Understanding the Impact of Threshold Voltage on MLC Flash Memory Performance and Reliability

Published in the 28th ACM International Conference on Supercomputing (ICS'2014)

Wei Wang[1], Tao Xie[2], and Deng Zhou[1], San Diego State University, San Diego, CA 92182

Introduction

MLC (multi-level cell) NAND flash memory based solid state drives (SSDs) have been increasingly used in supercomputing centres. However, as each cell starts to store two or more bits, a threshold voltage range employed to represent a state has to be continuously shrunk, and a narrowed thresh- old voltage range causes more bit errors.

In this research, we studied the the impact of threshold voltage on MLC flash performance and reliability.

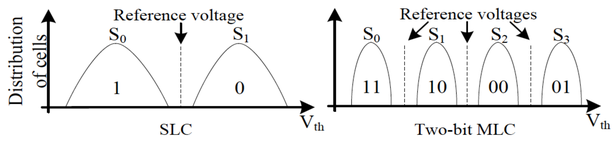

Fig. 1 Threshold voltage levels in SLC and MLC flash.

$$f(x) = \sum_{s=0}^{2^M-1} P(S_s)\frac{1}{\sqrt{2\pi}\delta_s}\exp\{\frac{-(x-\mu_s)^2}{2\delta_s^2}\}\ .$$

Testing Methodology

- Flash memory error collection scheme

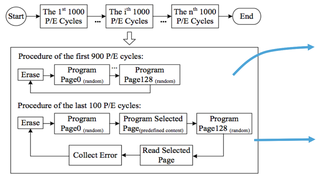

★ Wear out memory cells using constant threshold voltage.

★ Collect cell errors using pseudo-random data.

- Experimental environment

★ Xilinx FGPA evaluation board + flash daughter board.

★ Software stack:

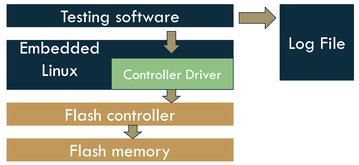

Reliability

★ Cells programmed to a higher voltage incur more errors.

★ The cell page programmed as '11' exhibits the most unreliable characteristic.

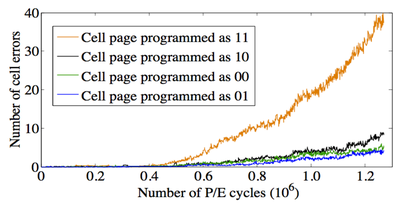

Fig. 2 Average number of cell errors.

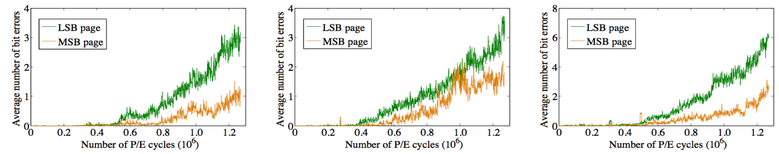

(a) Cell page programmed as 10 (b) Cell page programmed as 00 (c) Cell page programmed as 01

Fig. 3 Programming errors in LSB and MSB page within one cell page.

- **System implications:**
  - ❖ Continuously programming '11' to a cell page or erasing a block without any data programmed should be avoided.
  - ❖ Writing data patterns that are represented by a lower threshold voltage could prolong flash memory's lifetime.

Programming/Erase Performance

- Programming

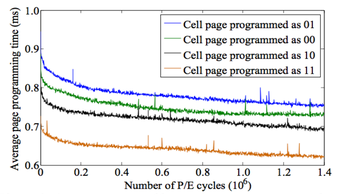

Fig. 4 Average cell page programming performance.

★ Pages programmed to a lower threshold voltage have a better programming performance.

The speed of programming a page to '11' is 15.5%, 23%, and 31% faster than that of programming a page to '10', '00', and '01', respectively.

★ The programming time decreases as the number of P/E cycles increases.

