Junchao CHEN, Guangyan ZHANG, Junyu WEI, 2023. A survey on design and application of open-channel solid-state drives. *Frontiers of Information Technology & Electronic Engineering*, 24(5):637-658.
https://doi.org/10.1631/FITEE.2200317

A survey on design and application of open-channel solid-state drives

Key words: Domain-specific storage; Flash translation layer; Garbage collection; Internal parallelism; Open-channel solid-state drives (OCSSDs)

Corresponding author: Guangyan ZHANG
E-mail: gyzh@tsinghua.edu.cn
ORCID: https://orcid.org/0000-0002-3480-5902

Motivation

1. Traditional solid-state drives (SSDs) have several well-documented shortcomings such as write amplification, long tail latency, and log-on-log problems because they are managed by the built-in flash translation layer (FTL).
2. Open-channel SSDs (OCSSDs) can achieve better performance due to the exposure of their physical layout and the host-side well-designed FTL.
3. The survey on design and application of open-channel SSDs can provide fundamental knowledge to researchers and promote better research on emerging flash memory.

Overview of open-channel SSDs

1. Move the FTL to the host side
2. A hierarchical address space for hardware accessing
3. The exposure of physical layout

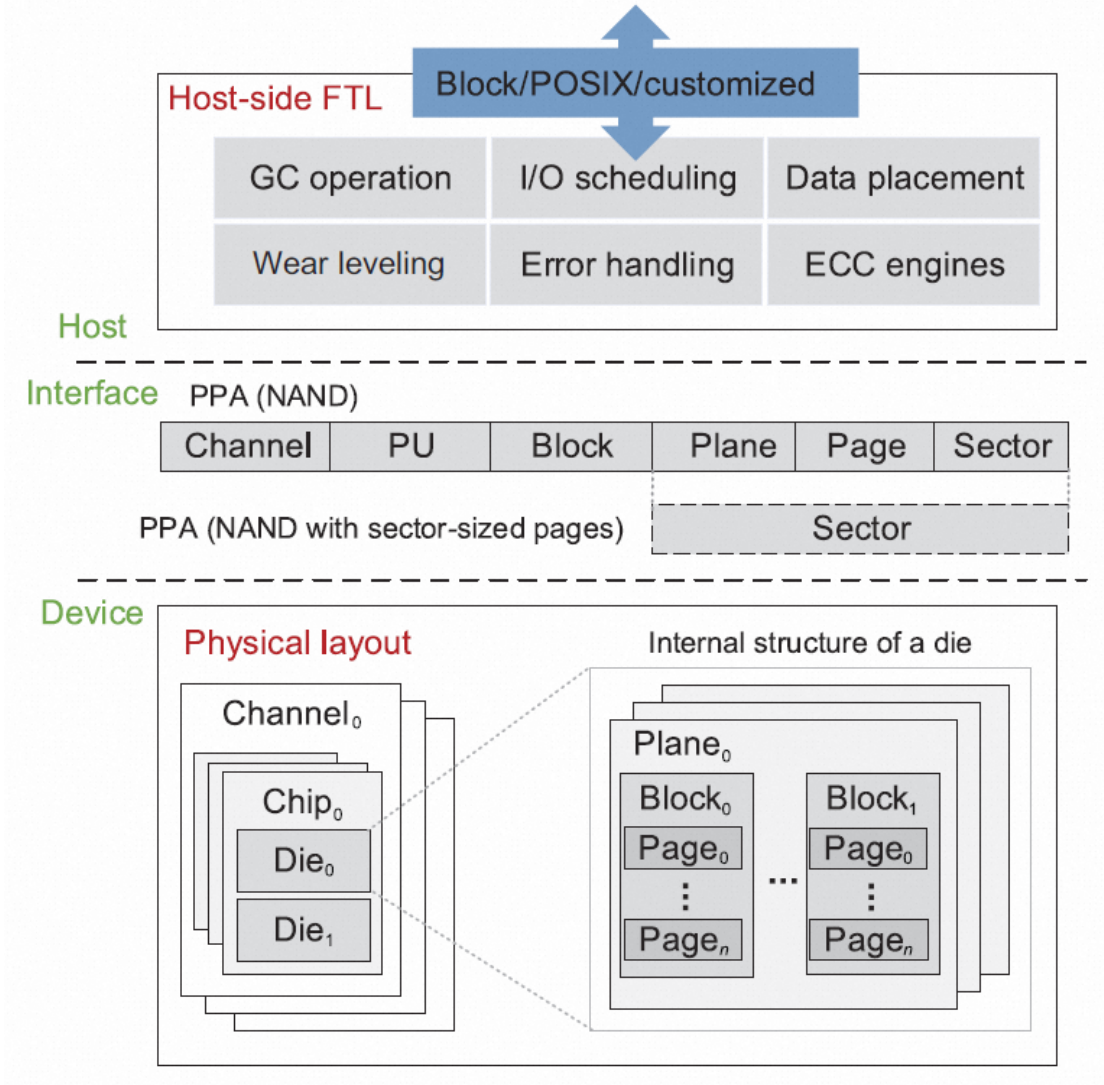

Opportunities

1. OCSSDs bring many opportunities for researchers in academia and industry to achieve specific performance goals or address practical problems in different scenarios.
2. Performance opportunities:
   - High throughput
   - Low latency
   - Long flash lifetime
   - Strong performance isolation
   - High resource utilization

