

Understanding **Issue Correlations**: A Case Study of the Hadoop System

Jian Huang

Xuechen Zhang[†]

Karsten Schwan



Why Issue Study Matters?

Scalable distributed systems are complex [Yuan et al., OSDI'14]



Complicated System

Why Issue Study Matters?

Scalable distributed systems are complex [Yuan et al., OSDI'14]



Complicated System



Error-prone

Why Issue Study Matters?


Scalable distributed systems are complex [Yuan et al., OSDI'14]



Complicated System



Error-prone



Hard to Debug

Why Issue Study Matters?

Scalable distributed systems are complex [Yuan et al., OSDI'14]




Complicated System



Error-prone



Hard to Debug



Issue Study



Issue Pattern

Why Issue Study Matters?

Scalable distributed systems are complex [Yuan et al., OSDI'14]




Complicated System



Error-prone



Hard to Debug



Issue Study



Issue Pattern

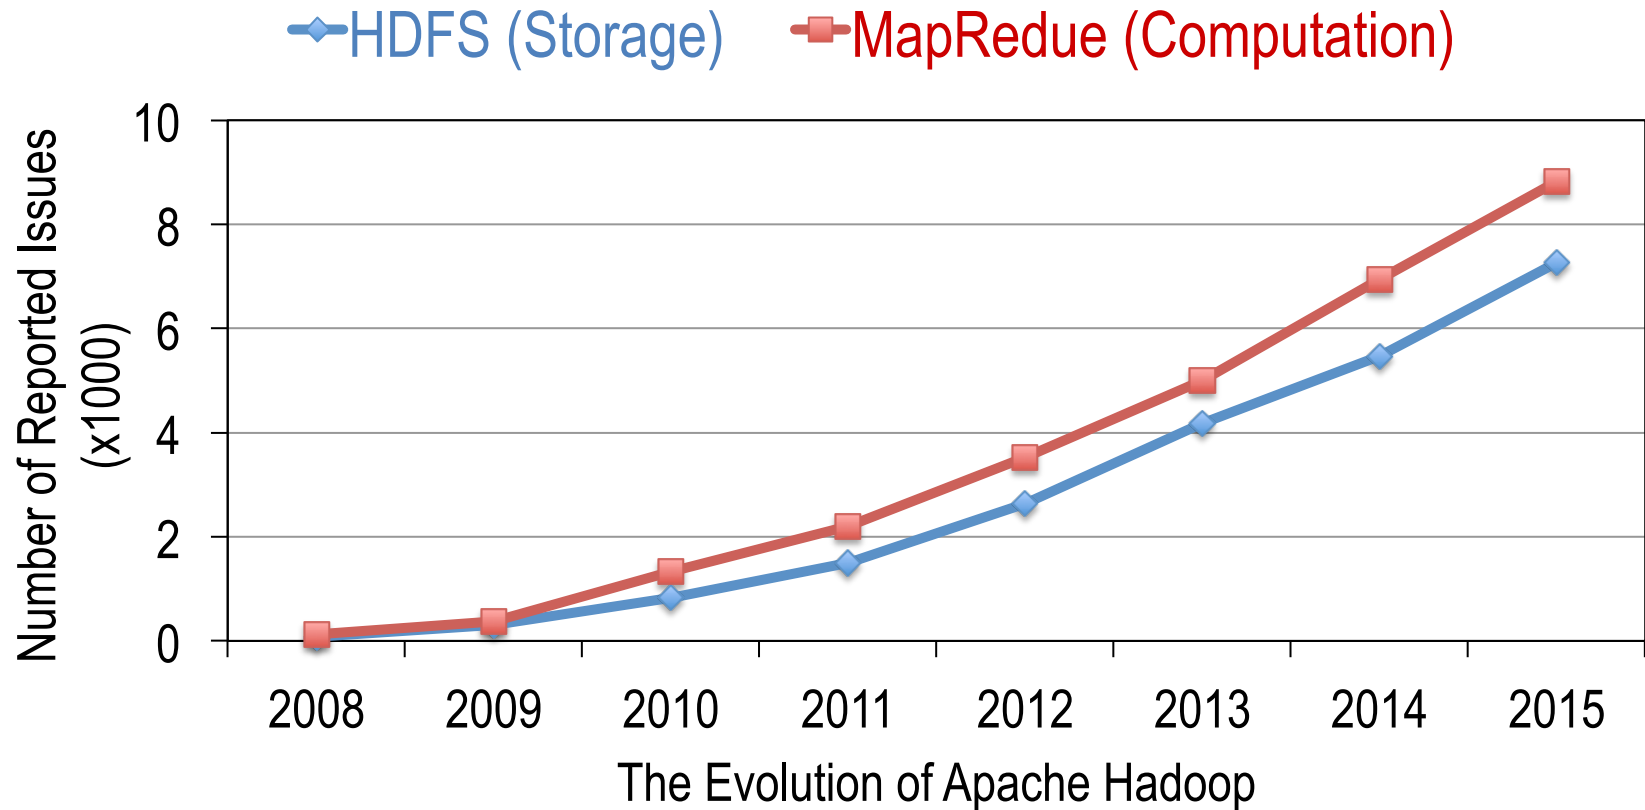


Better Software & Debugging Tools

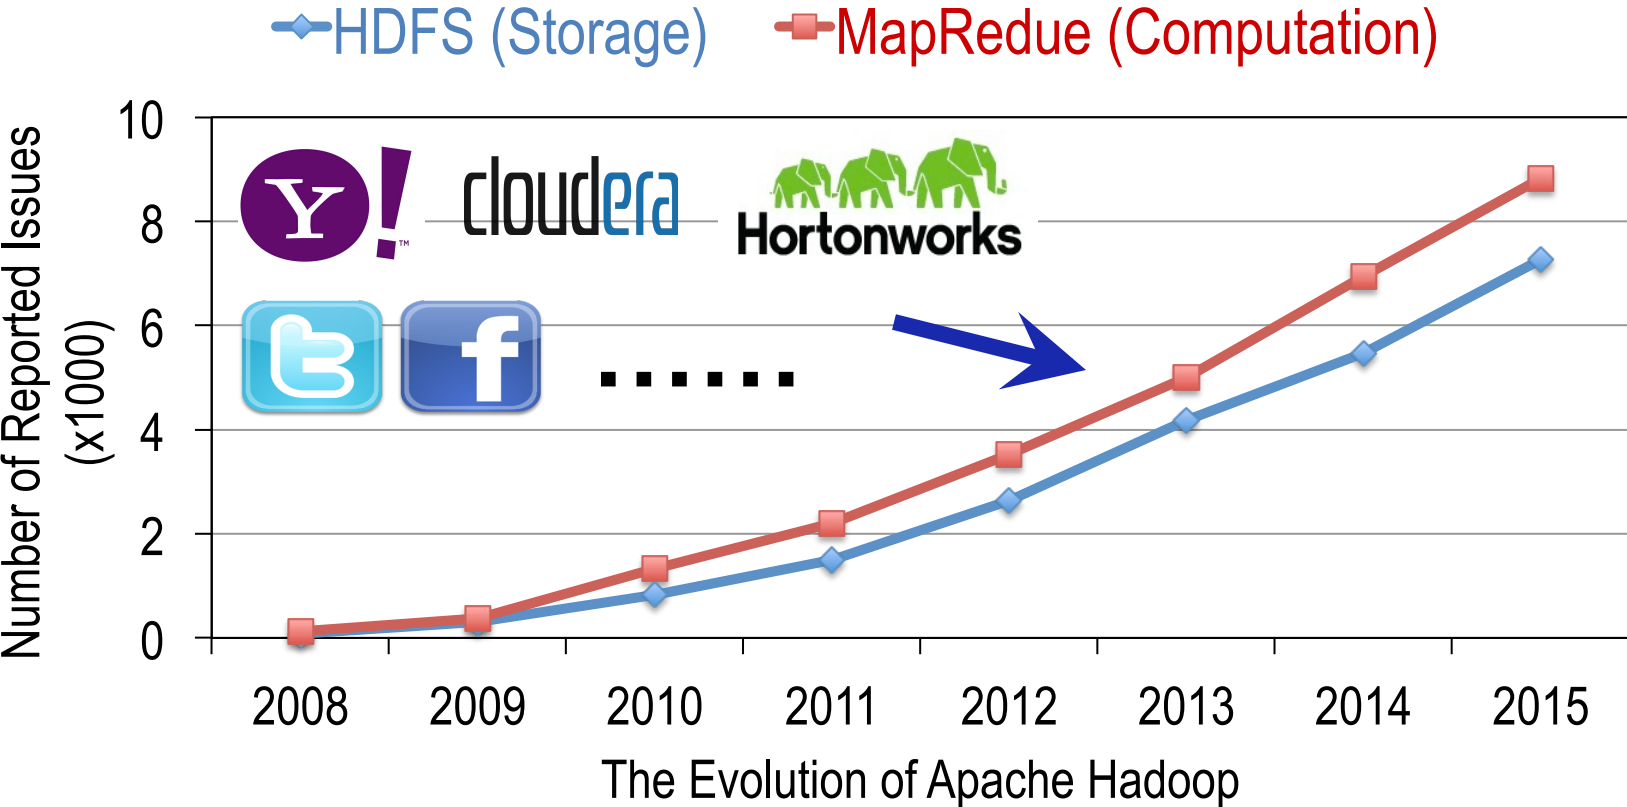


Hadoop: A Representative Distributed System

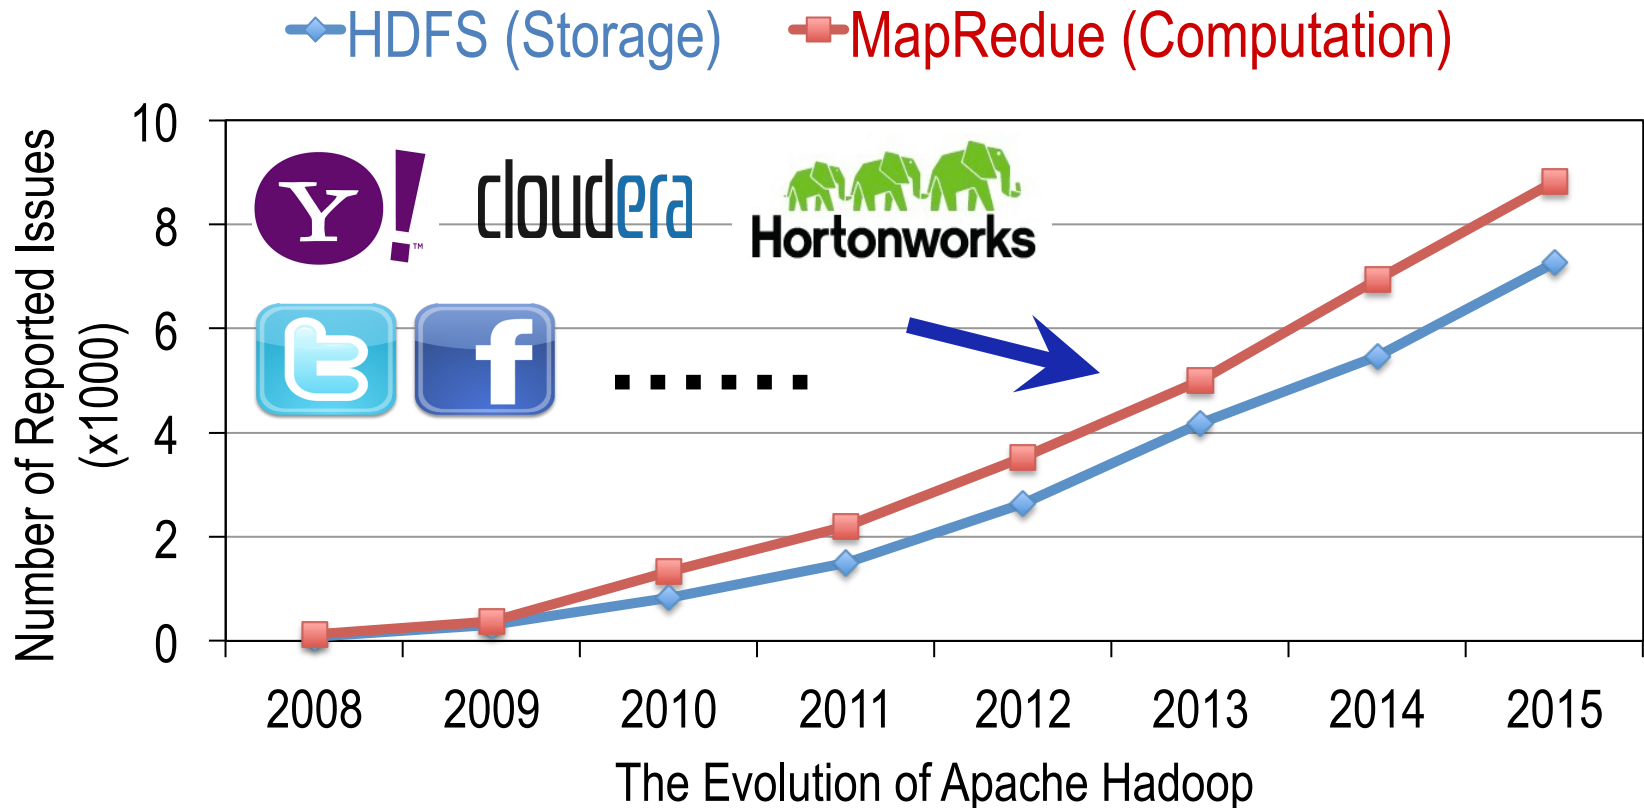
Hadoop: A Representative Distributed System



Hadoop: A Representative Distributed System



Hadoop: A Representative Distributed System



Learn from issues – more than 6 years of experience.

What Can We Learn From Issues?

Related Work



[Gunawi et al., SoCC'14]
What Bugs Live in the Cloud?

[Lu et al., FAST'13]
A Study of Linux File System Evolution

• • • • •

What Can We Learn From Issues?

Related Work

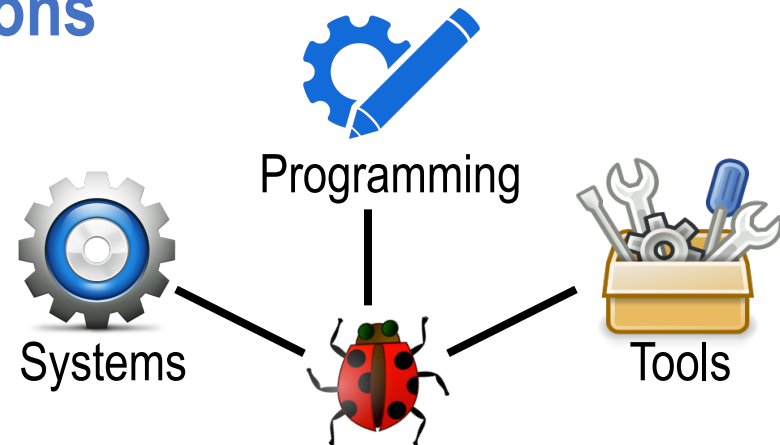



[Gunawi et al., SoCC'14]
What Bugs Live in the Cloud?

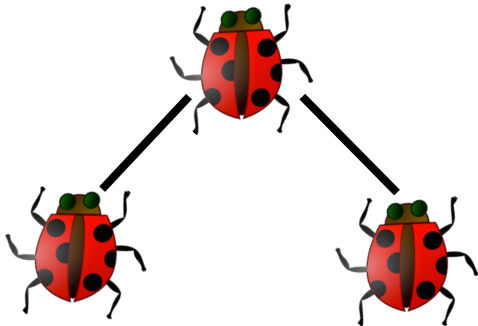
[Lu et al., FAST'13]
A Study of Linux File System Evolution

.....

