

Hot Random Off-loading: A Hybrid Storage System With Dynamic Data Migration

Lin Lin, Yifeng Zhu, Jianhui Yue, Zhao Cai and Bruce Segee

*Department of Electrical and Computer Engineering,
University of Maine, Orono, USA*

{llin, zhu, jyue, zcai, Segee}@eece.maine.edu

Outline

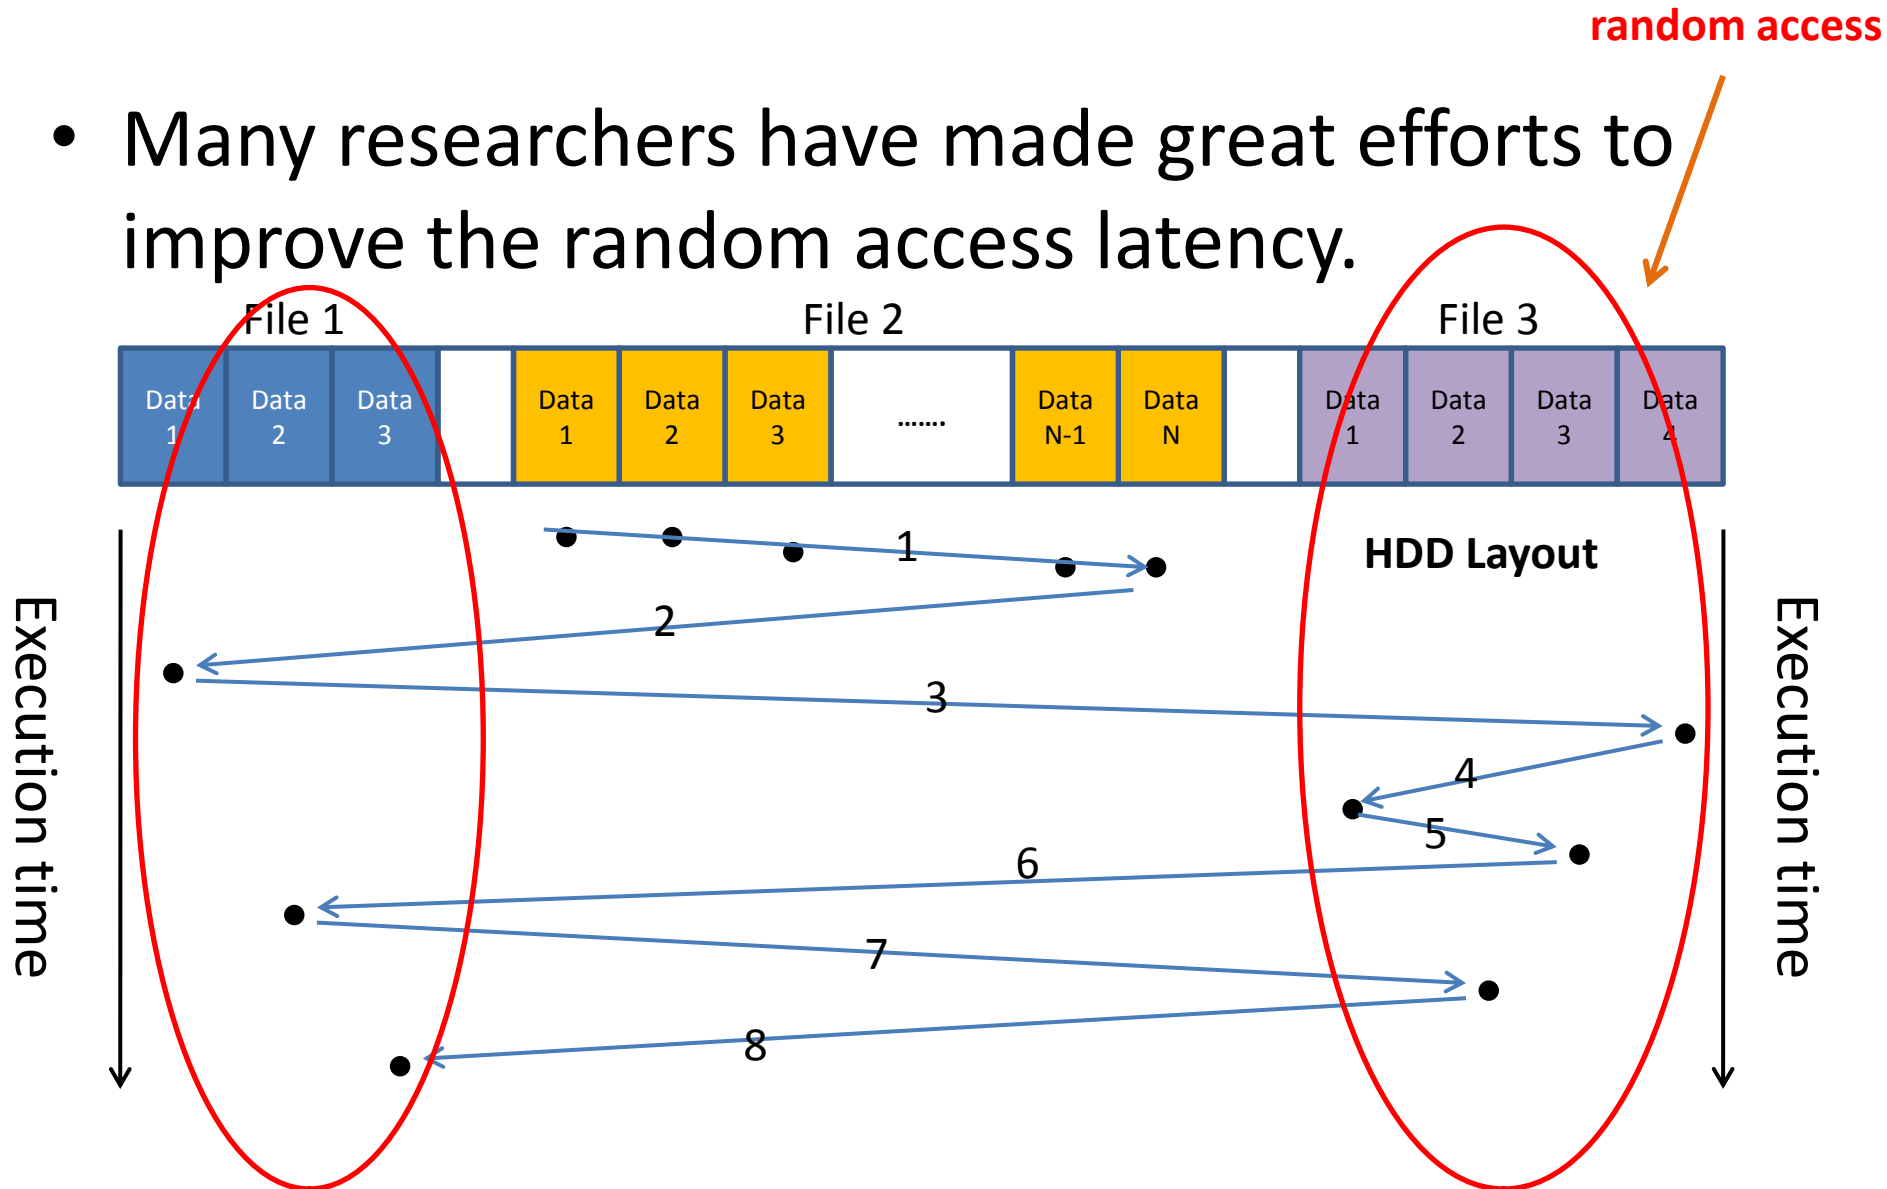
- Introduction
- Related work
- HRO scheme
- Evaluation
- Conclusions

Introduction

- The increasingly widening speed gap between hard disks and memory systems is a major performance bottleneck.
- Under random accesses, disk heads frequently move to different noncontiguous physical locations .

