



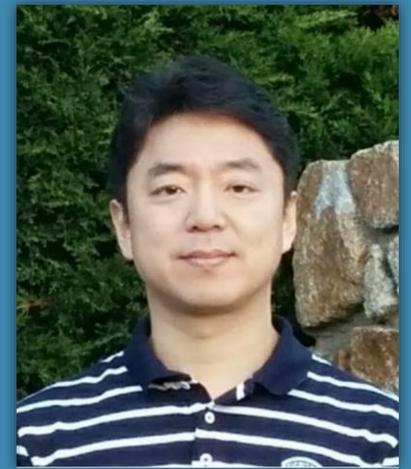
ANDROID IO STACK OPTIMIZATION

Prof. Youjip Won

Department of Computer Science
Hanyang University

2016年9月27日 星期二 2:00pm

理科五号楼410会议室



ABSTRACT: This work is dedicated to resolve the Journaling of Journal Anomaly in Android IO stack. We will overview the IO characteristics of the Android IO stack and will present recent effort on optimizing the Android IO stack, the WALDIO. We orchestrate SQLite and EXT4 filesystem so that SQLite's file-backed journaling activity can dispense with the expensive filesystem intervention, the journaling, without compromising the file integrity under unexpected filesystem failure. In storing the logs, we exploit the direct IO to suppress the filesystem interference. This work consists of three key ingredients: (i) Preallocation with Explicit Journaling, (ii) Header Embedding, and (iii) Group Synchronization. Preallocation with Explicit Journaling eliminates the filesystem journaling properly protecting the file metadata against the unexpected system crash. We redesign the SQLite B-tree structure with Header Embedding to make it direct IO compatible and block IO friendly. With Group Synch, we minimize the synchronization overhead of direct IO and make the SQLite operation NAND Flash friendly.

BIOGRAPHY: Youjip Won received his BS and MS from Dept. of Computer Science and Statistics in Seoul National University in 1990 and 1992, respectively. He received PhD in Computer Science from University of Minnesota, Minneapolis in 1997. Till 1999, he worked for Intel as Server Performance Analyst. Since 1999, he has been with Dept. of Computer Science, Hanyang University, Seoul, Korea. He is a director of Center for Filesystem and Storage Research at Hanyang University. His main research interests include Operating System, Storage System and Filesystem for new storage medium. He has lead the national consortium of five universities and three companies for developing an SSD controller for hyper scale high performance SSD. Currently, he is leading the national consortium of five universities and the four companies for developing OS stack for next generation byte addressable NVRAM. His work on Android received a best paper award on USENIX ATC, 2013. He serves as a reviewer and programming committee for major venues including IEEE ToC, ACM ToS, USENIX FAST, USENIX Hotstorage, ACM APSYS, IEEE ICPADS. He is currently an associate editor for ACM ToS. His group publishes widely open source tool to instrument and monitor the internal behavior of the Android IO: MOST, mobibench, androtrace, vssim.