University of Massachusetts Amherst

ECE697AA – Lecture 15

High-Performance Network Systems

Tilman Wolf
Department of Electrical and Computer Engineering
10/30/08

Network Systems

- Shift in objective of discussion
  - Focus so far: functionality
  - Focus now: performance
- Network systems
  - Devices that implement network protocols
- Layer of protocol processing determines functionality
  - Layer 1: repeater
  - Layer 2: bridge
  - Layer 3: router
  - Layer 4: transport-layer gateway
  - Layer 5: application-layer gateway

Protocol Implementations

- Protocols are standardized
  - Syntax: message/header format
  - Semantics: actions, order of messages
- Implementation of protocols is unspecified
  - Opportunity for different designs
- Examples:
  - Web browser
    - » Implements HTML rendering over HTTP
    - » Different implementations: Explorer, Firefox, etc.
  - IP router
    - » Home gateway
    - » Core router
  - Network Interface Card
    - » Standard desktop NIC
    - » Redundant, hot-swappable server NIC

Implementation Metrics

- What metrics should be considered when implementing a protocol on a network system?

Implementation Metrics

- What metrics should be considered when implementing a protocol on a network system?
  - Throughput
    - » Bandwidth (maximum or average)
    - » Delay (maximum or average)
  - Cost
    - » Initial cost
    - » Maintenance cost
  - Power consumption
    - » power consumption (peak or average)
    - » Standby power consumption
  - Quality
    - » Mean time to failure
  - Ease of use
    - » Configuration
  - Etc.

Aggregate Data Rate

- Definitions
  - From Douglas Comer: Network System Design using Network Processors
- Data rate:

  "The aggregate data rate is defined to be the sum of rates at which traffic enters or leaves a system. The maximum aggregate data rate of a system is important because it limits the type and number of network connections the system can handle"
- Packet rate:

  "For protocol processing tasks that have a fixed cost per packet, the number of packets processed is more important than the data aggregate rate"

Bandwidth Preservation

- Network systems often focus on bandwidth
  - Key performance metric in networks
  - Bandwidth bottlenecks cause problems
- Bandwidth metrics
  - Per-port peak bandwidth for max/min/avg packet sizes
  - Per-port peak packet rate for max/min/avg packet sizes
  - Aggregate bandwidth
- "Bandwidth preservation" drives system design
  - All components need to support target rate
  - Bottlenecks limit total system performance

Packet Size

- Packet rate depends on packet size:

Technology	Network Data Rate In Gbps	Packet Rate For Small Packets In Kpps	Packet Rate For Large Packets In Kpps
10Base-T	0.010	19.5	0.8
100Base-T	0.100	195.3	8.2
OC-3	0.156	303.8	12.8
OC-12	0.622	1,214.8	51.2
1000Base-T	1.000	1,953.1	82.3
OC-48	2.488	4,860.0	204.9
OC-192	9.953	19,440.0	819.6
OC-768	39.813	77,760.0	3,278.4

  - (Tables from Douglas Comer: Network System Design using Network Processors)

Time Per Packet

- Aggregate packet rate determines time per packet

Technology	Time per packet for small packets (in µs)	Time per packet for large packets (in µs)
10Base-T	51.20	1,214.40
100Base-T	5.12	121.44
OC-3	3.29	78.09
OC-12	0.82	19.52
1000Base-T	0.51	12.14
OC-48	0.21	4.88
OC-192	0.05	1.22
OC-768	0.01	0.31

- Packet processing requires in the order of 100s of instructions per packet

System Bottlenecks

- Where are potential bandwidth bottlenecks?
  - Example 1: Shopping portal on workstation computer
  - Example 2: Streaming video on cell phone
  - Example 3: FiOS service to home gateway
  - Example 4: Media Center via WiFi

System Bottlenecks

- Where are potential bandwidth bottlenecks?
  - Example 1: Shopping portal on workstation computer
    - » Processing power (web server, data base, etc.)
    - » Memory (if large number of connections)
  - Example 2: Streaming video on cell phone
    - » Data rate of network connection
    - » Processing power for video decoding
  - Example 3: FiOS service to home gateway
    - » Link bandwidth on external port of gateway
    - » NAT/firewall processing
  - Example 4: Media Center via WiFi
    - » Unreliability of wireless connections (limits average bandwidth)
    - » Possibly delay
- Network system design should avoid such problems

Scalability

- Designs evolve over time
  - Design process is expensive
  - Ability to easily extend to new requirements is important
  - Performance requirement increase really fast
    - » Moore's Law
  - Systems will eventually be used in a different context
- Scalability important design criterion
  - A system (design) is scalable if it can easily be extended in "size" and performance
    - » More ports
    - » Faster links

Scalability

- Example: simple software-based router

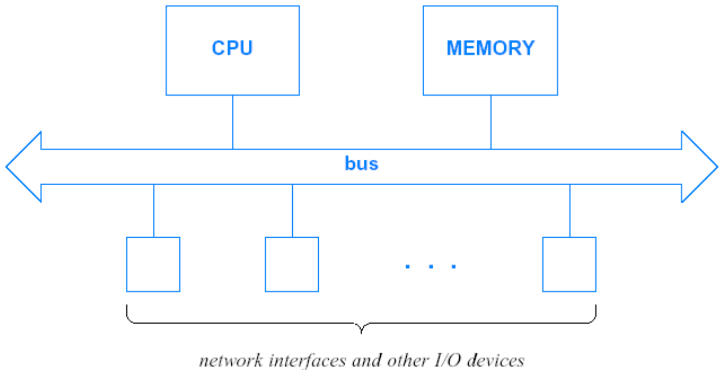

  - (Figure from Douglas Comer: Network System Design using Network Processors)
- How can we make this design more scalable?

Hardware-Based Routers

- Structure of hardware-based router:
  - Individual input and output ports with dedicated resources
  - Switch fabric as interconnect
  - Routing processor for control operations (e.g., routing)

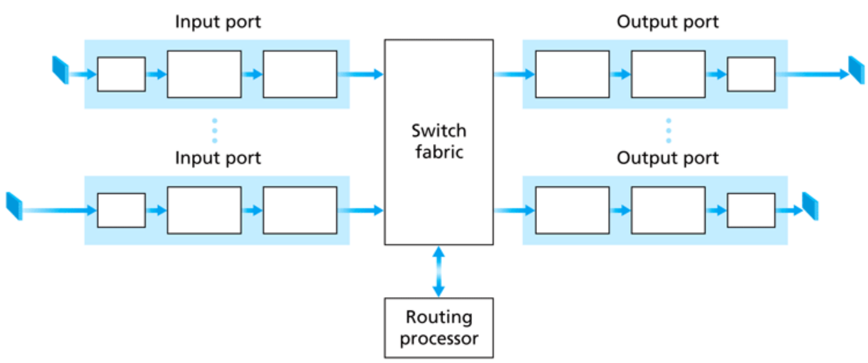

Upcoming Topics

- Queuing theory
  - Quantitative evaluation of statistical multiplexing
  - Basis for switching fabrics
- Switching fabrics
- Bridges
  - "Simplest" network devices
  - Spanning tree algorithm
- Routers
  - Prefix lookup
  - Packet classification
  - Scheduling
- Close relationship between system characteristics and algorithms

Assignments

- Read
  - SPARK Handout: Sections 2.1-2.4 from Leonard Kleinrock, *Queuing Systems - Volume I: Theory*, Wiley-Interscience, 1975.
- SPARK
  - Assessment quiz