

Aiman Priester

Review: Torturing Databases for Fun and Profit

## Summary

- The paper discusses ways to find faults in DB systems. It is implied that the research focuses on version control software as it does committing and transactional changes to test each DB implementation. This project also allows for black box testing every storage layer due to the iSCSI framework. Essentially, the paper goes into detail about what happens during a power outage / failure during a transaction. It then compares the output by logging timestamps and comparing it to keys given to it. If it violates the atomic principles set forth in the implementation, it reports back an invalid state.

## Strengths

- It is likely to be very effective in finding different errors. Due to its very journaled nature with various different checks, bugs were found in all of the DB's tested which was impressive. However, I would have liked to have seen Berkeley DB being tested as it is much more popular than TokyoCabinet.

## Weaknesses

- The flip side of being very journaled is its overhead. On checking for inconsistencies and atomic violations, it does store a significant amount of metadata. However, one could argue that it is unlikely to be run on a live system.

## Unresolved Issues

- As seen in past papers, storage systems can be stacked. The paper leaves the question (ultimately answered in future papers), for how big can you stack the DB's on top of each other before the complexity renders this project inoperable?

## Discussion

- Since this project requires the use of timestamps, how much will the implementation of EXPLODE affect this? This comes from the fact that EXPLODE does not store full stacks, but rather the choices made. Hence the timestamps will be different.