

Example: description of favourite paper

It is difficult to pick among the several interesting papers, so I decided to exclude the ones by my immediate collaborators. My pick is "Using Model Checking to Find Serious File System Errors" by Junfeng Yang, Paul Twohey, Dawson Engler, and Madanlal Musuvathi, from OSDI 2004.

This paper describes FiSC, a file system model checking tool based on CMC (another great paper from OSDI 2002). I like their work because, by applying formal verification techniques, they demonstrate how far reaching systems research is. But at the same time it has a pragmatic, systems approach: unlike traditional model checking that requires human effort to build an abstract specification of the system, called a model, they use the code itself as a model to check.

Furthermore, the FiSC paper involved real world systems. Their evaluation uses three widely-used, heavily-tested file systems: ext3, JFS, and ReiserFS, and they found serious bugs in all of them. For each file system, they found demonstrable events leading to the unrecoverable destruction of meta-data and entire directories, including the file system root. Notably some of these bugs were exploitable by malicious users. Most have led to patches.

To conclude, I like that they did not scale down to a toy problem. File systems are widely used, crucial code. Their errors can destroy persistent data and lead to unrecoverable corruption that cannot be fixed with a reboot. Another point is that they discuss the limitations of their solution, which include false positives and negatives. Finally, I appreciate that their research is not a one-off attack at the problem, and they have pursued commercial applications for their model checking real code research.