Characterizing SSDs with the SMART Monitor

Shuhei Tanakamaru and Erich F. Haratsch

Seagate Technology

Outline

- Overview of SMART
- SSD Characterization with SMART
- Summary

SMART

- Self-Monitoring, Analysis, and Reporting Technology (SMART) is a system that reports the drive's health statuses
- In SATA drives, the implementation is vendor specific but there are common SMART attributes as well as vendor specific attributes
- In NVMe drives, the NVMe specification defines what the contents are, in which log pages they are stored, as well as the bytes allocation
- There can be vendor specific SMART attributes in the vendor specific log pages

Common SMART Attributes

Reliability/Health

Flash related attributes

- Reallocated sector count (Bad blocks)
- Program fail count
- Erase fail count
- Program/Erase cycles (Lifetime used)
- Raw read error rate

Drive attributes

- Available reserved space
- Uncorrectable error count
- SATA downshift
- CRC error count

Performance/Others

- Host reads
- Host writes
- NAND writes
- Free space

- Power-on hours
- Power cycle count
- Unsafe shutdowns
- Temperature

A good summary can be found here:
https://en.wikipedia.org/wiki/S.M.A.R.T.

Details of the Attributes

- It is common in NAND Flash to have grown bad blocks as well as program/erase failures across the lifetime
- Therefore, non-zero values of Reallocated Sector Count and/or Program/Erase Failure count do not indicate the drive is in danger
- However, low Available Reserved Space *may* be a sign of potential issues (Vendors may indicate the threshold value for this)

- A Non-zero Uncorrectable Error Count itself might not immediately mean a catastrophic event, but *may* be a sign of potential severe issues

- Since NAND Flash has a predetermined endurance cycle, it is important to monitor the Lifetime Left attribute

Drive Testing / Drive Diagnosis

- It is beneficial to periodically read out the SMART information during testing and record it with the time stamp
- In case of any unwanted event, we can trace back the entire SMART history to see if there were any signs of the issue, such as …
  - Was the drive temperature within the expectation?
  - Any abnormal behavior from NAND? If yes, when? (Lots of bad blocks? Uncorrectable errors?)
  - Are write amplification and program/erase count expected?

Drive Health

- Below shows the SMART attributes that are potentially related to drive's health
- During the testing, there were zero occurrences of reallocated sector, program/erase fail, and uncorrectable ECC
- However, in case that there were these events, it will be very helpful to analyze later if these are recorded periodically

