NVMFS: A Hybrid File System for Improving Random Write in NAND-flash SSD

Sheng Qiu, A.L. Narasimha Reddy
Texas A&M University

Outline

- Introduction
- Design
- Implementation
- Evaluation
- Conclusion

Introduction

- Motivation
  - Random writes hurt lifetime & performance of flash SSD.
  - NVRAM has good random performance for both read/write.
- How to utilize NVRAM to improve performance and lifetime of SSD?

Design

- Hybrid file system and hybrid storage

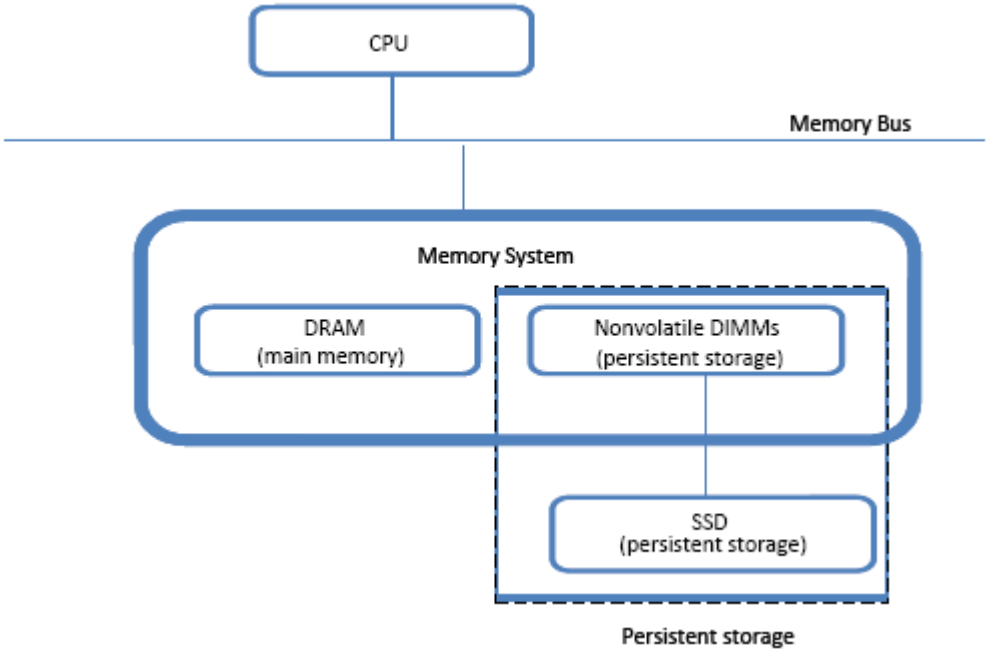

Design

- NVRAM as small cache and persistent storage
  - Newly allocated data
  - Hot file data and metadata
  - Temporarily accessed data
- Flash SSD as large persistent storage
  - Large relatively cold file data
- Writes to SSD are in units of 256KB

Implementation

- Track data hotness on NVRAM.
- Flush dirty LRU data to SSD when NVRAM is not sufficient.
- Do segment cleaning on SSD when SSD's space is too fragmented.

Implementation

- Two LRU lists for tracking data hotness

Fig. 2. Dirty and Clean LRU lists

Implementation

- Dirty LRU data are written back to SSD in units of 256KB.