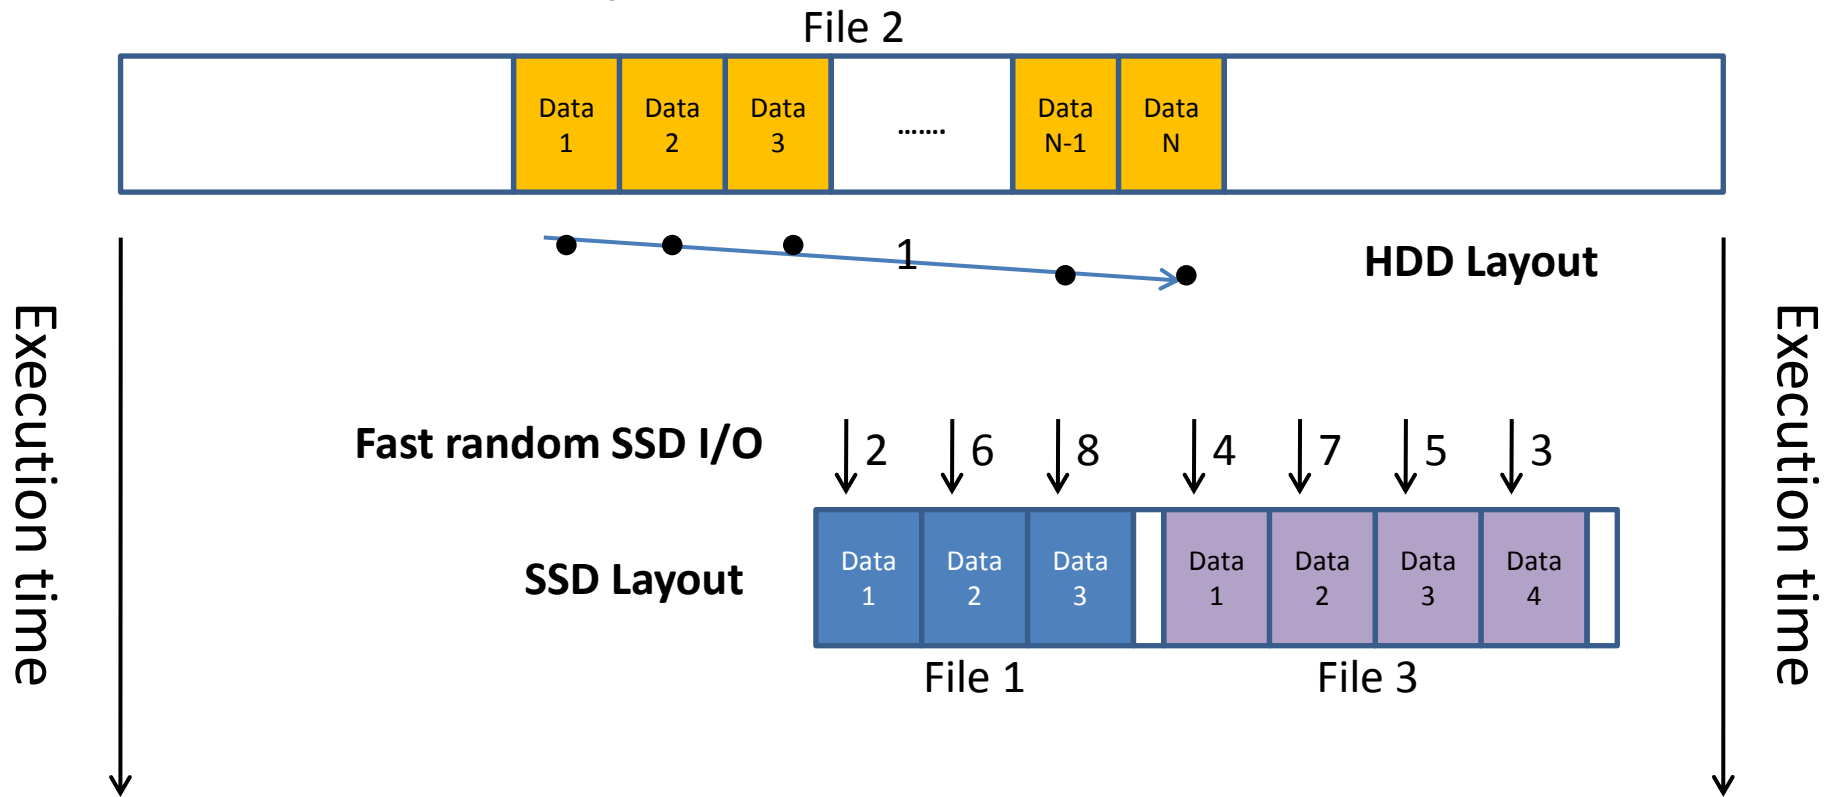
Introduction

- Many researchers have made great efforts to improve the random access latency.



Introduction

- The paper exploits solid state devices (SSDs) to address the performance issue.



