Flash Memory

- Non-volatile semiconductor storage (cells hold values even with no power applied).
- Different technology at the gate level from DRAM and SRAM, which are volatile (stored values persist only as long as external power is supplied).
- Two physical types: NAND and NOR
- Used in thumb drives for removable media, SD cards for smart phones and cameras, and more recently SSD (solid state drive) replacements for traditional hard drives.
- Somewhat slower and cheaper than RAM
- SSD is much faster but more expensive than a HD.
- Disadvantage: storage cells have a finite lifetime, shorter than both RAM and HD.
- Many trade articles on the differences between SSDs and HDDs, for example:
    - http://www.pcmag.com/article2/0,2817,2404258,00.asp

Organization

- Bits are organized into a row-and-column grid, like DRAM
- A single row in a grid is called a **page**, can be 8-16 kbytes per page, for example.
- Some number of rows makes a **block**, can be 64-512 pages per block, for example.

Operations

- The **write** operation is called **program**. Smallest programmable unit is a **page**.
- Pages are initially in the **erased** stage. After programming (writing) a page, it cannot be modified until it is cleared or **erased**. But the smallest erasable unit is a **block**. In other words, all pages in a block must be erased together.
- Is there a **modify** operation? Not directly, must do **read-modify-write.**
    - Copy existing page into RAM (read).
    - Modify the copy in RAM (modify).
    - Find a different free (already erased) page to program (can be in the same or different block) (write).
    - Update page mapping directory to point to new page.
    - Mark original page as invalid for later garbage collection.
- Garbage collection
    - As free pages are programmed and the modified, over time invalid pages will accumulate across different blocks.
    - Invalid pages cannot be reprogrammed until they are erased.
    - But an erase operation will erase the entire block, it can't be applied to individual pages within the block.
    - But if the block they are a part of contains other pages that are valid (programmed), the block is not ready to be erased.
    - Plus, we don't want to erase any one block too many times.
    - Garbage collection will find a block with mostly invalid pages, move and remap the remaining valid pages from that block to another block. Then the original block can be erased, making all of its pages ready to be reused.

Lifetime

- Finite number of PE (program-erase) cycles before storage becomes degraded/unreliable
- 5,000 cycles per block for consumer grade MLC (multilevel cell) flash
- 100,000 or 1M cycles per block for high end SLC (single level cell) flash
- Flash file systems include “wear leveling” algorithms that keep track of the number of erase cycles for each block and try to distribute erase operations evenly to all blocks.

Flash Transition Layer (FTL)

- Modified pages must usually be remapped to different physical pages.
- The FTL is a data structure to maintain logical-to-physical page/block pointers to simplify the remap operation.
- This is similar to a page table in a virtual memory system that maps logical page numbers to physical page frame numbers in RAM.
- External address references to a page use a logical page number, the FTL maps to a physical page, if the modified page is moved to a different physical page, the FTL can be easily modified to remap the logical page to the new physical page.

Flash File System

- Performs all the above tasks to manage the file system in flash.
- FTL for logical to physical page mapping
- Garbage collection to collect invalid pages together in one block for erasure.
- Wear leveling algorithm to ensure uniform numbers of erase operations across all blocks.
-