CSE398: Network Systems Design

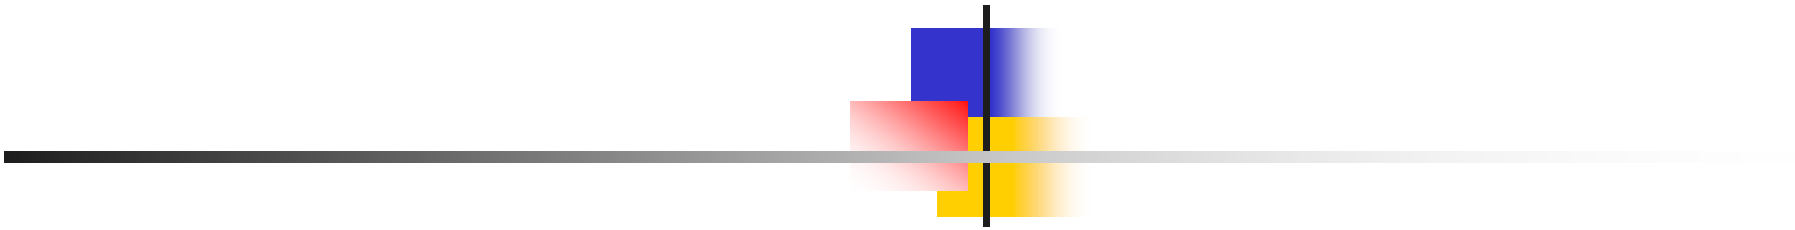

Instructor: Dr. Liang Cheng
Department of Computer Science and Engineering
P.C. Rossin College of Engineering & Applied Science
Lehigh University

March 28, 2005

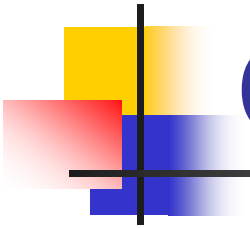

Outline

- Recap
  - Network processor architectures
- Issues in scaling a network processor
- Summary and homework

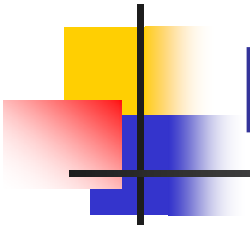

Recap

- Network processor architecture is characterized by
  - Processor hierarchy
  - Memory hierarchy
  - Internal buses
  - External interfaces
  - Special-purpose functional units
  - Support for concurrent or parallel execution
  - Programming model
  - Dispatch mechanisms

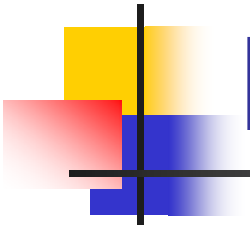

Design Questions

- Can we make network processors
  - Faster?
  - More powerful?
  - Cheaper?
  - More general?
  - Easier to use?
  - All of the above?
- Scalable design is fundamental

Scaling the Processor Hierarchy

- General purpose CPU
  - Highest level functionality
  - Administrative interface
  - System control
  - Overall management functions
  - Routing protocols
- Embedded processor
  - Intermediate functionality
  - Higher-layer protocols
  - Control of I/O processors
  - Exception and error handling
  - High-level ingress (e.g., reassembly)
  - High-level egress (e.g., traffic shaping)
- I/O processor
  - Basic packet processing
  - Classification
  - Forwarding
  - Low-level ingress operations
  - Low-level egress operations

Packet Flow through Hierarchy

- To maximize performance, packet processing tasks should be assigned to the lowest level processor capable of performing the task.

