

Performance and Concurrency Bug Detection Tools for Java Programs

Shan Lu

University of Chicago

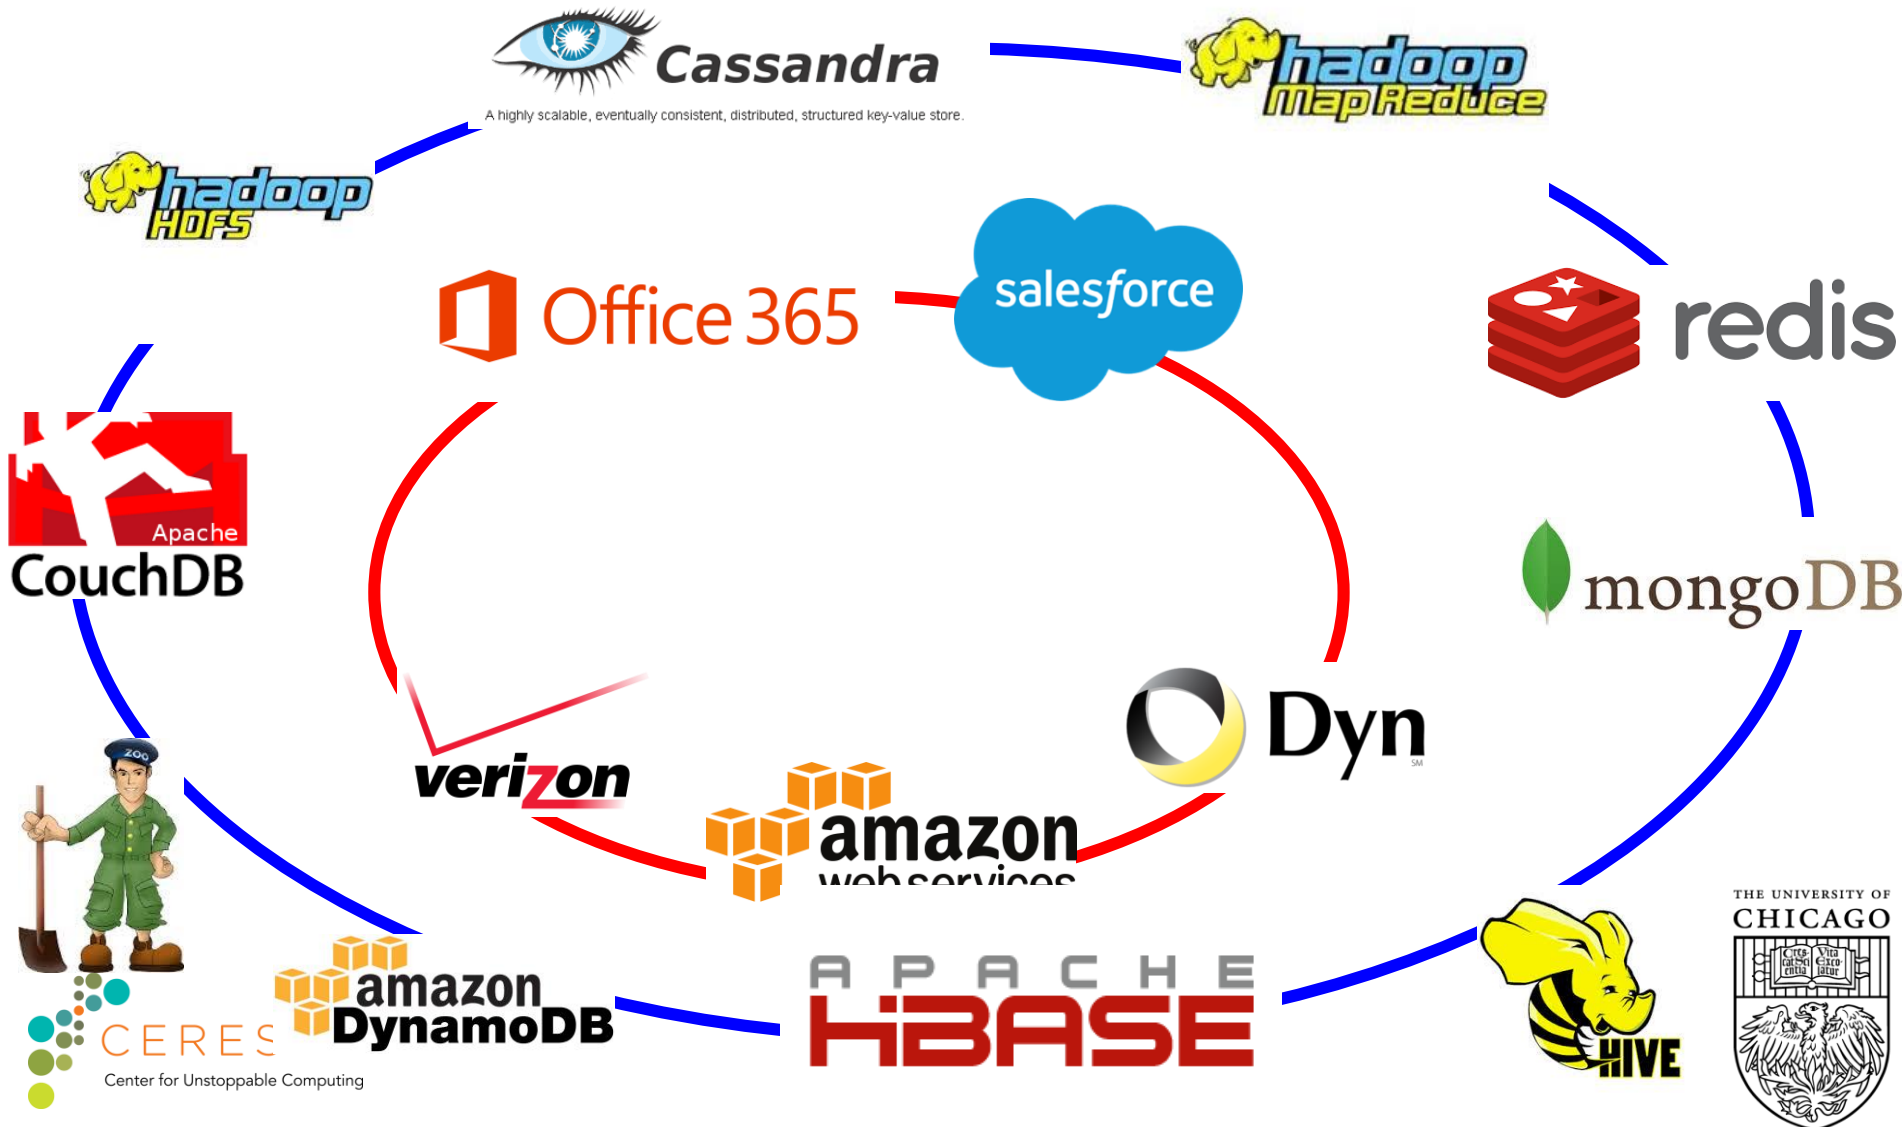


Fighting software bugs is crucial

- Software is everywhere
 - http://en.wikipedia.org/wiki/List_of_software_bugs
- Software bugs are widespread and costly
 - Lead to 40% system down time [Blueprints 2000]
 - Cost 312 Billion lost per year [Cambridge 2013]



Fighting bugs in cloud systems ...



