Hong-yan Li, Nai-xue Xiong, Ping Huang, Chao Gui, 2014. PASS: a simple, efficient parallelism-aware solid state drive I/O scheduler. *Journal of Zhejiang University-SCIENCE C (Computers & Electronics)*, **15**(5):321-336. [doi:10.1631/jzus.C1300258]

PASS: a simple, efficient parallelism-aware solid state drive I/O scheduler

Key words: Solid state drives (SSD), I/O scheduler, Parallelism

Corresponding author: Hongyan Li E-mail: hongyanli78@aliyun.com

Introduction

- Propose PASS, an optimized I/O scheduler at the Linux block layer, to accommodate the changing trend of underlying storage devices toward flash-based SSDs.
- Take the rich internal parallelism in SSDs into account when dispatching requests to the device driver in order to achieve high performance.
- Partition the logical storage space into fixed-size regions (preferably the component package sizes) as scheduling units.
- The experimental results have shown that PASS outperforms the Linux four off-the-shelf I/O schedulers by a degree of 3% up to 41%, while at the same time it improves the lifetime significantly.

PASS design and implementation

1. Space partition

Divide the entire logical space into different regions of continuous logical space and dispatch requests to those individual regions in a parallel and interleaved manner.

Each region is associated with a dispatch sub-queue and each sub-queue has its own data structures to track requests that visit locations within the same region. Incoming requests are forwarded to respective subqueues according to their visiting addresses. Subqueues are serviced in a round-robin manner.

PASS design and implementation

2. Request management



PASS services the sub-queues in a round-robin manner. For each sub-queue's turn, it consecutively dispatches a batch of read or write requests and a batch of the other type requests in its next turn.

PASS design and implementation

3. Parallel dispatching requests



Major results



The number of erase operations of the workloads under different schedulers

Conclusions

- PASS groups requests targeting the same disk area together in dedicated queues and dispatches those queues in a round-robin manner.
- Within each queue, it sorts requests to create sequentiality and separately dispatches read and write requests to reduce interference.
- Experiments with a wide variety of workloads and SSDs have shown that PASS can not only improve the performance but also extend SSD's lifetime.