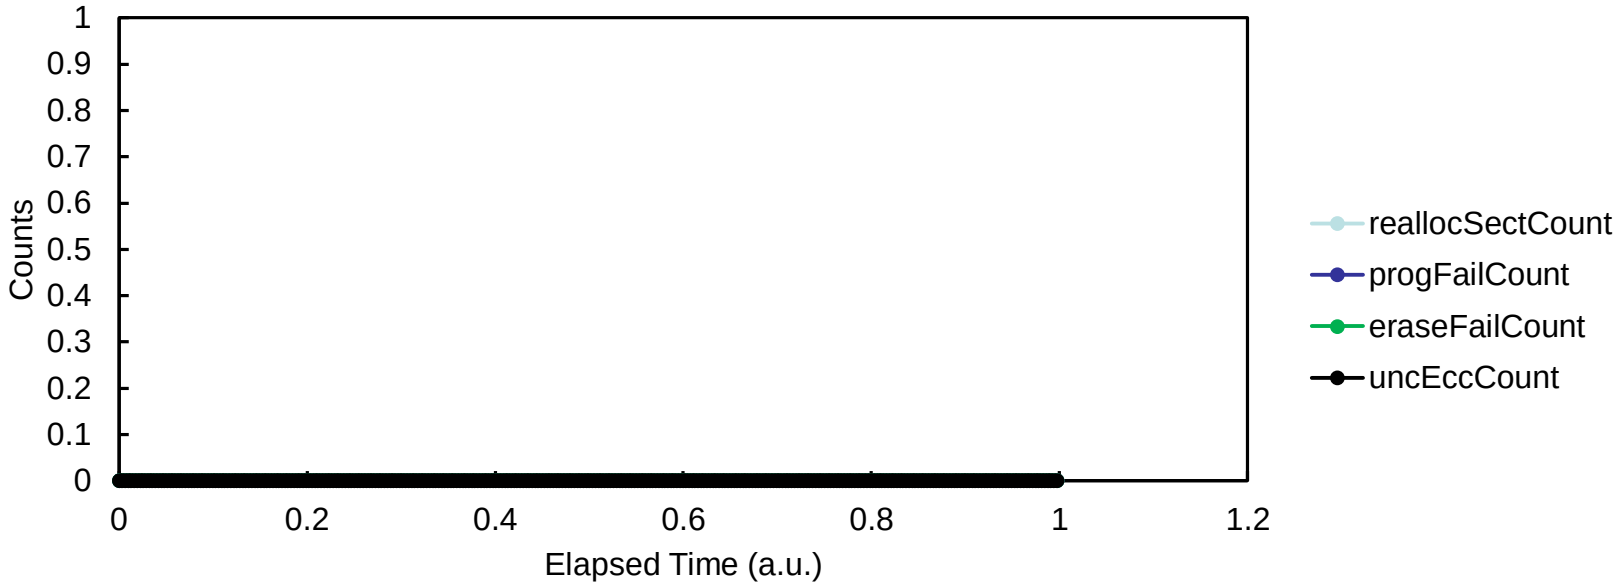

Write Amplification

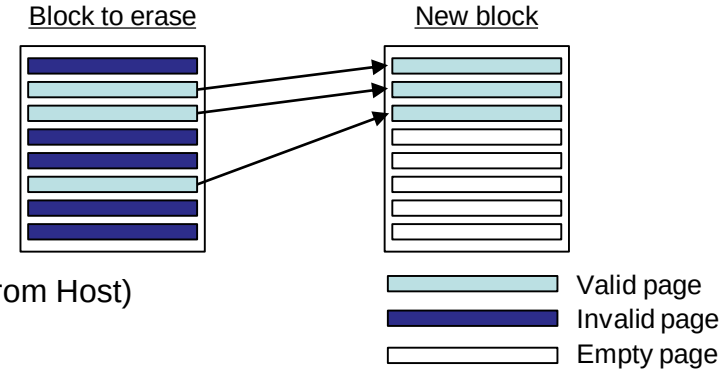

- Write Amplification (WA): Garbage collection, ECC, metadata, etc.

- WA := Amount of Writes to Flash / Amount of Writes from Host
- Instantaneous WA: = Δ(Lifetime Writes to Flash) / Δ (Lifetime Writes from Host)

