Hui SUN, Xiao QIN, Chang-sheng, XIE, 2014. Exploring optimal combination of a file system and an IO scheduler for underlying solid state disks. *Journal of Zhejiang University-SCIENCE C (Computers & Electronics)*, **15**(8):607-621. [doi:10.1631/jzus.C1300314]

Exploring optimal combination of a file system and an IO scheduler for underlying solid state disks

Key words: Solid state disk (SSD), Performance, Energy consumption, File system, I/O scheduler

Corresponding author: Chang-sheng XIE
E-mail: cs_xie@hust.edu.cn

Current issue

- ✓ Performance and energy consumption of a solid state disk (SSD)  File systems and I/O schedulers in operating systems

highly depend on

- ✓ File system + I/O scheduler  Performance + energy consumption

Optimize

Our solution

- ✓ **Aggregative indicator (AI)**

AI is the ratio of SSD performance (e.g., MB/s or IOPS) to energy consumption

This metric is aimed to evaluate SSD performance per energy consumption and to study the SSD with high performance at low energy consumption in a combination of a file system and an I/O scheduler.

✓ Cemp

Cemp is the change of energy consumption and mean performance for an Intel SSD (i.e., SSD-I) when it provides the largest AI, lowest power, and highest performance, respectively.

Using Cemp, we attempt to find the combination of a file system and an I/O scheduler to make SSD-I deliver a smooth change in energy consumption.

Our measurement system (I)

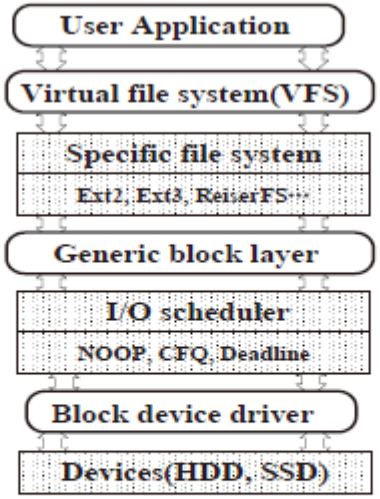

Fig. 1 The storage level based on SSDs and HDDs: We concentrate on specific file systems and I/O schedulers layer in this paper.

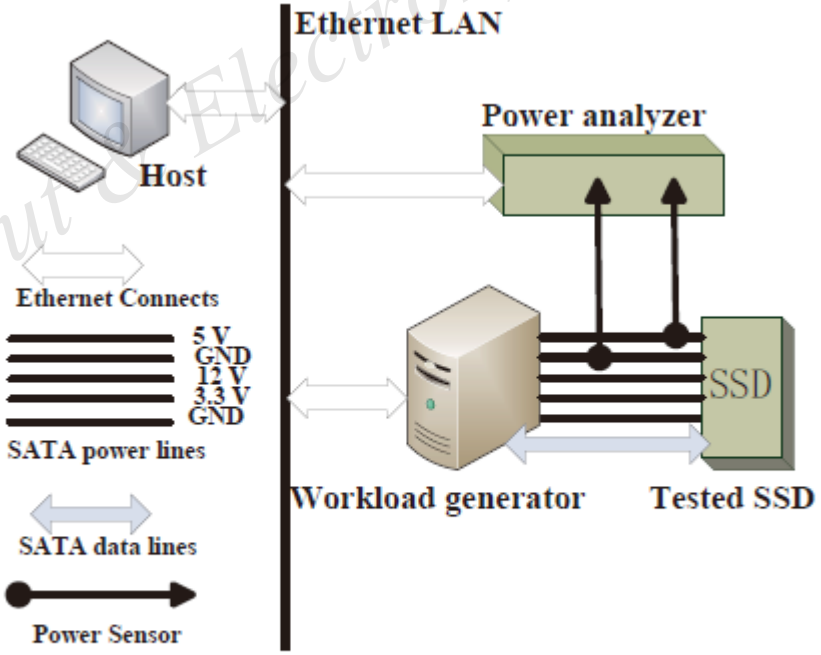

Fig. 2 The measurement system is to measure performance and energy consumption of tested SSDs

Our measurement system (II)

Table 1 Characteristics of Macro-benchmarks in Filebench

Workload	File size (KB)	Amounts of files	Amounts of threads	I/O size (Bytes)	Time (second)	N_{read} : N_{write}
Varmail	16	1000	16	16K(M), 1M (R)	600	1 :1
Fileserver	128	10000	50	16K(M), 1M (F)	600	1:2
Webserver	16	1000	100	1M	600	10:1

In the column of I/O size, F, R and M stand for fixed size, read size and mean size of I/O, respectively

Table 2 Characteristics of Tested SSDs

Tested SSDs	SSD-I	SSD-S	SSD-K
Device Capacity	80GB	64GB	64GB
NAND flash Type	MLC	MLC	MLC
Controller	Intel	Xilinx	Toshiba
	PC29AS21AA0	XC3S1600E	T6UG1XBG
DRAM	16MB	128MB	128MB
Performance (R)	201.3MB/s	98.8MB/s	220MB/s
(W)	73.8MB/s	103.6MB/s	140MB/s
Power(work/idle)	2.6W/0.15W	4.1W/1.5W	5.9W/3.3W
Support Secure Erase	Yes	Yes	Yes

Conclusions

- ✓ Confirm that SSD performance and energy consumption largely depend on the combination of a file system and an I/O scheduler
- ✓ SSDs coupled with the extended file systems + any I/O scheduler provide better AI than the SSD counterparts under varmail
- ✓ SSDs based on NILFS2 and any I/O scheduler offer high AI under fileserver
- ✓ SSD-I based on the extended file systems delivers high performance at low energy efficiency in most cases
- ✓ NILFS2-based SSD-I performs with lower energy consumption and smoother changes in energy consumption