CSE398: Network Systems Design

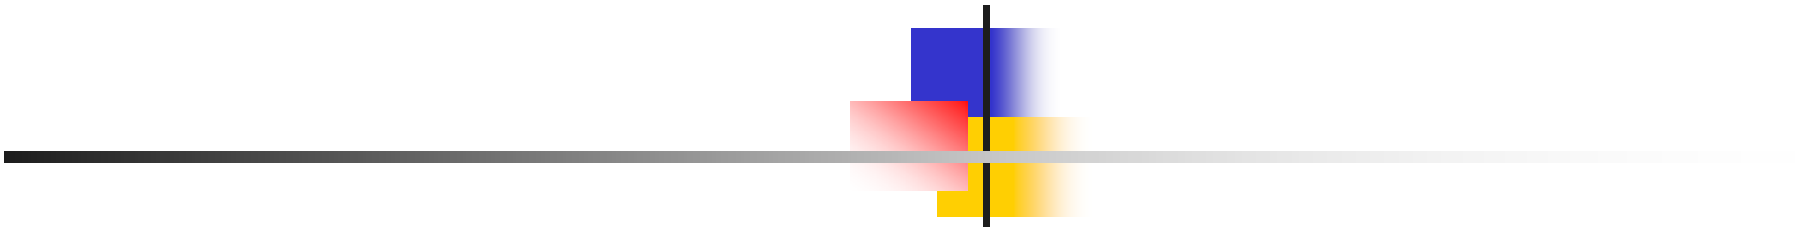

Instructor: Dr. Liang Cheng
Department of Computer Science and Engineering
P.C. Rossin College of Engineering & Applied Science
Lehigh University

February 23, 2005

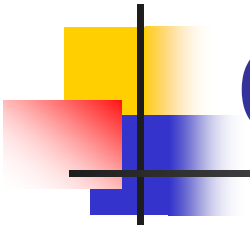

Outline

- Recap
  - Protocol software
- Hardware architecture
- Summary and homework

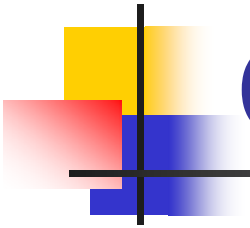

Outline

- Recap
- Hardware architecture for protocol processing
- Summary and homework

A Brief History of Computer Hardware

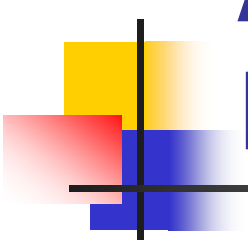

- 1940s
  - Beginnings
- 1950s
  - Consolidation on von Neumann architecture
  - I/O controlled by CPU
- 1960s
  - I/O becomes important
  - Evolution of third generation architecture with interrupts

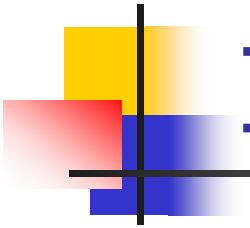

I/O Processing

- Low-end systems (e.g., microcontrollers)
  - Dumb I/O interfaces
  - CPU does all the work (polls devices)
  - Single, shared memory
  - Low cost, but low speed
- Mid-range systems (e.g., minicomputers)
  - Single, shared memory
  - I/O interfaces contain logic for transfer and status
  - Operations: CPU starts device then resumes processing
  - Device: transfers data to / from memory; interrupts when operation complete
- High-end systems (e.g., mainframes)
  - Separate, programmable I/O processor
  - OS downloads code to be run
  - Device has private on-board buffer memory

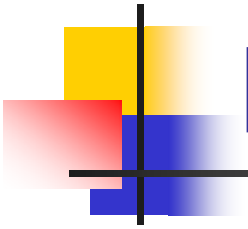

Networking Systems Evolution

- ~ Thirty year history
- Same trend as computer architecture
  - Began with central CPU
  - Shift to emphasis on I/O
- Three main generations

First Generation Network Systems

- Traditional software-based router
- Used conventional (minicomputer) hardware
  - Other generations later
  - Single general-purpose processor handles most tasks
  - Single shared memory, I/O over a bus
  - Network interface cards use same design as other I/O devices

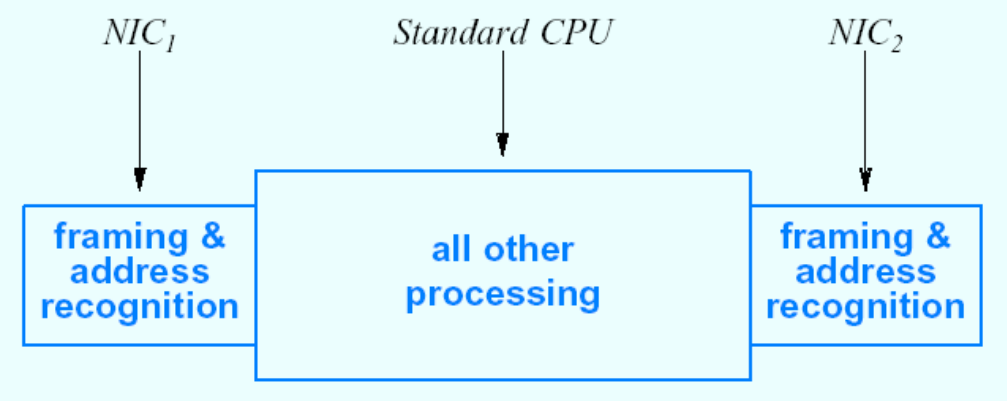

How Fast a CPU Need to Be?