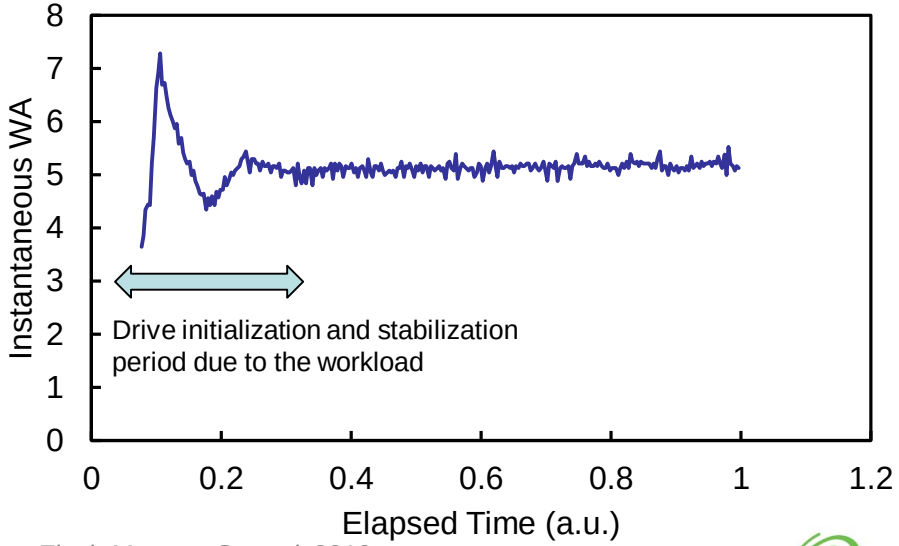

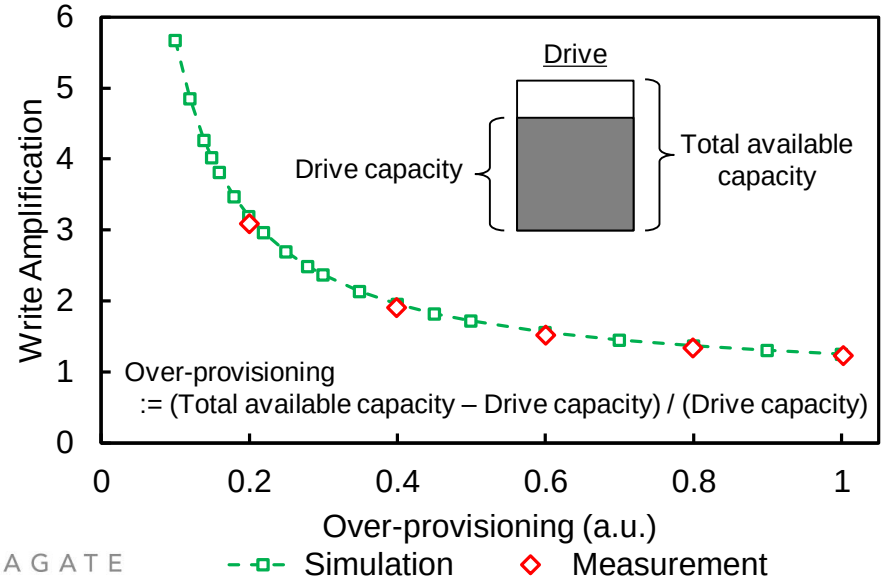

Measuring Throughputs

- Host Write Throughput := Δ (Lifetime Writes from Host) / (Time between two SMART reads)
- Benefits of using SMART
  - No other tools necessary
  - Good accuracy (though less time resolution)
  - Drive's internal throughput can also be measured