... is crucial



Cassandra

A highly scalable, eventually consistent, distributed, structured key-value store.



The 10 Biggest Cloud Outages of 2016

www.crn.com/slide-shows/cloud/.../the-10-bigges

Dec 29, 2016 - The biggest cloud outages of 2016 incl large, are increasingly vulnerable from downtime.



redis



mongoDB



THE UNIVERSITY OF
CHICAGO



**APACHE
HBASE**



Different aspects of fighting bugs

In-house
bug detection



In-field
failure recovery



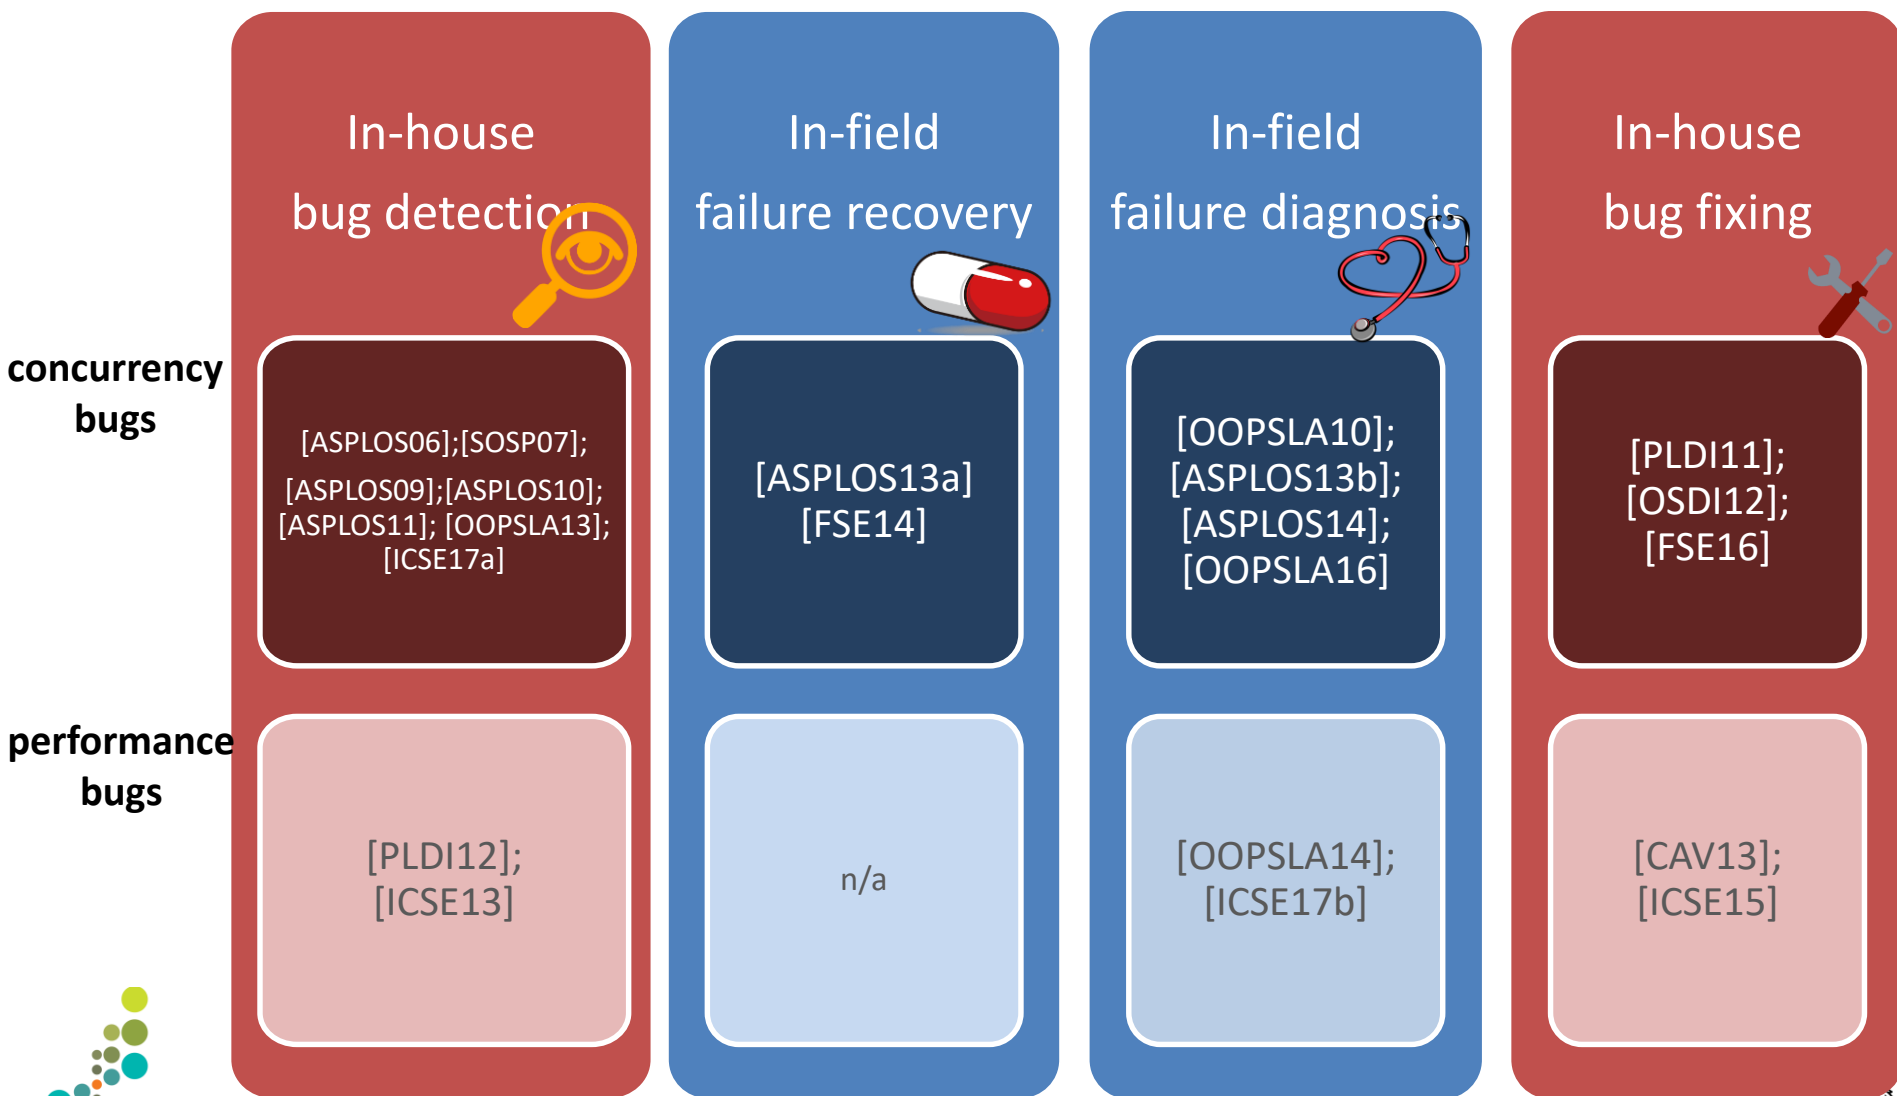
In-field
failure diagnosis



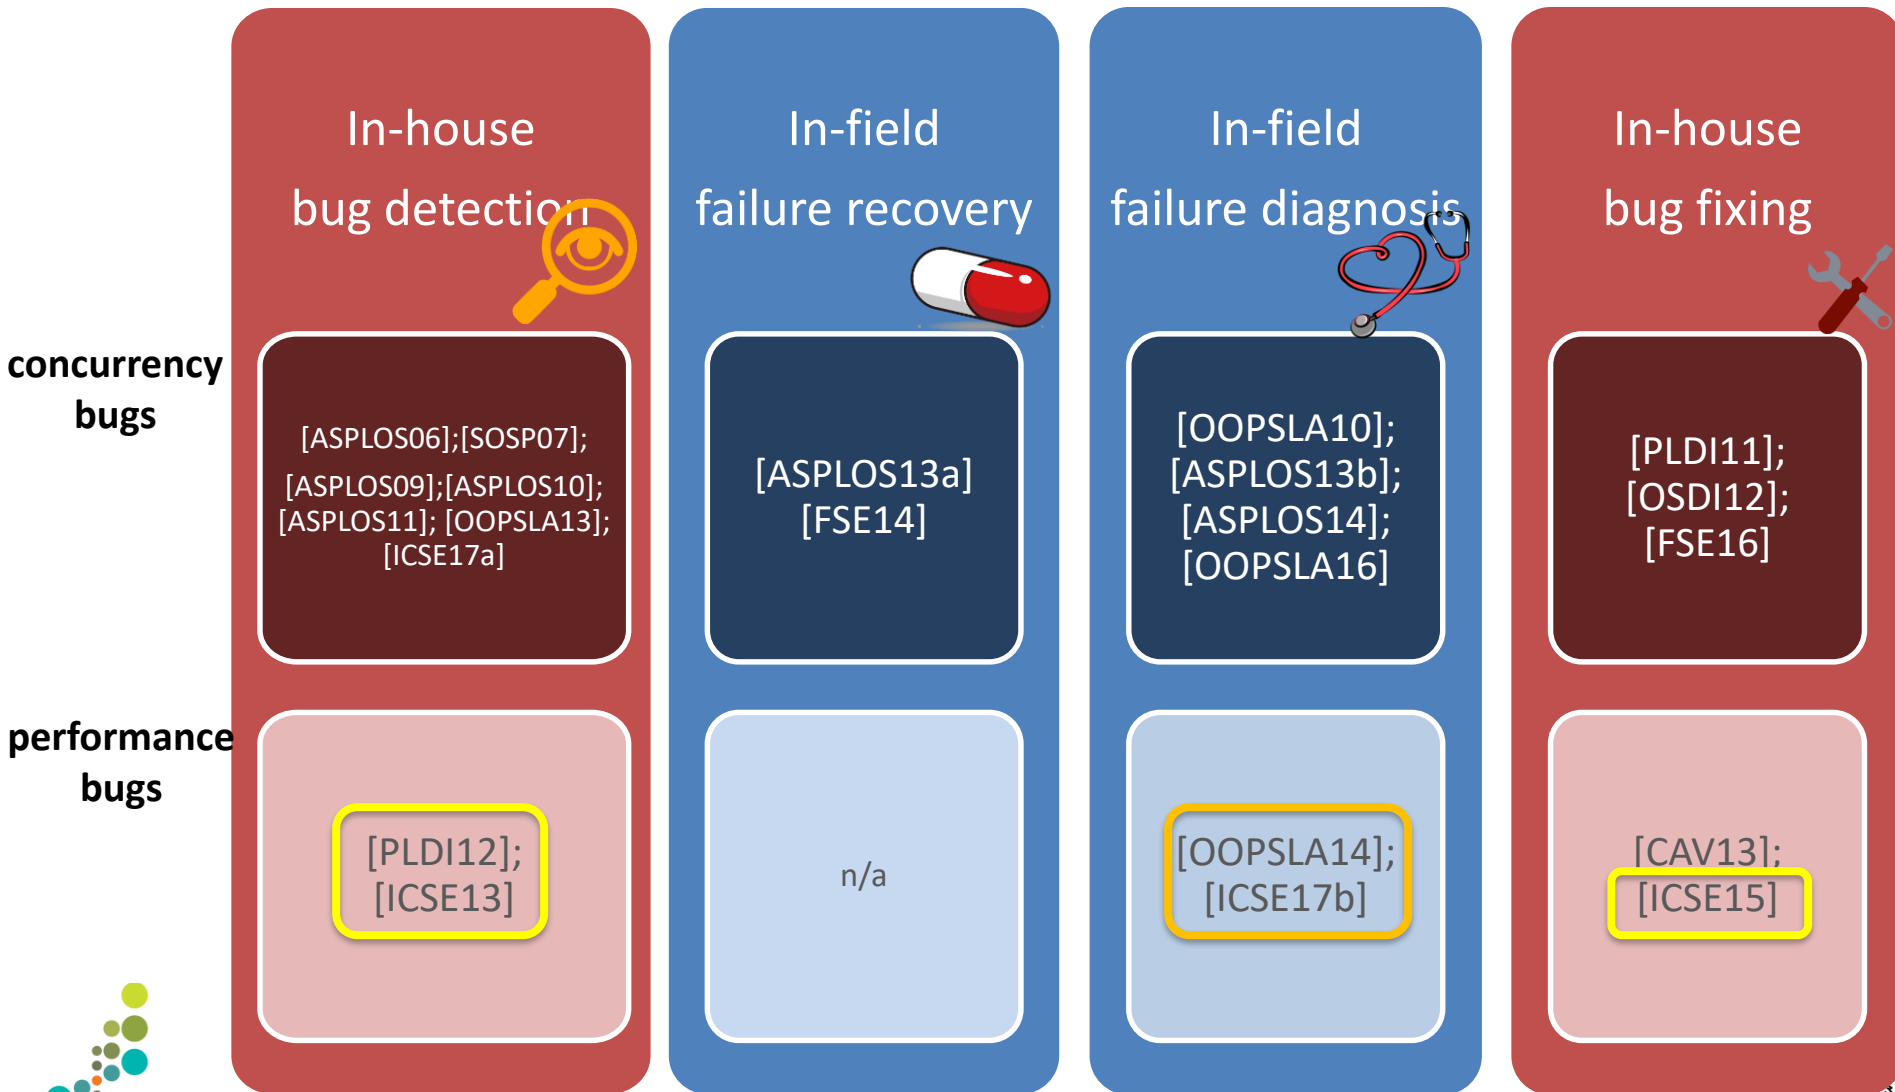
In-house
bug fixing



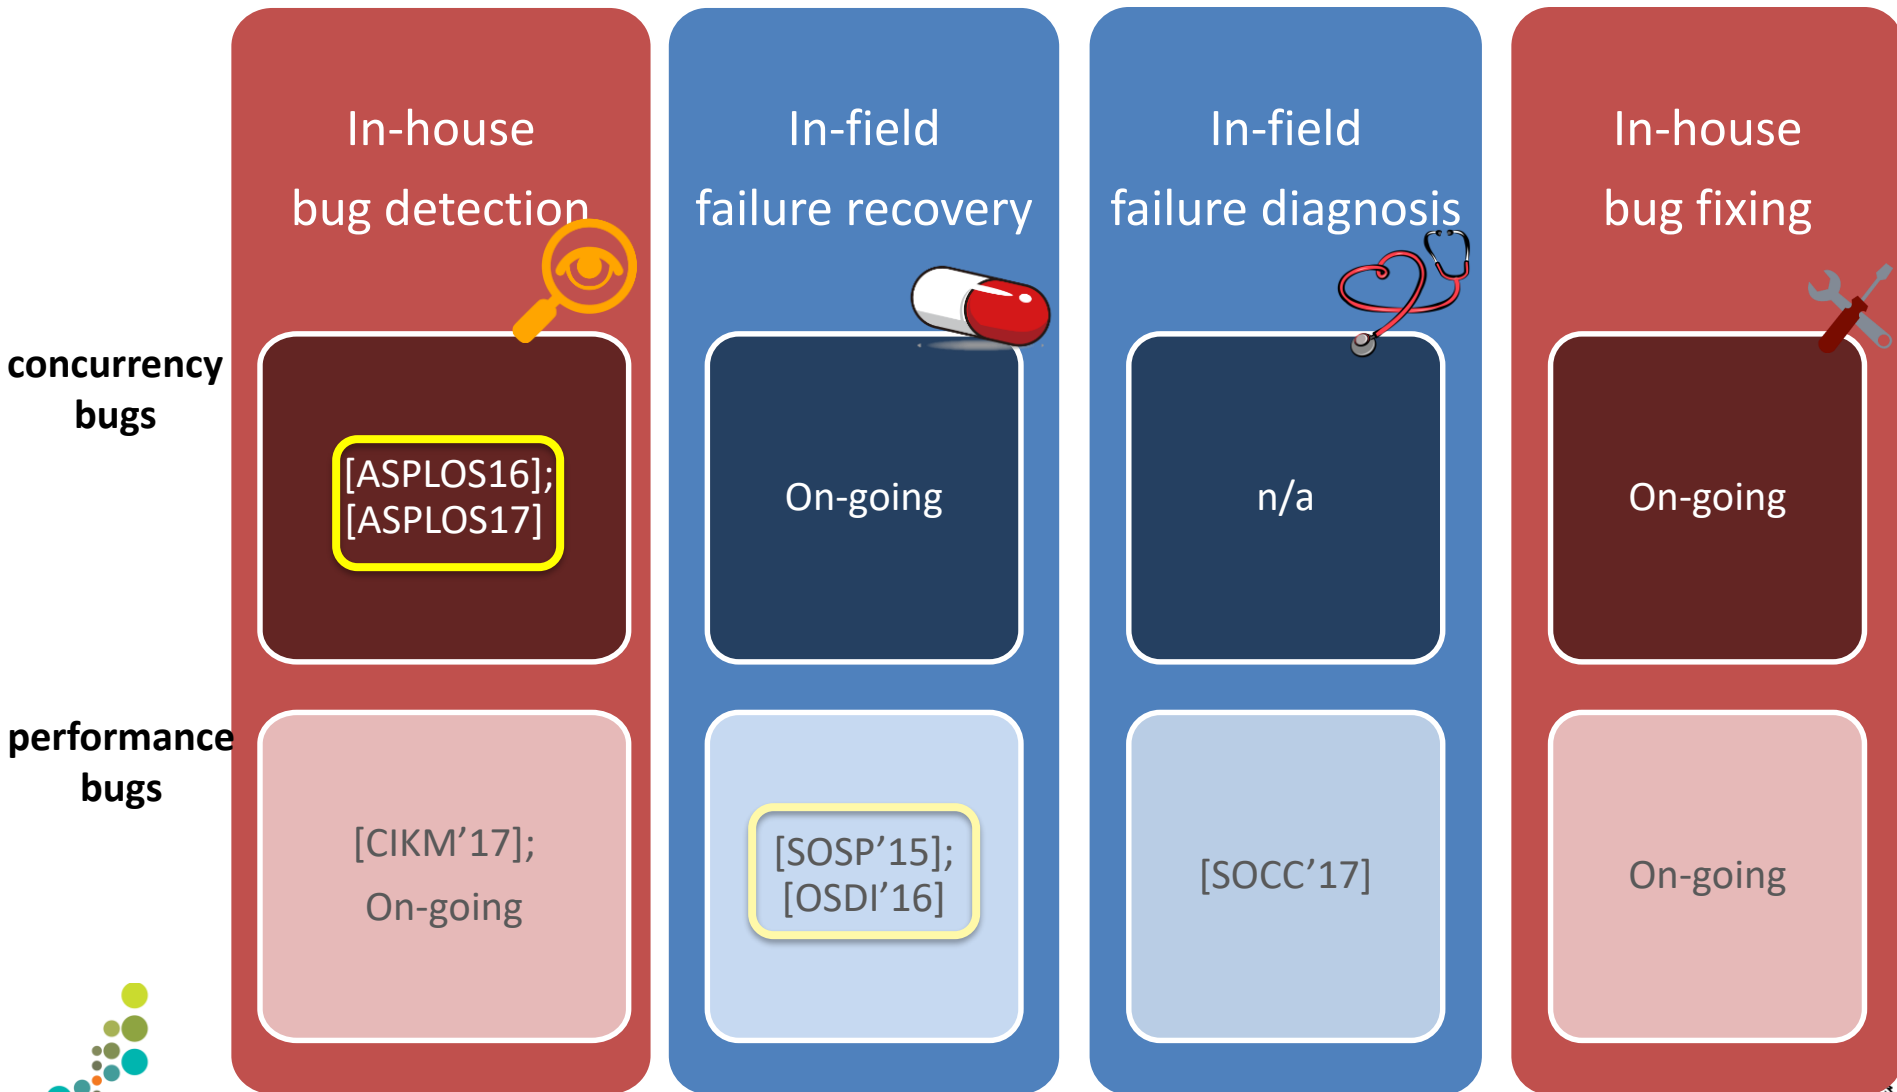
Work from my group (local systems)



