CS561 Spring 2025 - Research Project

Title: Concurrency-aware Algorithm Design for ZNS SSDs

Background: Modern storage devices like ZNS SSDs offer high throughput and stable performance, by giving the host more control over data placement and management. However, the applications need to design algorithms that take advantage of the device's internal parallelism and achieve maximum bandwidth. Recent research has developed a parametric I/O(PIO) model for characterizing the device based on read/write asymmetry and access concurrency [1]. PIO was designed for conventional block interface. However, to adapt this model for the ZNS interface there are few key differences to consider:

(1) Zones have to be written sequentially
(2) Data placement strategy has impact on concurrency
(3) ZNS interface limits the number of zones that can be written in parallel

In ZNS SSDs zone size and striping ways can impact the parallelism. For instance, if the ZNS SSD has small zones, the host should maximize parallelism across the zones to take advantage of the device's concurrency [3]. In large zone SSDs intra-zone parallelism might be more optimal.

For this study we propose using a state of the art SSD simulator ConfZNS++[2]. Commercial SSDs often keep the internal design details hidden from users. However, the emulator allows fine-grained control over SSD configuration such as number of chips, channels, block mapping and striping mechanisms.

Objective: The objective of the project is to understand how to design systems for the ZNS interface to maximize the concurrency.

Steps:

- Explore the SSD architecture and configuration options through the emulator
- Understand the ZNS interface, I/O and zone management utilities
- As an initial experiment use ConfZNS++ to emulate ZNS SSDs with different sizes and striping mechanisms currently supported by the emulator. Then create a multi-threaded workload that writes to each zone sequentially and compare the inter and intra zone performance for each configuration.

Responsible mentor*: Teona Bagashvili*

[1]Papon, Tarikul Islam, and Manos Athanassoulis. "A parametric I/O model for modern storage devices." *Proceedings of the 17th International Workshop on Data Management on New Hardware*. 2021

[2]Doekemeijer, Krijn, et al. "Exploring I/O Management Performance in ZNS with ConfZNS++." *Proceedings of the 17th ACM International Systems and Storage Conference*. 2024.

[3] Im, Minwoo, Kyungsu Kang, and Heonyoung Yeom. "Accelerating RocksDB for small-zone ZNS SSDs by parallel I/O mechanism." *Proceedings of the 23rd International Middleware Conference Industrial Track*. 2022.