In-Depth Reliability Characterization of NAND Flash based Solid State Drives in High Performance Computing Systems

Shuwen Liang, Zhi Qiao, Song Fu and Weisong Shi

UNT University of North Texas

Wayne State University

Motivations

NAND flash based solid state drives (SSD) have been widely used in high performance computing (HPC) systems due to their better performance compared with the traditional hard disk drives. However, little is known about the reliability characteristics of SSDs in production systems.

Methods & Process

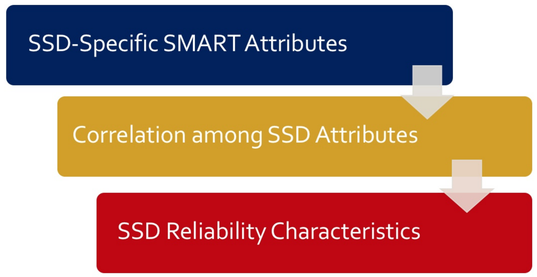

Data

SSD-SMART data files are pulled from a active production system.

- Window size is 6 months
- Hourly recored
- Over 500,000 record
- Contains 20 SMART attributes
- All data are raw values

ATTRIBUTE ID	ATTRIBUTE NAME	DESCRIPTION
Environmental Attributes		
9	POH	Power On Hour
12	PCC	Power Cycle Count
194	TC	Temperature Celsius
Performance Attributes		
5	RSC	Reallocated Sector Count
166	MWEC	Min Write/Erase Count
167	MBB	Min Bad Block/Die
168	MEC	Max Erase Count
169	TBB	Total Bad Block
171	PFC	Program Fail Count
172	EFC	Erase Fail Count
173	AWEC	Average Write/Erase Count
174	UPLC	Unexpected Power Lost Count
187	RU	Reported Uncorrect
212	SPE	SATA PHY Error
230	PWEC	Percentage Write/Erase Count
232	PARS	Percentage Avaliable Reserved Space
233	TNWG	Total NAND Write(GB)
241	TWG	Total Write(GB)
242	TRG	Total Read(GB)
243	Unknown	N/A

Results 1

Pearson's coefficient reveals that write and erase operations appear strong linear correlations to each other.

Figure 4: Pearson Correlation Coefficient

Other results includes:

- I/O operations are not distributed evenly in the system.
- Environmental attributes do not directly affect SSD health.
- No evidence shows that I/O operations lead to increase of back blocks in SSDs.

Results 2

- SSDs in Cluster 1 experience I/O intensive operations.
- SSDs in Cluster 3 experience the highest number of write and erase operations, while the number of read operations is the average.
- Clusters 0, 2 and 4 include the majority of SSDs which experience the average number of I/O operations.
- SSDs in the three clusters have reliability degradation following the same transition pattern

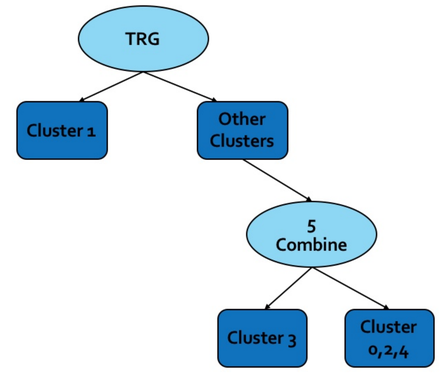

Figure 1: Cluster Distinguished by Wear Level

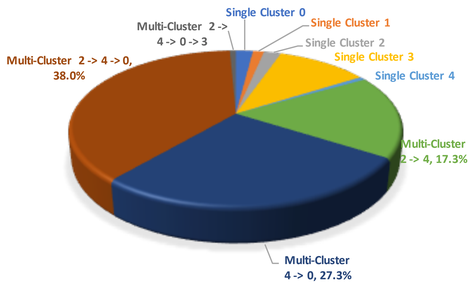

Figure 2: Disk distribution on Clusters

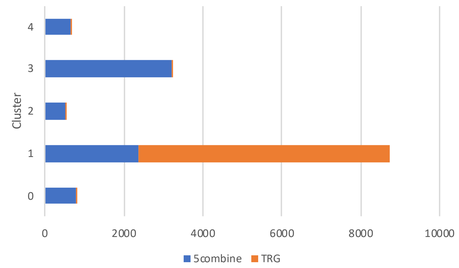

Figure 3: I/O operations on Clusters

Conclusion

We explore SSDs SMART data from a real world production system. Our research found that:

- The volume of I/O operations and cycles has a major influence on the wear level of SSD.
- Write and erase related attributes display a strong correlation.
- Read operation is relatively independent and is not evenly distributed in among the drives.
- Observe many health status transitions.

Future Research

- Investigate the reliability degradation process of SSDs
- Accurate model this process for a better characterization of SSD's reliability

References

[1] Song Huang, Song Fu, Quan Zhang, and Weisong Shi. Characterizing disk failures with quantified disk degradation signatures: An early experience. In *Proceeding of IISWC*, pages 150–159. IEEE, 2015.

Acknowledgment

This work was supported in part by NSF grants CCF-1563728 and CCF-1563750.