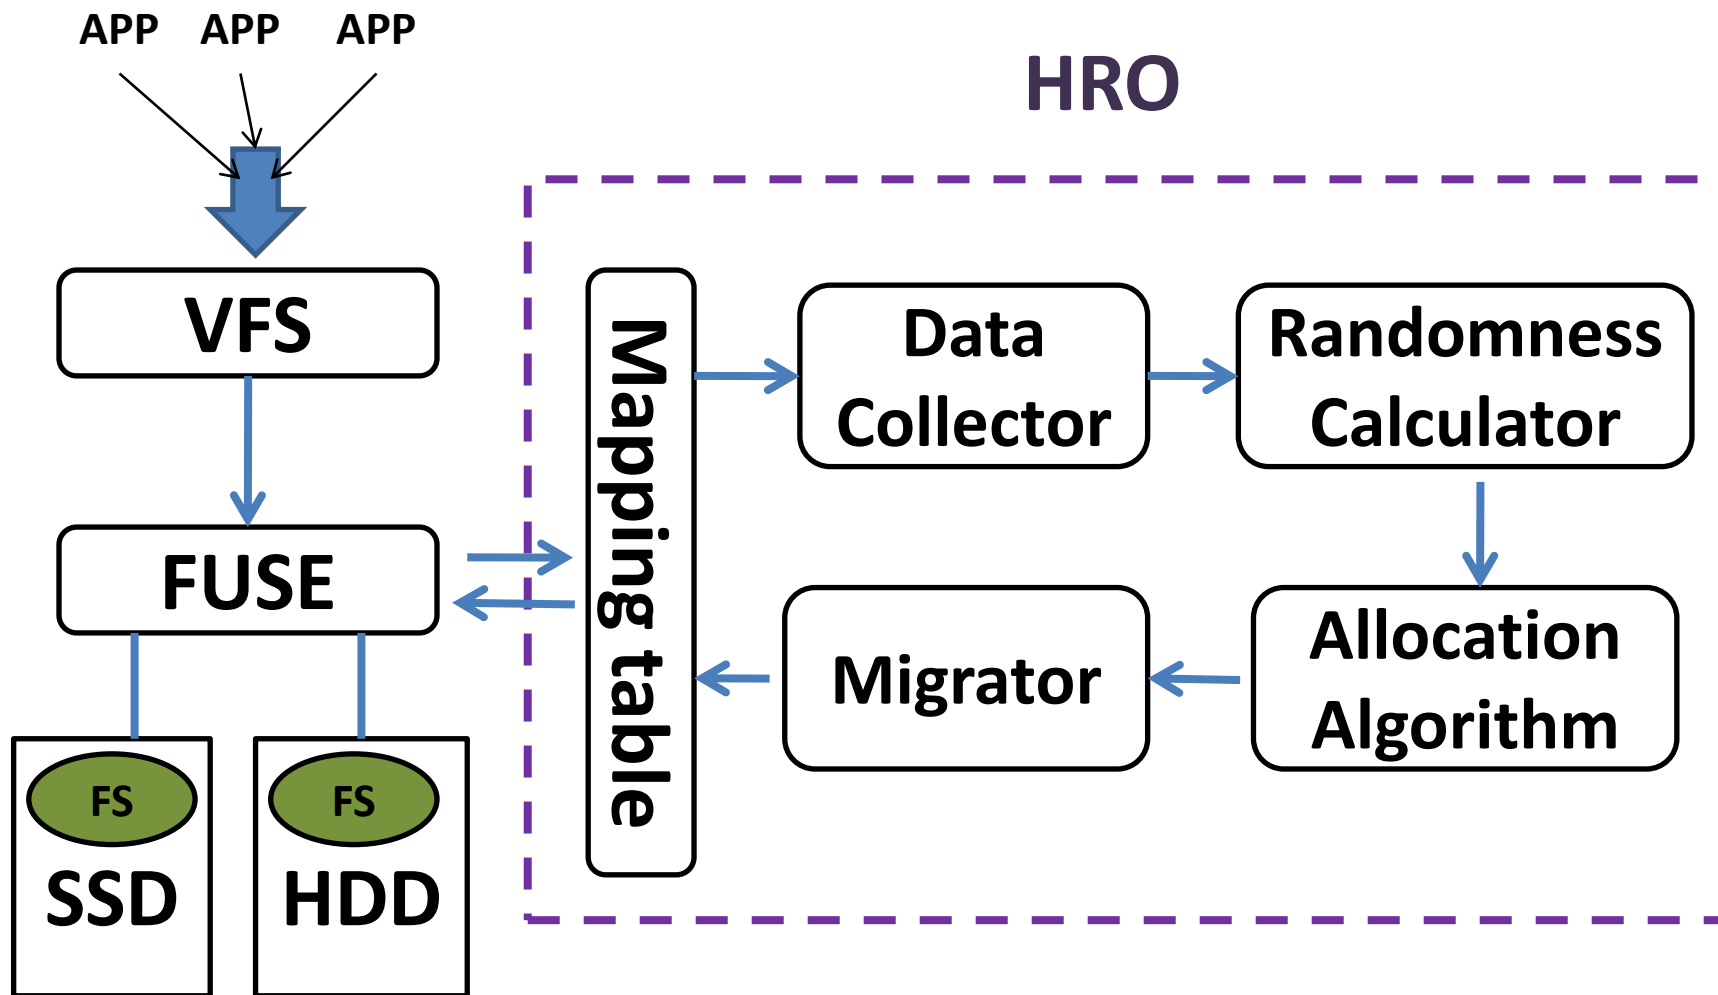
Introduction

- This key advantage gives us an exciting opportunity to build a hybrid system.
- Our basic principle is to place hot and randomly accessed data on the SSD, and other data on hard disk.

Related work


- They are something difference with RAF.
 - The unit.
 - File management for SSD and HDD.

HRO Architecture



HRO scheme

- **Mapping table structure**



Field	Explanation
Device ID	File located device ID (8 bits)
Inode	File inode number in the file system (48 bits)
Count	File randomness count number (32 bits)
Size	File size (64 bits)
Frequency	File access frequency (32 bits)
Type	File last access type (8 bits)
Time	File last access time (32 bits)
Offset	File last access offset (64 bits)
End	File last access end (64 bits)

HRO scheme

- **DATA COLLECTOR** captures all access requests on the system call I/O path .
- **RANDOMNESS CALCULATOR** mainly evaluates the randomness of each individual file.
 - Threshold = 0.5s

HRO scheme

- **ALLOCATION ALGORITHM** is specially designed to maximize the utilization of SSD in our hybrid storage systems.
 - Benefit value : $V_i = (f_i/s_i) * R_i$
- **MIGRATOR** handles data movement between SSD and disks.