Methodologies

1. Interface design
   - High-level I/O interface
   - Co-design with host side
   - Physical page addresses (PPAs)
2. FTL design
   - Design dedicated FTLs with different placement strategies to achieve specific performance.

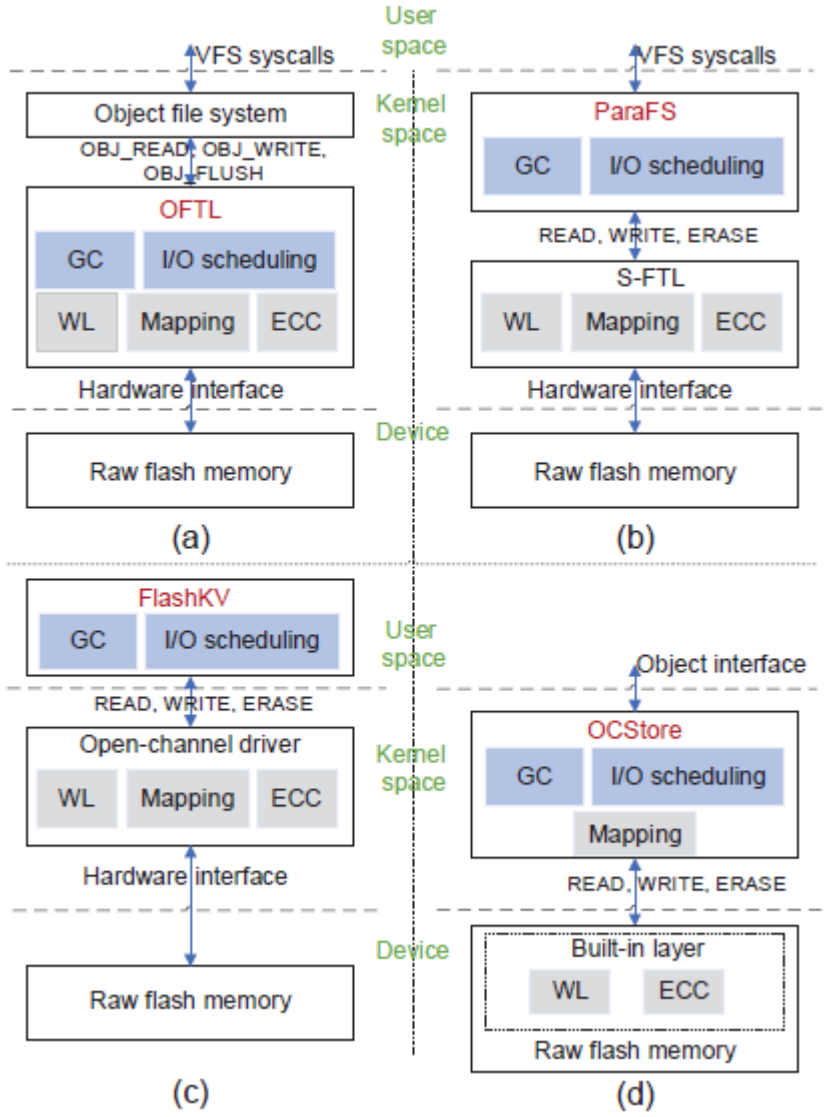

Methodologies (Cont'd)

1. Interface design
2. FTL design
3. Internal parallelism exploitation
   - Independent channel parallelism
   - Coarse-grained striping
   - Channel-level parallelism with hotness grouping
   - Fine-grained parallelism

