中国科学院软件研究所学术年会’2019
暨计算机科学国家重点实验室开放周

学术论文

大规模分布式系统中节点失效恢复缺陷实证研究

高钰, 窦文生, 秦锋, 高楚舒, 王栋, 魏峻, 黄瑞瑞, 周利, 吴永明

An Empirical Study on Crash Recovery Bugs in Large-Scale Distributed Systems, ESEC/FSE 2018, pp. 539-550.
ACM SIGSOFT Distinguished Paper Award

联系方式：窦文生，wensheng@iscas.ac.cn，http://www.tcse.cn/~wsdou

Crash Recovery Bugs in Distributed Systems and Study Methodology

Crash recovery processes in cloud systems

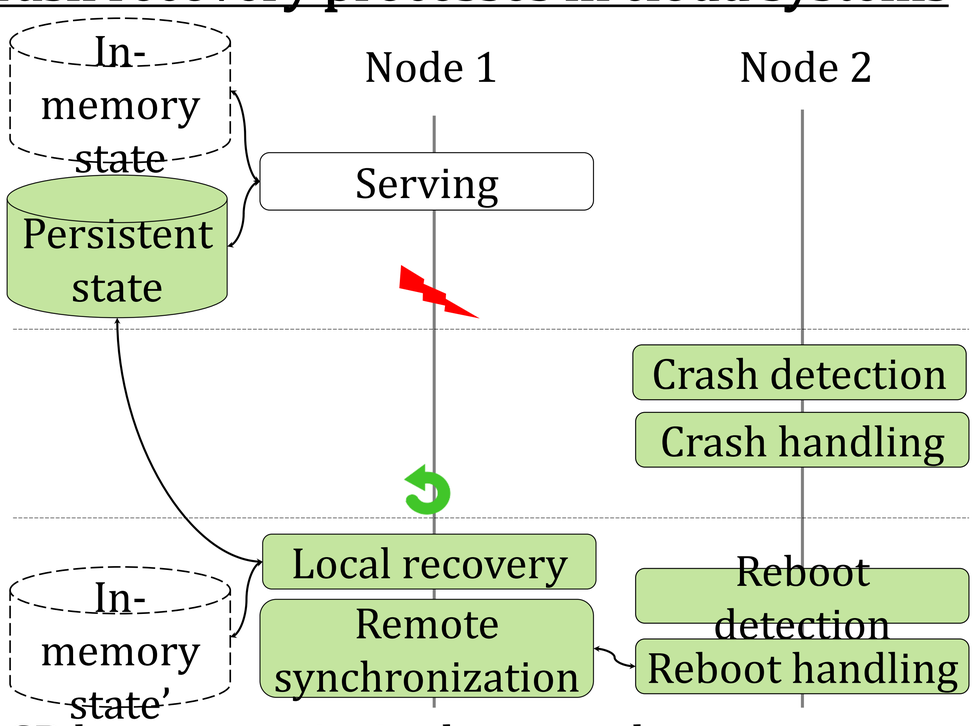

◆ CR bugs can exist in these crash recovery

Main research questions

✓RQ1: What are the root causes for crash recovery bugs?
✓RQ2: How is a crash recovery bug triggered?
✓RQ3: What impacts do crash recovery bugs have?

Methodology

✓We study **4** distributed systems:
- ZooKeeper, Hadoop MapReduce, HBase, Cassandra

✓We select CBS [1] as our study base:
- 3,655 vital cloud issues reported from Jan 2011 to Jan 2014

✓We collect **103** crash recovery bugs from CBS and analyze each bug to answer our four research questions by studying:
- Developer comments, patches and source code

Key Findings – Root Causes

We got five bug patterns and each bug pattern contains multiple subcategories.

No backup

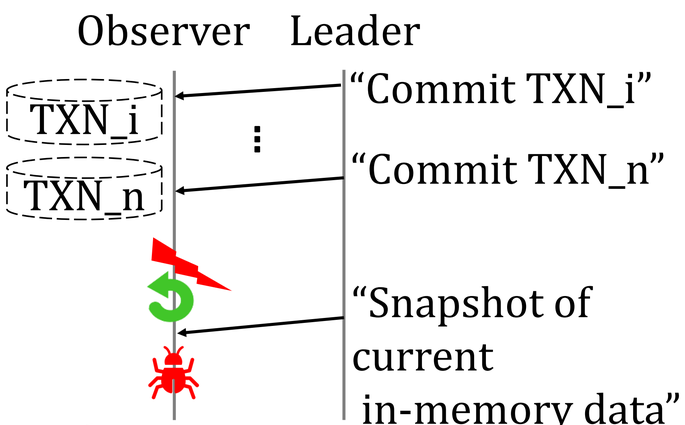

Take a long time for recovery!