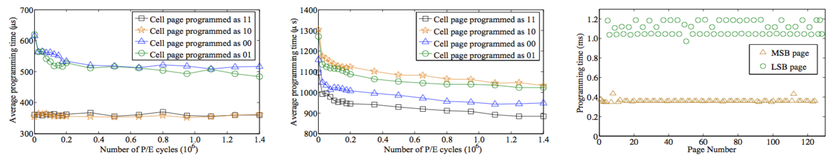

(a) MSB page programming time (b) LSB page programming time (c) Page programming time in a block

Fig. 5 Programming time of LSB and MSB page in a block.

★ Programming speed of an MSB page is much faster than that of an LSB page.

★ In the case of MSB page programming, the time of programming a cell page to '11' state is almost the same as that of programming it to '10' state, whereas the same observation can be made in '00' state and '01' state.

- Erase

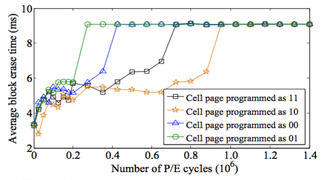

Fig. 6 Average block erase time.

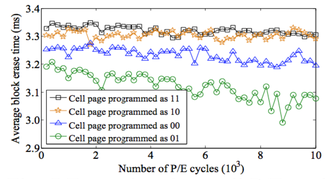

Fig. 7 Erase time in small P/E cycles.

- **System implications:**
  - ❖ Programming content that is represented by a lower threshold voltage can have a higher P/E performance.
  - ❖ Judiciously rearranging the programming order could improve SSD's performance.

Threshold Voltage Reduction

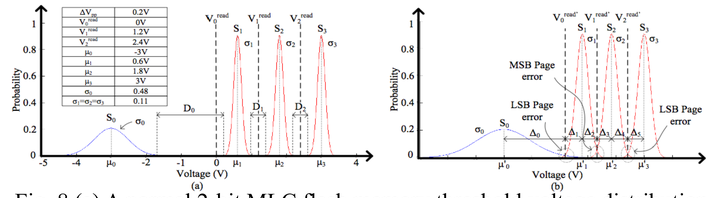

Fig. 8 (a) A normal 2-bit MLC flash memory threshold voltage distribution, (b) an example of a retention free threshold voltage distribution.

★ In typical data center workloads like proxy and MapReduce, most data are overwritten frequently, which suggests a short data retention in only days or even hours.

★ Data retention capability offered by the worst-case oriented design is always under-utilized.

Retention Time ⟷ tradeoff ⟷ **Reliability** / **Performance**

- **In the retention free case:**
  - ❖ Programming speed can be improved by 50%.
  - ❖ The number of P/E cycles can be improved by 7.1%.
- Impact on SSDs

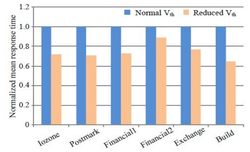

★ DiskSim 4.0 and Microsoft SSD module

★ TVR can reduce SSD's overall mean response time by 11% to 35%.

Conclusions

The P/E performance and reliability of MLC flash are highly correlated to threshold voltages. The retention time can be judiciously reduced to gain a higher reliability and performance flash memory.

Acknowledgement

We would like to thank Steven Swanson and his Non-volatile Systems Lab for the support of hardware platform. This work was supported by the US National Science Foundation under Grant CNS(CAREER)-0845105.

References

[1] R.-S. Liu, et al. Optimizing NAND flash-based SSDs via retention relaxation. In UESNIX FAST, 2012

[2] T. Bunker, M. Wei and S. Swanson. Ming II: A flexible platform for NAND flash-based research. Tech. report, UCSD, 2012.

[3] L. Grupp, A. Caulfield, et al. Characterizing flash memory: anomalies, observations, and applications. In IEEE MICRO, 2009

1. Computational Science Research Center
2. Computer Science Department