- Depends on
  - Rate at which data arrives (pp. 104)
    - Data rate (bits per second): per interface rate, aggregate rate
    - Packet rate (packets per second): per interface & aggregate rate
    - Definition of fast data rate keeps changing
      - 1960: 10 Kbps
      - 1970: 1 Mbps
      - 1980: 10 Mbps
      - 1990: 100 Mbps
      - 2000: 1000 Mbps (1 Gbps)
      - 2003: 2400 Mbps
  - Amount of processing to be performed

Aggregate Rate vs. Per-Interface Rate

- Interface rate
  - Rate at which data enters / leaves
- Aggregate
  - Sum of interface rates (assumptions)
  - Measure of total data rate system can handle
  - Note: aggregate rate crucial if CPU handles traffic from all interfaces
- Packet rate vs. data rate
  - Sources of CPU overhead
    - Per-bit processing
    - Per-packet processing
  - Interface hardware handles much of per-bit processing

Possible Ways to Solve the CPU Bottleneck

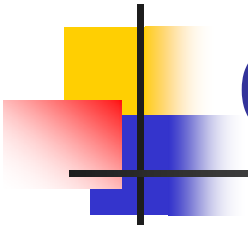

- Special-purpose coprocessors
- NICs with onboard processing
- Smart NICs with onboard stacks
- General principle
  - To optimize computation, move operations that account for the most CPU time from software into hardware.
- Fine-grain parallelism
- Symmetric coarse-grain parallelism
- Asymmetric coarse-grain parallelism
- Cell switching
- Data pipelines

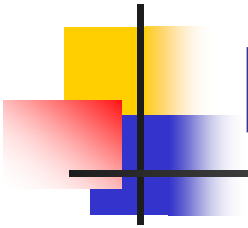

Fine-Grain Parallelism

- Multiple processors
- Instruction-level parallelism
- Example: parallel checksum
  - Add values of eight consecutive memory locations at the same time
- Assessment: insignificant advantages for packet processing

Symmetric Coarse-Grain Parallelism

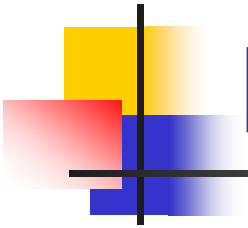

- Symmetric multiprocessor hardware
  - Multiple, identical processors
- Typical design: each CPU operates on one packet
- Requires coordination
- Assessment: coordination and data access means N processors cannot handle N times more packets than one processor

Asymmetric Coarse-Grain Parallelism

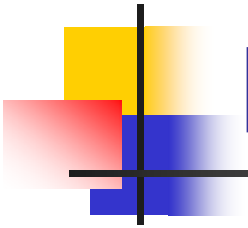

- Multiple processors
- Each processor
  - Optimized for specific task
  - Includes generic instructions for control
- Assessment
  - Same problems of coordination and data access as symmetric case
  - Designer much choose how many copies of each processor type

Parallel & Pipelined Hardware

- Cell switching as an alternative to new hardware: ATM
  - Fixed-size packets
    - Allows fixed-size buffers
    - Guaranteed time to transmit/receive
  - Relative (connection-oriented) addressing
    - Smaller address size
    - Label on packet changes at each switch
    - Requires connection setup
- Data pipeline
  - Move each packet through series of processors
  - Each processor handles some tasks
  - Assessment
    - Well-suited to many protocol processing tasks
    - Individual processor can be fast

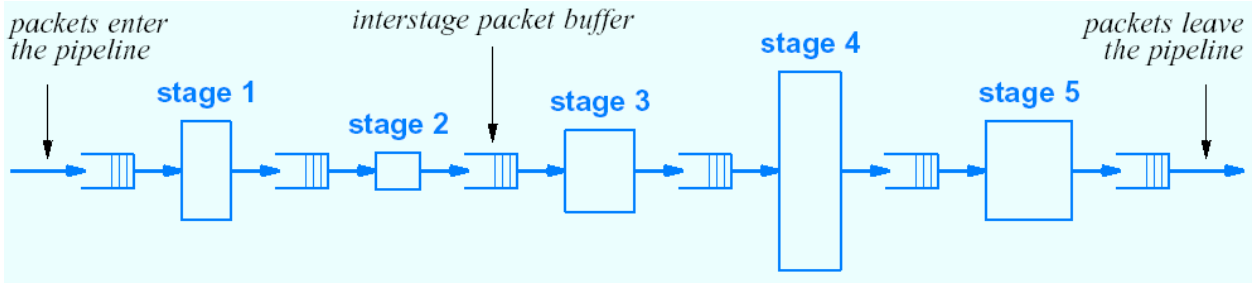

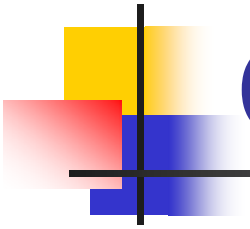

Outline

- Recap
- Hardware architecture
- Summary
  - No homework today
  - Midterm on Wed.