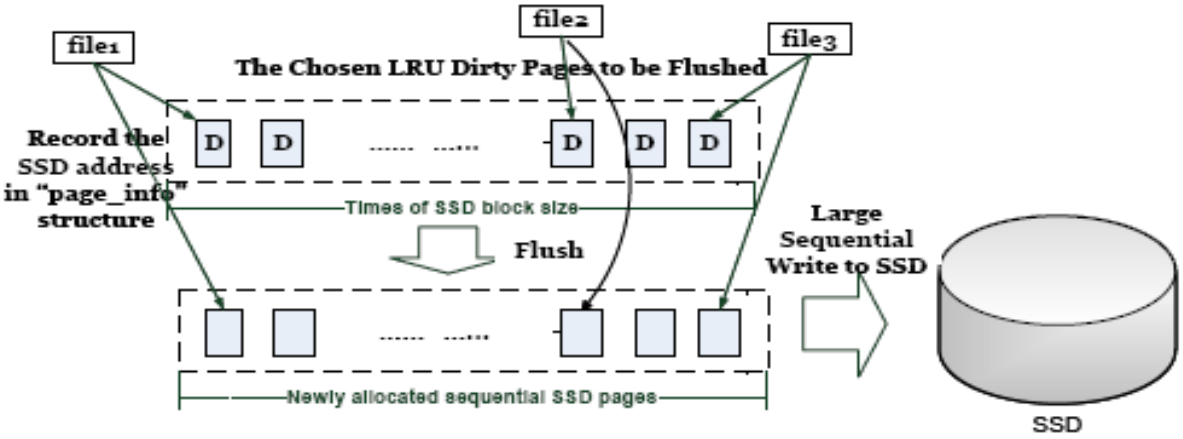

Fig. 3. Migrate Dirty NVRAM Pages to SSD

Implementation

- Segment Cleaning on SSD

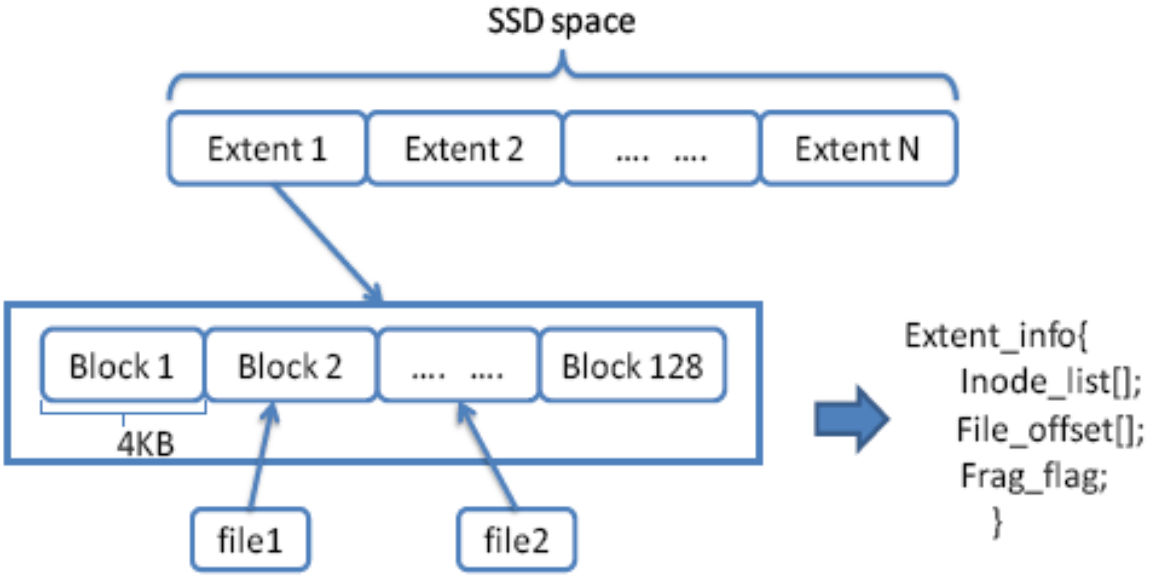

Evaluation

- Reduced writes to SSD
- Reduced erase overhead on SSD
- Improved performance on SSD

