Computer Systems and Networks

ECPE 170 – Jeff Shafer – University of the Pacific

Memory Hierarchy

(Performance Optimization)

Lab Schedule

Activities

- **Today**
  - Background discussion
  - **Lab 6 – Performance Optimization (Memory)**
- **Thursday**
  - **Lab 6 – Performance Optimization (Memory)**

Assignments Due

- **Wednesday Oct 10th**
  - **Lab Report for Lab 5 due by 11:59pm**
    - **Includes post-lab**
- **Next Tuesday**
  - ***** Midterm Exam *****
- **Sunday Oct 21st**
  - **Lab Report for Lab 6 due by 11:59pm**
    - **Includes post-lab**

Midterm Exam

Midterm Exam

Part 1

- Paper & Pencil only
  - No computer, no notes, no Internet, etc…
- When finished, turn in paper and obtain part 2

Part 2

- With Computer
  - Open notes, open Internet
  - Not allowed: communication with classmates!

You can choose how much time to allocate to Part 1 versus Part 2, but a safe goal is to leave 1 hour for Part 2

Midterm Exam – Part 1

- Format: **Short answer questions**
  - A few words up to a few sentences
- Content
  - All class material up through end of Lab 5
  - Lab 6 is **not** on the midterm
- Study Tips
  - Review lecture slides
  - Review questions asked in lab reports

Midterm Exam – Part 2

- Format: **Perform "lab-like" activities**
  - Utilize version control?
  - Create a Makefile?
  - Profile a program for CPU / memory utilization?
  - Use Valgrind to identify a memory leak, and then fix it?
- **You will only be given a text description of the problem**
  - You determine the right command(s) to use

Midterm Exam

- **Note:** I will be off-campus Oct 15th and 16th *(all day, both days, including the exam time)*
  - Interviewing new SOECS Dean candidates since Dr. Jain is retiring
- **Proctor:** Mike Lagomarsino (ECPE grad student)
  - Mike will have a copy of the solutions and is authorized to make *official decisions* in the event of questions

Bring Laptop!

**** Bring your laptop to the midterm! ****

Memory Hierarchy

Memory Hierarchy

Goal as system designers:

Fast Performance **and** Low Cost

Tradeoff: Faster memory is
more expensive than slower memory

Memory Hierarchy

- To provide the best performance at the lowest cost, memory is organized in a hierarchical fashion
  - **Small**, **fast** storage elements are kept **in the CPU**
  - **Larger**, **slower** main memory are **outside the CPU** (and accessed by a data bus)
  - **Largest**, **slowest**, permanent storage (disks, etc…) is **even further** from the CPU