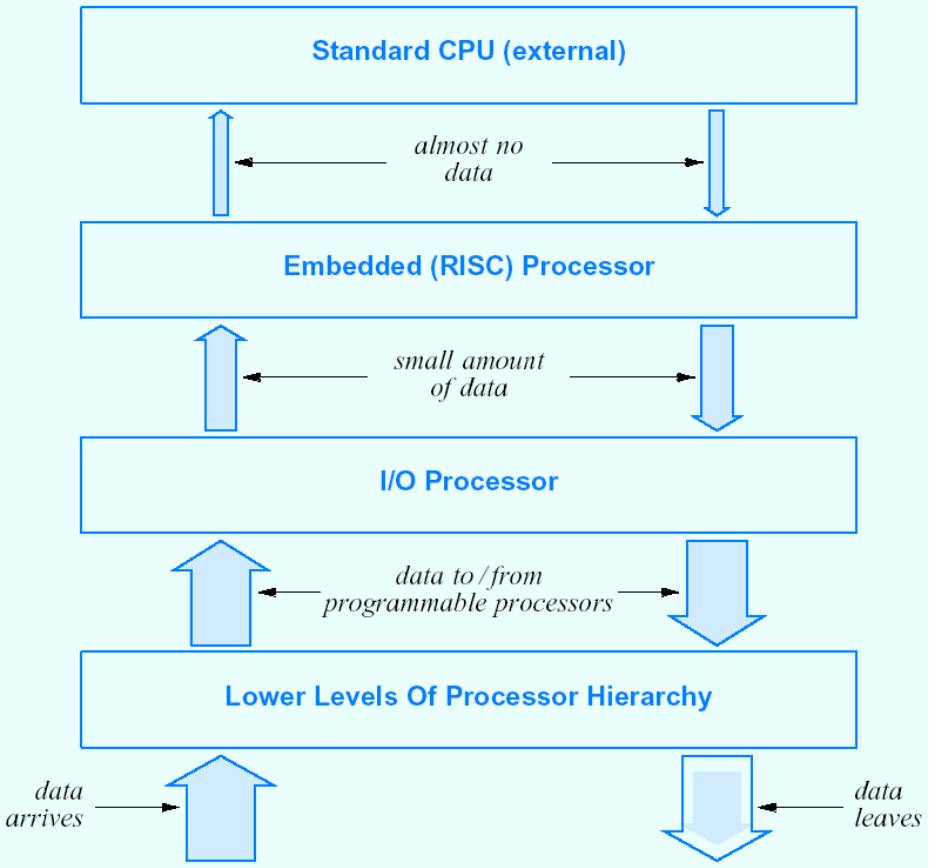

The Pyramid of Processor Scale

- Lower levels need the most increase

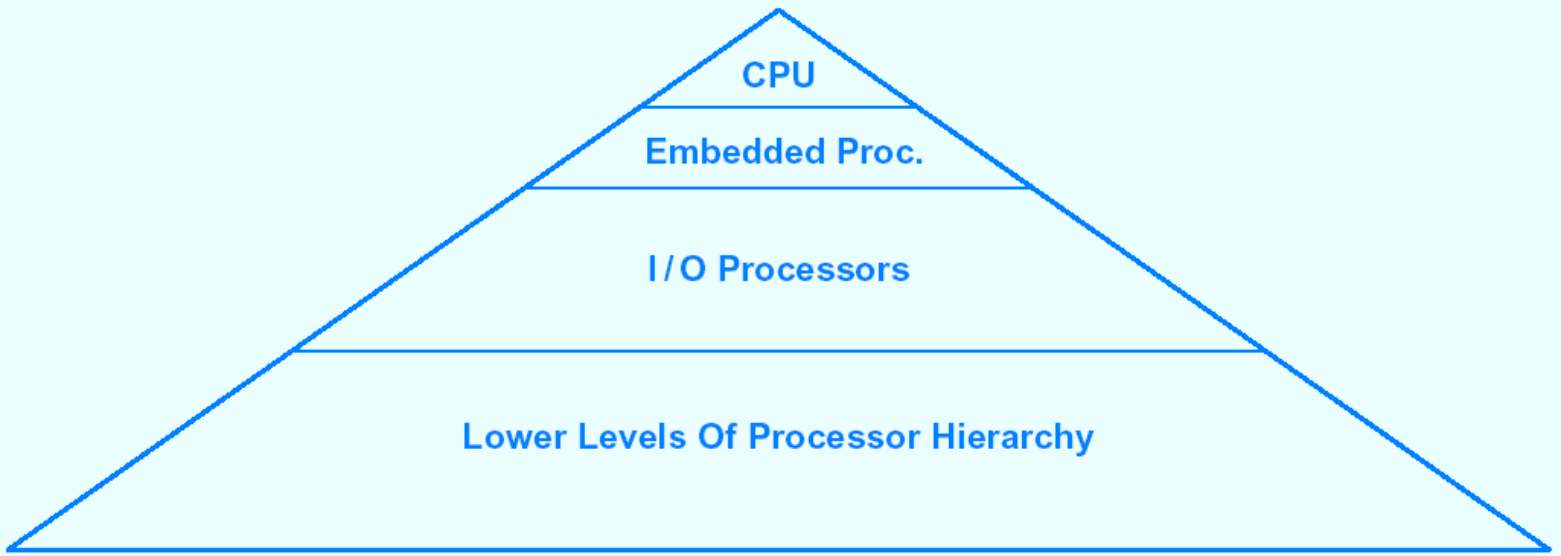

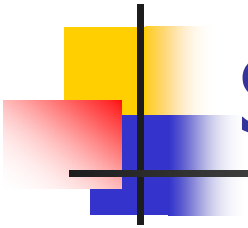

Scaling the Processor Hierarchy

- Make processors faster
- Use more concurrent threads
- Increase numbers of processors
- Increase processor types

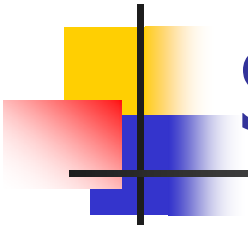

Scaling the Memory Hierarchy

- Size
- Speed
- Throughput
- Cost

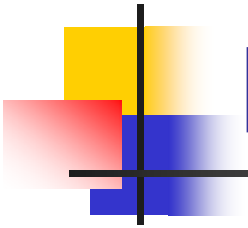

Memory Speed

- Access latency
  - Raw read/write access speed
  - SRAM 2 - 10 ns: flip-flop, 4-6 transistors
  - DRAM 50 - 70 ns: capacitor under 1 transistor
  - External memory takes order of magnitude longer than onboard
- Memory cycle time
