

# Homework 1 – Algorithms

Due via dropbox by September 8, 2021

**Submissions only in pdf format, Latex documents preferred**

In preparation for looking at data structures for Flash Drives, read the following articles.

Narayanan, I., Wang, D., Jeon, M., Sharma, B., Caulfield, L., Sivasubramanian, A., ... & Vaid, K. (2016, June). SSD failures in datacenters: What? when? and why?. In *Proceedings of the 9th ACM International Conference on Systems and Storage* (pp. 1-11).

Zhang, T., Zuck, A., Porter, D.E. and Tsafir, D., 2017, May. Flash drive lifespan\* is\* a problem. In *Proceedings of the 16th Workshop on Hot Topics in Operating Systems* (pp. 42-49).

- (1) Explain Write Amplification and its causes for SSD.
- (2) Explain the various modes of flash failures:
  - (1) data retention
  - (2) program disturb
  - (3) read disturb
  - (4) endurance
  - (5) power faults.
- (3) Use your search skills on the web in order to verify the current cost ratio per GB between HDD and SSD (for servers). Show your work.
- (4) Derive a "back of the envelope" formula that gives the lifespan of an SSD in years assuming
  - no write amplification
  - a constant write rate of  $x$  MB/min
  - pages wear out after exactly  $L$  writes
  - Perfect distribution of the stored data: no page is written more than any other page.