ZOOKEEPER-1552

No crash detection

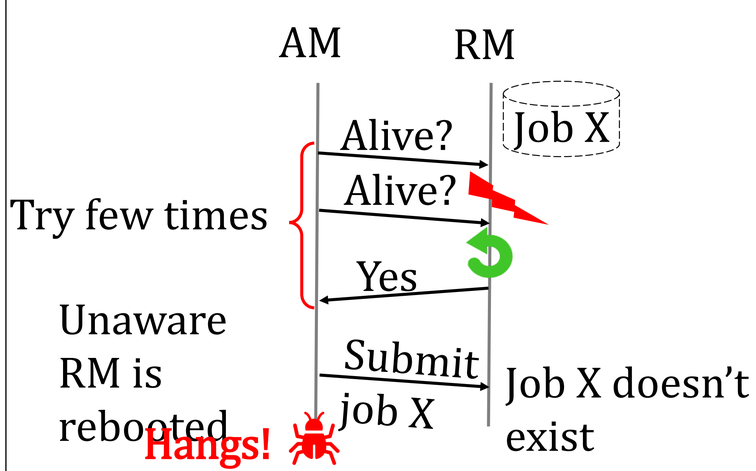

MAPREDUCE-3186

No handling of leftovers

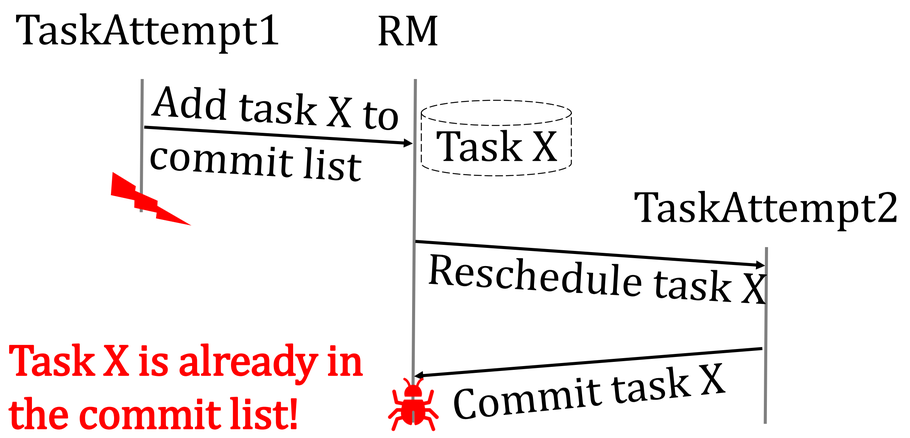

MAPREDUCE-3858

✓ Important data should be backed up in all cases.
✓ Crashes and reboots can happen at any time. Missing/untimely crash/reboot detection indicates CR bugs.
✓ Unhandled leftovers of a crash node indicates CR bugs.

Key Findings – Triggering Conditions

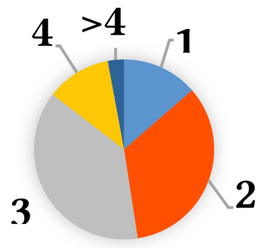

◆ 97% of CR bugs involve **four** nodes or fewer

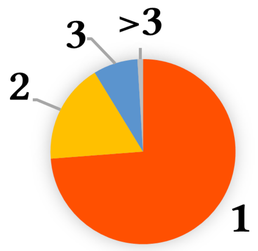

◆ No more than **three** crashes can trigger 99% of CR bugs

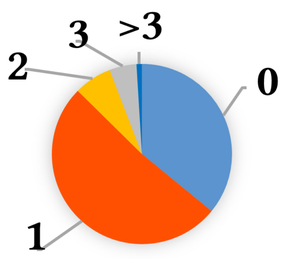

◆ No more than **one** reboot can trigger 87% of CR bugs

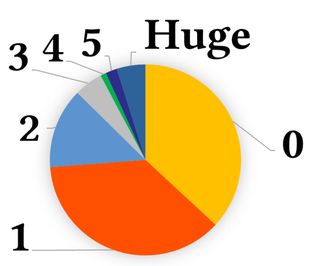

◆ 92% of CR bugs require no more than **three** user requests.

✓ CR bugs can be detected in a small cluster, with injection of limited crashes and reboots.

Key Findings – Impacts

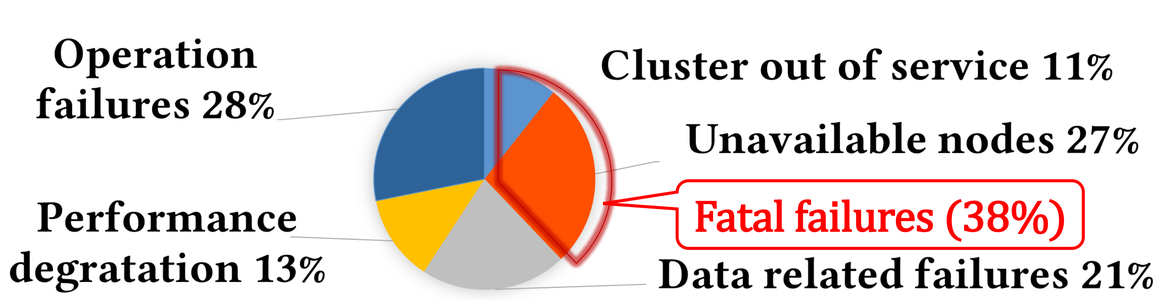

✓ Compared with the bugs in CBS [1] (18%) and TaxDC [2] (17%), crash recovery bugs are more likely to cause fatal failures.

References
[1] H. Gunawi et al., “What Bugs Live in the Cloud? A Study of 3000+ Issues in Cloud Systems”, SoCC 2014.
[2] T. Leesatapornwongsa et al., “TaxDC: A Taxonomy of Non-Deterministic Concurrency Bugs in Datacenter Distributed Systems”, ASPLOS 2016.