  - Measure of successive read/write operations
  - Important for networking because packets are large
  - Read Cycle time (tRC) is time for successive fetch operations
  - Write Cycle time (tWC) is time for successive store operations

The Pyramid of Memory Scale

- Largest memory is least expensive

Memory Type	Rel. Speed	Approx. Size	On Chip?
Control store	100	10^3	yes
G.P. Registers†	90	10^2	yes
Onboard Cache	40	10^3	yes
Onboard RAM	7	10^3	yes
Static RAM	2	10^7	no
Dynamic RAM	1	10^8	no

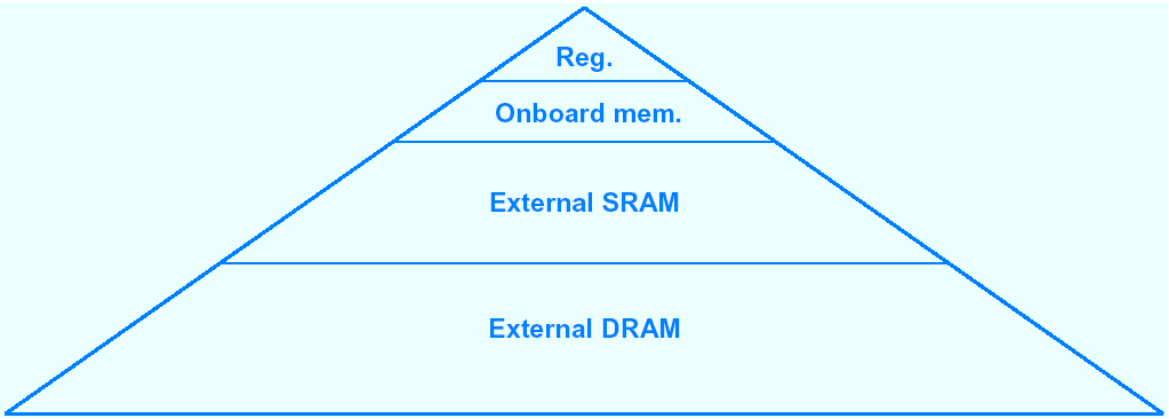

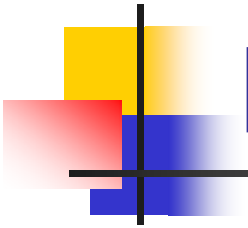

Memory Bandwidth

- General measure of throughput
- More parallelism in access path yields more throughput
- Cannot scale arbitrarily
  - Pinout limits
  - Processor must have addresses as wide as bus