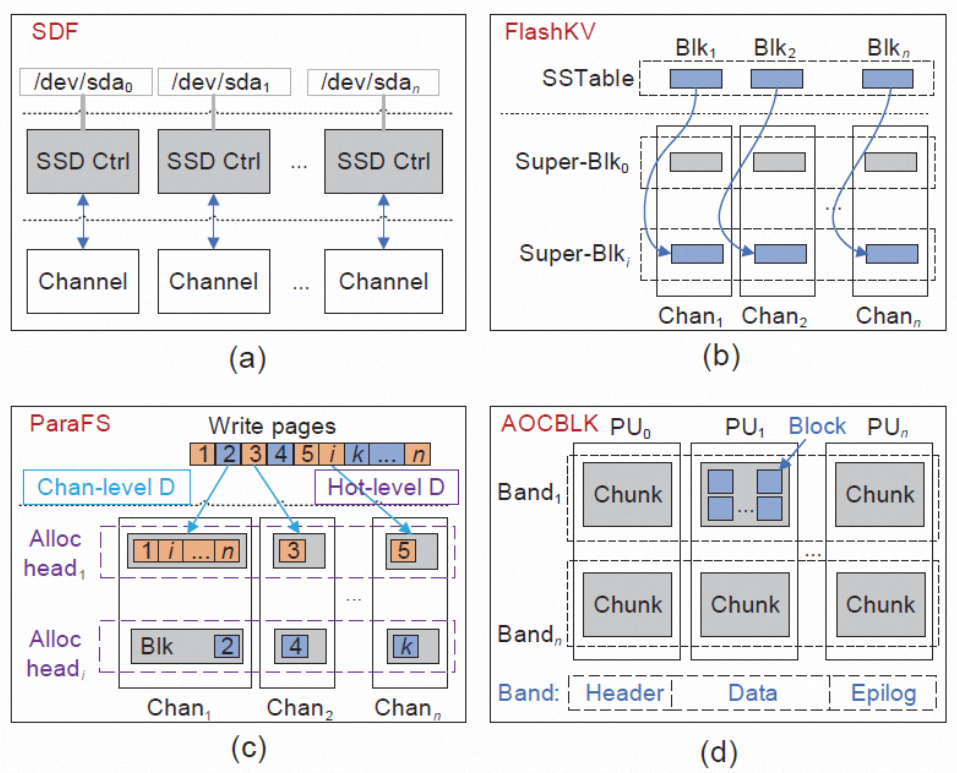

Fig. 6 Four instances of exploiting OCSSD's internal parallelism: (a) SDF; (b) FlashKV; (c) ParaFS; (d) AOCBLK. Chan: channel; Blk: block; D: dimension; PU: parallel unit; OCSSD: open-channel solid-state drive (References to color refer to the online version of this figure)

Methodologies (Cont'd)

1. Interface design
2. FTL design
3. Internal parallelism exploitation
4. I/O scheduling optimization
   - For overall throughput
   - For read throughput
   - For latency