File 1

File 2

File 3

File 4

File 5

$f_1=6, s_1=3, R_1=5, f_2=8, s_2=2, R_2=4, f_3=6, s_3=2, R_3=6, f_4=2, s_4=2, R_4=7, f_5=2, s_5=4, R_5=6,$

Step1 : To Calculate benefit values

$$V_1 = (6/3)*2 = 4, \quad V_2 = 16, \quad V_3 = 18, \quad V_4 = 7, \quad V_5 = 3$$

Step2 : Sort all items in the decreasing order of V_i / s_i

File 1 = $4/3$, File 2 = 8 , File 3 = 9 , File 4 = 3.5 . File 5 = $3/4$

File 3, File 2 , File 4, File1 , File 5

CAPACITY_{ssd} = 8

Step3 :

File 4 ,

MoveSet \leftarrow File 3, File 2 ,

TotalSize \leftarrow 6

for $i \leftarrow 1$ to n and TotalSize < CAPACITY_{ssd} do

{

if (TotalSize + size_i \leq CAPACITY_{ssd})

{

TotalSize \leftarrow TotalSize + size_i;

$x_i \leftarrow 1$;

Add i to MoveSet;

}

Step4 :

FileSet_{ssd} : File1,File5

FileSet_{disk} : File2,File3,File4

SET_{SSD} \leftarrow File 2 , File 3 , File4

SET_{HDD} \leftarrow File 1 , File 5

Experimental study

- Configuration of storage server

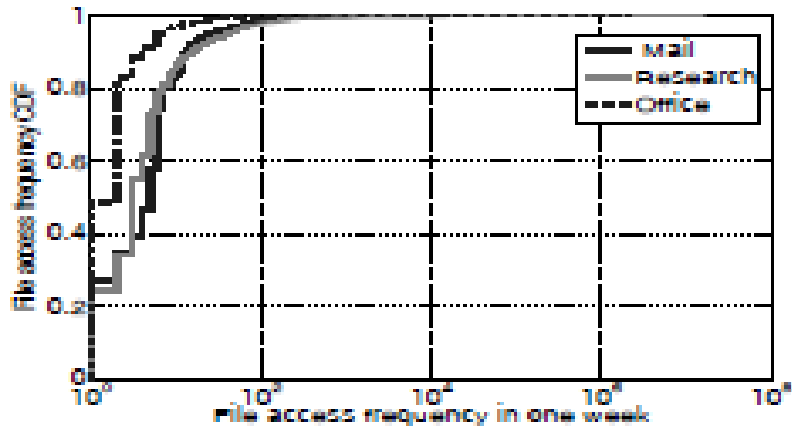
Components	Specification
Operating system	Ubuntu 10.04 with kernel 2.6.31
File system	Ext3
CPU	AMD Opteron dual core 1000Hz
Memory	1G DDR2 667 Hz
SSD	OCZ-AGILITY2
Capacity	60GB
Sequential Read/Write	20us/70us
Random Read/Write	270us/375us
Hard Disk *2	3*WDC WD7500AAKS
Capacity	750GB
Rotational speed	7200RPM
Read Seek	8.9ms
Track-to –track Seek	2ms

Experimental study


- Statistic of three workloads studied

	Office	Mail	Research
Trace period	One week	One week	One week
Total ops	211,308,494	187,974,468	29,550,778
Metadata ops	66%	14%	75%
Read	24%	65%	10%
Write	10%	21%	15%
Read (MB)	833,135	845,123	32,498
Write (MB)	242,376	313,987	61,488
R/W Ratio	3.4	2.5	0.5


About three workloads



(a) File access frequency CDF




(b) File size CDF




(c) Random hot access.

