

CSE 123b Graduate Networking

Spring 2003

Last Lecture: Potpourri

Stefan Savage

Final

- Check Studentlink to be sure, but I believe that its June 11th 3-6pm in this room.
- Closed book
- You can bring on 8.5x11 sheet of paper and you can do just about anything to that piece of paper you want (i.e. write on both sides, print on it, etc)
 - You can't tape your textbook to the sheet of paper though
 - Or your laptop
- Same style as midterm
 - FYI: You may be asked questions about project

June 3, 2003

CSE 222A – Lecture 16 – Potpourri

Thanks to Nick Weaver for some slides

Review session

- Monday, June 9th, 4:00-5:30pm in AP&M 4301
- I will also be available on the 10th from 2-3:30 (normal class hours)
 - Or by appointment

June 3, 2003

CSE 222A – Lecture 16 – Potpourri

Thanks to Nick Weaver for some slides

Stuff we'll look at today

- Skim
 - IPv6
 - IPSec
 - Active Networks
 - Sensor Networks
- A bit more detail:
 - Network measurement
 - Router design

June 3, 2003

CSE 222A – Lecture 16 – Potpourri

Thanks to Nick Weaver for some slides

IPSEC

- Security at lowest layer: IP
- Framework for security
 - Select different encryption algorithms, security protocols
 - Select security services (e.g., integrity, authentication, etc.)
 - Select granularity (e.g., connection, all end-point flows)
- Two parts
 - Authentication Header (AH)
 - » Access control, integrity, authentication, anti-replay
 - Encapsulating Security Payload (ESP)
 - » Confidentiality in addition to above
 - » Relies upon an encryption algorithm

June 3, 2003

CSE 222A – Lecture 16 – Potpourri

Thanks to Nick Weaver for some slides

IPSEC (2)

- Security Association (SA) binds AH and ESP
 - The "association" defines a set of security services between end-points (security gateways)
 - Simplex – only defined for a single direction
 - » Need two for both directions of a connection
- Negotiation
 - Internet Security Association and Key Management (ISAKMP)
 - » Procedures and formats to establish, negotiate, modify, delete security associations
 - » Again, a framework
 - Internet Key Exchange (IKE)
 - » Specific protocol for exchanging keys

June 3, 2003

CSE 222A – Lecture 16 – Potpourri

Thanks to Nick Weaver for some slides

IPv6

- **Originally motivated by crisis in IP address space**
 - 32 bits not enough if everything gets an IP address
 - Subnetting and CIDR helped alot, but not a long-term solution (mobile phones)
- **Solution: Increase the size of IP addresses**
 - Originally to 64, then to 128 bits
 - Requires changes to IP header
- **If we're going to change the header, might as well change other aspects of IP**
- **Rule of thumb: IPv6 = IPv4 + big addresses + IPSEC, Mobile IP + DHCP + ARP + improvements to these**

June 3, 2003

CSE 222A – Lecture 16 – Potpourri

Thanks to Nick Weaver for some slides

IPv6 (2)

- **IP addresses are 128 bits**
 - Unicast, multicast, local, etc.
 - 1234:ABCD:1234:ABCD:1234:ABCD:1234:ABCD
 - Registry, provider, subscriber, subnetwork, interface
 - Interface ID used as lowest 6 bytes
- **Simplified headers**
 - 40 bytes (20 for IPv4)
 - Use header extensions for options
- **Autoconfiguration**
 - Multiple addresses per host (link-local address)
- **FlowLabel, priority, security, mobility**

June 3, 2003

CSE 222A – Lecture 16 – Potpourri

Thanks to Nick Weaver for some slides

IPv6 and IPv4 Interoperability

- **Dual-stack operation**
 - IPv6 nodes that support both IPv4 and IPv6
- **Tunneling**
 - IPv6 packets encapsulated within IPv4 packets
 - End-points speak IPv6, but use IPv4 packet to use standard Internet routing
 - Easy mapping if IPv6 embeds IPv4 address, otherwise need to configure a table

