Announcements

- Program is due on Tuesday at 10:30 AM
- This is the final regular lecture
  - the next lecture will be a review for the final
- Course evaluations were handed out at the end of class

Disks

- Several types:
  - Hard Disks - rigid surface with magnetic coating
  - Floppy disks - flexible surface with magnetic coating
  - Optical (read only, write once, multi-write)
- Hard Disk Drives:
  - collection of platters
  - platters contain concentric rings called tracks
  - tracks are divided into fixed sized units called sectors
  - a cylinder is a collection of all tracks equal distant from the center of disk
  - a disk head floats on a cushion of air above the platter
    - when it hits the disk its called a head crash
      - metal flies everywhere: time to buy a new disk

Disks (cont)

- Factors in disk performance
  - seek time
    - one head reads all tracks on a platter and it needs to move (seek) to the correct track for the requested data
  - rotational latency
    - head is at a fixed spot on the disk and has to wait for the sector to be used to rotate (spin) under the head
- Current Performance:
  - capacity: megabytes to tens of gigabytes
  - throughput: sustained < 10 megabytes/sec
  - latency: mili-seconds

Networks

- Two major types
  - bus based (Ethernet)
  - switched (ATM)
  - break data into packets

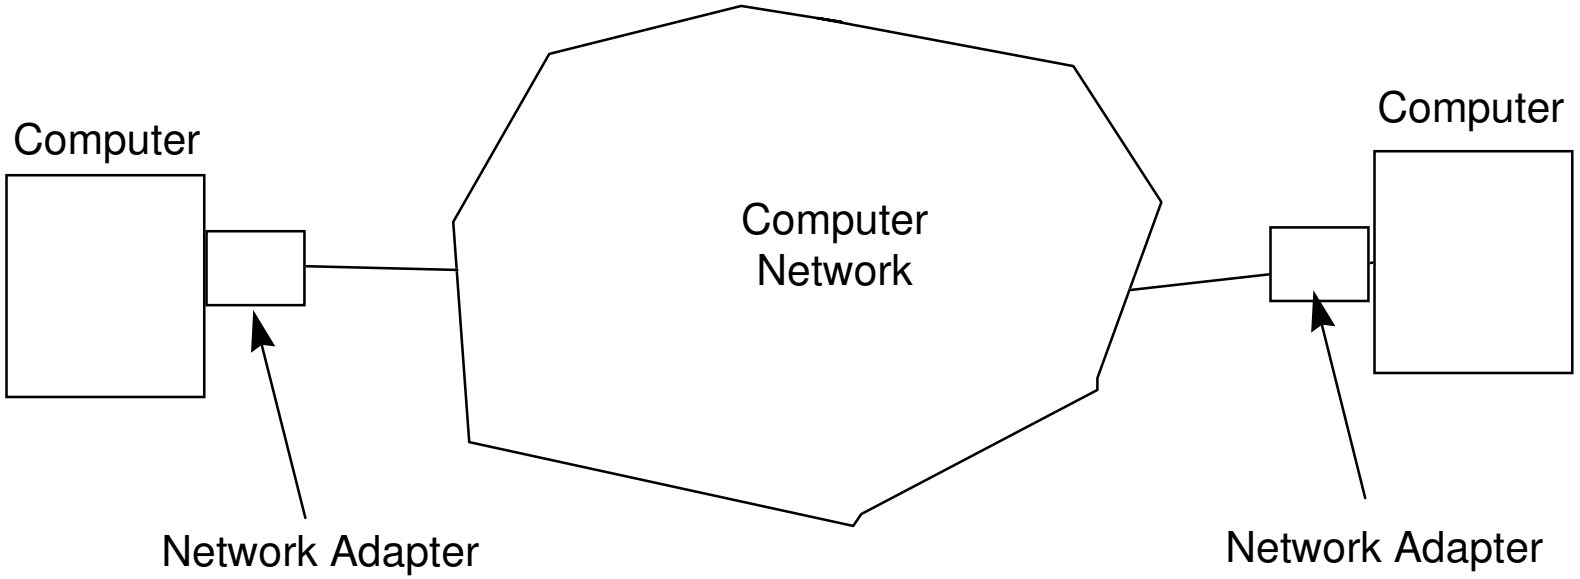

Ethernet

- 10 Mbps (to 100 Mbps)
- mili-second latency
- limited to several kilometers in distance
- variable sized units of transmission
- bus based protocol
  - requests to use the network can collide