Types of Memory

Memory Technology	Abbreviation	Purpose
Synchronized DRAM	SDRAM	Synchronized with CPU for lower latency
Quad Data Rate SRAM	QDR-SRAM	Optimized for low latency and multiple access
Zero Bus Turnaround SRAM	ZBT-SRAM	Optimized for random access
Fast Cycle RAM	FCRAM	Low cycle time optimized for block transfer
Double Data Rate DRAM	DDR-DRAM	Optimized for low latency
Reduced Latency DRAM	RLDRAM	Low cycle time and low power requirements

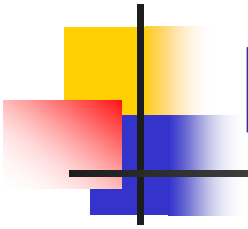

Memory Cache

- General-purpose technique better than
  - External memory
    - Memory controller
    - Magnitude longer to access than on-chip memory
- May not work well in network systems
  - Low temporal locality
  - Large cache size (either more entries or larger granularity of access)

Content Addressable Memory (CAM)

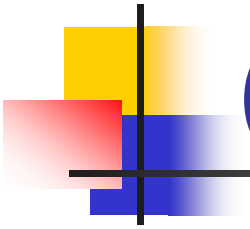

- Combination of mechanisms
  - Random access storage
  - Exact-match pattern search
- Rapid search enabled with parallel hardware

Arrangement of CAM

- Organized as array of slots

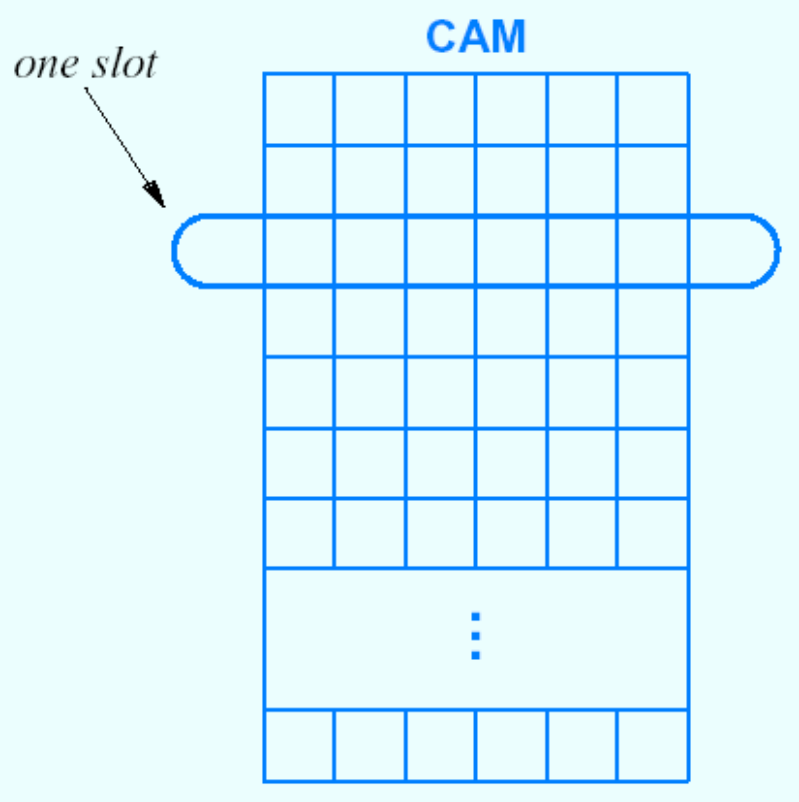

Lookup in Conventional CAM

- Given
  - Pattern for which to search
  - Known as *key*
- CAM returns
  - First slot that matches key, or
  - All slots that match key
- Algorithm

```
for each slot do {
    if ( key == slot ) {
      declare key matches slot;
    } else {
      declare key does not match slot;
    }
}
```