Our Focus: Issue Correlations




Our Findings

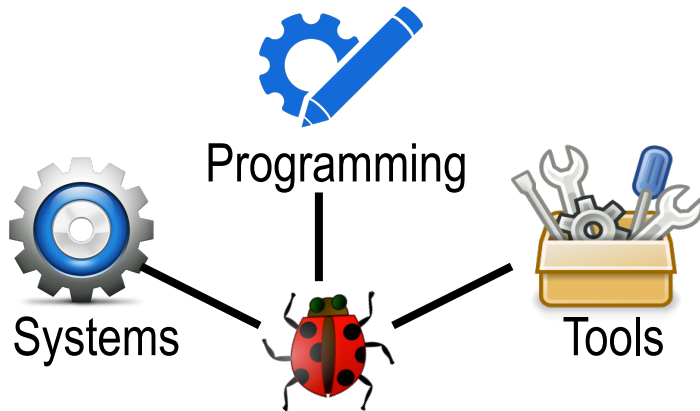


- Half of the issues are independent
- MapReduce issues tend to relate to YARN
- One third of the issues have similar causes
-

Our Findings

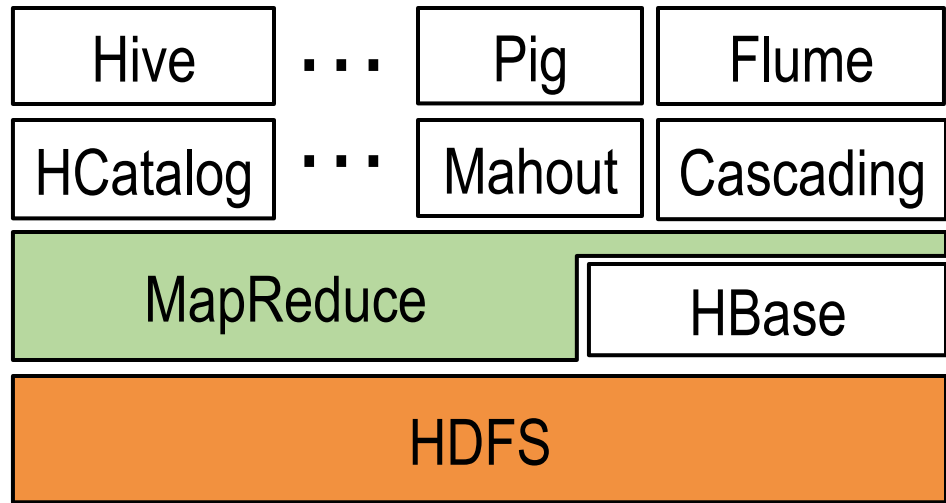


- Half of the issues are independent
- MapReduce issues tend to relate to YARN
- One third of the issues have similar causes
-



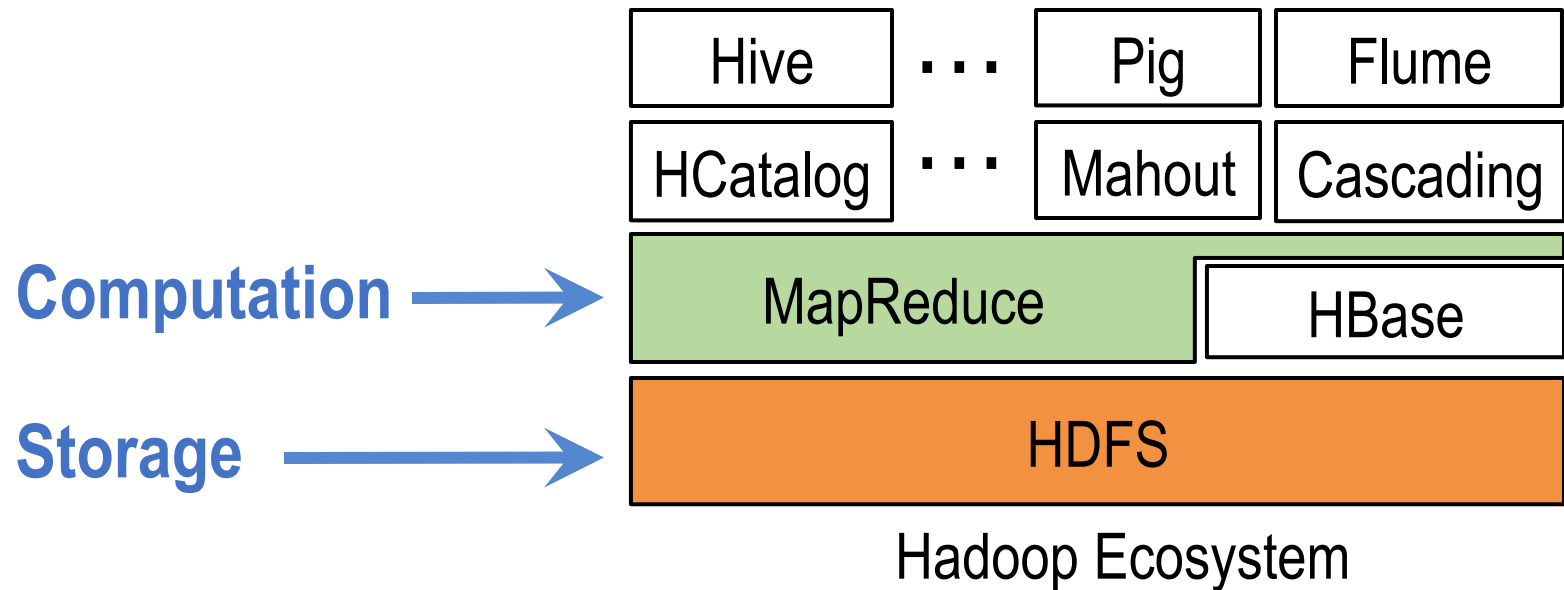
- **Memory:** GC is still the No. 1 concern
- **Storage:** “99.99% of data reliability” is challenged
- **Programming:** one third of them relate to interfaces
- **Tools:** the logging in Hadoop is error-prone
-