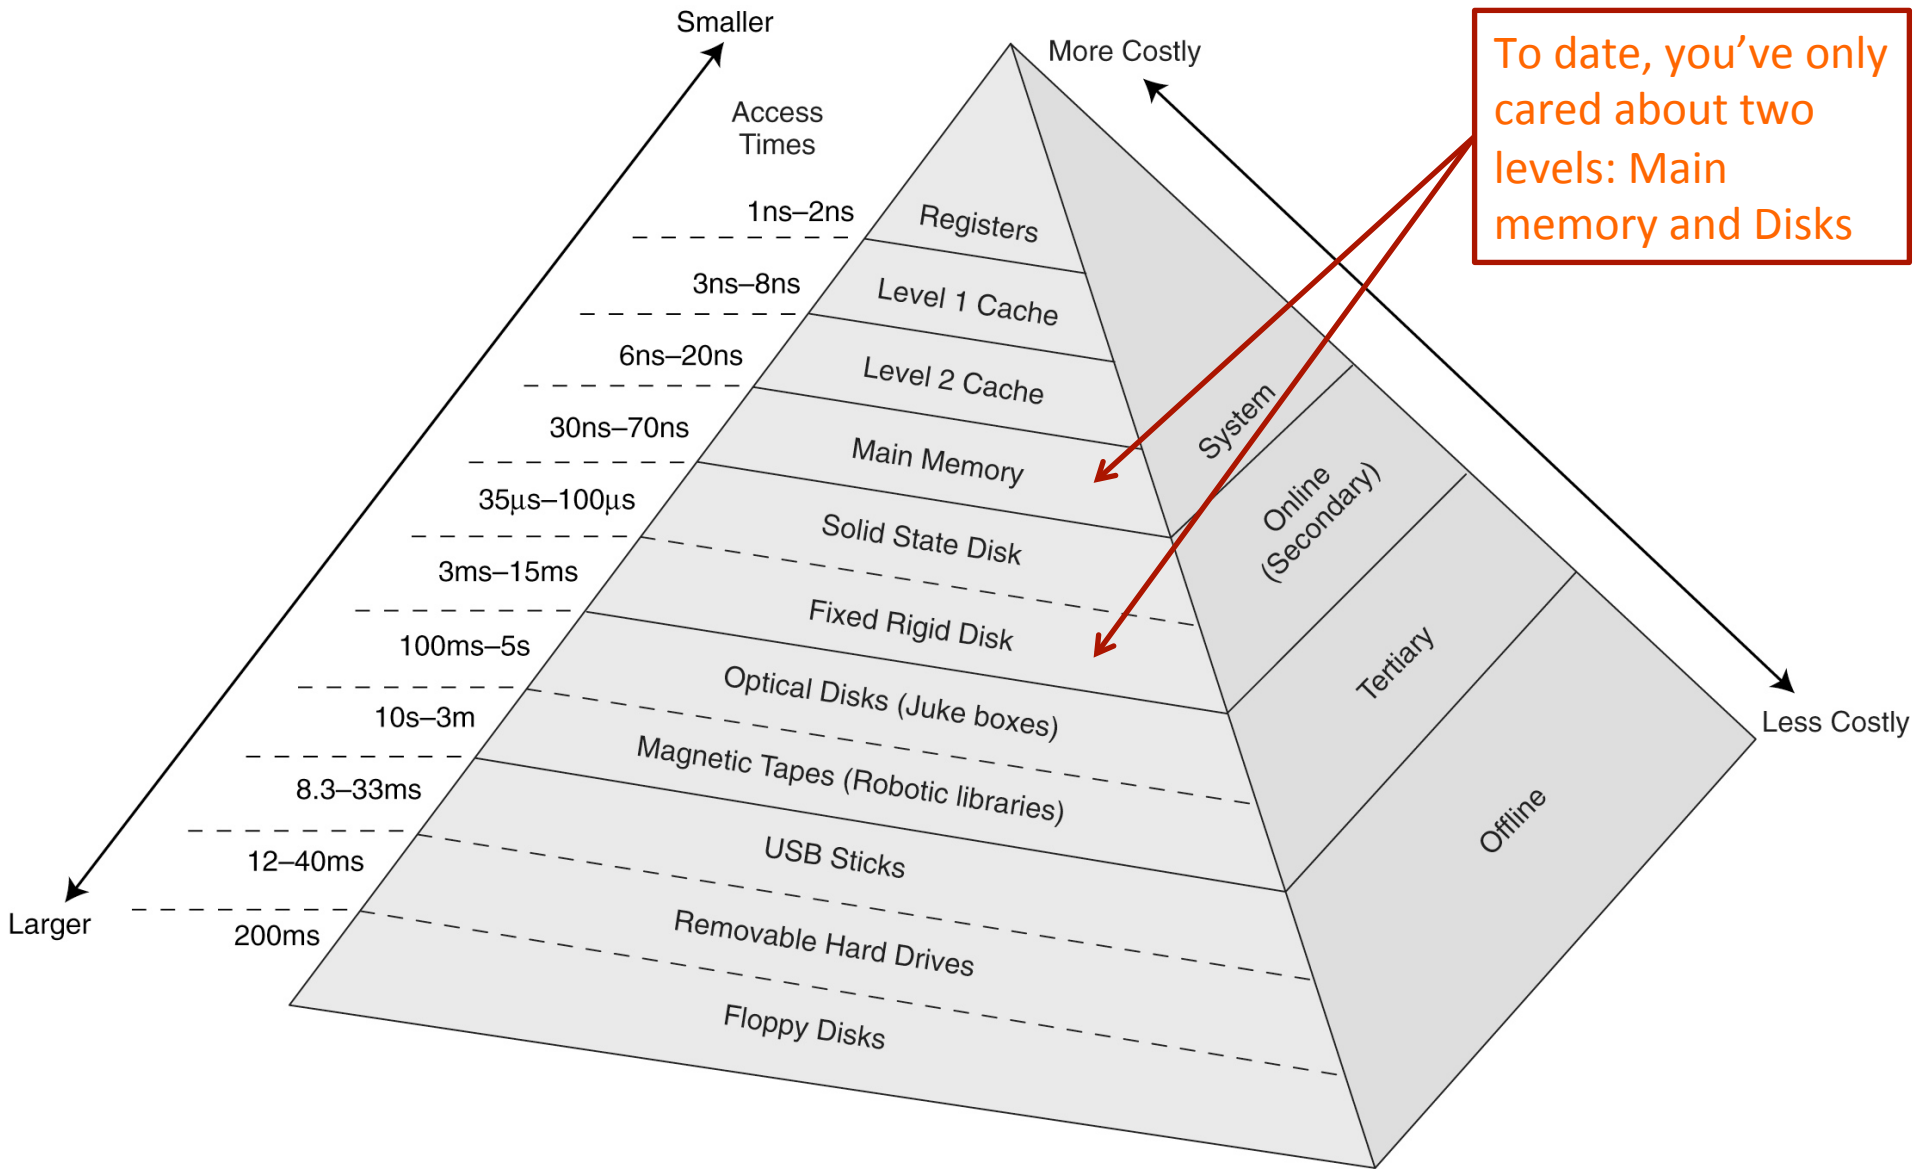

To date, you've only cared about two levels: Main memory and Disks

Memory Hierarchy
– Registers and Cache

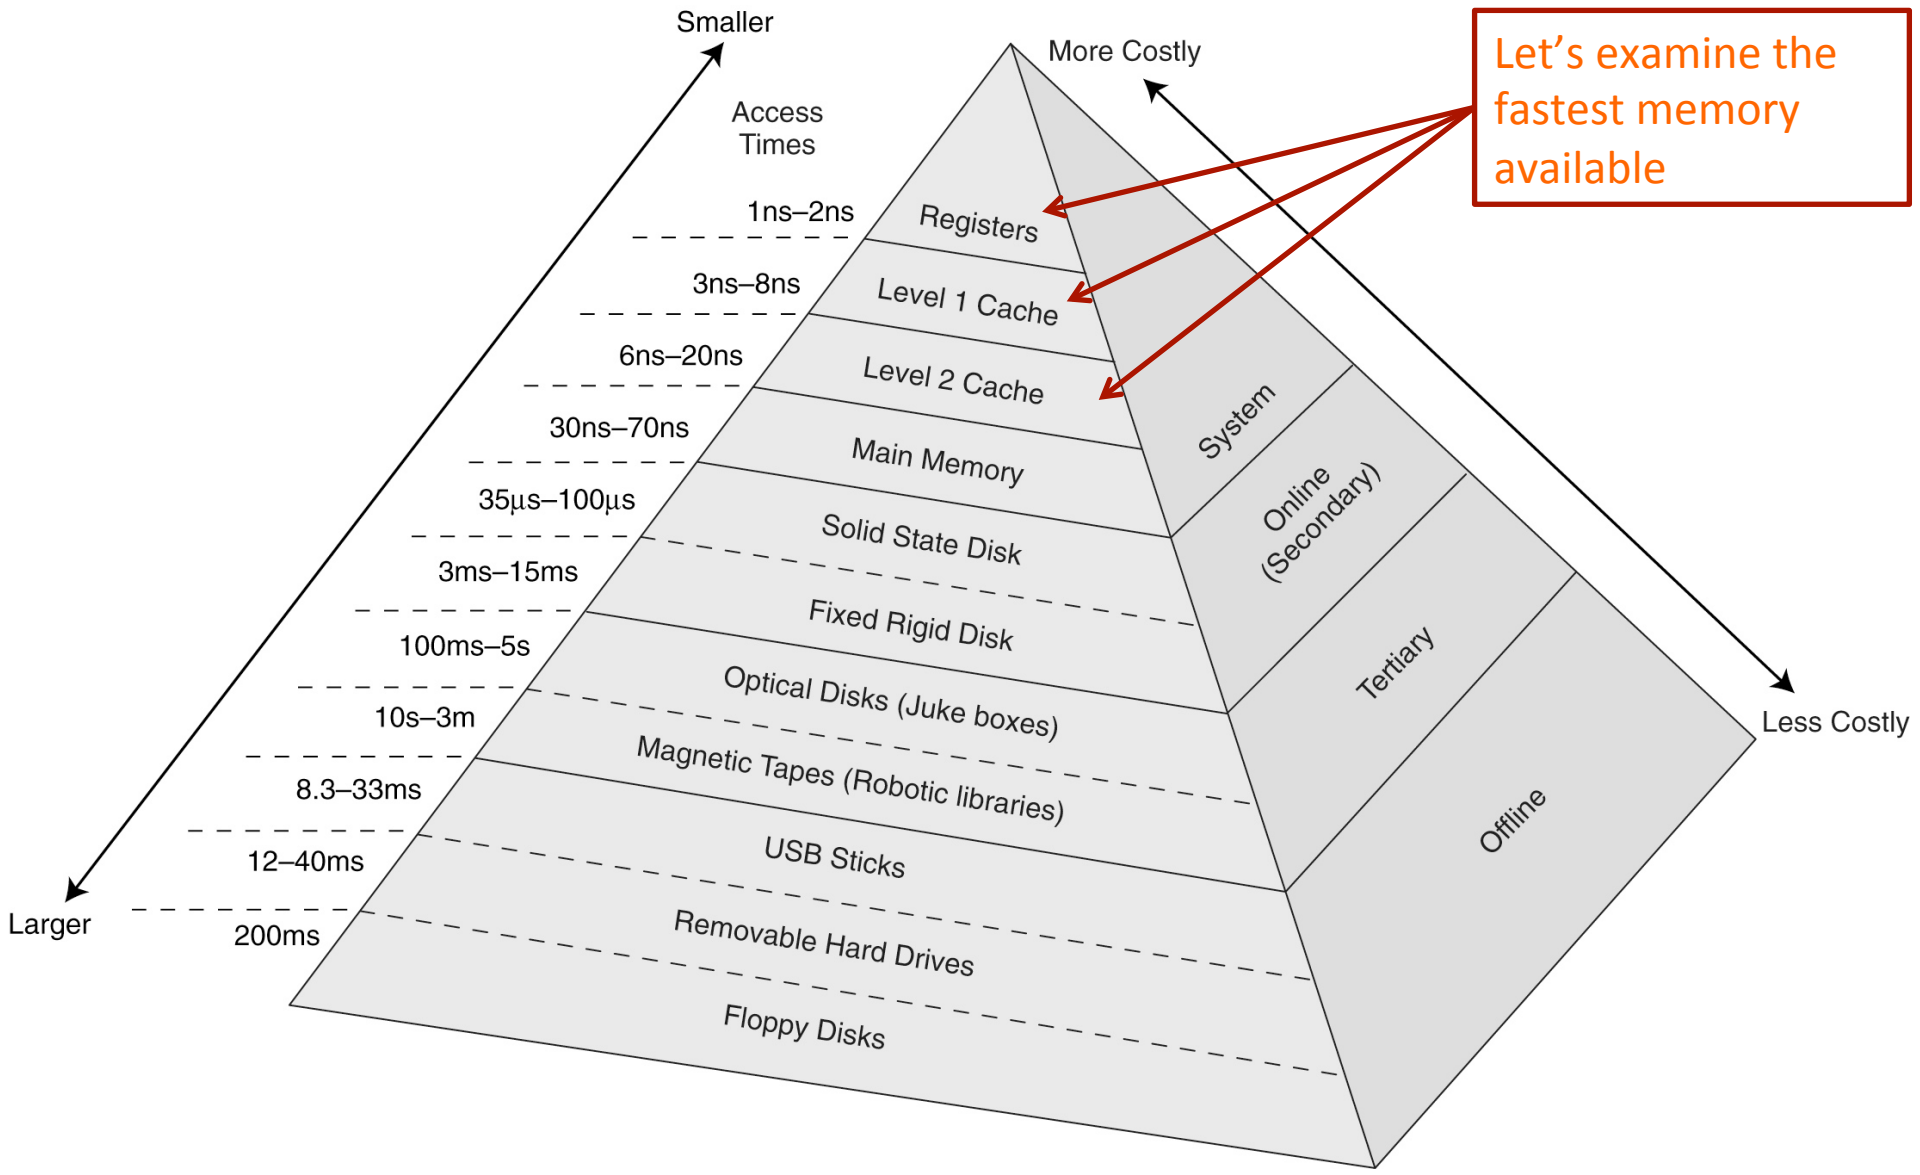
Smaller
Access Times
Larger
More Costly
Less Costly
1ns–2ns
3ns–8ns
6ns–20ns
30ns–70ns
35µs–100µs
3ms–15ms
100ms–5s
10s–3m
8.3–33ms
12–40ms
200ms
Registers
Level 1 Cache
Level 2 Cache
Main Memory
Solid State Disk
Fixed Rigid Disk
Optical Disks (Juke boxes)
Magnetic Tapes (Robotic libraries)
USB Sticks
Removable Hard Drives
Floppy Disks
System
Online (Secondary)
Tertiary
Offline
Let's examine the fastest memory available