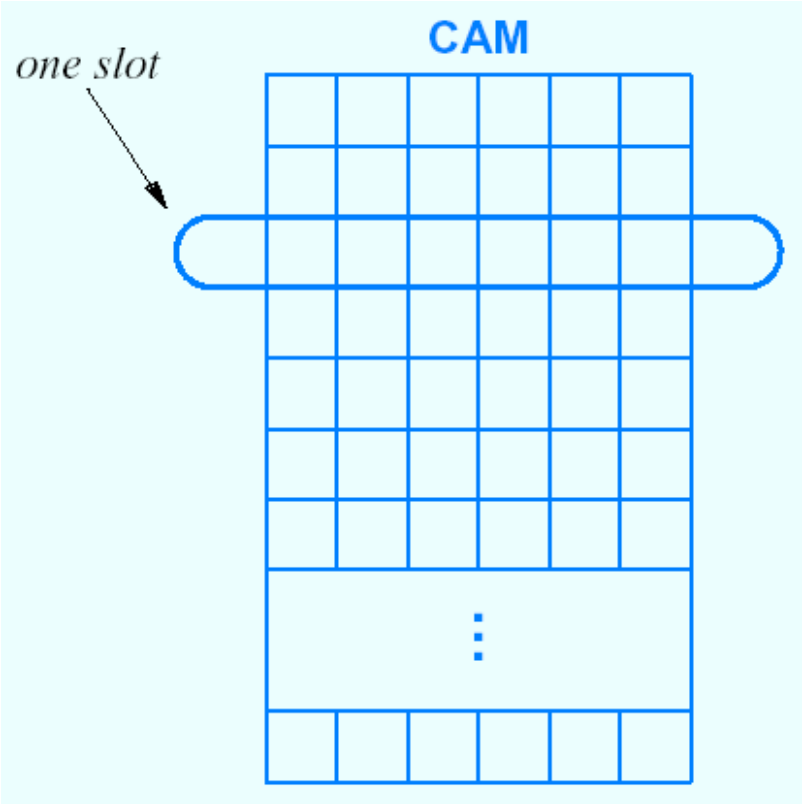

Ternary CAM (T-CAM)

- Allows masking of entries
- Good for network processor
- One, zero, don't care

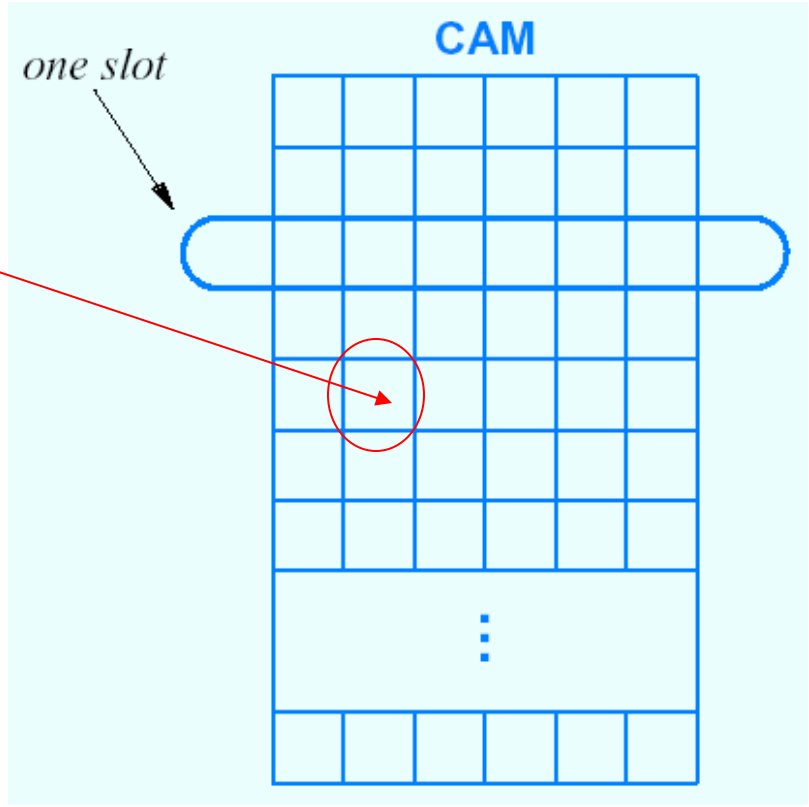

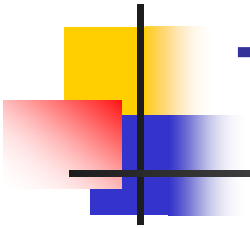

T-CAM Lookup

- Each slot has bit mask
- Hardware uses mask to decide which bits to test
- Algorithm

```
for each slot do {
    if ( ( key & mask ) == ( slot & mask ) ) {
      declare key matches slot;
    } else {
      declare key does not match slot;
    }
}
```