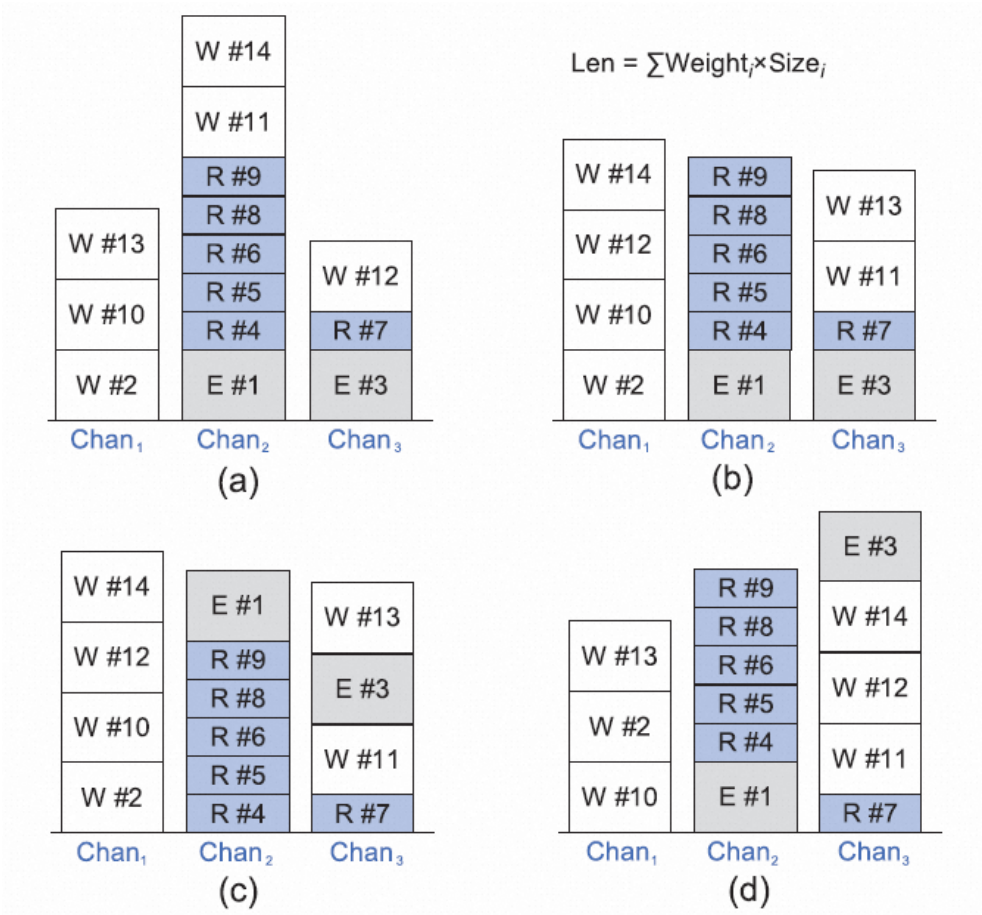

Fig. 7 I/O scheduling policies in LOCS and FlashKV: (a) round-robin dispatching; (b) least weighted-queue-length; (c) scheduling optimization for erase; (d) priority-based scheduling. I/O: input/output; Op: operation; Chan: channel; Pri: priority

Methodologies (Cont'd)

1. Interface design
2. FTL design
3. Internal parallelism exploitation
4. I/O scheduling optimization
5. Garbage collection (GC) optimization
   - Data locality
   - Hotness grouping
   - Write granularity adjustment
   - Coordinated GC mechanism
   - …

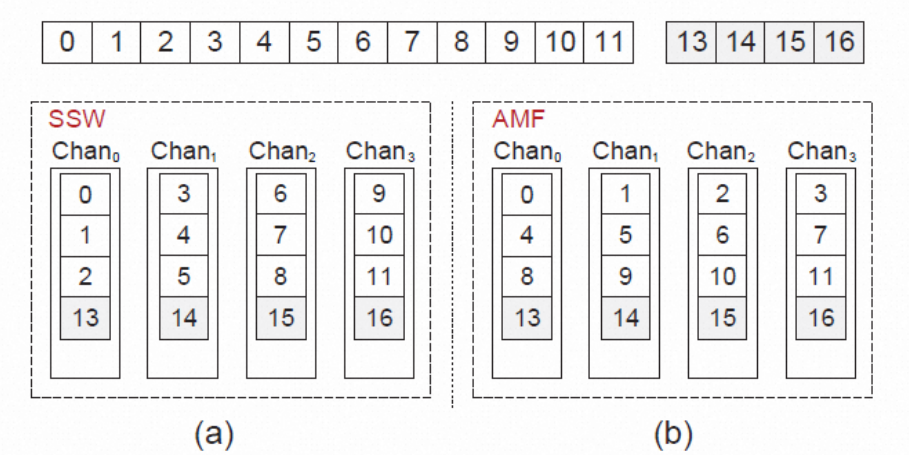

Fig. 9 Data placement strategies in SSW (a) and AMF (b). Chan: channel

Challenges and future directions

1. Challenges
   - Complexity and nonversatility of the FTL design
   - Lack of standards
   - Flexibility, security, and universality
2. Future development
   - NVMe SSDs and multi-stream SSDs earn a great application prospect for their standard and semi-open design
   - Despite facing many huge challenges, users can have more choices to leverage OCSSD’s features for optimal performance and specific goals
   - OCSSDs should form their own standards and norms to better promote their research and application

Conclusions

1. By moving the FTL to the host side and co-designing with host-side software, open-channel SSDs can achieve better performance such as higher throughput, lower latency, and better resource utilization.
2. Existing studies give us many insightful ideas such as rich FTL co-design at different software layers, full parallelism exploitation at different parallel levels, rational I/O scheduling policies to improve performance, and efficient GC mechanism to reduce performance overhead.
3. Some research issues are still open, and future research efforts especially in domain-specific areas, need to take all factors into consideration, including standards, universality, products, and performance.

Junchao CHEN received his BS degree from the National University of Defense Technology (NUDT), Changsha, China in 2009. He is pursuing his MS degree in the Department of Computer Science and Technology at Tsinghua University. His research interests include data compression, big data storage and management.

Guangyan ZHANG received the PhD degree from the Department of Computer Science and Technology, Tsinghua University, Beijing, China, in 2008. He is now an associate professor with the Department of Computer Science and Technology, Tsinghua University. His current research interests include big data computing, network storage, and distributed systems.

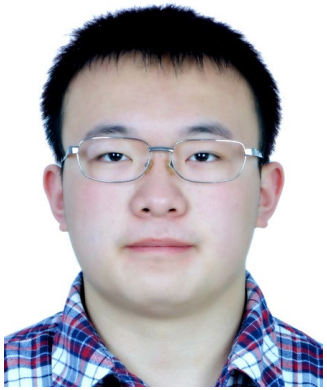

Junyu WEI received his BS degree from the Department of Computer Science and Technology, Tsinghua University, China in 2019. He is pursuing his PhD degree in the Department of Computer Science and Technology, Tsinghua University, China. His research interests include data compression, storage and cloud log management.