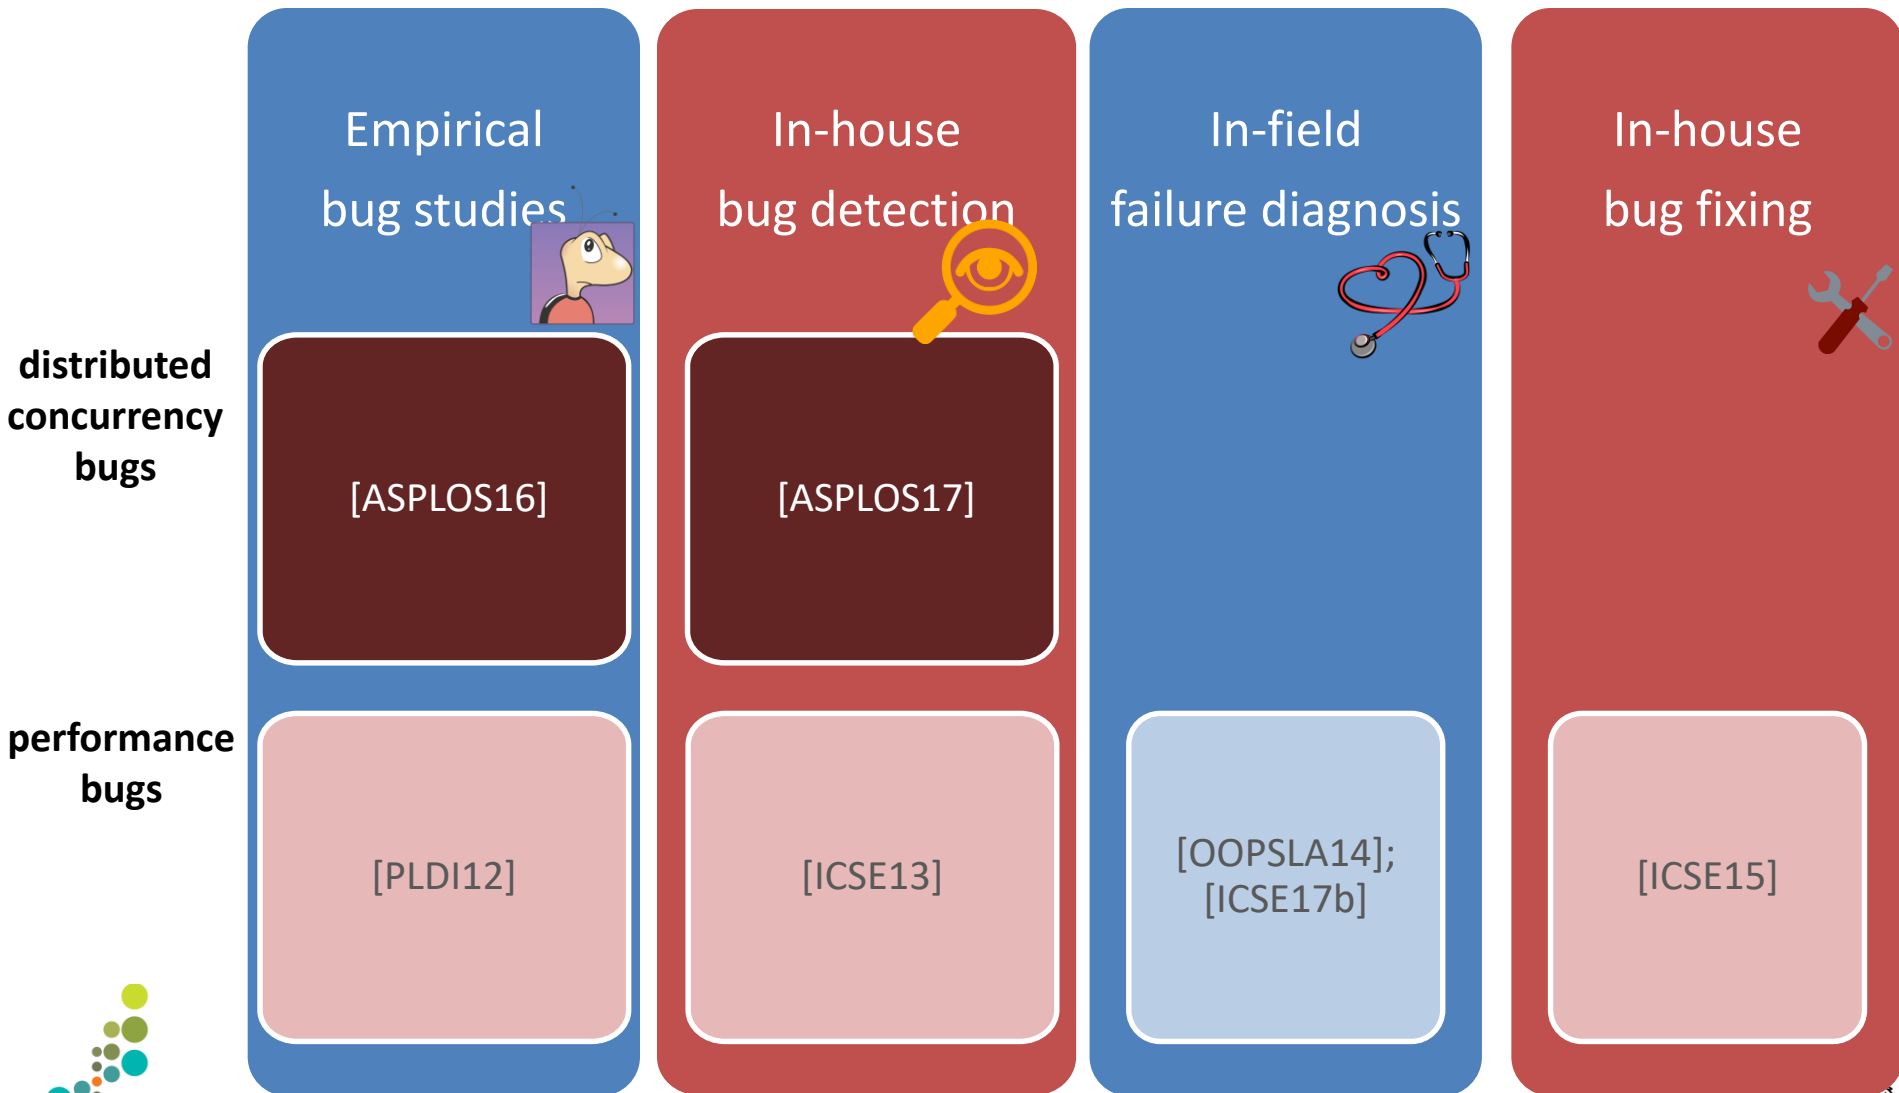
Work from my group (local systems)