Partial Matching with a T-CAM

- Key matches slot #1

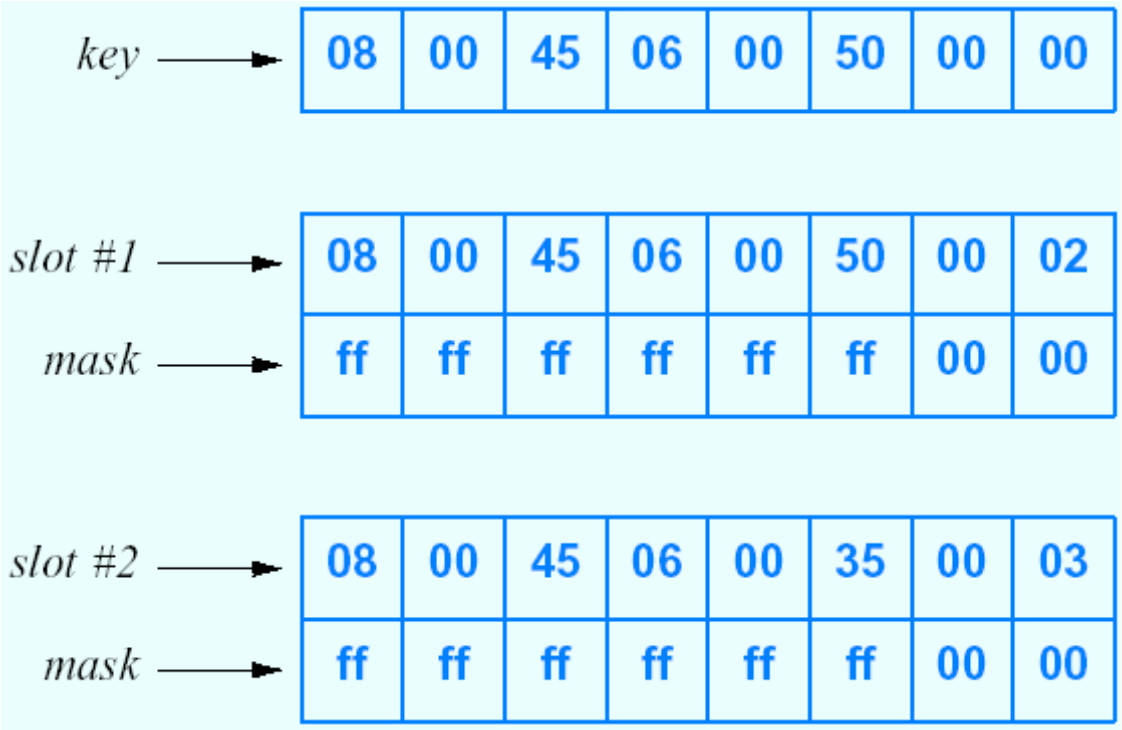

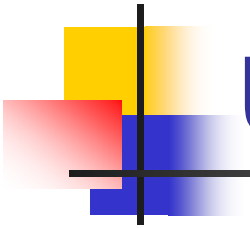

Using a T-CAM for Classification

- Extract values from fields in headers
- Form values in contiguous string
- Use a key for T-CAM lookup
- Store classification in slot

