
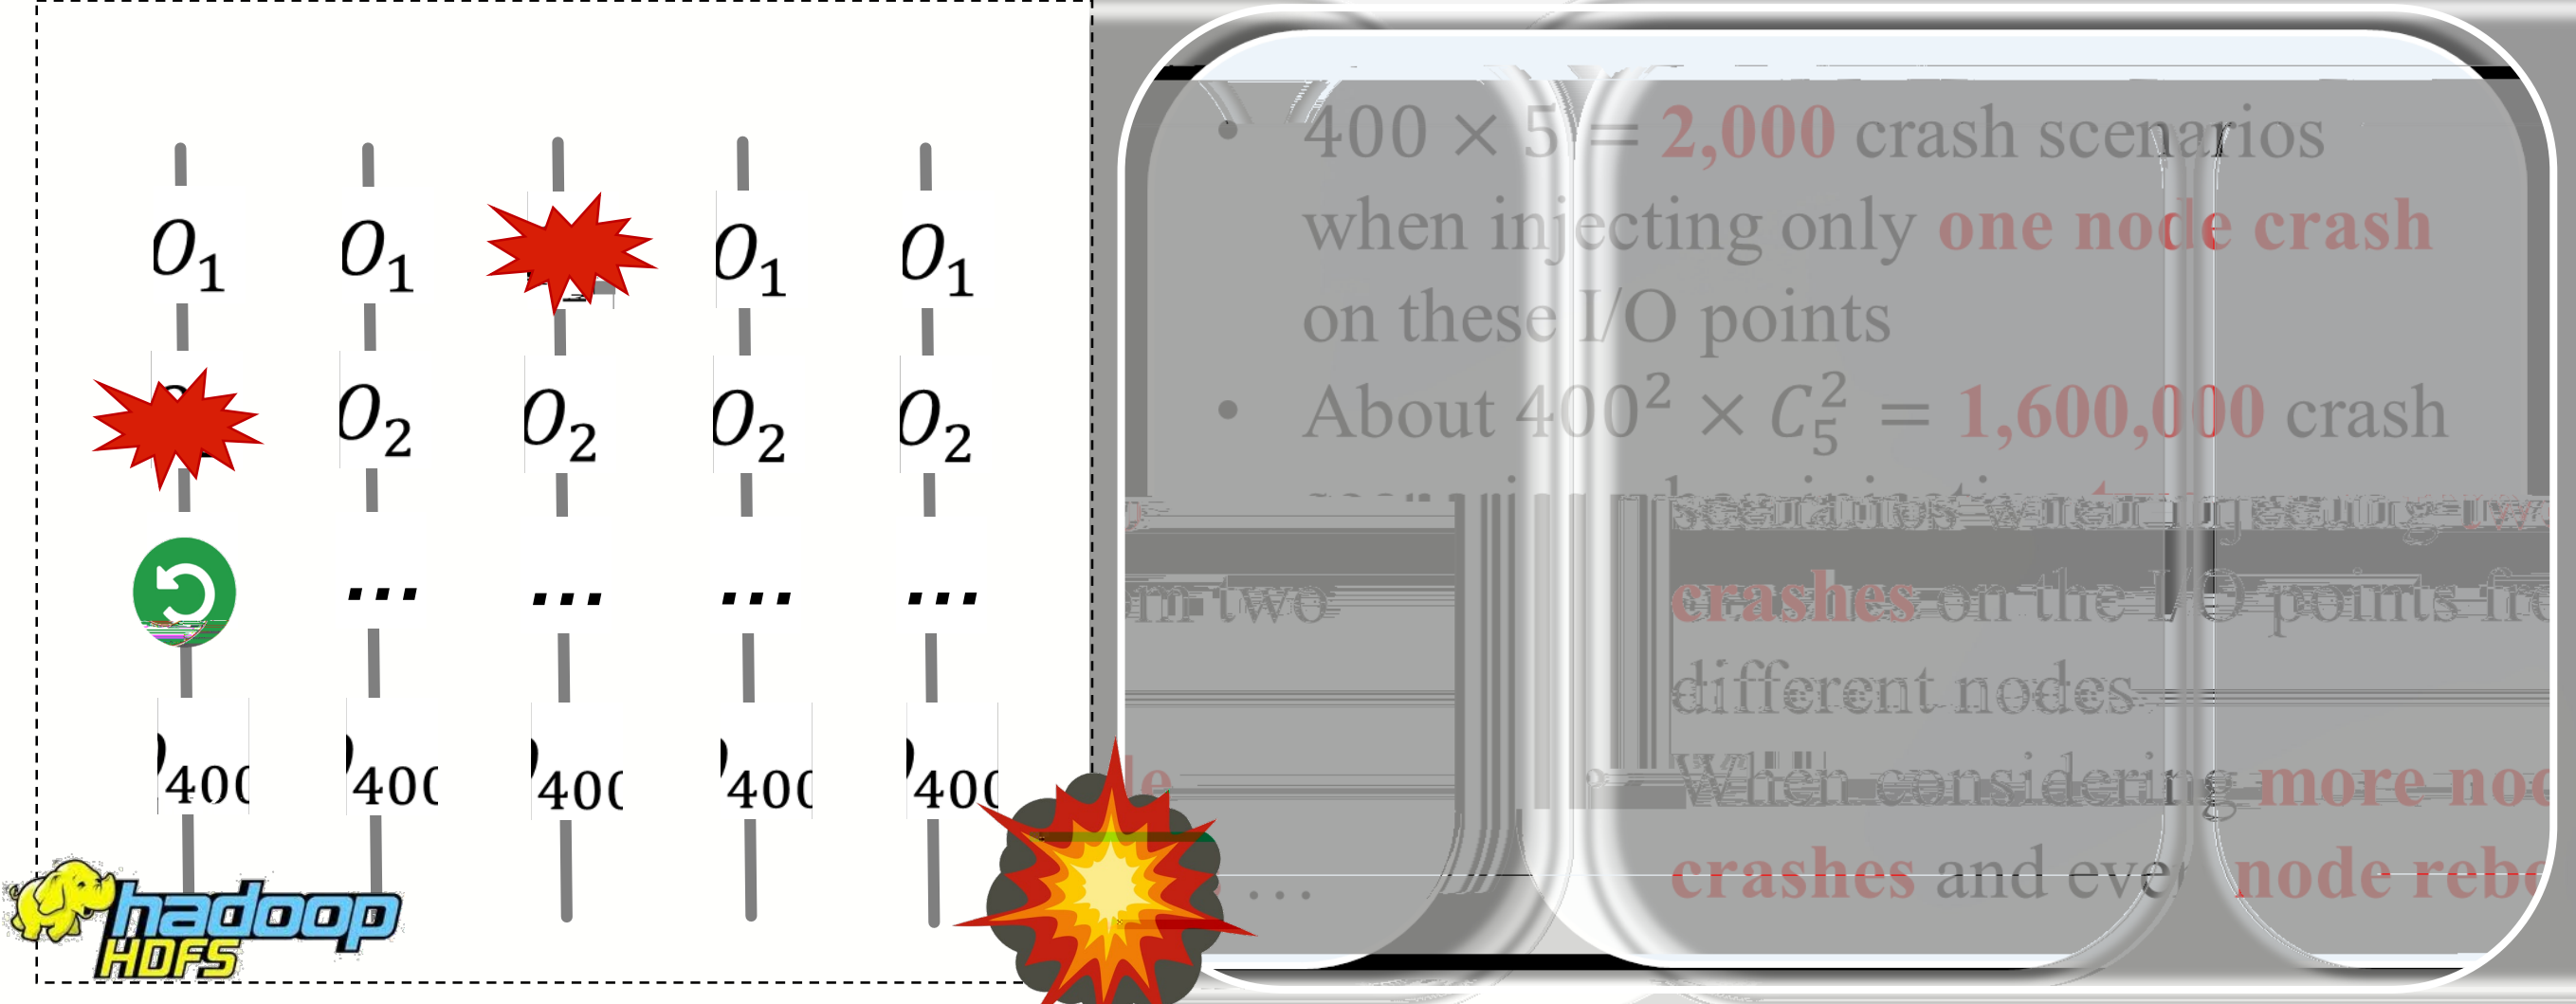


The 45th IEEE/ACM International Conference on Software Engineering (ICSE'23)

13552557320 gaoyu15@otcaix.iscas.ac.cn

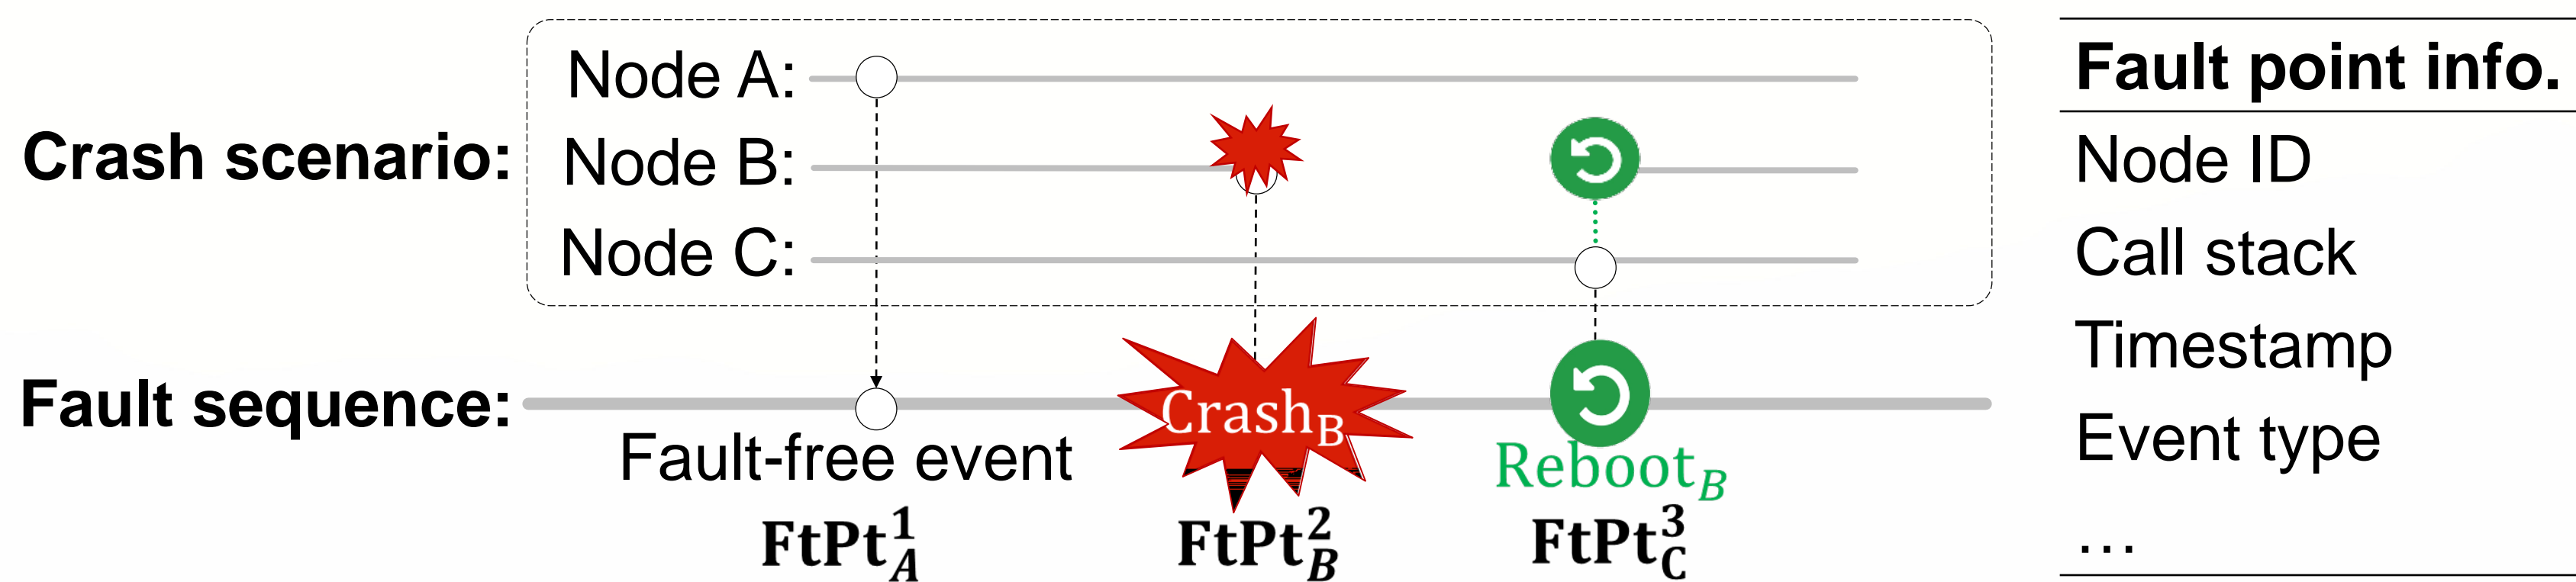