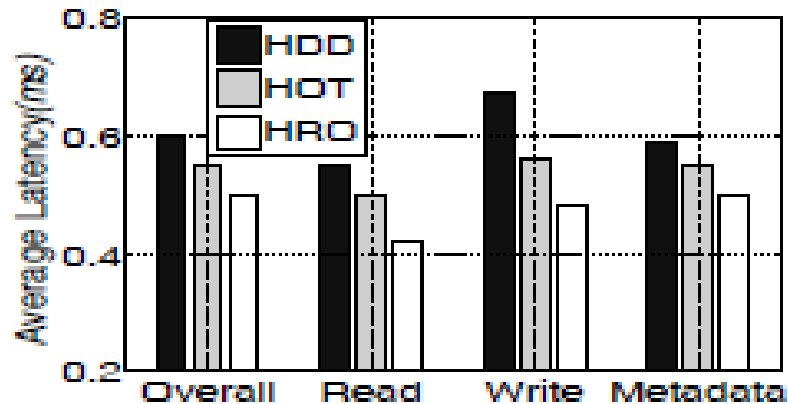
Average Latency



(a) Mail server workload




(b) Office workload

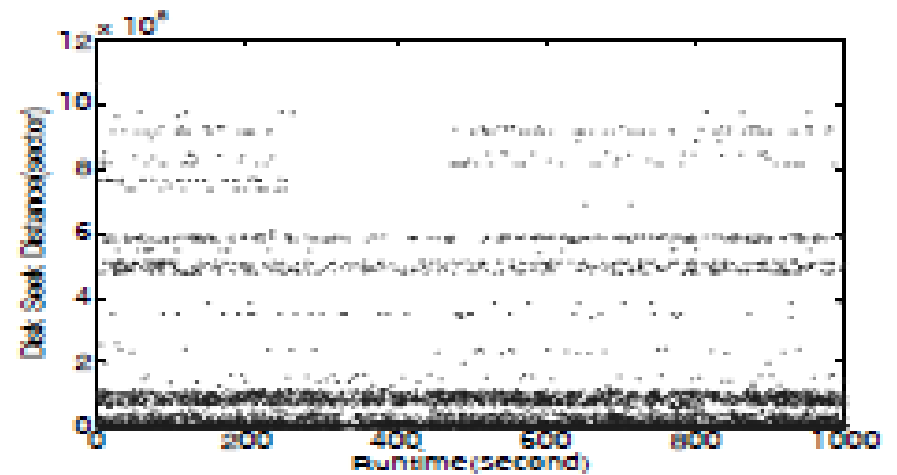


(c) Research workload

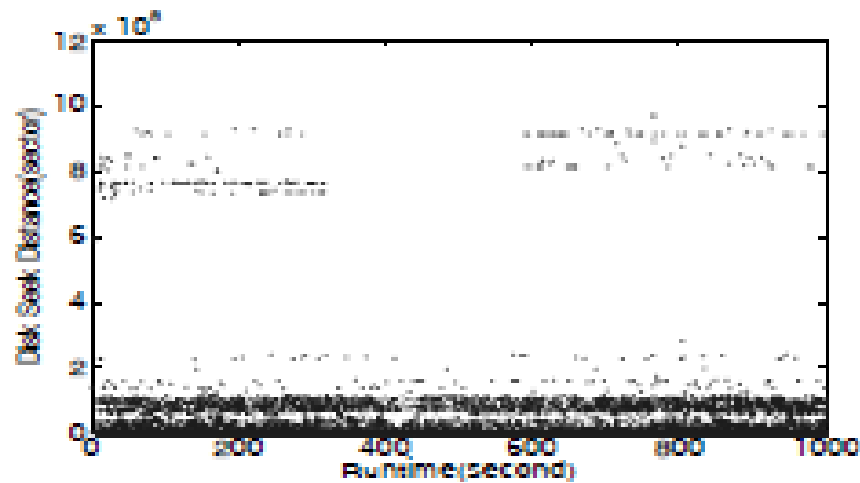
Disk seek distance



(a) HDD




(b) HOT



(c) HRO

Average bandwidth



Conclusions

- In our experiments, the usable storage capacity of SSD is set less than 1% of the hard disk capacity.
- It can effectively improve the overall I/O performance of disks by up to 39% and the latency up to 23%.