Memory Hierarchy – Registers

- Storage locations available **on the processor** itself
- **Manually** managed by the assembly programmer or compiler
- *You'll become good friends with registers when we do MIPS assembly programming*

Memory Hierarchy – Caches

- **What is a cache?**
  - Speed up memory accesses by storing recently used data closer to the CPU
  - **Closer** that main memory – on the CPU itself!
  - Although cache is much smaller than main memory, its access time is much faster!
  - Cache is **automatically** managed by the hardware memory system
    - *Clever programmers can help the hardware use the cache more effectively*

Memory Hierarchy – Caches

- **How does the cache work?**
  - Not going to discuss how caches work internally
    - If you want that, take ECPE 173!
  - This class is focused on *what does the programmer need to know about the underlying system*

Memory Hierarchy – Access

- CPU wishes to **read data** (needed for an instruction)
  1. Does the instruction say it is in a register or memory?
     - If register, go get it!
  2. If in memory, send request to nearest memory (the cache)
  3. If not in cache, send request to main memory
  4. If not in main memory, send request to virtual memory (the disk)

(Cache) Hits versus Misses

Hit

- When data is found at a given memory level (e.g. a cache)

Miss

- When data is **not** found at a given memory level (e.g. a cache)

You want to write programs that produce a lot of **hits**, not misses!

Memory Hierarchy – Cache

- Once the data is located and delivered to the CPU, it will also be saved into cache memory for future access
  - We often save more than just the specific byte(s) requested
  - Typical: Neighboring 64 bytes (called the **cache line size**)
- **Hardware Prefetcher**
  - Look at the memory addresses used by your program
  - Guess pattern…
  - Load cache lines into cache **before** your program actually needs them
  - **Cache appears "magical!"**

Cache Locality

Principle of Locality

Once a byte is accessed, it is likely that a nearby data element will be needed soon, or perhaps even the same element again

Cache Locality

- **Temporal locality** – Recently-accessed data elements tend to be accessed again
  - Imagine a *loop counter*…
- **Spatial locality** - Accesses tend to cluster in memory
  - Imagine scanning through all elements in an array, or running several sequential instructions in a program

Programs with good locality **run faster** than programs with poor locality

A program that randomly accesses memory addresses (but never repeats) will gain no benefit from a cache

Recap – Cache

- **Which is bigger – a cache or main memory?**
  - Main memory
- **Which is faster to access – the cache or main memory?**
  - Cache – It is **smaller** (which is faster to search) and **closer** to the processor (signals take less time to propagate to/ from the cache)
- **Why do we add a cache between the processor and main memory?**
  - Performance – hopefully frequently-accessed data will be in the faster cache (so we don't have to access slower main memory)

Recap – Cache

- **Which is manually controlled – a cache or a register?**
  - Cache is automatically controlled by hardware
  - Registers are manually controlled by the assembly language program (or the compiler)
- **Suppose a program wishes to read from a particular memory address. Which is searched first – the cache or main memory?**
  - Search the cache first – otherwise, there's no performance gain

Recap – Cache

- **Suppose there is a cache miss (data not found) during a 1 byte memory read operation. How much data is loaded into the cache?**
  - Trick question – we always load data into the cache **1 "line" at a time**.
  - Cache line size varies – 64 bytes on a Core i7 processor

Cache Example – Intel Core i7 980x

- High-end 6 core processor with a sophisticated multi-level cache hierarchy
- 3.5GHz, 1.17 billion transistors (!!!)