Methodology Used in Our Study

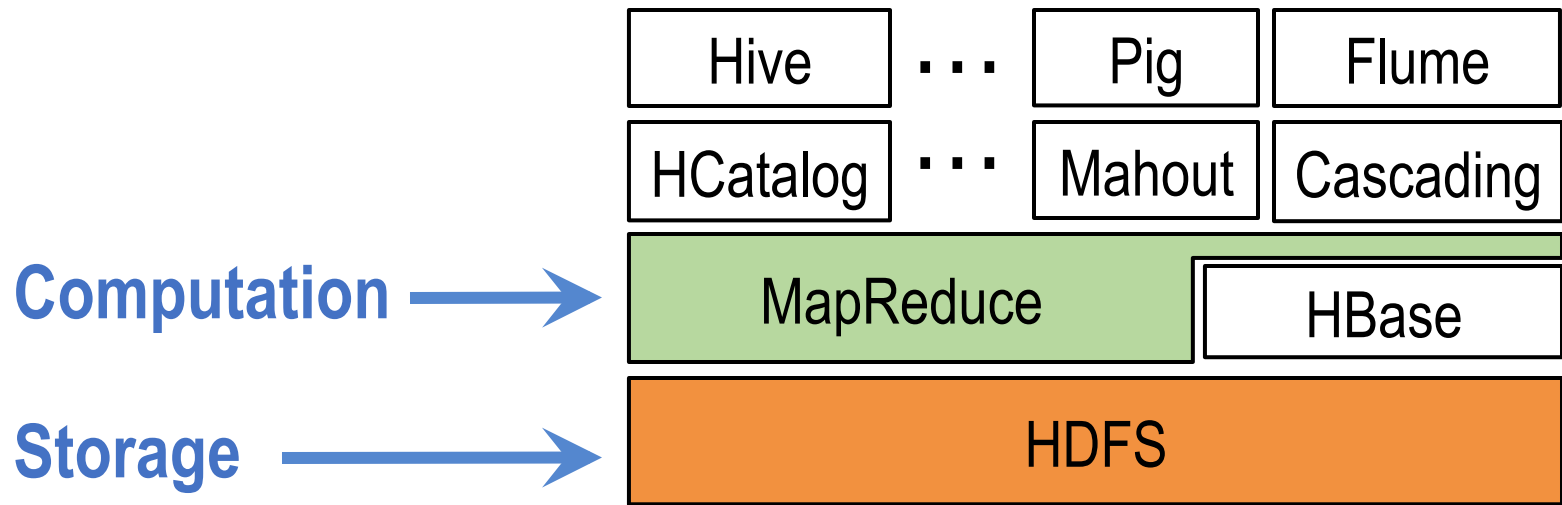


Hadoop Ecosystem

Methodology Used in Our Study



Methodology Used in Our Study



Hadoop Ecosystem



Closed Issues

2359

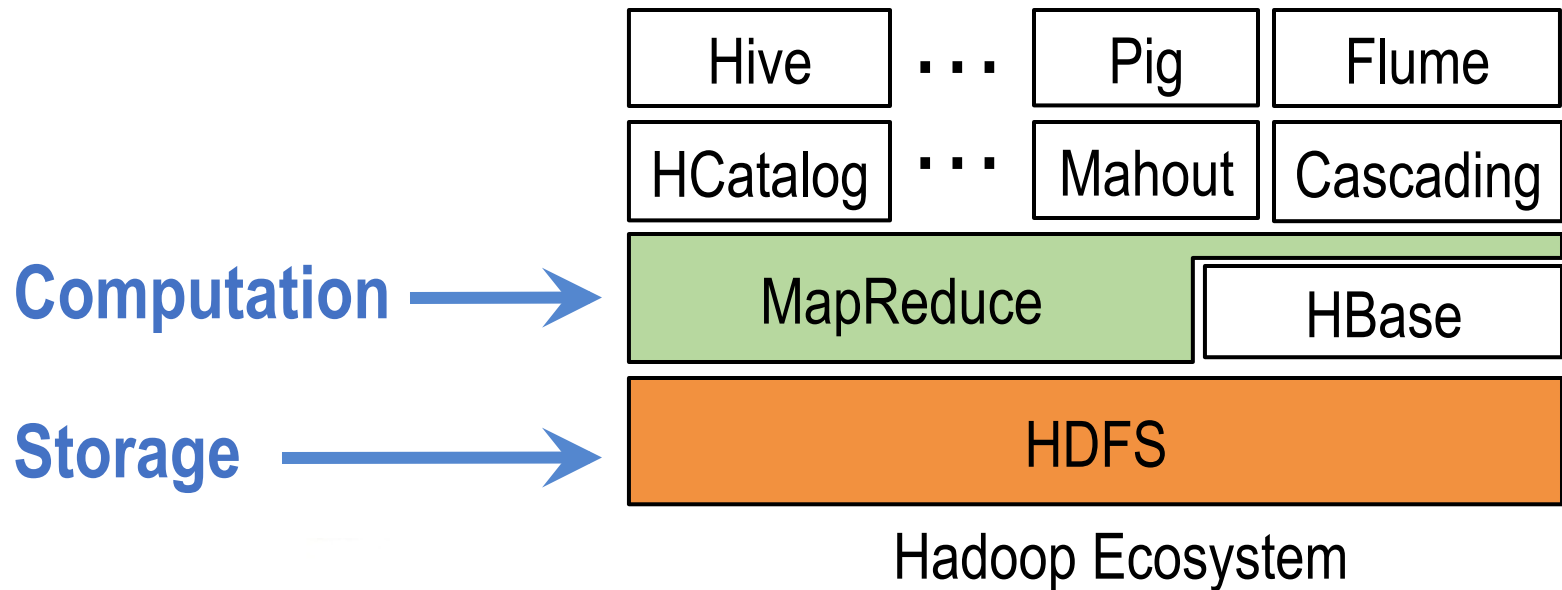
2340

Examined Issues

2180


2038

Methodology Used in Our Study



Closed Issues	2359	2340	Sampling Rate 89.8%
Examined Issues	2180	2038	

Methodology Used in Our Study

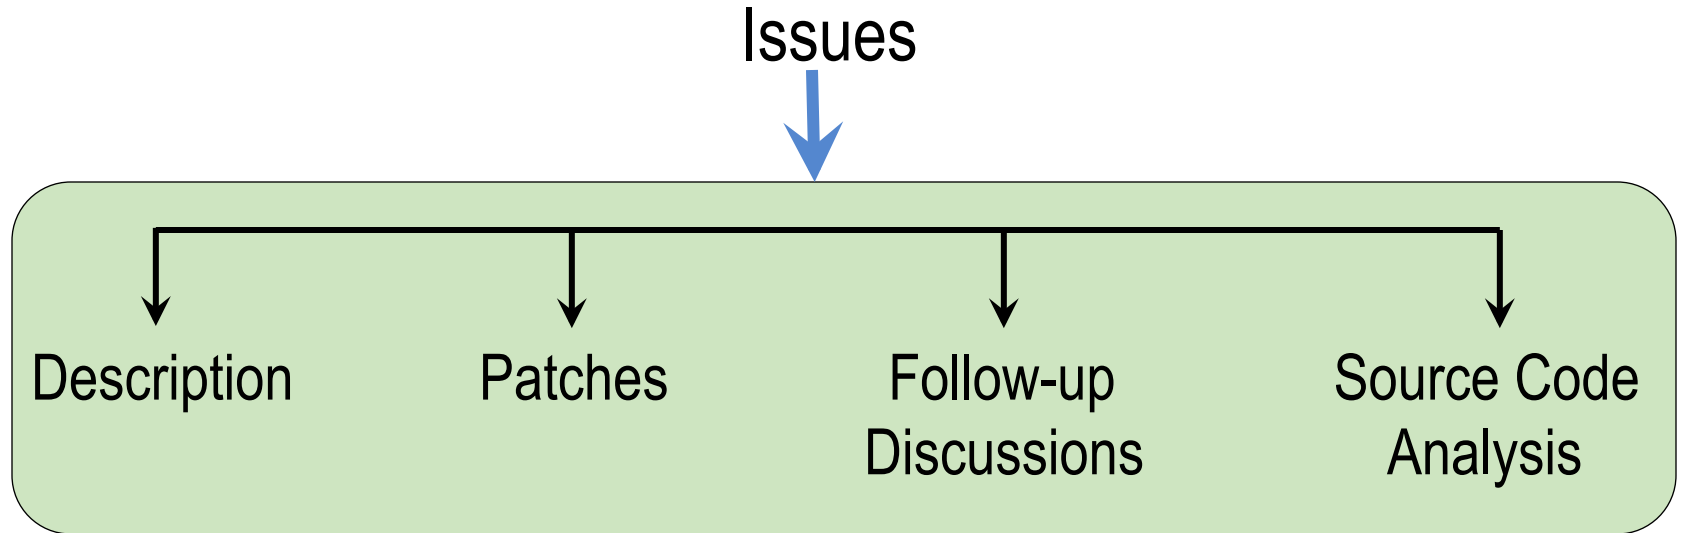


Hadoop Ecosystem

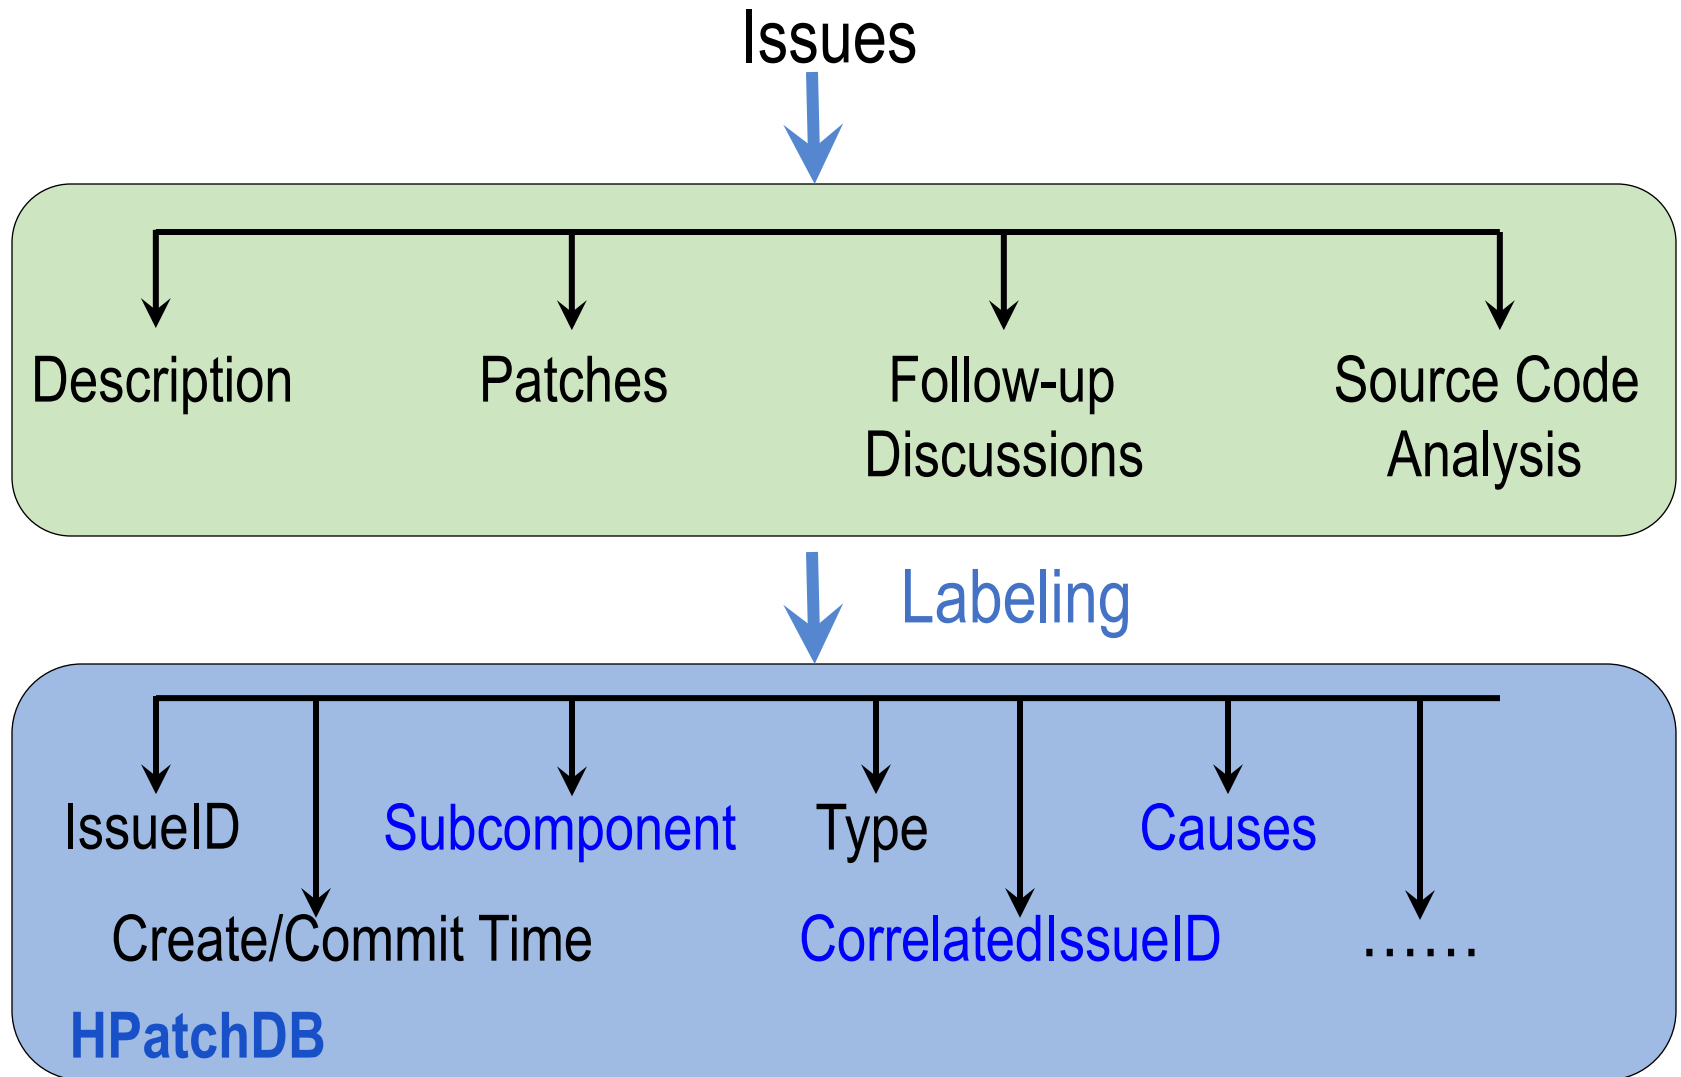


Closed Issues	2359	2340	Sampling Rate 89.8%
Examined Issues	2180	2038	
Sampling Period	~6 years	5 years	

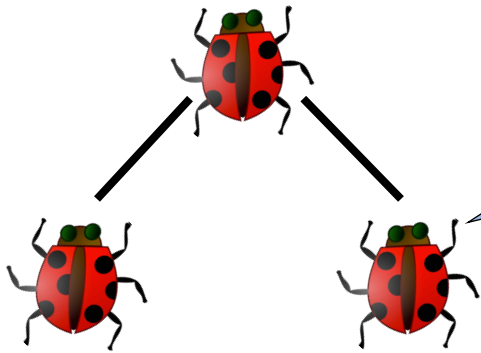
Methodology Used in Our Study



Methodology Used in Our Study




Where Are the Correlated Issues From?



Do you know where I'm from?

Where Are the Correlated Issues From?




External Correlation

correlated issues appear in other systems

