[TG3-R-3]

FTL Optimization with Linear Regression-Based Mapping Management

Seokyoon Choi, Dongkun Shin
Department of Electrical and Computer Engineering, Sungkyunkwan University; Korea
E-mail : dongkun@skku.edu

The Flash Translation Layer(FTL), a key component managing the flash memory of an SSD, performs essential tasks such as address translation, garbage collection, and wear leveling. To enable fast address translation, FTL stores a mapping table in the SSD's internal DRAM. However, as SSD capacity increases, the size of the mapping table grows proportionally, making it challenging to store the entire table in memory. LearnedFTL [1], which combines DFTL and learned indexes, addresses this issue but is constrained by low model capacity—the proportion of Logical Page Numbers whose Physical Page Numbers can be directly utilized by the model. Low model capacity increases double-read occurrences, where a flash read is performed twice: first to access the mapping table and then to access the actual data. To address these challenges, we propose an on-demand model loading technique. It enhances model capacity using mapping models for finer-grained address ranges and manages larger-sized model information by utilizing both memory cache and flash storage. Frequently accessed models are cached in memory to optimize performance, while less frequently accessed models are stored in flash to conserve memory resources. Fig1. illustrates the overall structure of our proposed idea. This method not only enhances model capacity but also reduces double-read occurrences and read latency compared to previous methods.

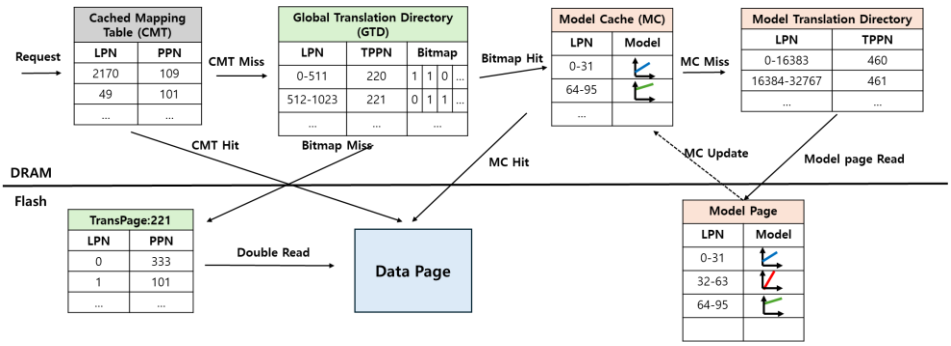

Fig 1. Architecture of proposed technique

Acknowledgments This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (No. NRF-2021R1A2C2095125)

References [1] S. Wang et al., "LearnedFTL: A Learning-based Page-level FTL for Reducing Double Reads in Flash-based SSDs," HPCA, 2024