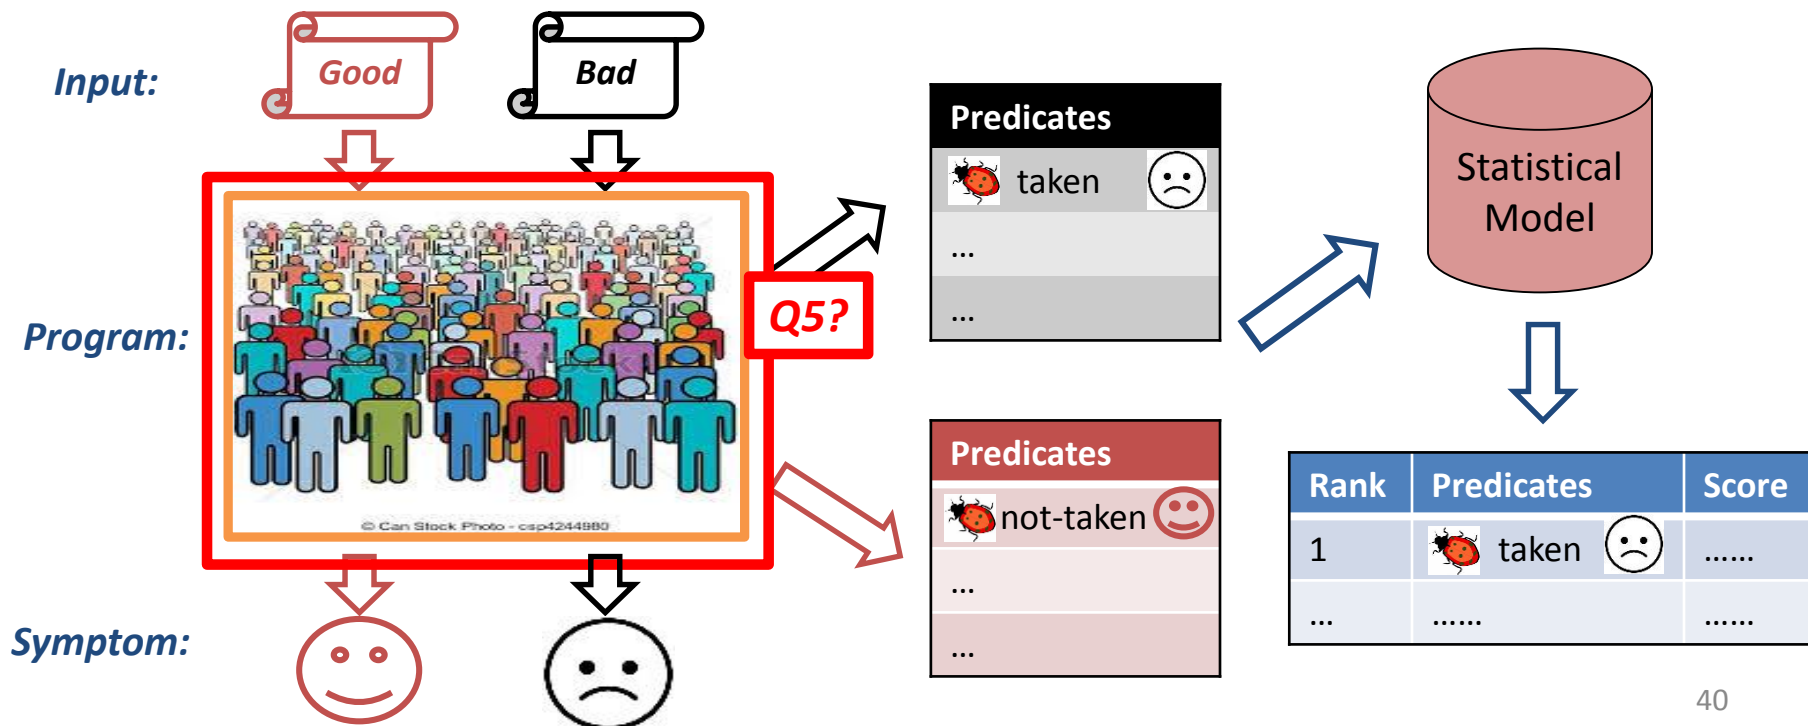
BugID	Basic Model			Branch	ΔLDA		Profiler
	Branch	Return	S-pair		Return	S-pair	
Mozilla258793	√1	-	/	-	-	/	-
Mozilla299742	√1	-	/	-	-	/	-
Mozilla347306	-	-	-	√1	√1	√1	√1
Mozilla416628	-	-	-	√1	-	√1	√1
MySQL15811	-	-	/	√1	√1	/	√1
MySQL26527	√1	-	/	-	-	/	-
MySQL27287	-	-	/	√1	-	/	√1
MySQL40337	√1	-	/	-	-	/	-
MySQL42649	√1	-	/	-	-	/	-
MySQL44723	√1	-	/	-	-	/	-
Apache3278	√1	√1	√1	-	-	-	-
Apache34464	-	-	-	√3	√1	-	√5
...