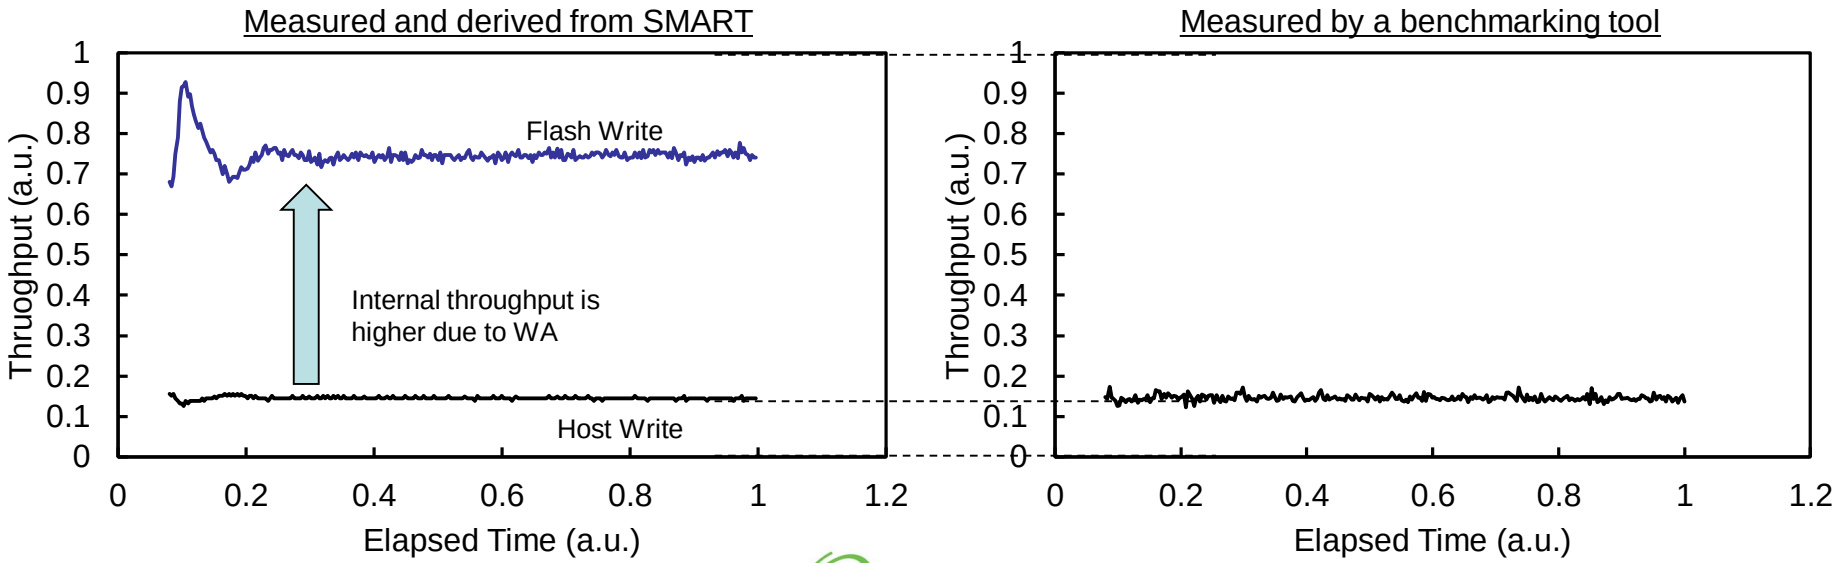

Compression

- When a drive supports compression, there will be more free space remaining when the incoming data is compressible
  - Due to higher over-provisioning, there will be less write amplification and therefore better performance / endurance
  - The extra free space can be utilized as the additional capacity that the host can use
- When the data is compressible, the write amplification can be below 1

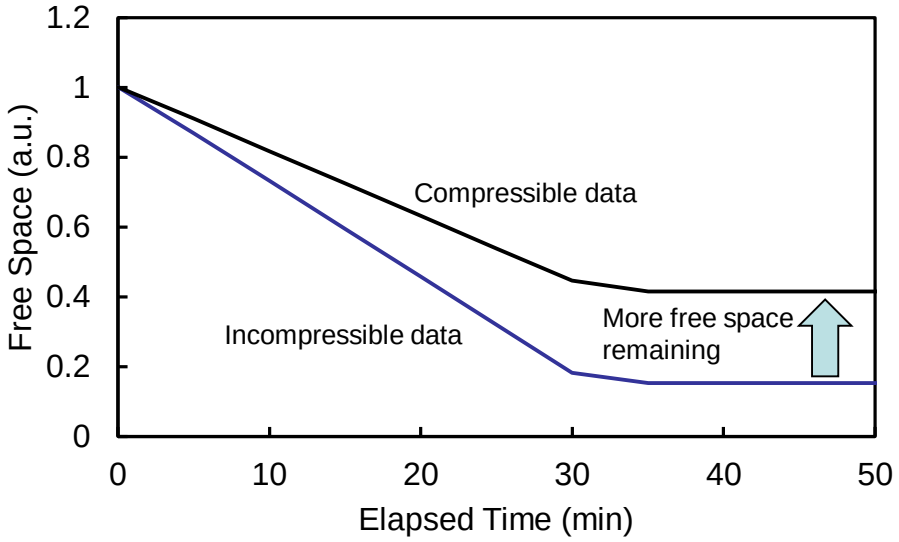

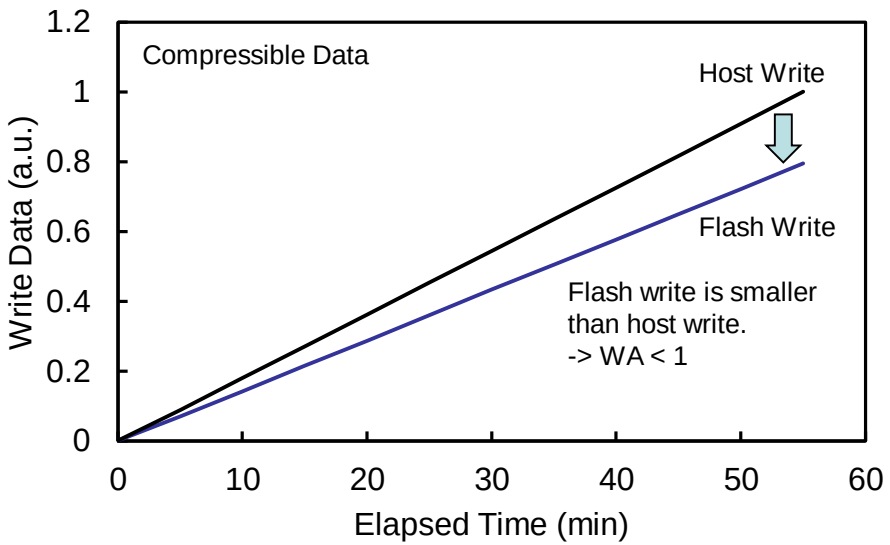

Temperature and Average Program/Erase Count

- Users can also monitor temperature and possibly correlate with the workload
- SSDs have endurance limitations. Average program/erase count or SSD Life Left is a good metric to see if the SSDs are excessively cycled or not

