CS 61 Scribe Notes, Tuesday, 10/30/12

Caching (Performance Optimization) & Memory Hierarchy

Brian Zhang, Hikari Senju, Muhammad Ibrahim, and William Chambers

CS 61 Server

- CS 61 server went down over the weekend, let's look at the graph
- Red line = 4 GB, total in-core memory, crashed when we exceeded this limit

IBM 7090

- In 1963, IBM gave the MIT Computation Center a computer called the IBM 7090 which had 36 bits of memory per word, could store 32 kilowords (128 KB), and cost $3.5 million in 1963 dollars
- Inputs were fed in via cards, read into volatile memory and used for computation, then punched on output cards
- Usage of the machine was split into three 8-hour shifts per day (1 for MIT, 1 for other universities, 1 for IBM)

Motivation for Swapping

- At the time, IBM had a yacht handicapping program that took priority whenever the president of IBM called, all other processes would be cancelled
- One day, MIT is using the IBM 7090 and is in hour 5.5 of a 6-hour computation when IBM's president calls, the process is killed to run the yacht program
- The CPU utilization that day can be at most 18.5 / 24 = 77%, since 5.5 hours of computation goes to waste
- We need a way to save the state of a computation in the middle
- Swapping: Storing process state in secondary storage so that primary memory is available for other use

Two types of memory

- Volatile memory / primary storage: unstable, values preserved only when powered but much faster for computation
- Stable memory / secondary storage: hard disk, flash memory, tape, cheaper per byte
- Hard disk and flash are optimized for bytes per dollar
- Memory is optimized for computing speed
- These differences affect all of computer systems

Cost Comparisons

$/MB, constant 2012 dollars	Memory	Flash	Hard Disk
1957	$411,000,000	Didn't exist	$6,200
2012	$0.0047	$0.00078	$0.000038

Observations:

- The difference in memory cost from 1957 to 2012 is 10^{11}
- The jet engine only increased our speed by 2 orders of magnitude from 8 mph to

about 400 mph
- The gap between flash memory and hard disk is narrowing (230x in 2003 compared to 21x in 2012)

Hard Disk Design
- While memory is made of transistors, the hard disk is a spinning disk with an arm that reads coded magnetic bits
- Important vocabulary: platter, spindle, arm, head, tracks
- Example image: http://en.wikipedia.org/w/index.php?title=File:Hard_drive-en.svg
- Not labeled on this image are the tracks, which are the individual circles of bytes on each platter that the head reads from
- Swinging the arm and waiting for the rotation are the two primary contributions to latency

Speed Breakdown for Hard Disks
1. Seek time (avg = 4 ms) is the time for the arm to swing and settle into place
2. Rotational latency (avg = 2 ms) comes from waiting for the platter to spin around
3. Sustained transfer speed of 58 – 96 MB / s
    a. (The reason why there is a range is because more data can be read when the head is near the outside of the platter, where the circumference is longer)
- So it takes about 4 ms + 2 ms + 2 ms = 8 ms to read 4 pages of memory (16 KB) from the hard disk to primary memory (disk read)
- Then it takes 15 microseconds to read those 16 KB from primary memory (memory read)

Speed Comparisons
- A CPU cycle is the time it takes to execute one instruction, about 1 ns
- Can execute 6 million arithmetic instructions in the time it takes to access 1 byte from the hard disk (6 ms versus 1 ns)
- Disk read (8 ms) takes about 500x the time as memory read (15 microseconds)

Improving Swapping
- In virtual memory, fixed size allocations = less fragmentation = better utilization
- For swapping, instead of using units of processes (low utilization), we can use units of memory pages (higher utilization)
- Locality of Reference: data accessed at time t is close to data accessed at time t + 1, property of well-written programs
- Contiguous allocation has good locality