



SDSU
presents
a thesis defense for
Master of Science
degree in
Computer Science

Wednesday,
October 24, 2012

5:00pm
GMCS 405

Xinhua Fahy

A New RAID Linux Flash File

Abstract

NAND flash memory has been widely used in the embedded market because of its desirable properties such as non-volatility, shock resistance, low power and low latency for reading and writing. However, hardware limitations like the need to erase before writing again, erase block sizes span multiple write pages, and limited memory cell lifespan, make file system requirements for flash memory very different from those for conventional block devices. As such, conventional file system cannot be applied directly to flash memory.

To address these issues, two different approaches have been proposed: (1) add a software layer that emulates a block device using the flash device, and (2) create special-purpose file systems that are targeted to flash devices. The first approach, which is based on a Flash Translation Layer (FTL) concept, succeeds in making flash devices universally accessible, but hiding flash device characteristics from the file system inevitably leads to reduced performance compared to the second option, of creating “flash-aware” file systems.

In contrast, the flash-specific file system is designed for direct use on raw flash chips. It proves to be a better solution on flash memory. However, there is currently no support for flash-based RAID file systems in standard Linux distributions. RAID support can provide performance enhancements that enable video rate recording of data, which has already been demonstrated in some high performance video cameras. RAID support is particularly important for NAND flash devices because they have a high rate of cell failure as they age compared to disk drives. As NAND flash devices continue to fall in price at a faster pace than disk drivers, price barriers will be reduced or eliminated. As such it is anticipated that RAID-capable flash file systems will eventually become a fundamental requirement for Linux operating systems.

There are currently three dominant flash file systems for Linux: JFFS2, YAFFS2, and UBIFS. JFFS2 has issues scaling to large sizes and does not write the index data to flash as it goes, but instead reconstructs it every time the system mounts. This means mounting a large JFFS2 file system can take a long time, which would only be worsened in the context of RAID. YAFFS2 and UBIFS do write their indexes to flash as they go but have other issues. UBIFS is the most advanced of the three but is somewhat immature, very complex (three times the code size of JFFS2), and has high storage overhead for the index (over 10%). YFFS2 is not stable, consumes a lot more memory and has the worst performance in many cases.

This thesis adds enhancements to JFFS2 by eliminating its main deficiencies, adding RAID capability, making all the IO operations asynchronous to reap performance. The implemented RAID-based file system, JFFS2R, is compared to JFFS2, YAFFS2, and UBIFS and analysis is implemented.

Keywords: File System, Flash memory, RAID

Thesis Committee

Tao Xie, Thesis Chair, Department of Computer Science
Joseph Lewis, Department of Computer Science
Jianwei Chen, Department of Mathematics & Statistics