



Flash Memory Summit

Memory Expansion Technology Using Software-Controlled SSD

○ S. Kazama*, S. Gokita*, S. Kuwamura*, E. Yoshida*, J. Ogawa*, Y. Honda**

*Fujitsu Laboratories Ltd.

**Fujitsu Ltd.

Contact: sc-ssd-fms2017@ml.labs.fujitsu.com



Challenge

- Flash memory has about 1000 times slower performances than that of DRAM, so that it usually makes any system performance greatly degraded

Challenges to the performance gap:

- To avoid to store data in slower flash as much as possible?
- To have highly parallelism in hardware?
- To step away from FTL housekeeping operations?
- To wipe off frequent data transfer between flash and CPU?

The answer is ... “Memory Expansion Technology,” which we have developed

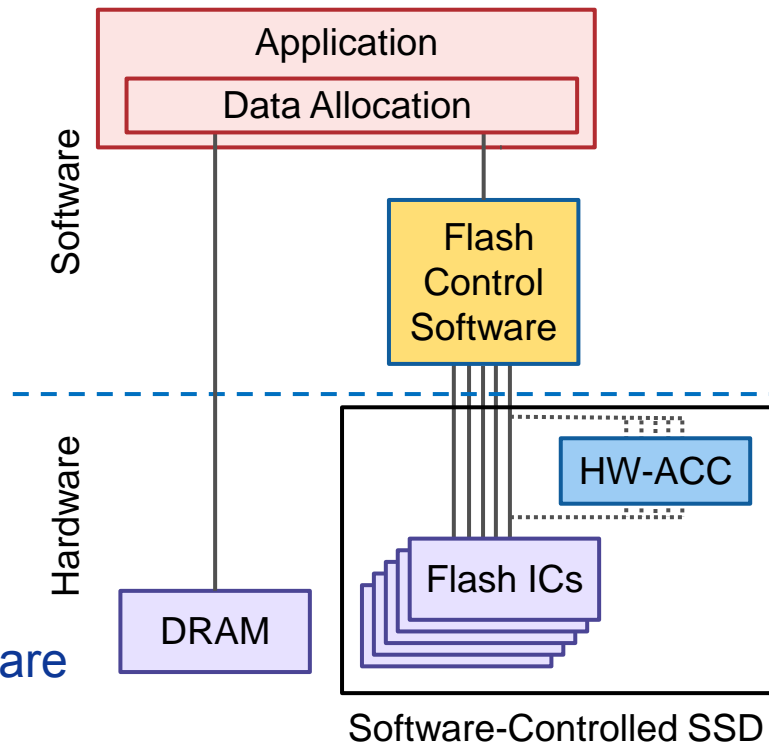


Memory Expansion Technology Architecture

Flash Memory Summit

- Data allocation control
 - Software control to allocate data to either DRAM or flash according to application data characteristics
- Flash IC manipulation
 - Dedicated highly parallel flash IC hardware in order to realize high BW and to control the data allocation
 - SW takes over GC/WL-like housekeeping operation
- Hardware accelerator (HW-ACC)
 - Near data processing, with accelerators close to flash, can easily wipe off data transfers btw flash and CPU

We developed “Software-Controlled SSD” hardware to realize the memory expansion technology



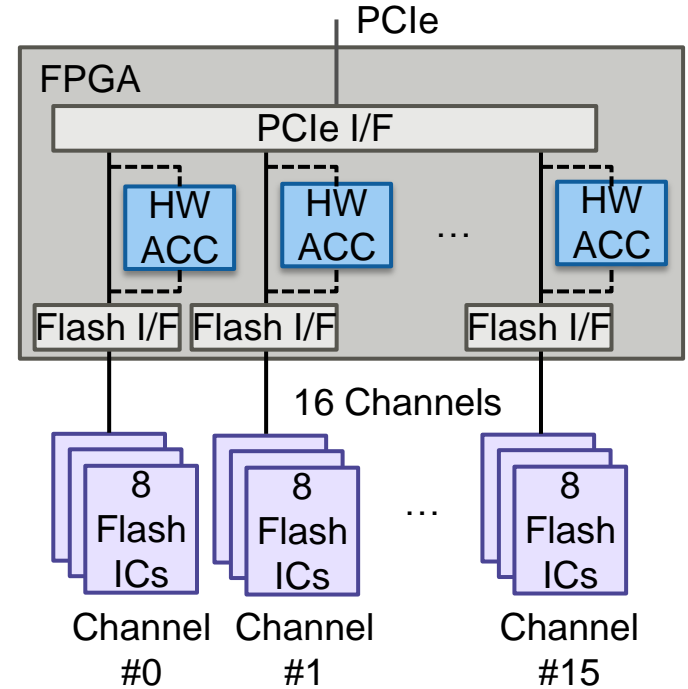


Software-Controlled SSD

- 128 parallel flash-ICs controlled by software
- 16 parallel hardware accelerator