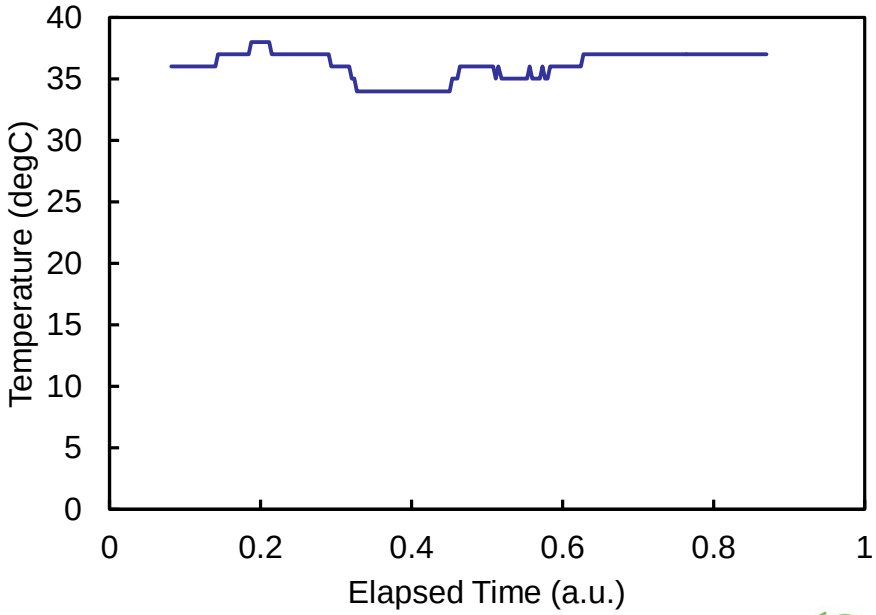

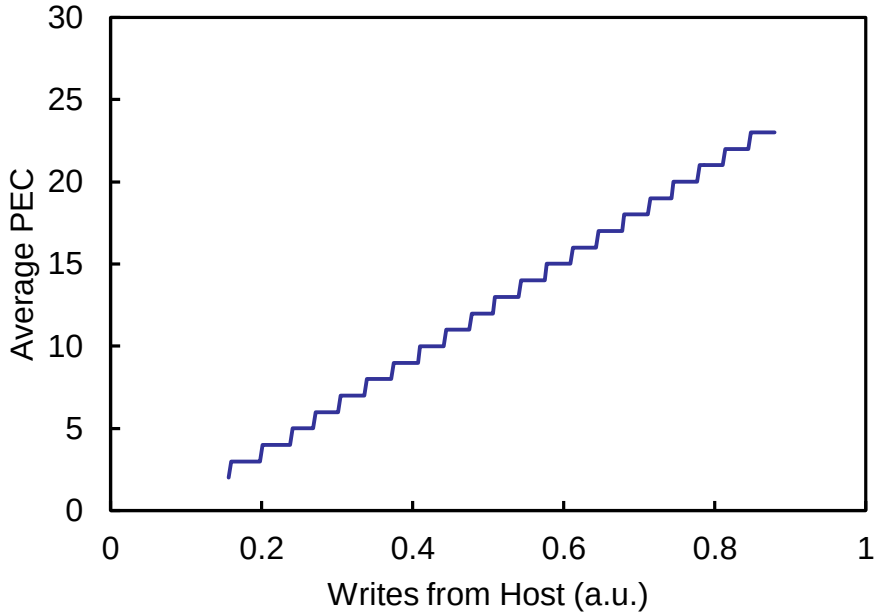

Summary

- SMART is an extremely useful tool not only to monitor the SSD's health but also to broadly characterize the SSD

- Write amplification is an important metric in SSD characterization
  - Using Flash Writes and Host Writes will give the WA value
- Host throughput as well as internal throughput can be measured
- Benefits of compression can be quantitatively analyzed
- Other attributes like temperature and average program/erase cycles are also beneficial to see if there were any unexpected conditions occurring