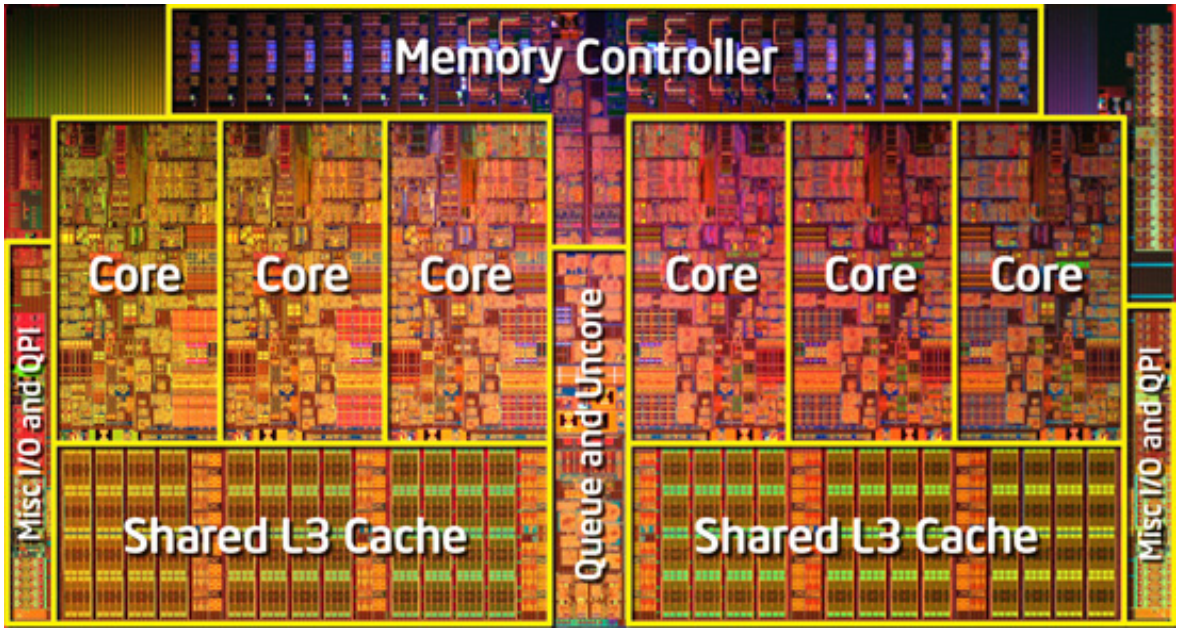

Cache Example – Intel Core i7 980x

- Each processor core has its own a L1 and L2 cache
  - 32kB Level 1 (L1) data cache
  - 32kB Level 1 (L1) instruction cache
  - 256kB Level 2 (L2) cache (both instruction and data)
- The entire chip (all 6 cores) **share** a single 12MB Level 3 (L3) cache

Cache Example – Intel Core i7 980x

- Access time? (Measured in 3.5GHz clock cycles)
  - 4 cycles to access L1 cache
  - 9-10 cycles to access L2 cache
  - 30-40 cycles to access L3 cache
- Smaller caches are faster to search
  - And can also fit closer to the processor core
- Larger caches are slower to search
  - Plus we have to place them further away

Caching is Ubiquitous!

Many types of "cache" in computer science, with different meanings

Type	What Cached	Where Cached	Managed By
CPU Registers	4-byte or 8-byte word	On-chip CPU registers	Compiler or assembly programmer
TLB	Address Translation (Virtual->Physical Memory Address)	On-chip TLB	Hardware MMU (Memory Management Unit)
Buffer cache	Parts of files on disk	Main memory	Operating Systems
Disk cache	Disk sectors	Disk controller	Controller firmware
Browser cache	Web pages	Local Disk	Web browser

Memory Hierarchy – Virtual Memory

Virtual Memory

Virtual Memory is a BIG LIE!

- We **lie** to your application and tell it that the system is simple:
  - Physical memory is infinite! (or at least huge)
  - You can access *all* of physical memory
  - Your program starts at *memory address zero*
  - Your memory address is *contiguous* and *in-order*
  - Your memory is *only RAM* (main memory)

What the System Really Does

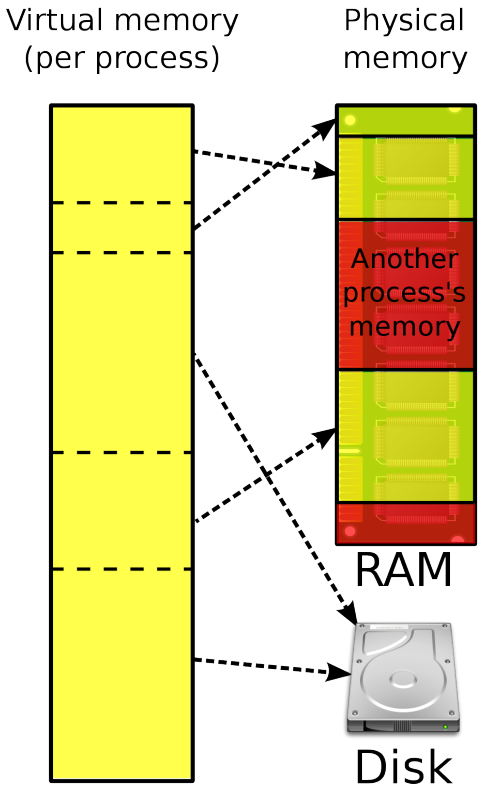

Why use Virtual Memory?