Classification using a T-CAM

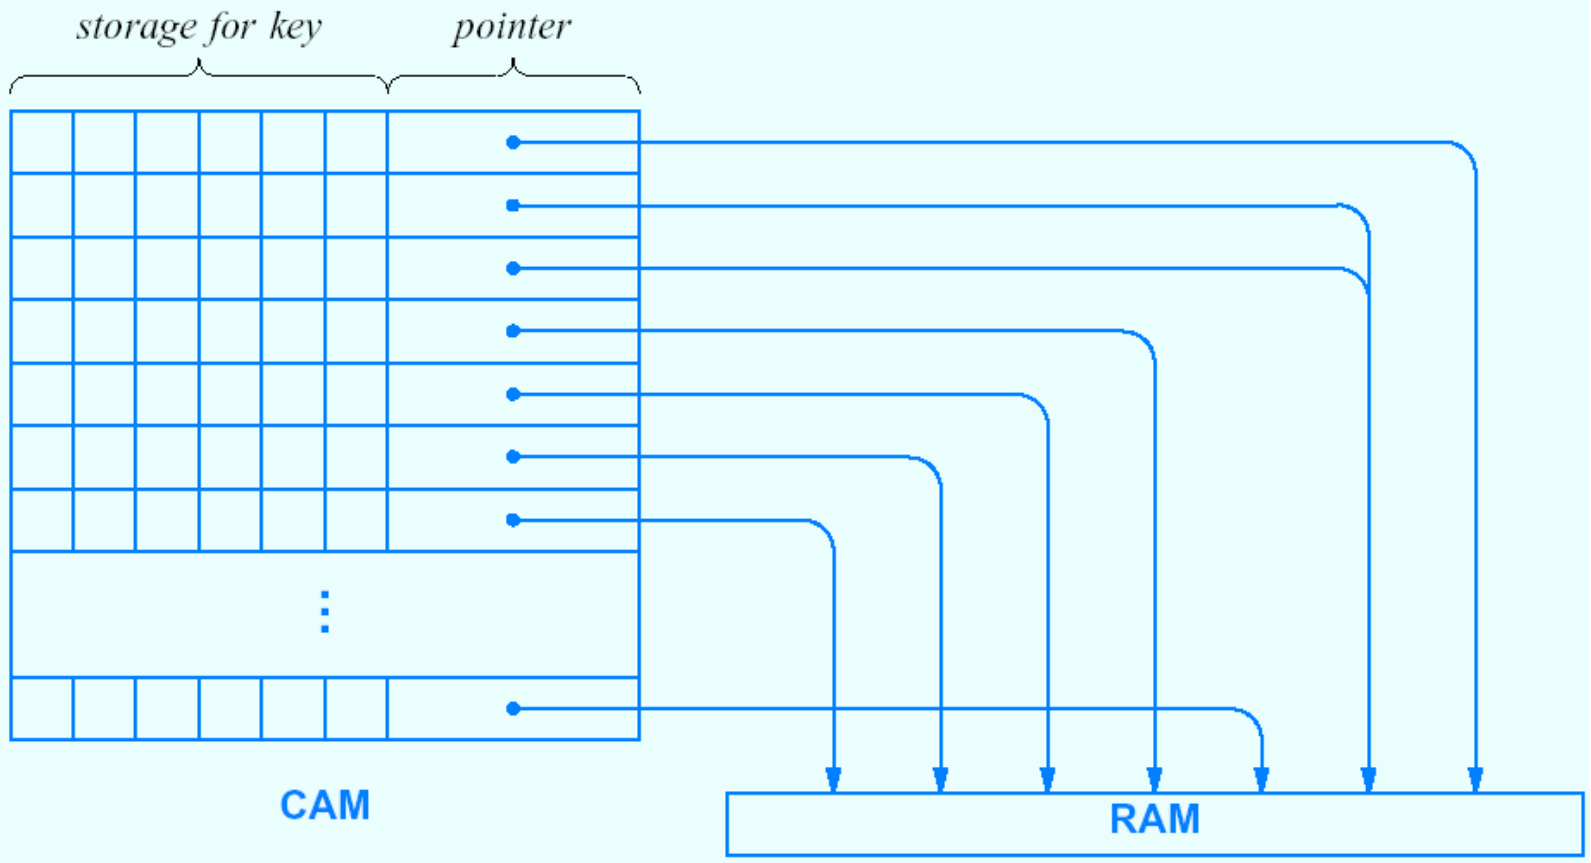

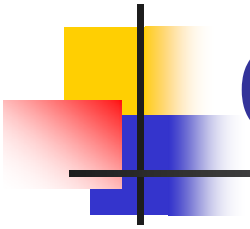

Outline

- Recap
- Network processor architectures
- Summary and homework

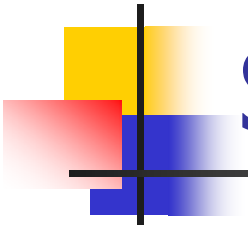

Summary

- Scalability key issue
- Primary subsystems affecting scale
  - Processor hierarchy
  - Memory hierarchy
- Many memory types available
  - SRAM
  - SDRAM
  - CAM
- T-CAM useful for classification

Homework (due on 04/05)

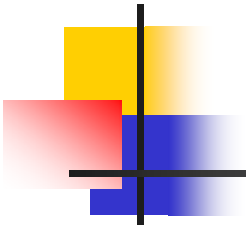

- 10.1. (a) Problem 8 of Chapter 13 (Page 200); (b) Problem 1 of Chapter 14 (Page 218).
- 10.2 Find out what are SDRM, QDR-SRAM, ZBT-SRAM, FCRAM, DDR-DRAM, RLDRAM.