Internal Correlation


correlated issues appear in the same system

Where Are the Correlated Issues From?



#Correlated Issues		0	1	2	3	≥ 4
External	HDFS	94.7%	4.8%	0.5%	-	-
	MapReduce	79.3%	17.1%	2.8%	0.5%	0.3%


Where Are the Correlated Issues From?



A significant number of issues are independent.

#Correlated Issues		0	1	2	3	>=4
External	HDFS	94.7%	4.8%	0.5%	-	-
	MapReduce	79.3%	17.1%	2.8%	0.5%	0.3%


Where Are the Correlated Issues From?



Half of them are from YARN.


#Correlated Issues		0	1	2	3	>=4
External	HDFS	94.7%	4.8%	0.5%	-	-
	MapReduce	79.3%	17.1%	2.8%	0.5%	0.3%

Where Are the Correlated Issues From?



#Correlated Issues		0	1	2	3	>=4
External	HDFS	94.7%	4.8%	0.5%	-	-
	MapReduce	79.3%	17.1%	2.8%	0.5%	0.3%
Internal	HDFS	52.7%	32.8%	9.1%	3.1%	2.3%
	MapReduce	59.3%	32.7%	5.6%	1.3%	1.0%

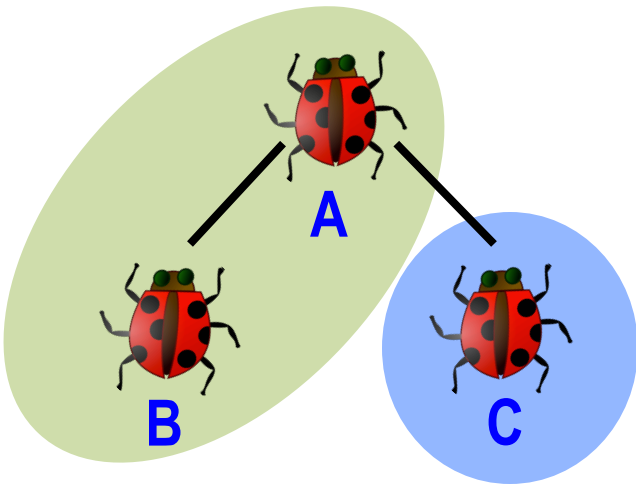
Where Are the Correlated Issues From?