Reduced writes to SSD

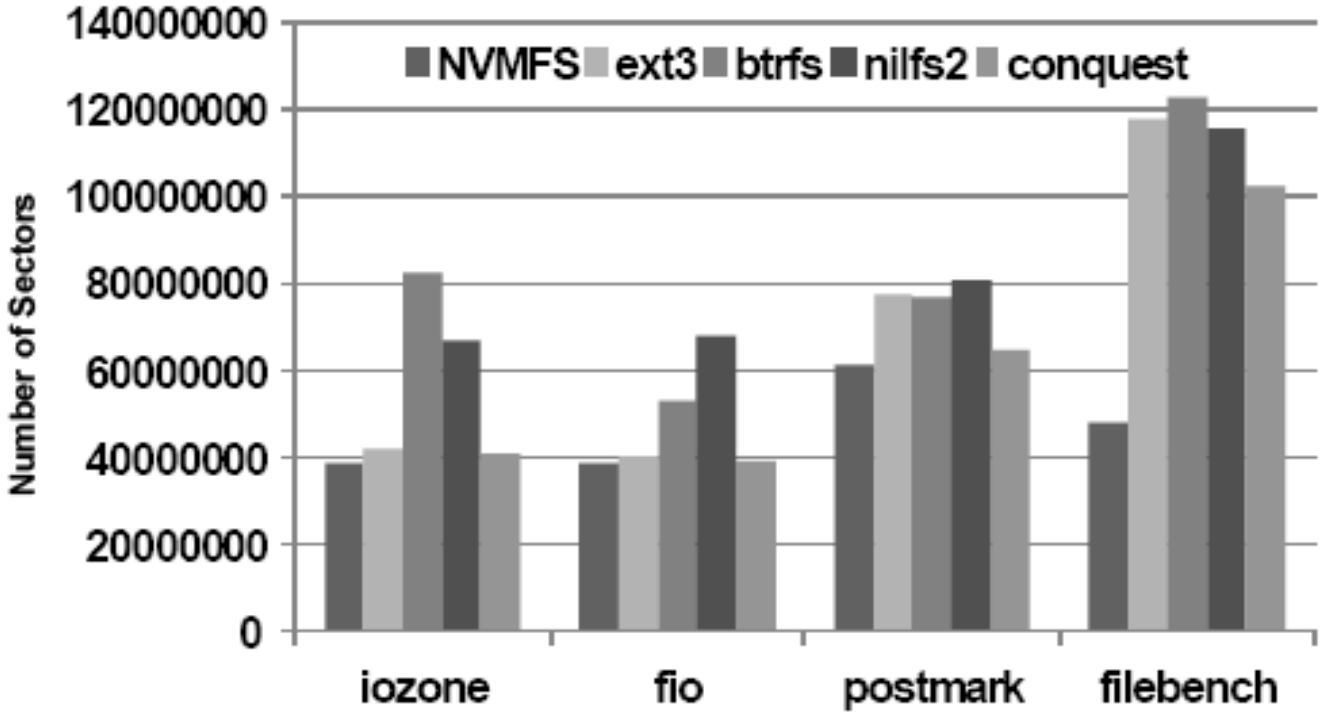

Reduced GC overhead

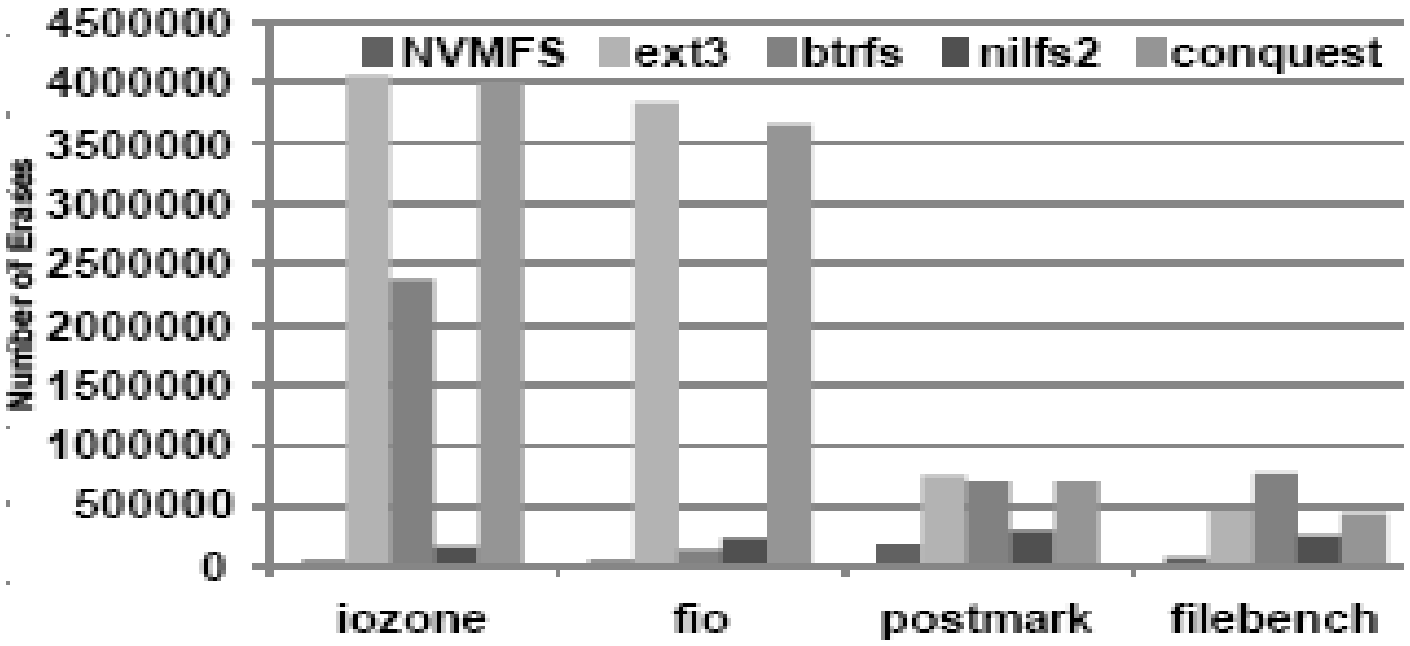

Improved IO throughput

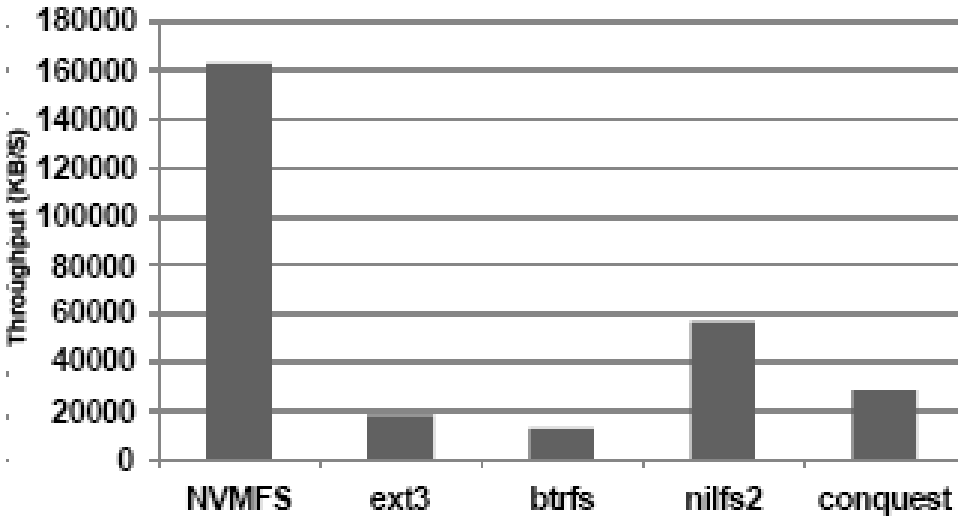

<a> 50-70% FIO

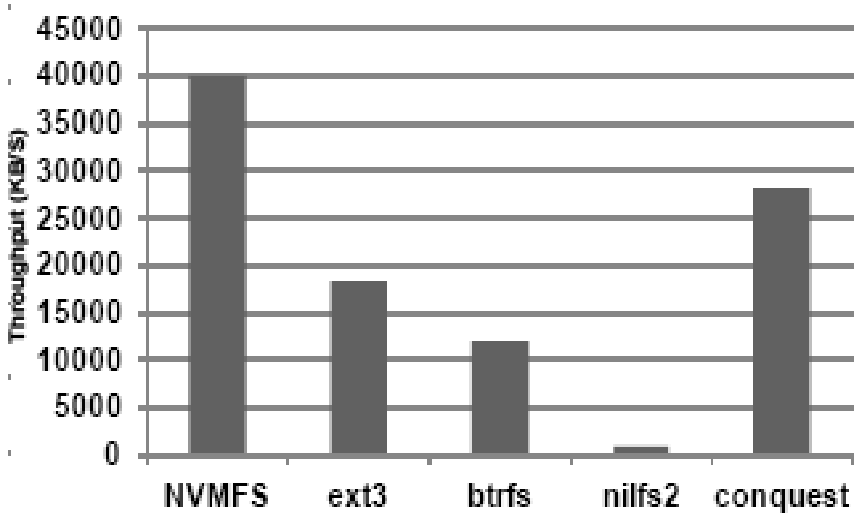

<b> over 85% FIO

Conclusion

- Proposed a hybrid FS for NVRAM+SSD based hybrid storage
- Efficiently utilized NVRAM to reduce random writes to flash SSD
- Demonstrated reduced erase overhead and improved FS performance on SSD