Outline

- Overview
- Is it feasible to apply SD for PPs?
- **How to conduct SD for PPs?**
 - In-house diagnosis scenario
 - On-line diagnosis scenario
- Conclusion

On-line Statistical Debugging

- Q5: How to do on-line performance diagnosis?
 - Less information from one single run
 - Diagnosis capability relies on multiple runs info.



Experimental Methodology

- Tool design
 - CBI in sampling mode for C benchmarks
 - LBR for C++ benchmarks
- Benchmarks
 - Reuse benchmarks from in-house experiments
 - Most effective predicate and model
- Experiment design
 - Default sampling rate is roughly 1/10000
 - 1000 success runs and 1000 failure runs

Experimental Results

BugID	Diagnosis Capability	Run-time Overhead	Requested Failure Runs
Mozilla258793	√1	2.39%	100
Mozilla299742	√1	4.27%	500
Mozilla347306	√1	1.42%	10
Mozilla416628	√1	2.03%	10
MySQL15811	√1	2.25%	10
MySQL26527	√1	6.05%	500
MySQL27287	√1	3.02%	10
MySQL40337	√1	2.69%	100
MySQL42649	√2	6.10%	500
MySQL44723	√1	3.16%	100
Apache3278	-	0.23%	>1000
Apache34464	√3	0.18%	10
...

Experimental Results

BugID	Diagnosis Capability	Run-time Overhead	Requested Failure Runs
	√1	2.39%	
	√1	4.27%	500
Mozilla37500	√1	1.42%	10
Mozilla416628	√1	2.03%	10
MySQL15811	√1	2.25%	10
MySQL26527	√1	6.05%	500
MySQL27287	√1	3.02%	10
MySQL40337	√1	2.69%	100
MySQL42649	√2	6.10%	500
MySQL44723	√1	3.16%	100
Apache3278	-	0.23%	>1000
Apache34464	√3	0.18%	10
...

Same Diagnosis Capability

< 10%

Experimental Results

BugID	Diagnosis Capability	Run-time Overhead	Requested Failure Runs
Mozilla258793	√1	2.39%	
Mozilla299742	√1	4.27%	
Mozilla347306	√1	1.42%	
Mozilla416628	√1	10 X 10000 ??	
MySQL15811	√1	2.25%	
MySQL26527	√1	6.05%	
MySQL27287	√1	3.02%	
MySQL40337	√1	2.69%	
MySQL42649	√2	6.10%	
MySQL44723	√1	3.16%	
Apache3278	-	0.23%	
Apache34464	√3	0.18%	
...

Experimental Results

BugID	Diagnosis Capability	Run-time Overhead	Requested Failure Runs
Mozilla258793	√1	2.39%	100
Mozilla299742	√1	4.27%	500
Mozilla347306	√1	1.42%	10
Mozilla416628	√1	2.03%	10
MySQL15811	√1	2.25%	10
MySQL26527	√1	6.05%	500
MySQL27287	√1	3.02%	10
MySQL40337	√1	2.69%	100
MySQL42649	√2	6.10%	500
MySQL44723	√1	3.16%	100
Apache3278	-	0.23%	>1000
Apache34464	√3	0.18%	10
...

Conclusion and Future Works

- Study diagnosis process for PPs
 - Statistical debugging is a natural fit
- Study statistical debugging on PPs
 - Branch predicates + two statistical models
- Future works
 - Analyze inefficient loops
 - Provide detailed fix hints

Thanks a lot!