Half of them are independent.

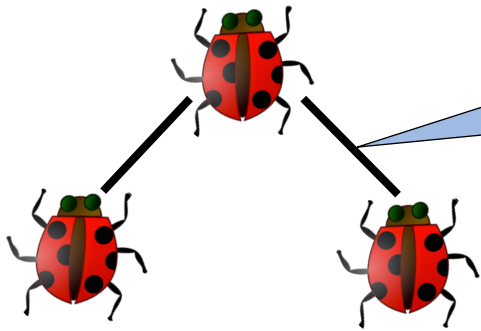
#Correlated Issues		0	1	2	3	>=4
External	HDFS	94.7%	4.8%	0.5%	-	-
	MapReduce	79.3%	17.1%	2.8%	0.5%	0.3%
Internal	HDFS	52.7%	32.8%	9.1%	3.1%	2.3%
	MapReduce	59.3%	32.7%	5.6%	1.3%	1.0%

Where Are the Correlated Issues From?




#Correlated Issues		0	1	2	3	>=4
External	HDFS	94.7%	4.8%	0.5%	-	-
	MapReduce	79.3%	17.1%	2.8%	0.5%	0.3%
Internal	HDFS	52.7%	32.8%	9.1%	3.1%	2.3%
	MapReduce	59.3%	32.7%	5.6%	1.3%	1.0%

How the Issues Are Correlated?



Do you know our relationship?

How the Issues Are Correlated?



Similar Causes

Issues have similar causes


Blocking Other Issues

Issues need to be fixed before fixing other issues

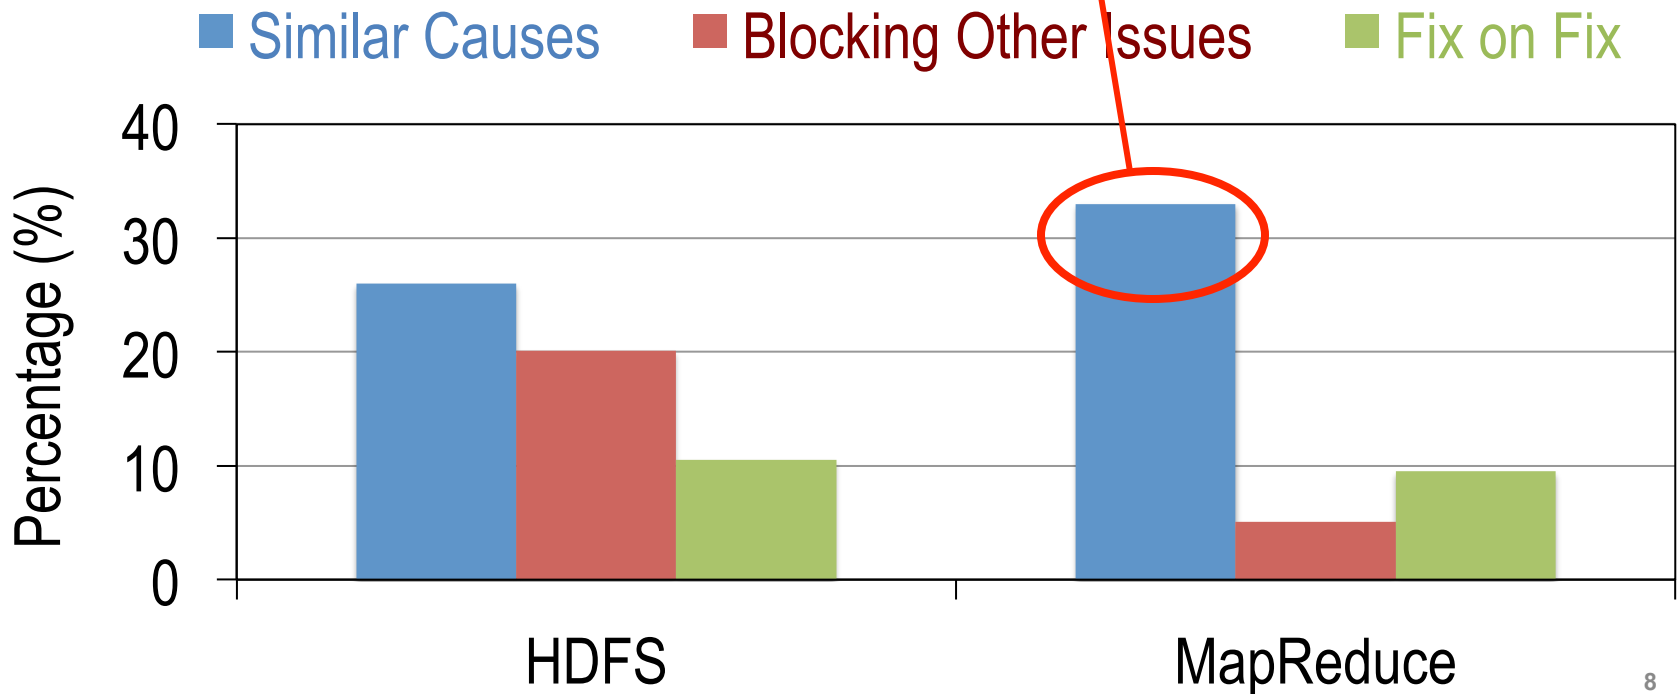
Fix on Fix

Issues are caused by fixing other issues


How the Issues Are Correlated?



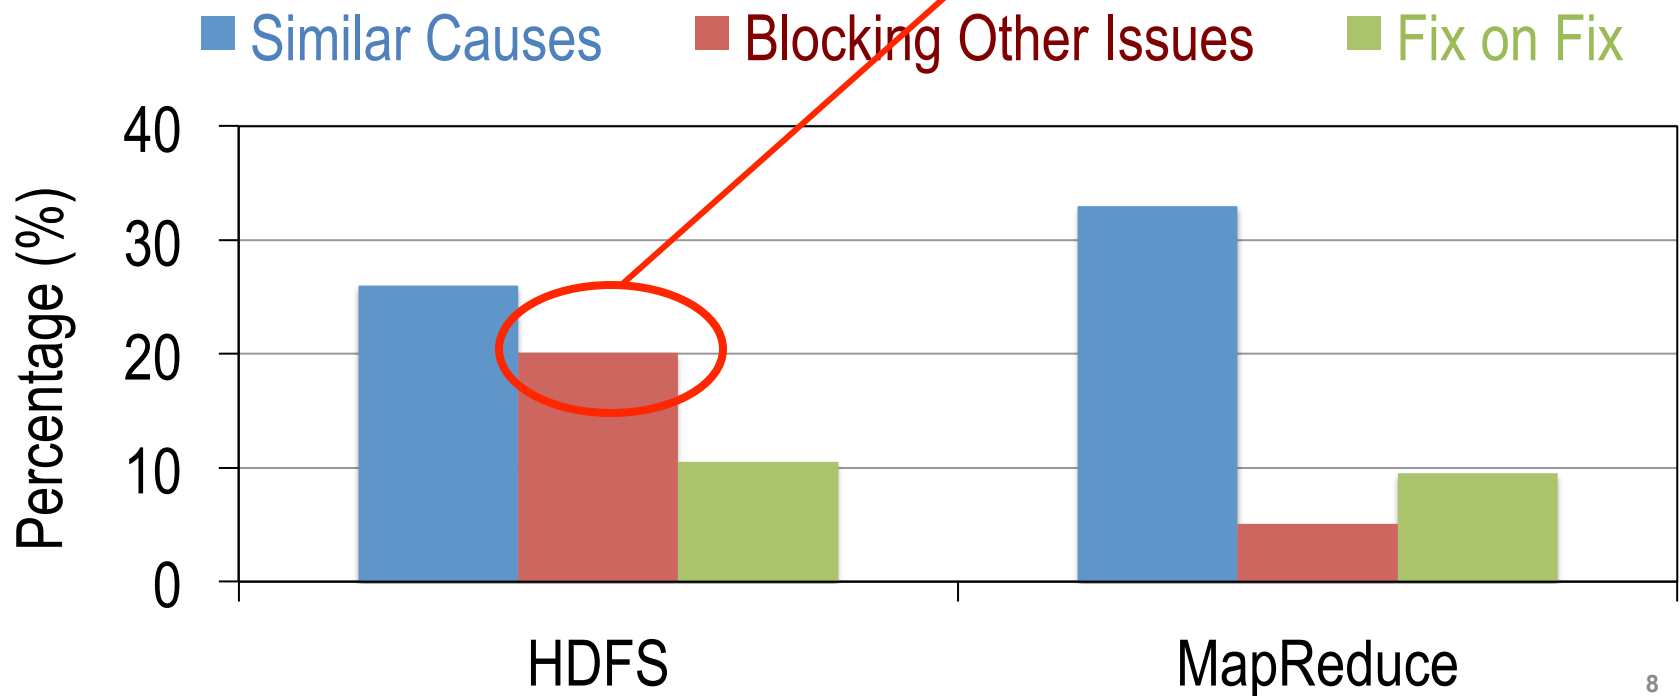
26-33% of the issues have similar causes.