June 3, 2003

CSE 222A – Lecture 16 – Potpourri

Thanks to Nick Weaver for some slides

Active Networks

- **Problem: How can we change the network without replacing the network?**
 - Routing algorithms (multicast), queueing disciplines (WFQ, RED), mobility, measurement, etc.
- **Approach: Active Networks**
 - Make routers **programmable**
 - Example:
 - » Packets carry programs or pointers to programs\
 - » Programs executed on arrival of packets
 - Need platform, execution environment, resource control, security, storage, etc. – essentially, a whole new OS
 - Many different incarnations, some small testbeds
 - Controversial (security, management, overhead)
 - A point in the design space of non-“client/server” models

June 3, 2003

CSE 222A – Lecture 16 – Potpourri

Thanks to Nick Weaver for some slides

Sensor Networks

- **Scenario**
 - Hundreds or thousands of low-powered wireless sensors distributed haphazardly over an area
 - Applications to environmental monitoring, object tracking, etc.
- **Research problems**
 - How to communicate information through network efficiently and with lower overall power (**application-specific** multihop)
 - How to locate nodes and provide time sync among them
 - How to write an application for such a network

June 3, 2003

CSE 222A – Lecture 16 – Potpourri

Thanks to Nick Weaver for some slides

Network Measurement

- **The Internet is an artifact of sufficient complexity that we don't understand it**
 - Difficult to measure many things directly -> inference
 - Too big to measure all of it
 - No such thing as “typical” (sampling difficult)
 - Constantly changing
- **A whole community focused on these problems**
 - Packet characteristics: application usage; dynamic behavior of existing protocols;
 - Path characteristics: delay, loss rate, queuing, bandwidth [end-to-end vs hop-by-hop]
 - Network topology: connections between routers, between ISPs; graph analysis of Internet topologies
- **Big challenge: no support for measurement**

June 3, 2003

CSE 222A – Lecture 16 – Potpourri

Thanks to Nick Weaver for some slides

For example:

- How to measure one-way **network path measurements (e.g. packet loss rate)**?
 - Critical for network performance analysis
 - Requires measurement from **both** endpoints
- **Remote hosts are indifferent to problem**
 - **Unrealistic** to deploy new software
 - Remote hosts may block measurements

June 3, 2003

CSE 222A – Lecture 16 – Polpourri

Thanks to Nick Weaver for some slides

Key idea

- **Exploit services needed by remote host**
 - Web, E-mail, news, file transfer, etc.
 - All use *Transmission Control Protocol (TCP)*
 - » Reliable, in-order, data transfer protocol
- **Exploit standard protocol behavior**
 - TCP has rich specification
 - Can be leveraged to perform measurements

June 3, 2003

CSE 222A – Lecture 16 – Polpourri

Thanks to Nick Weaver for some slides

Simple example: Deducing one-way loss rate

- Send TCP data packets to remote host
- What we know
 - Number of data packets sent
 - Number of acknowledgments received
- What we need to know
 - How many data packets were received?
 - » Extract from ACKs; TCP is a reliable protocol
 - How many acknowledgments were sent?
 - » Arrange that **one** ACK sent for each data packet


June 3, 2003

CSE 222A – Lecture 16 – Polpourri

Thanks to Nick Weaver for some slides

How TCP reveals packet loss


- Data packets ordered by sequence #'s
- ACK packets specify next seq# expected



CSE 222A – Lecture 16 – Polpourri

Thanks to Nick Weaver for some slides

Loss deduction example



$$\text{Loss}_{data} = 1 - (\text{dataReceived}/\text{dataSent}) = 33\%$$

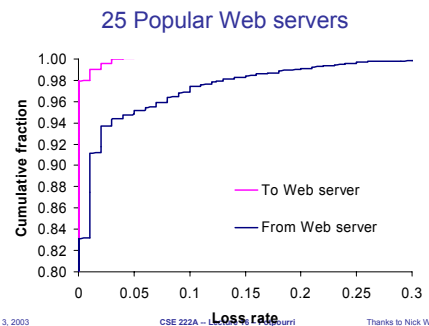
$$\text{Loss}_{ack} = 1 - (\text{ackReceived}/\text{ackSent}) = 50\%$$