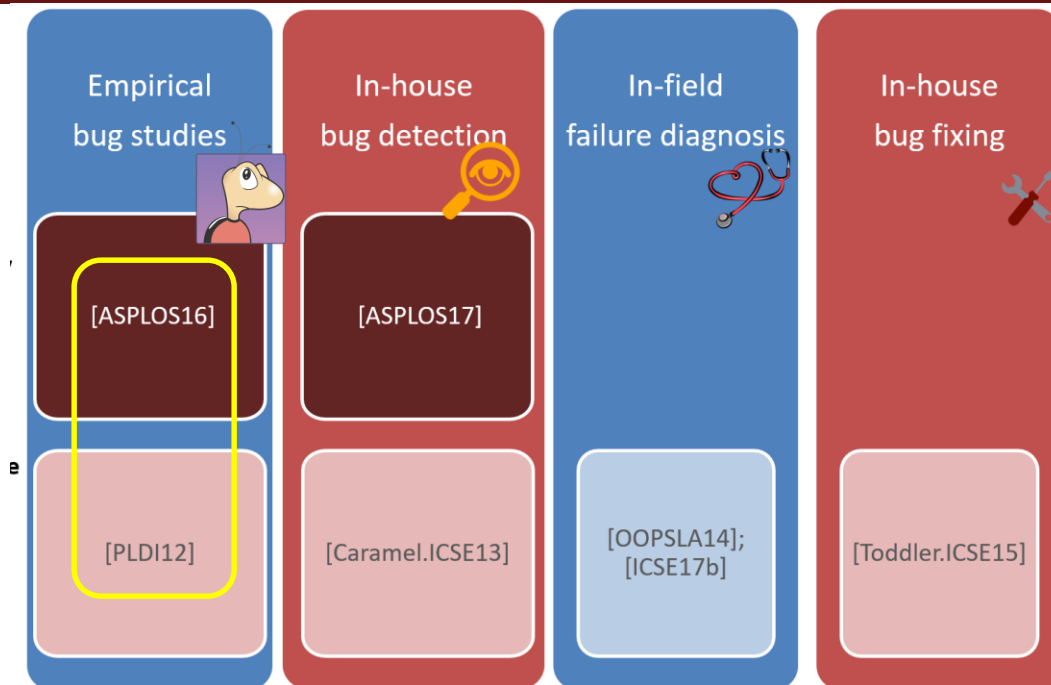
Work from my group (cloud systems)



Our bug-tools for Java programs



Empirical Bug Studies



Performance bug studies



- Why did we do this study?

Benchmark Suite

Application	Software	Language	MLOC	Bug DB History	Tags	# Bugs
Apache	Command-line Server + Library	C/Java	0.45	13 y	N/A	25
Chrome	GUI Application	C/C++	14.0	4 y	N/A	10
GCC	Compiler	C/C++	5.7	13 y	Compile-time-hog	10
Mozilla	GUI Application	C++/JS	4.7	14 y	perf	36
MySQL	Server Software	C/C++/C#	1.3	10 y	S5	28

Ant, Tomcat

What/How did we study?

- Read on-line discussion
- Check patches
- Check source code
- Bug root causes
- Bug locations
- Bug triggering conditions
- Bug fix strategies
- Bug symptoms
- Bug-related inputs

All Manual

What could have been better?

- Read on-line discussion
- Check patches
- Check source code
- Bug root causes
- Bug locations
- Bug triggering conditions
- Bug fix strategies
- Bug symptoms
- Bug-related inputs

All Manual

Distributed concurrency bug studies

- Why did we do this study?



• TaxDC: A Comprehensive Taxonomy of Non-Deterministic Concurrency Bugs in Cloud Distributed Systems [[ASPLOS '16](#)]



Benchmark Suite

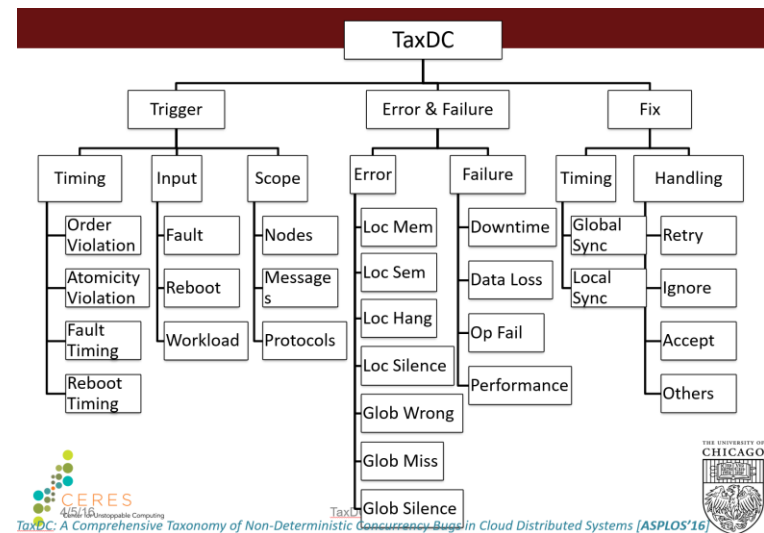
Application	Software Type		MLOC	Bug DB History	# Bugs
<i>Cassandra</i>	Distributed Key-Value Store	Java	0.06	9 y	19
<i>Hadoop</i>	Distributed computing	Java	1.2	12 y	36
<i>HBase</i>	Distributed Key-Value Store	Java	0.2	10 y	30
<i>Zookeeper</i>	Distributed Synch. Service	Java	0.1	10 y	19



What/How did we study?

- Read on-line discussion
- Check patches
- Check source code
- **Read software tutorial**
- Triggering conditions
- Errors & Failures
- Fix strategies