How the Issues Are Correlated?




These issues that block others appear more frequently in HDFS.




How the Issues Are Correlated?



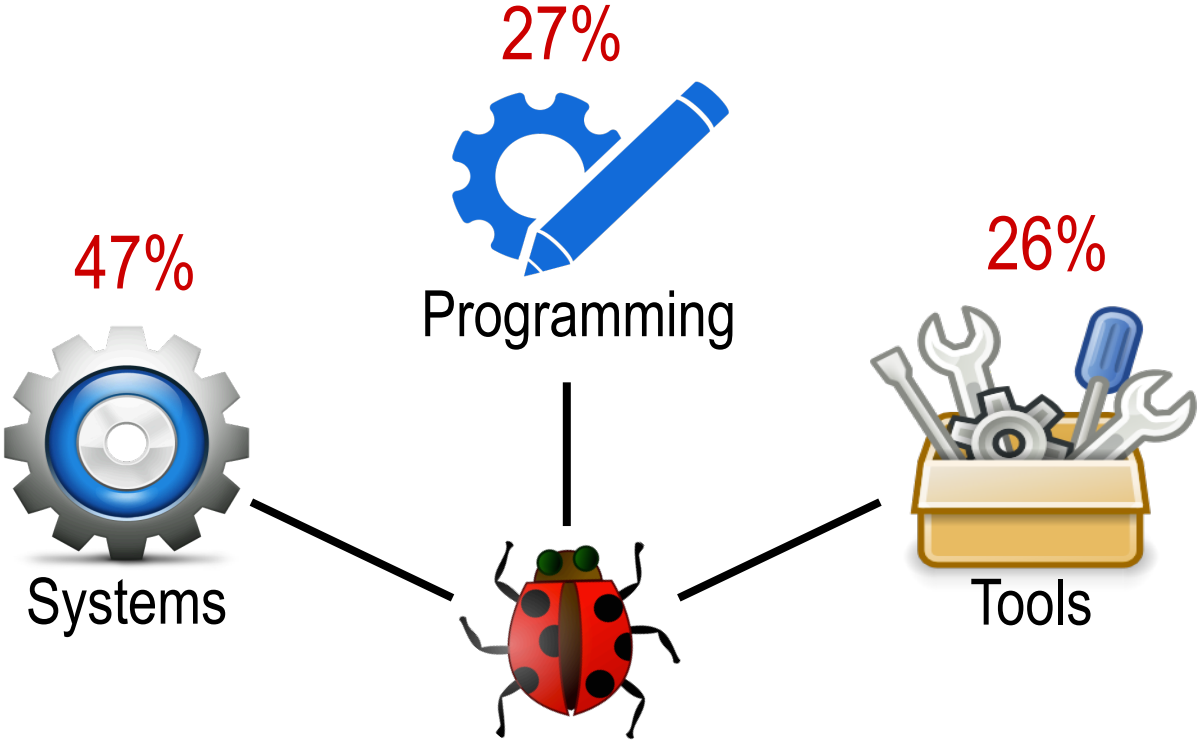
Mostly due to functional dependency.



On the Issue Correlations with System Characteristics




On the Issue Correlations with System Characteristics




How Issues Relate to Systems?

How Issues Relate to Systems?




How Issues Relate to Systems?



GC is still the No.1 concern,
memory-friendly objects are preferred.

- LightWeightGSet Vs. java.util structure
- Object cache for long lived object:
ReplicasMap, ReplicasInfo

How Issues Relate to Systems?