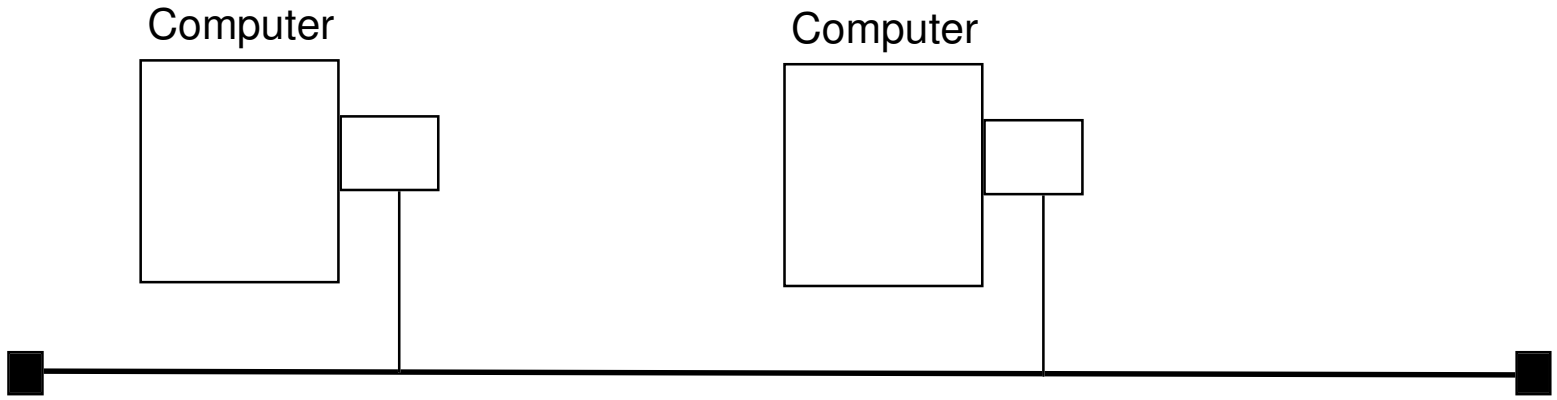

ATM (Asynchronous Transfer Mode)

- 155Mbps and up
- fixed sized unit of transmission called a cell
  - cells are 48 bytes plus 5 bytes header
- switch based protocol
- for both local area and wide area networking
- I-way
  - coast to coast ATM network
  - minimum speed is 155Mps
  - see http://www.iway.org/ for more info