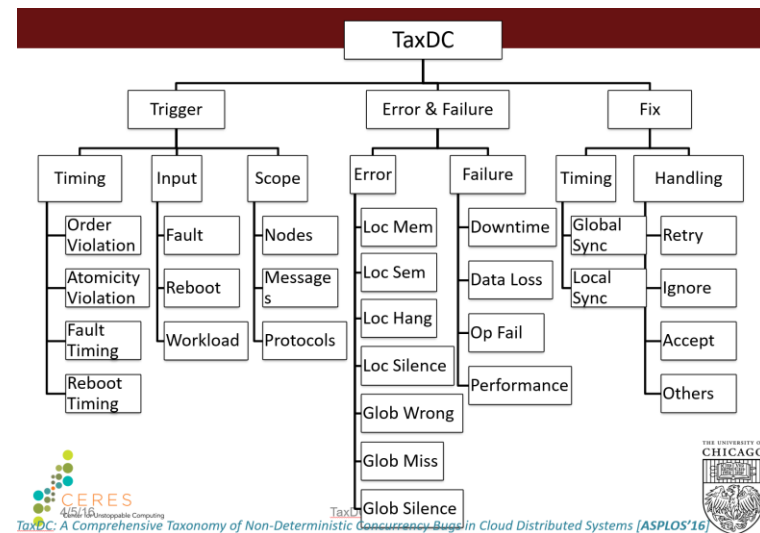
All Manual



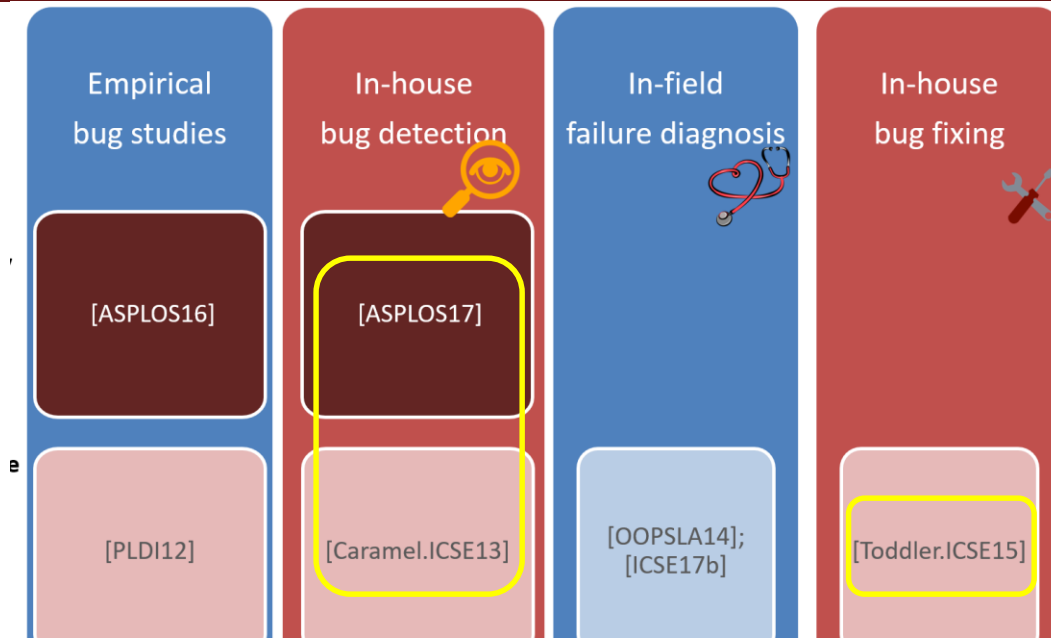
What could have been better?

- Read on-line discussion
 - Check patches
 - Check source code
 - Read software tutorial
- Triggering conditions
 - **Errors & Failures**
 - **Fix strategies**

All Manual



Bug Detection & Fixing



DCatch: Automatically Detecting Distributed Concurrency Bugs in Cloud Systems [[ASPLOS'17](#)]

Toddler: Detecting Performance Problems via Similar Memory-Access Patterns [[ICSE '13](#)]


CAMEL: Detecting and Fixing Performance Problems That Have Non-Intrusive Fixes [[ICSE'15](#)]

Dynamic DCbug Detection




- DCbugs
Bugs caused by improper timing among distributed operations
- Why? Why not applying single-machine detectors?
 - Distributed triggering
 - Different happens-before model
 - Distributed error propagation
 - Much larger scales

DCatch Tool



DCatch Tool Implementation



Benchmark Suite

Application	Software Type	Workload	# Bugs
<i>Cassandra</i>	Distributed Key-Value Store	startup	1
<i>Hadoop</i>	Distributed Computing	wordcount	2
<i>HBase</i>	Distributed Key-Value Store	enable, split, alter	2
<i>Zookeeper</i>	Distributed Synch. Service	startup	2

Dynamic PerfBug Detection

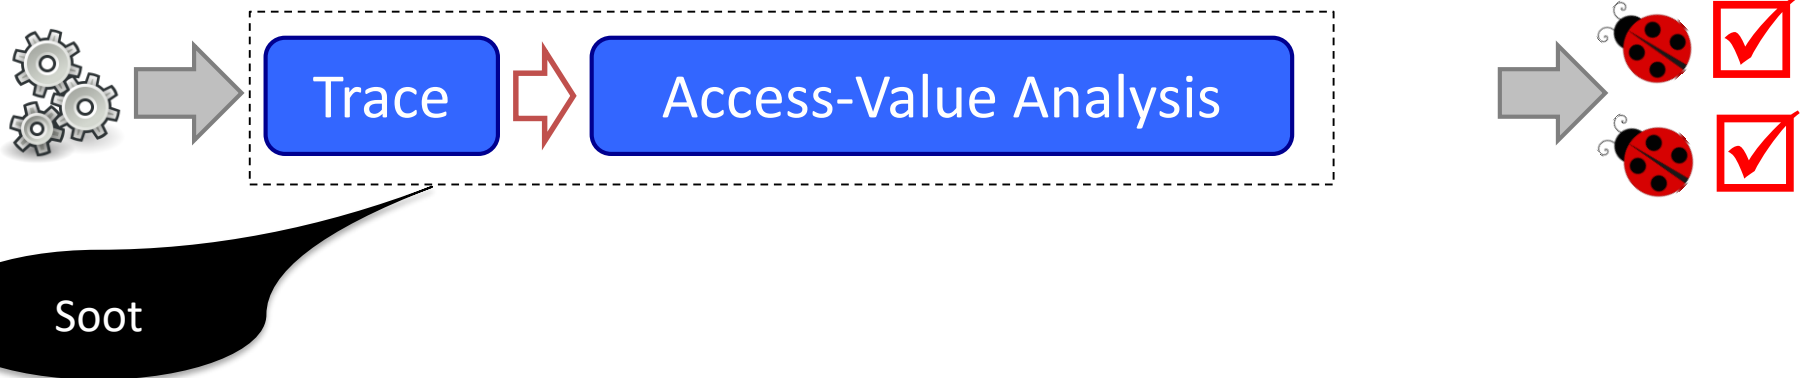


- Loop inefficiency bugs
 - Inefficient data structure
 - Redundant computation
- Why?

Toddler Tool



Toddler Tool



Benchmark Suite

Application	Software Type		KLOC	Known Bugs	New Bugs
Ant	Build tool	Java	110	1	8
Apache Col.	Collections library	Java	51	1	20
Groovy	Dynamic language	Java	137	1	0
Ggl Core Lib	Collections library	Java	156	2	10
JFreeChart	Chart framework	Java	64	1	8
JMeter	Load testing tool	Java	86	1	1
Lucene	Text search engine	Java	321	2	0
PDFBox	PDF framework	Java	78	1	0
Solr	Search server	Java	373	1	0

How to get inputs?

Static PerfBug Detection & Fixing

- Missing-break bugs

How Is
Computation
Wasted?

Where Is Computation Wasted?

	Every Iteration	Late Iterations	Early Iterations
No-Result	Type 1	Type 2	Type Y
Useless-Result	Type X	Type 3	Type 4


- Why?