- Cloud systems face huge crash scenario space.
- Some crash scenarios may trigger the same recovery code.

A 5-node HDFS system that produces around 400 I/O operations for each node

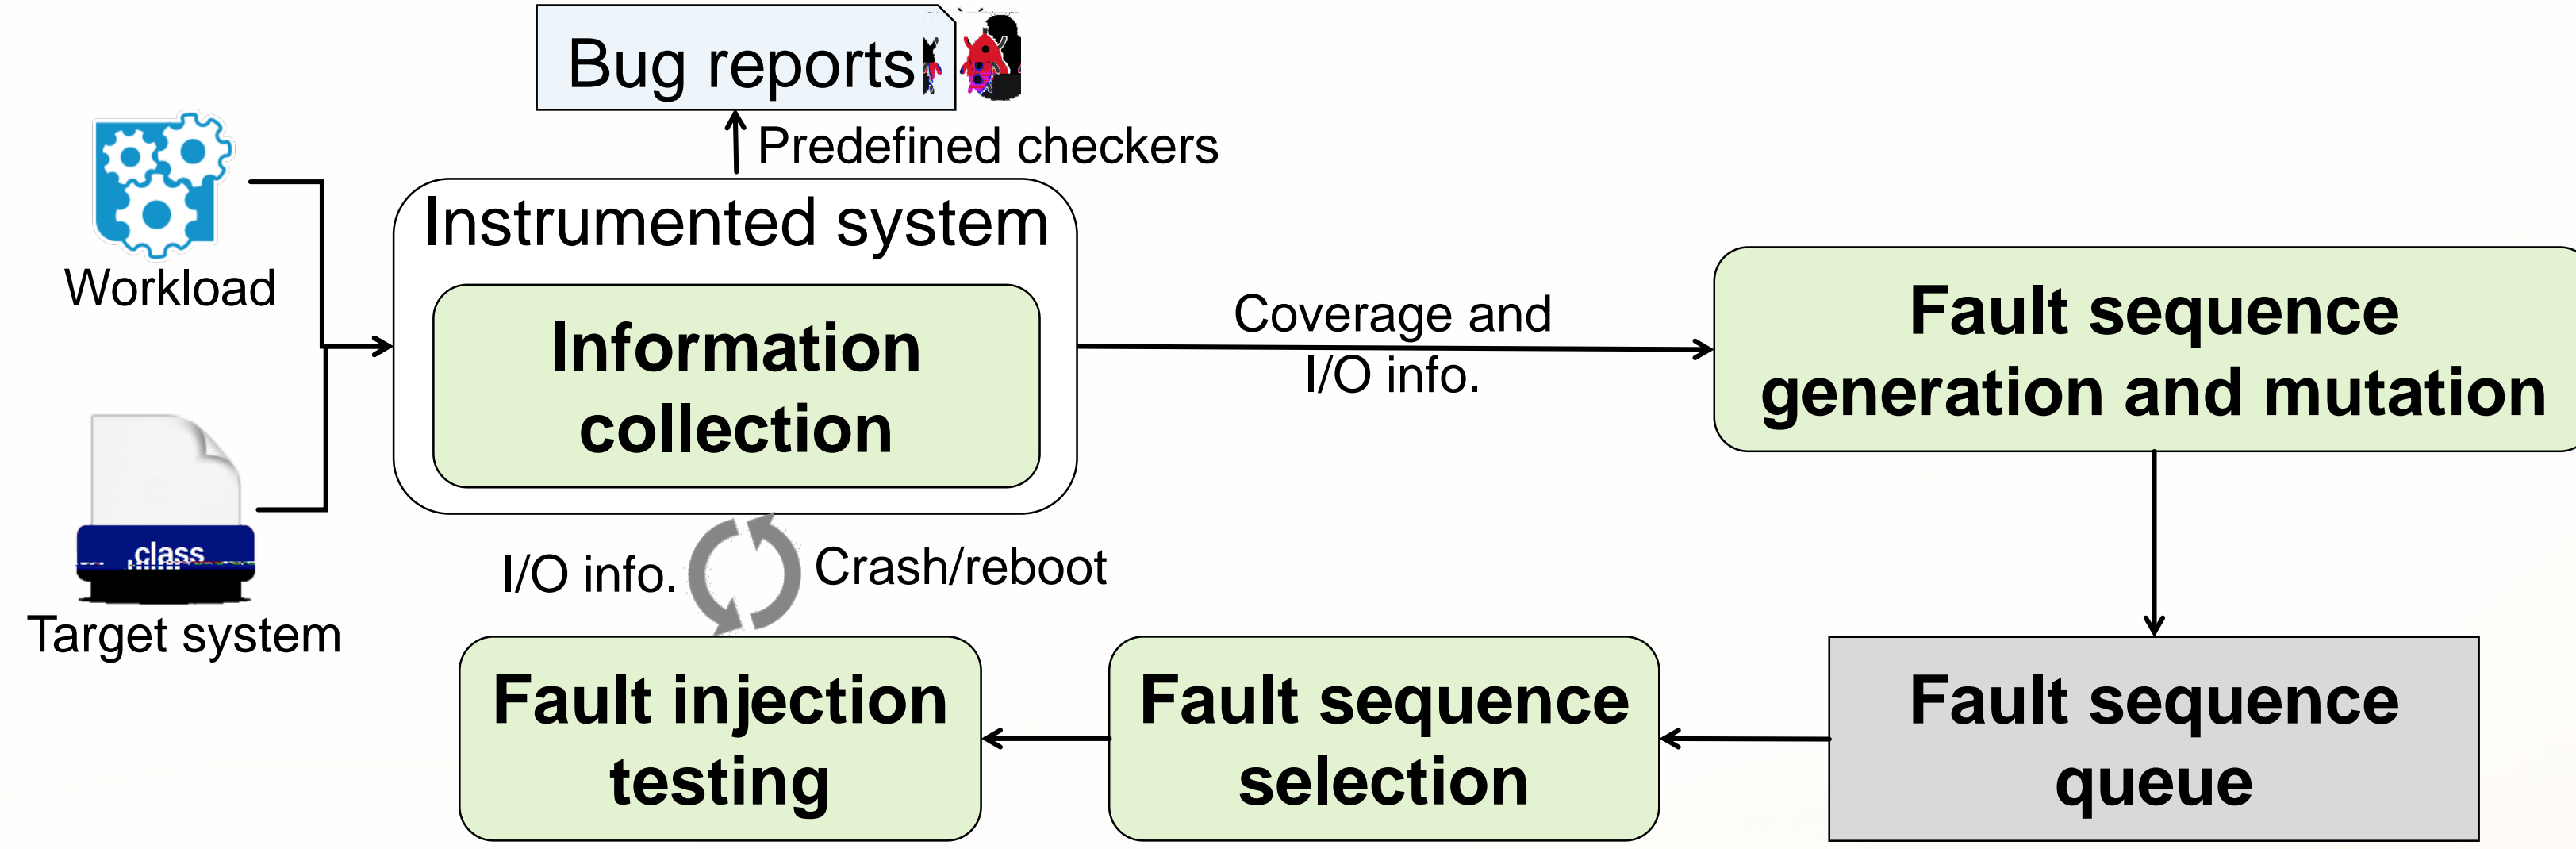


- In cloud systems, any node can crash or reboot at any time. Node crashes/reboots can trigger crash recovery procedures.