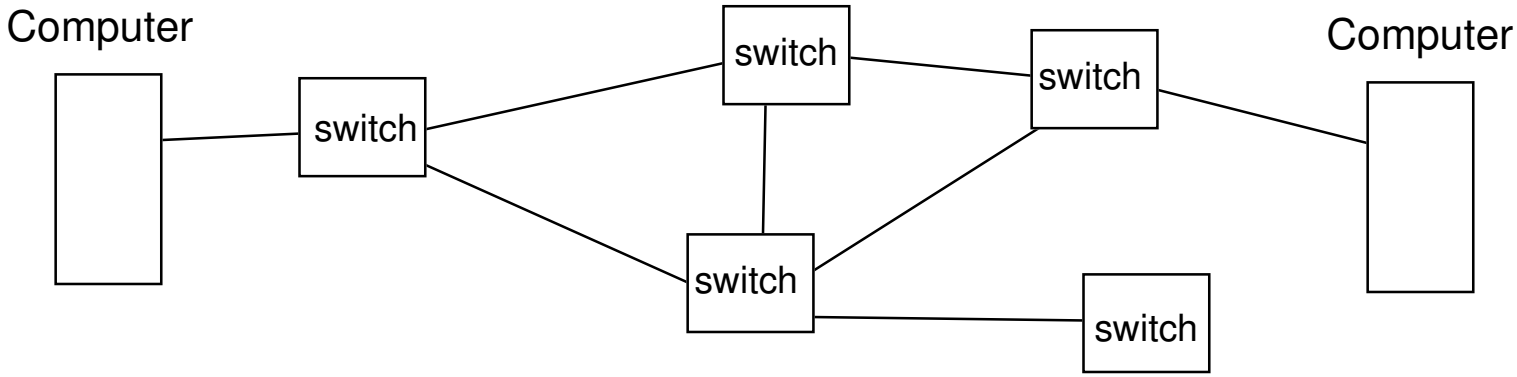

How to Make Computations go faster?

- Do more than one thing at once
  - called parallel computing
- Processors can work on more than one instruction at once
- Have more than one processor

Pipelining

- different parts of the CPU are dedicated to parts of executing an instruction
  - one part does the fetch
  - another the decode
  - while yet a third does the memory access
- Sort of like a washer/dryer pair
- Each stage in the pipeline works on a different instruction

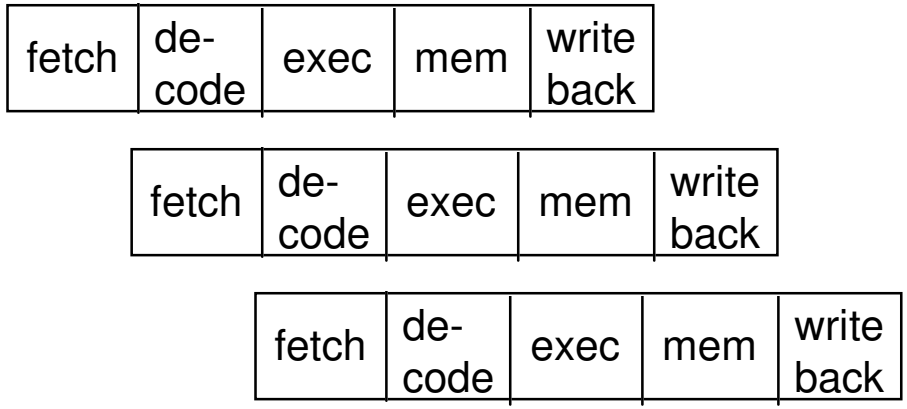

Superscalar

- Two separate copies of the parts of the processor
  - registers are shared
- Each part works on a different instruction
- Sort of like having two combined washer/dryer units

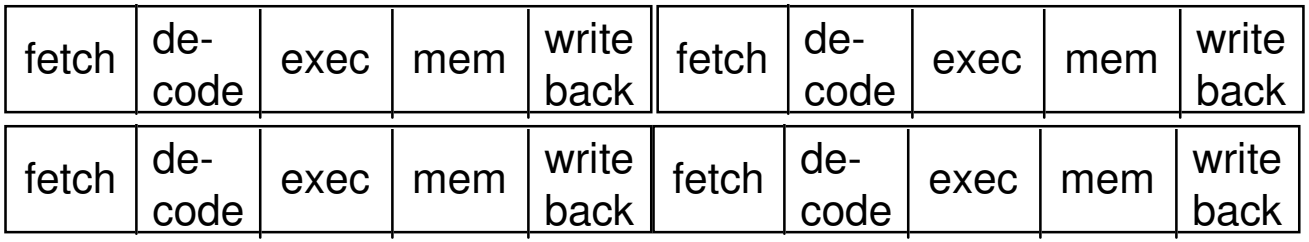

Problems with these Schemes

- Instructions are not (necessarily) independent
  - might have to wait for the result of one instruction before you can execute the next one
- Hardware is more complex
- Performance is less predictable