- We want to run multiple programs on the computer concurrently *(multitasking)*
  - Each program needs its own separate memory region, so physical resources must be divided
  - The amount of memory each program takes could vary dynamically over time (and the user could run a different mix of apps at once)
- We want to use multiple types of storage (main memory, disk) to increase performance and capacity
- We don't want the programmer to worry about this
  - Make the processor architect handle these details

Pages and Virtual Memory

- Main memory is divided into **pages** for virtual memory
  - Pages size = 4kB
  - Data is moved between main memory and disk at a page granularity
    - i.e. we don't move single bytes around, but rather big groups of bytes

Pages and Virtual Memory

- Main memory and virtual memory are divided into equal sized pages
- The entire address space required by a process need not be in memory at once
  - Some pages can be on disk
    - Push the unneeded parts out to slow disk
  - Other pages can be in main memory
    - Keep the frequently accessed pages in faster main memory
- The pages allocated to a process do not need to be stored contiguously-- either on disk or in memory

Virtual Memory Terms

- **Physical address** – the actual memory address in the *real* main memory
- **Virtual address** – the memory address that is seen in your program
  - We need some special hardware/software to map between virtual addresses and physical addresses!
- **Page faults** – a program accesses a virtual address that is not currently resident in main memory (at a physical address)
  - The data must be loaded from disk!
- **Pagefile –** The file on disk that holds memory pages
  - Usually twice the size of main memory

Cache Memory vs Virtual Memory

- Goal of **cache memory**
  - Faster memory access speed (**performance**)
- Goal of **virtual memory**
  - Increase memory **capacity** without actually adding more main memory
    - Data is written to disk
    - If done carefully, this can improve performance
    - If overused, performance suffers greatly!
  - Increase system flexibility when running multiple user programs (as previously discussed)

Memory Hierarchy – Magnetic Disks

Magnetic Disk Technology

- Hard disk platters are mounted on spindles
- Read/write heads are mounted on a comb that swings radially to read the disk
  - All heads move **together**!