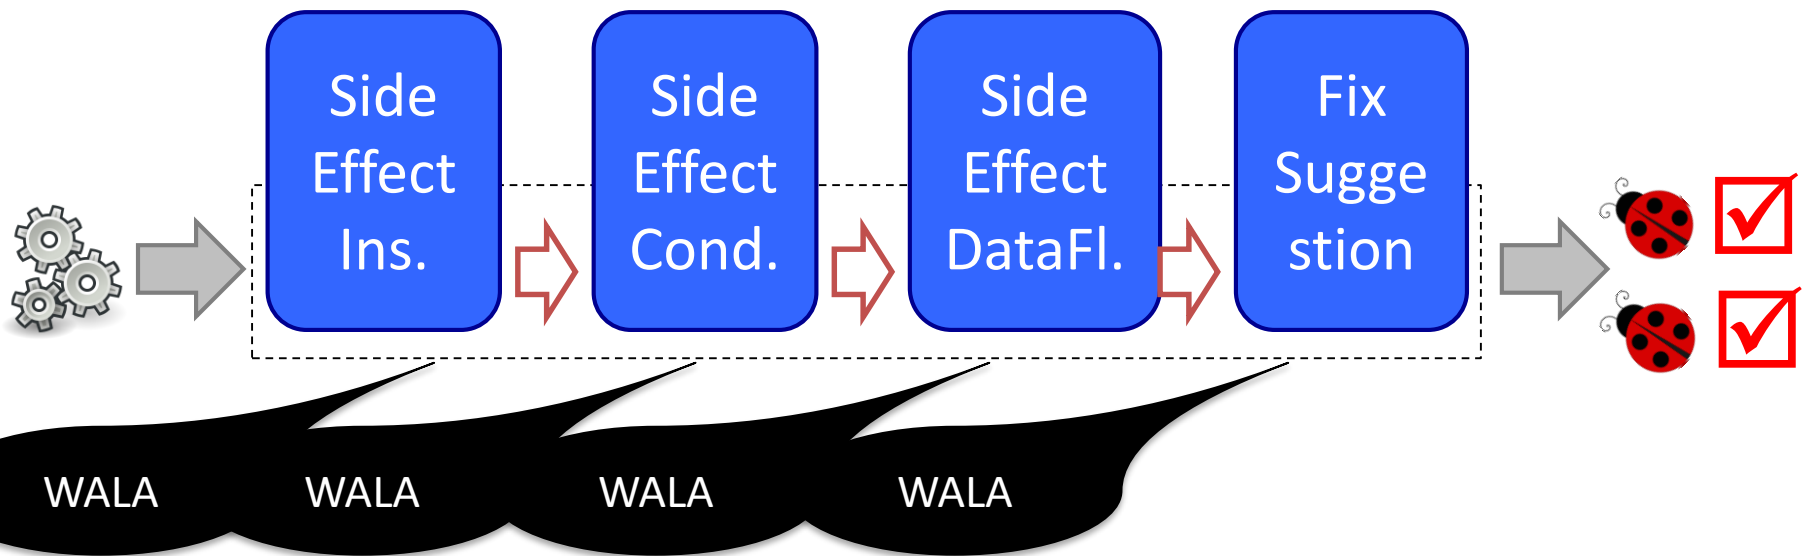
[CARMEL: Detecting and Fixing Performance Problems That Have Non-Intrusive Fixes \[ICSE'15\]](#)



Caramel Tool



Caramel Tool Implementation



Benchmark Suite (1)

Application	Software Type		KLOC	New Bugs
Ant	Build tool	Java	110	1
Apache Col.	Collections library	Java	51	20
Groovy	Dynamic language	Java	137	9
GgI Core Lib	Collections library	Java	156	10
JFreeChart	Chart framework	Java	64	8
JMeter	Load testing tool	Java	86	4
Lucene	Text search engine	Java	321	14
PDFBox	PDF framework	Java	78	10
Solr	Search server	Java	373	2

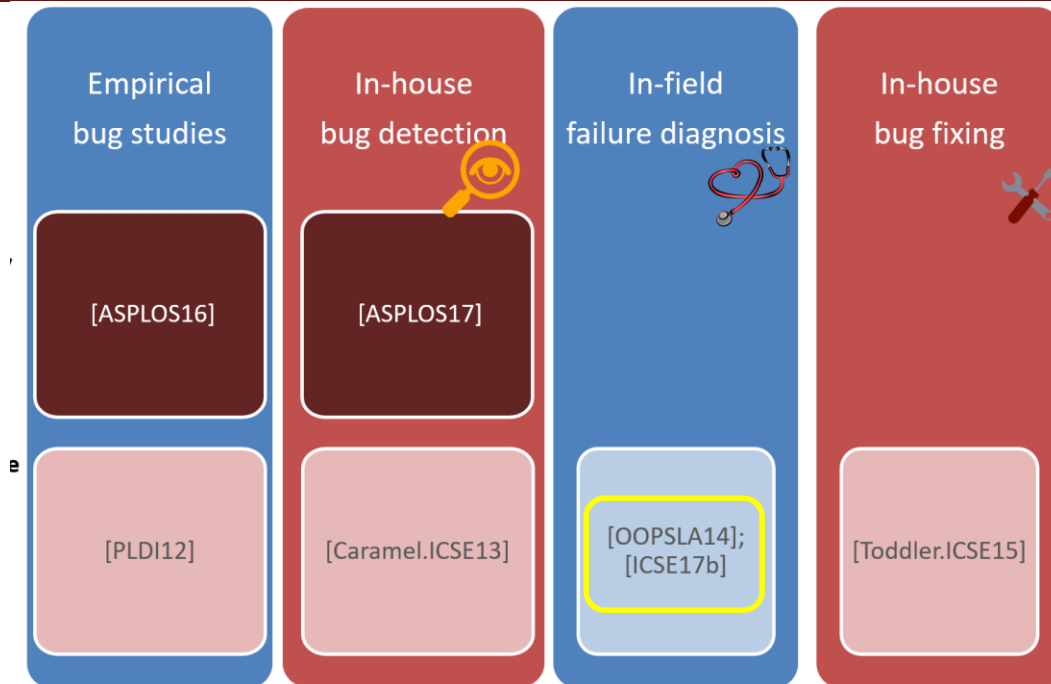


Benchmark Suite (2)

Application	Software Type		KLOC	New Bugs
<i>Log4J</i>	Logging framework	Java	52	6
<i>Sling</i>	Web app. framework	Java	202	6
<i>Struts</i>	Web app. framework	Java	175	4
<i>Tika</i>	Content extraction	Java	50	1
<i>Tomcat</i>	Web server	Java	295	4

No input requirements

Failure Diagnosis



Conclusion

- 2 bug studies, 3 bug detection tools
- 18 Java benchmark software (4 distributed)
- Workloads
 - Functionality & Performance
- Analysis mechanisms
 - Human
 - WALA, Soot, Javassist
- Others
 - Bug-tracking systems

Thanks!

Shan Lu

University of Chicago

shanlu@uchicago.edu