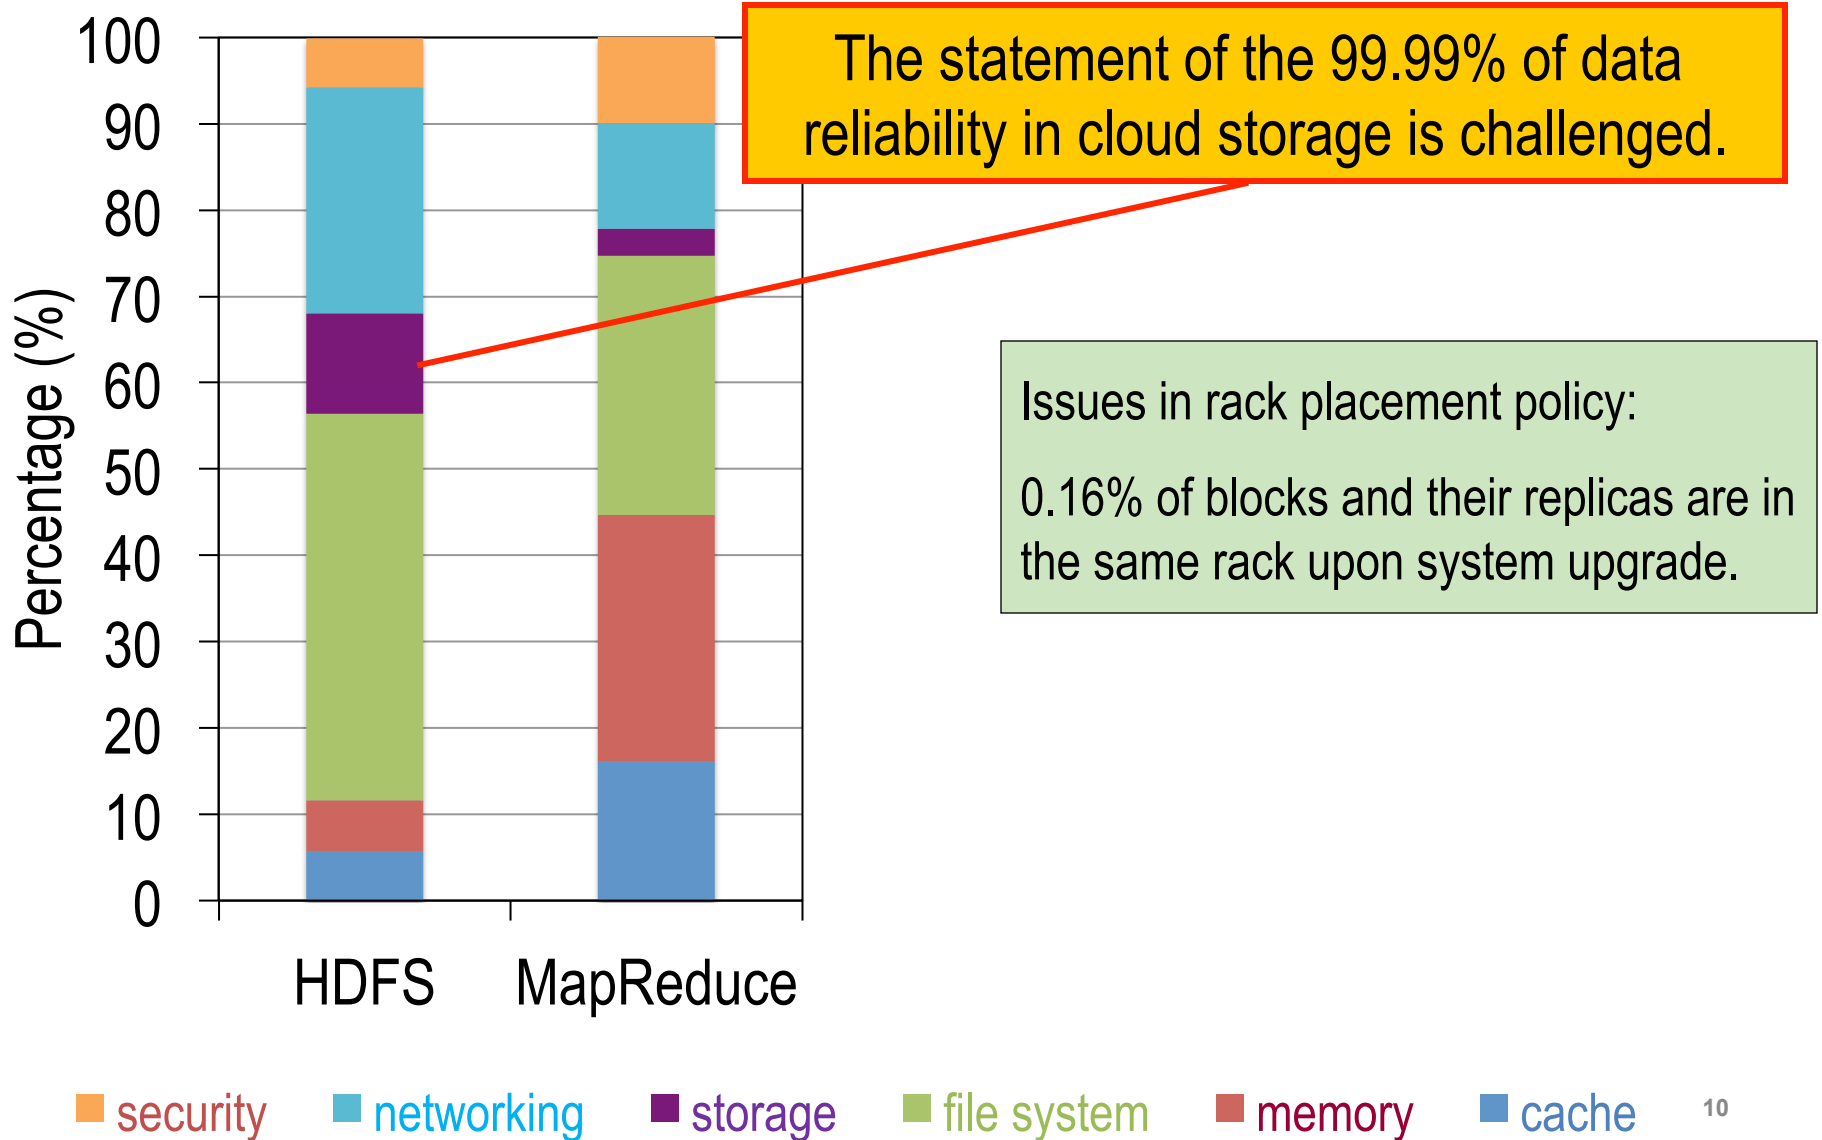


Many issues happened in file system like EXT4 appear in Hadoop.


File system semantic:

namespace management, file permission, consistency (e.g., fsck), etc.

How Issues Relate to Systems?



How Issues Relate to Systems?




One quarter of networking issues cause resource wastage.

Socket leak!


Read a block:

```
Peer peer = newTcpPeer(dnAddr),
- return newBlockReader(...)
+ try{
+   reader = newBlockReader(...)
+   return reader
+ } catch (IOException ex) {
+   throw ex;
+ } finally {
+   if(reader == null) closeQuietly(peer);
+ }
```

How Issues Relate to Programming?




How Issues Relate to Programming?



Half of them relate to code maintenance.

■ typo ■ lock ■ interface ■ maintenance


How Issues Relate to Programming?



Mainly caused by interface changes.

typo lock interface maintenance

How Issues Relate to Programming?




5.6% of programming issues are caused by typos !

A fsimage cannot be accessed due to:

- elif ["COMMAND" = "oiv_legacy"] then
- + elif ["\$COMMAND" = "oiv_legacy"] then


■ typo ■ lock ■ interface ■ maintenance

How Issues Relate to Tools?




■ configuration ■ debugging ■ documents ■ testing

How Issues Relate to Tools?

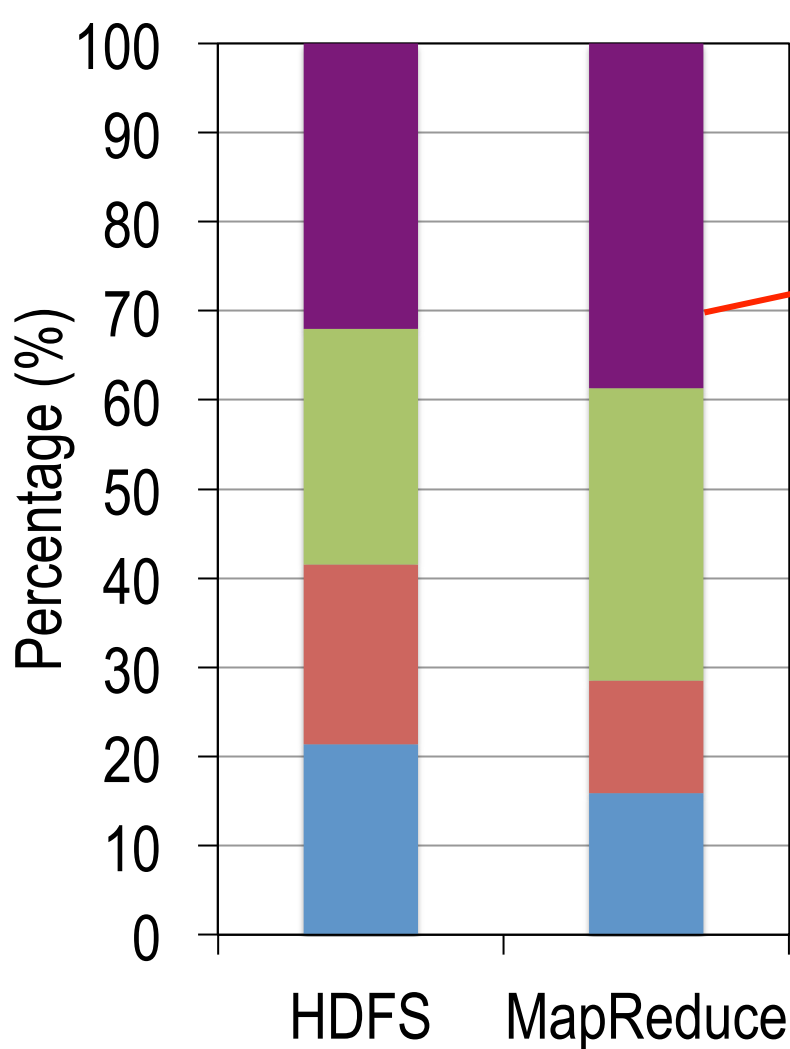


Logs are misleading:
incorrect, incomplete, indistinct output.

How Issues Relate to Tools?




How Issues Relate to Tools?



A majority of configuration issues are related to system performance.

59% of the 219 configuration parameters in MapReduce are performance related.

Conclusion



1 Correlations Between Issues

Issues are independent; 33% of issues have similar causes, etc.

2 Correlations With System Characteristics

More efforts are required to achieve highly reliable distributed system

Thanks!

Jian Huang

jian.huang@gatech.edu

Xuechen Zhang[†]

Karsten Schwan



Q&A