June 3, 2003

CSE 222A – Lecture 16 – Polpourri

Thanks to Nick Weaver for some slides

Experimental finding: Packet loss is highly asymmetric



June 3, 2003

CSE 222A – Lecture 16 – Polpourri

Thanks to Nick Weaver for some slides

What's in a router? (how to make a very fast router?)

- **Physical components**
 - One or more **input interfaces** that receive packets
 - One or more **output interfaces** that transmit packets
 - A chassis (box + power) to hold it all
- **Functions**
 - **Forward** packets
 - **Drop** packets (congestion, security, QoS)
 - **Delay** packets (QoS)
 - **Transform** packets? (Encapsulation, Tunneling)

June 3, 2003

CSE 222A – Lecture 16 – Potpourri

Thanks to Nick Weaver for some slides

What a router does: the normal case


- **Receive incoming packet** from link input interface
- **Lookup packet destination** in forwarding table
 - (destination, output port(s))
- **Validate checksum, decrement ttl, update checksum**
- **Buffer packet** in input queue
- **Send packet to output interface** (interfaces? Mcast)
- **Buffer packet** in output queue
- **Send packet to output interface link**

June 3, 2003

CSE 222A – Lecture 16 – Potpourri

Thanks to Nick Weaver for some slides

What a router looks like?



June 3, 2003

CSE 222A – Lecture 16 – Potpourri

Thanks to Nick Weaver for some slides


What a router looks like (2)



CSE 222A – Lecture 16 – Potpourri

Thanks to Nick Weaver for some slides

What a router looks like (3)



June 3, 2003

CSE 222A – Lecture 16 – Potpourri

Thanks to Nick Weaver for some slides

High-performance routers

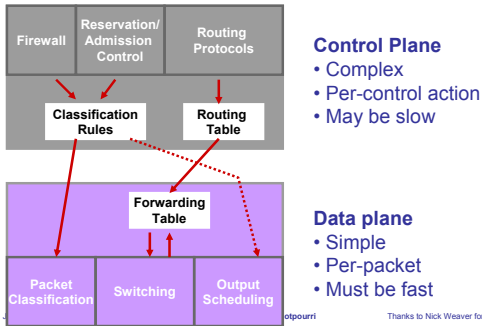
- **Geared to core and distribution service needs**
 - Requirements: high speed & high density
- **Why do we care?**
 - Moore's Law slower than link speed growth (and BW demand)
 - » OC48c (2.5Gbps), now, 128ns/packet
 - » OC192c (10Gbps), in deployment, 33ns/packet
 - » OC768c (40Gbps), 2002-3, 8ns/packet
 - Need high density/low power to manage POP complexity
 - » \$20-100k & 2-400W per port, 50% ports frequently for internal connectivity
 - » DWDM can help with the former, but requires more interfaces

June 3, 2003

CSE 222A – Lecture 16 – Potpourri

Thanks to Nick Weaver for some slides

Functional architecture




- Control Plane**
- Complex
 - Per-control action
 - May be slow
- Data plane**
- Simple
 - Per-packet
 - Must be fast