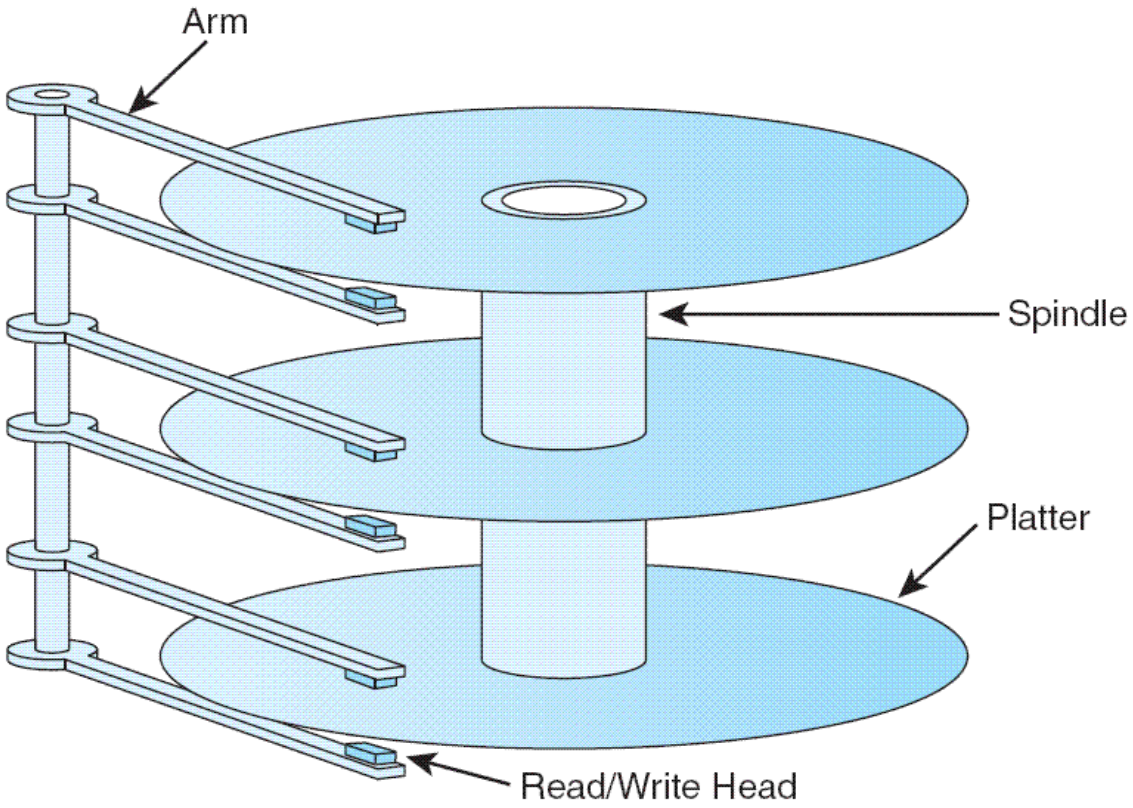

Magnetic Disk Technology

- There are a number of *electromechanical* properties of hard disk drives that determine how fast its data can be accessed
- **Seek time** – time that it takes for a disk arm to move into position over the desired cylinder
- **Rotational delay** – time that it takes for the desired sector to move into position beneath the read/ write head
- Seek time + rotational delay = **access time**

How Big Will Hard Drives Get?

- Advances in technology have defied all efforts to define the ultimate upper limit for magnetic disk storage
  - In the 1970s, the upper limit was thought to be around 2Mb/in^2
- As data densities increase, bit cells consist of proportionately fewer magnetic grains
  - There is a point at which there are too few grains to hold a value, and a 1 might spontaneously change to a 0, or vice versa
  - This point is called the **superparamagnetic limit**

How Big Will Hard Drives Get?

- **When will the limit be reached?**
- In 2006, the limit was thought to lie between 150Gb/in^2 and 200Gb/in^2 *(with longitudinal recording technology)*
- 2010: Commercial drives have densities up to 667Gb/in^2
- 2012: Seagate demos drive with 1 Tbit/in^2 density
  - *With heat-assisted magnetic recording* – they use a laser to heat bits before writing
  - Each bit is ~12.7nm in length (a dozen atoms)

Memory Hierarchy – SSDs

Emergence of Solid State Disks (SSD)

- **Hard drive advantages?**
  - Low cost per bits
- **Hard drive disadvantages?**
  - Very slow compared to main memory
  - Fragile (ever dropped one?)
  - Moving parts wear out
- Reductions in flash memory cost is opening another possibility: **solid state drives** (SSDs)
  - SSDs appear like hard drives to the computer, but they store data in non-volatile **flash memory** circuits
  - Flash is **quirky!** Physical limitations pose engineering challenges…

Flash Memory

- Typical flash chips are built from dense arrays of NAND gates
- Different from hard drives – we **can't** read/write a single bit (or byte)
  - **Reading or writing?** Data must be read from an entire **flash page** (2kB-8kB)
    - Reading much faster than writing a page
    - It takes some time before the cell charge reaches a stable state
  - **Erasing?** An entire **erasure block** (32-128 pages) must be erased (set to all 1's) first before individual bits can be written (set to 0)
    - Erasing takes two orders of magnitude more time than reading

Flash-based Solid State Drives (SSDs)

Advantages

- Same block-addressable I/O interface as hard drives
- No mechanical latency
  - Access latency is independent of the access pattern
    - Compare this to hard drives
- Energy efficient (no disk to spin)
- Resistant to extreme shock, vibration, temperature, altitude
- Near-instant start-up time

Challenges

- Limited endurance and the need for **wear leveling**
- Very slow to erase blocks (needed before reprogramming)
  - Erase-before-write
- Read/write asymmetry
  - Reads are faster than writes

Flash Translation Layer

- **Flash Translation Layer (FTL)**
  - Necessary for flash reliability and performance
  - **"Virtual" addresses** seen by the OS and computer
  - **"Physical" addresses** used by the flash memory
- Perform writes out-of-place
  - Amortize block erasures over many write operations
- Wear-leveling
  - Writing the same "virtual" address repeatedly won't write to the same physical flash location repeatedly!