PCIe Gen3 x 8
8 TB
FHHL form factor





Case 1: Memcached Server

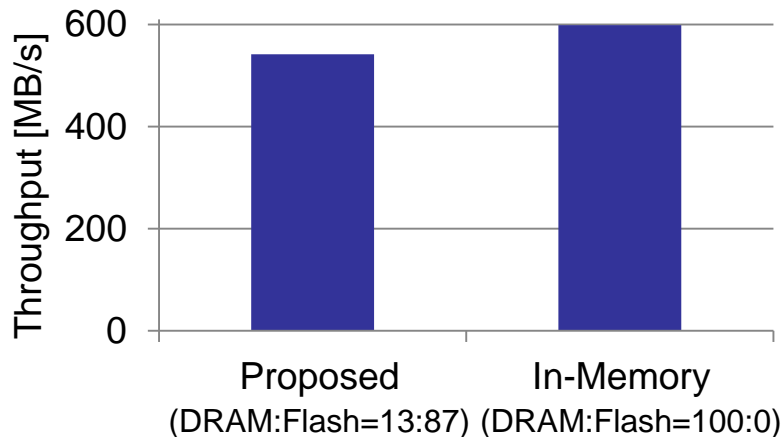
- Memcached server is in-memory KVS, and usually used as a cache server for web services
- In terms of large-scale web services, several thousand cache server units operates at a time in order to increase cache capacity
- The memory expansion technology can reduce memory cost of servers





Memcached Performance Evaluation

- Even when 87% of DRAM is replaced with flash, evaluation results show that performance degradation can be practically suppressed to less than -10%
- Estimated memory cost goes around 1/6



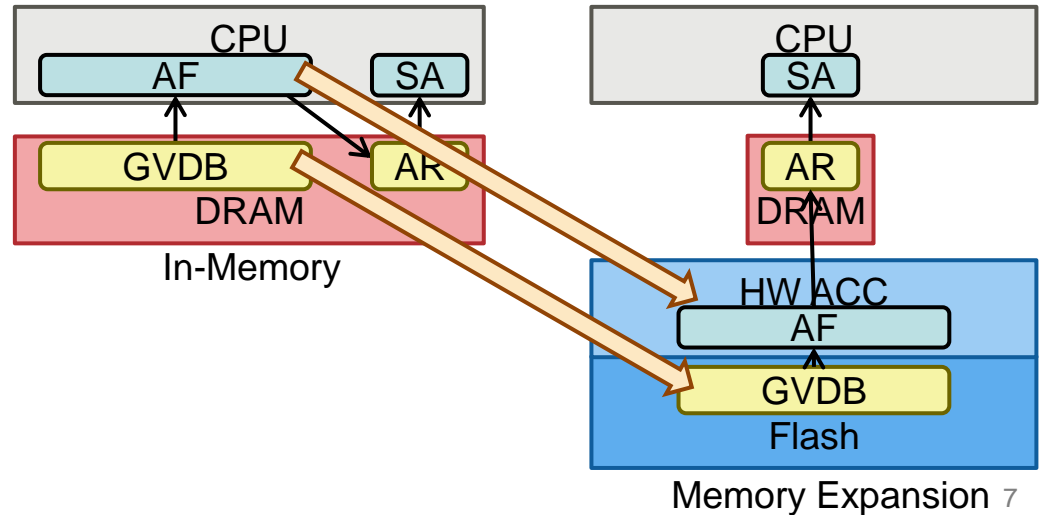


Case 2: Genome Analysis

- Genome-wide association study
 - Search and find out correlations between disease and gene variant (GV)
 - Aggregation function (AF) for GV database (GVDB): 500GB per 100K people
 - Statistical analysis (SA) for aggregation result (AR) : 640MB

- Approach

- Place GVDB in flash
- Process AF with hardware accelerator

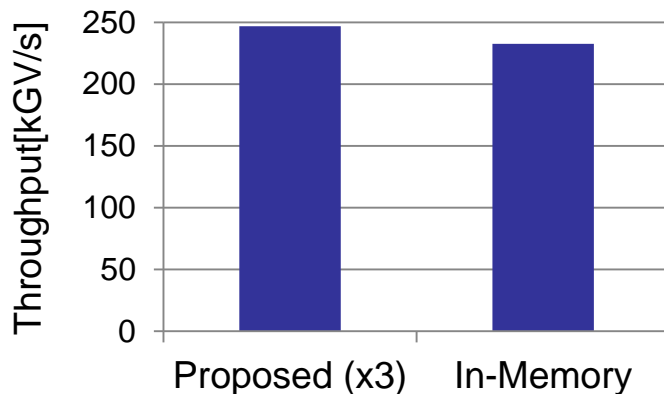




Genome Analysis Performance Evaluation

Flash Memory Summit

- Using three software-controlled SSDs in parallel
- Evaluated throughput is almost equivalent to “in-memory processing” results
- Estimated memory cost goes around 1/7
- Resultantly, DRAM can be well replaced with the memory expansion technology





Summary

- Developed “Memory Expansion Technology” using software-controlled SSD
 - Try to fill performance gap between DRAM and flash
- Applied to both memcached server and genome analysis applications
 - Proposed “Memory Expansion Technology” can practically reduce memory cost of servers, at the expense of around -10% performance degradation



Flash Memory Summit

Thank you!

FUJITSU