

An Experimental Study on Basic Performance of Flash SSDs with Micro Benchmarks and Real Access Traces

Yongkun Wang[†] Kazuo Goda[‡] Miyuki Nakano[‡] Masaru Kitsuregawa[‡]

Graduate School of Information Science and Technology[†] Institute of Industrial Science[‡]

The University of Tokyo

Abstract

Flash SSDs demonstrate outstanding performance improvement over the traditional hard disk drives. This paper provides an experimental study on basic performance of flash SSDs, by using micro benchmarks and real access traces. We present our intensive experimental measurements and then derive the basic model from micro benchmark results. This model has been utilized by us to implement a performance simulator that can simulate storage access behavior on flash SSDs. The simulator is expected to be effective to design flash-based database systems.


1 Introduction

Flash SSDs (Solid State Drives) demonstrate outstanding performance improvement over the traditional hard disk drives, and have been considered for enterprise storage platform. As for the performance of flash SSD, there are some literatures disclosing it, such as [1]. However, the benchmark results may deviate from the actual performance when running with the real workload in the enterprise storage system. We studied the basic performance by the real IO traces from OLTP applications running at two large financial institutions. Based on the micro benchmarks and the results of real access traces, we have derived a basic performance model and implemented as a performance simulator. Unlike some ideal simulator such as [2], our simulator is expected to be effective to design large-scale flash-based database systems.

The rest of this paper is organized as follows: Section 2 describes the experiment setup. Section 3 presents the experimental study. Our conclusion is provided in Section 4.

2 Experiment Setup

We have built a test system shown in Figure 1. High-end SLC flash SSDs are connected to the computer system with SATA 3.0Gbps hard drive controller.



3 Experimental Study

3.1 Micro Benchmark and Performance Model

We have developed a micro benchmark in C. The benchmark result of Mtron SSD is shown in Figure 2. We do not show the similar result of other SSDs here for brevity.

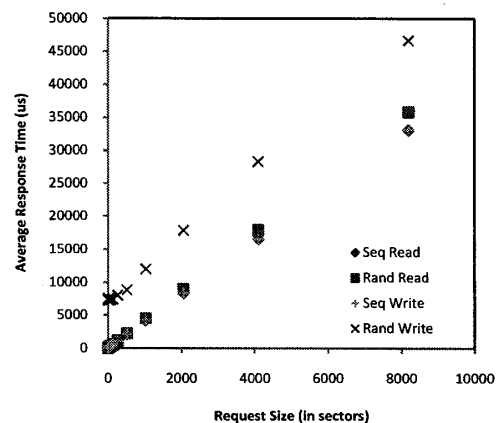


Figure 2 Micro benchmark result of Mtron SSD

Figure 2 shows that the response time of the request increases by size in a almost linear fashion. We have calculated the trend line for each data series. The general trend line equation is $y=ax+b$, where y stands by the response time and x is the request size in sectors. The a and b are constants decided by four access patterns (sequential read, random read, sequential write, random write). Figure 3 gives an

example of the equation of random read.

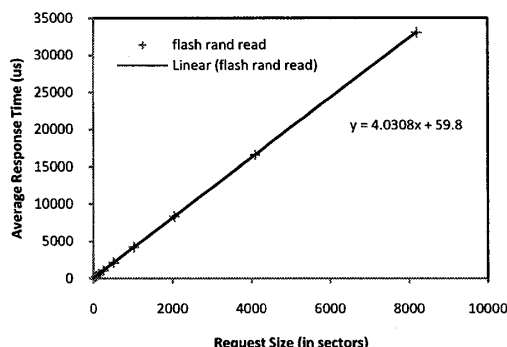


Figure 3 Micro benchmark results with trend line equation for random read on Mtron SSD

We have derived a basic performance model by all the trend lines of different access patterns. A performance simulator of SSD has been implemented on the basis of this performance model.

3.2 Test of the Performance Simulator

We have derived the performance simulator based on the results of the micro benchmark. The accuracy of this simulator has been studied with IO traces generated by a trace generator we have developed, as well as the real IO traces from OLTP applications running at two large financial institutions. Here we focus on the results of the real traces.

The IO traces we used were made by Ken Bates from HP and Bruce McNutt from IBM, following the SPC trace file format [3]. We describe the test by the *Financial1* trace file, which contains more than 5 million records, of which more than 80% are write requests, and more than 98% of the writes are random write requests.

Figure 4 confirms that the gap between the real SSD and our simulator is large on random write. In Figure 4, the line of “Cumulative Probability (Model on Trace File)” shows the response time distribution of our model-based simulator feeding with the trace file. As a comparison, the response time distribution of the real SSD indicates that the number of long response time is much less than that of our simulator.

We believe that the main difference comes from the influence of the write-back cache. Therefore, we have added a cache model as an enhancement in our simulator. 16MB cache buffer is managed in a LRU-like manner. The result is shown in Figure 5, which indicates that the gap between the real SSD and our model is small. The overall error ratio of response

time is 0.08156 for more than 4 million random write requests.

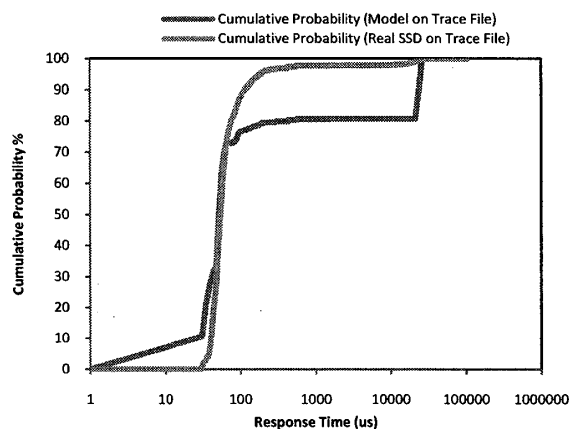


Figure 4 Response time distribution of random write of the simulation result by *Financial1* trace file on Mtron SSD




Figure 5 Response time distribution of random write of the simulation result by *Financial1* trace file on Mtron SSD, after added cache enhancement

4 Conclusion

We gave a performance study on flash SSD with the performance model derived by our micro benchmark. Based on the performance model, we developed and tested a simulator which is expected to be effective to design flash-based database systems.

References

1. Bouganim, L., Jonsson, B.T., Bonnet, P.: uFLIP: Understanding Flash IO Patterns. In: CIDR.(2009)
2. SSD Extension for DiskSim Simulation Environment, <http://www.pdl.cmu.edu/DiskSim/>
3. Storage Performance Council (SPC), SPC TRACE FILE FORMAT SPECIFICATION, Revision 1.0.1, www.storageperformance.org