- Specific node crashes/reboots can trigger crash recovery bugs hidden in incorrect crash recovery mechanisms and implementations.

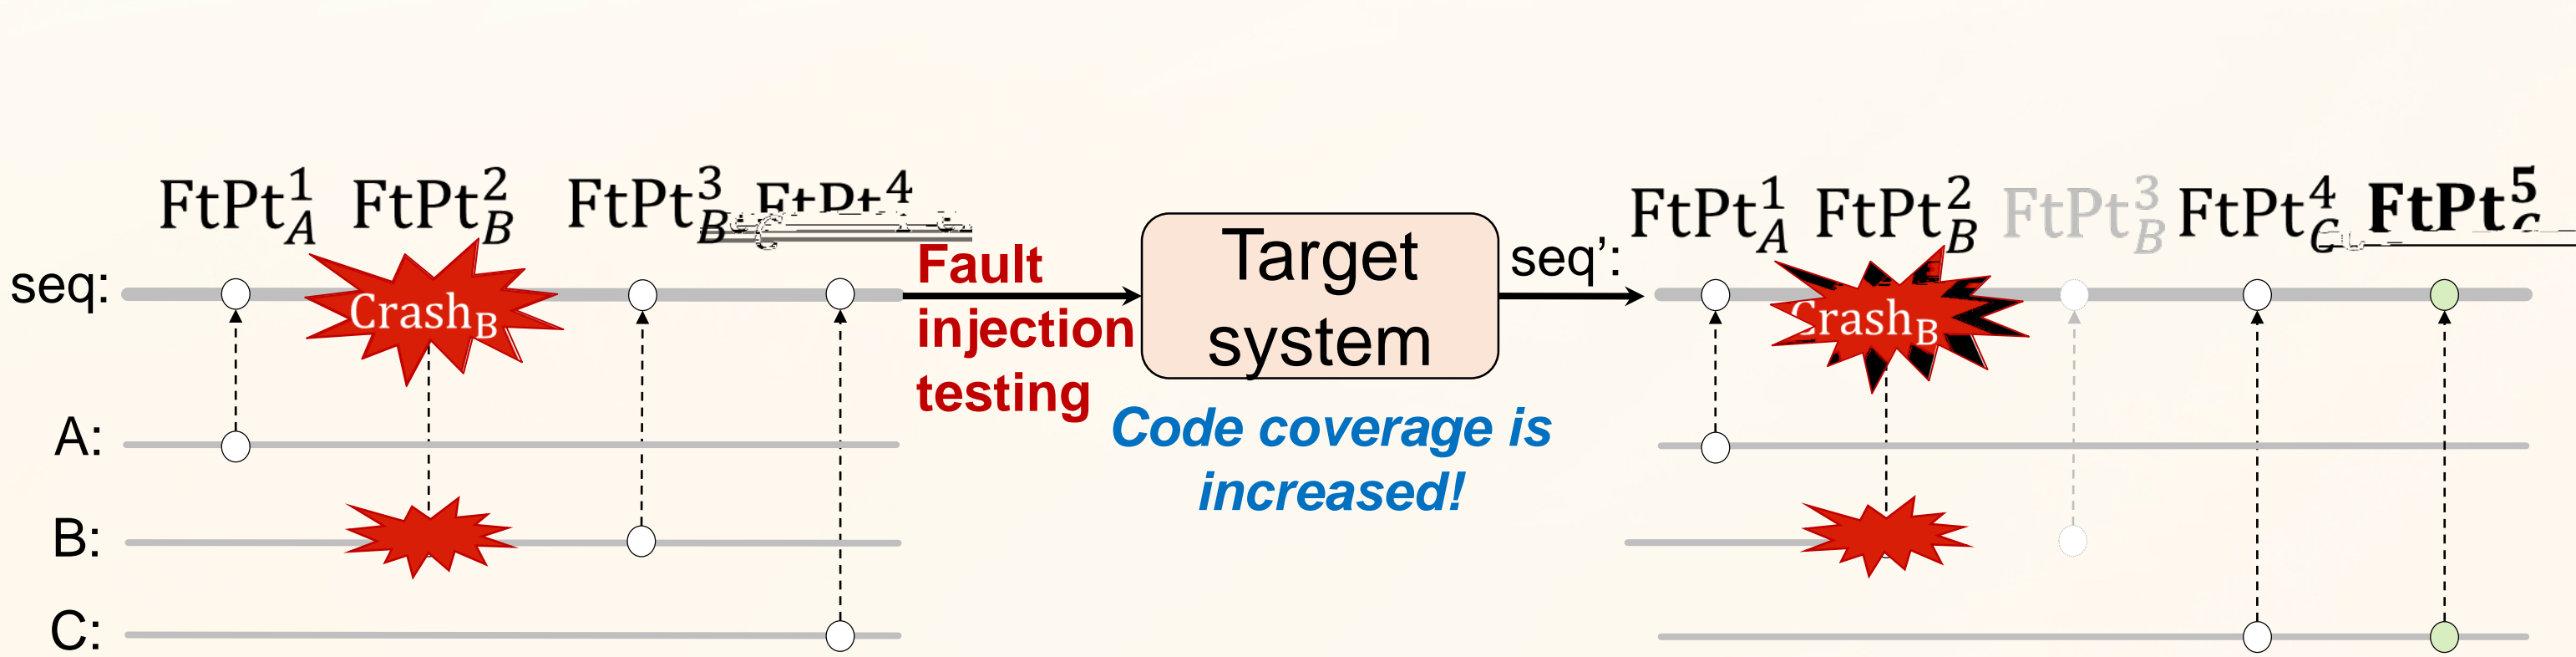
- We use fault sequences to represent various crash scenarios
- Fault sequence: all the I/O points executed in a run and their corresponding events.



and explore the crash scenario space of cloud systems.



- Guide a cloud system to cover new crash recovery code and increase the chance of triggering crash recovery bugs.
- Take a fault sequence as a special system input.
- Adjust fault sequences according to system feedbacks.



Bug ID	Failure Symptoms
HBASE-26883	Data loss
ZOOKEEPER-4503	Data staleness
HBASE-26897	Cluster out of service
HBASE-26370	Misleading error message
HDFS-16508	Operation failure