Packet classification

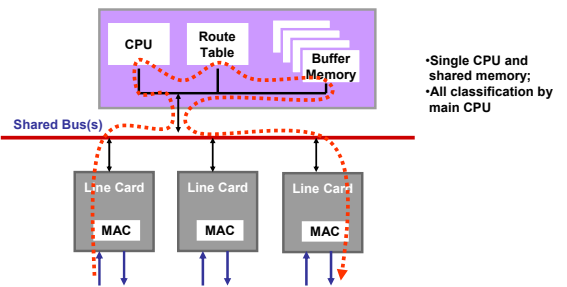
- **Forwarding**
 - Longest prefix match of destination against forwarding table
 - Returns (output port, Next-hop MAC header) tuple
 - Key issue: forwarding table growth
 - George will talk about this next time
- **QoS tagging**
 - Certain traffic tagged with higher priority
 - Per flow (src ip, src port, dst ip, dst port), pre source or dest prefix, per protocol (Napster, etc...)
- **Firewall rules**
 - Block access to TCP packets with dst port != 80

Interconnect architecture

- Input & output connected via switch fabric
- Kinds of switch fabric
 - Bus
 - Crossbar
 - Shared Memory
- How to deal with transient contention?
 - Input queuing
 - Output queuing
 - Combination

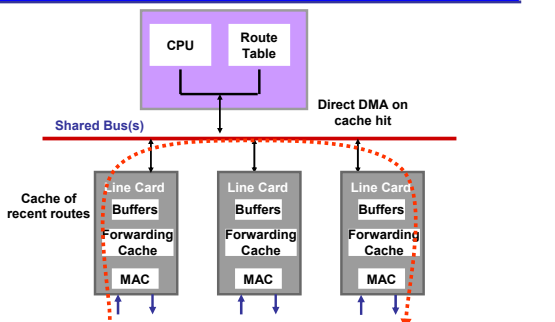


First Generation Routers

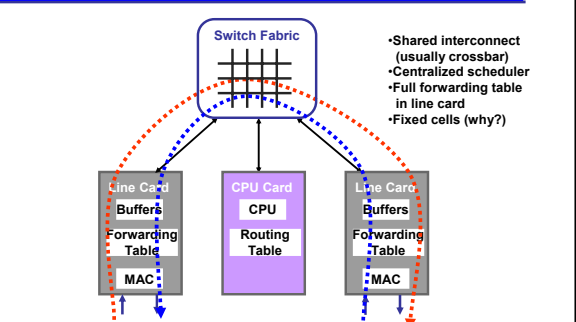


- Single CPU and shared memory;
- All classification by main CPU

Second Generation Routers



Third Generation Routers




- Shared interconnect (usually crossbar)
- Centralized scheduler
- Full forwarding table in line card
- Fixed cells (why?)

Output queuing

- Output interfaces buffer packets

- Pro
 - Simple algorithms
 - Single congestion point
- Con
 - N inputs may send to the same output
 - Requires speedup of N



June 3, 2003


CSE 222A – Lecture 16 – Polpourri

Thanks to Nick Weaver for some slides

Input queuing

- Input interfaces buffer packets

- Pro
 - Single congestion point
 - Simple to design algorithms
- Con
 - Must implement flow control
 - Low utilization due to Head-of-Line (HoL) Blocking
 - » Util limited to $2 - 2^{\wedge} .5 = 58\%$




June 3, 2003

CSE 222A – Lecture 16 – Polpourri

Thanks to Nick Weaver for some slides

Head-of-Line Blocking




June 3, 2003

CSE 222A – Lecture 16 – Polpourri

Thanks to Nick Weaver for some slides

Virtual Output Queues

(courtesy Nick McKeown)




some slides

IQ + Virtual Output Queuing

- Input interfaces buffer packets in per-output virtual queues

- Pro
 - Solves blocking problem
- Con
 - More resources per port
 - Complex arbiter
 - Still limited by input/output contention (scheduler).
 - RR: $1 = 1/e = 63\%$



June 3, 2003

CSE 222A – Lecture 16 – Polpourri

Thanks to Nick Weaver for some slides

Switch scheduling

- Problem
 - Match inputs and outputs
 - Resolve contentions
 - No output packet drops
 - Maximize throughput
 - Do it in constant time...
- Many algorithms for uniform traffic assumption
 - E.g. TDM, Maximum size bipartite match
 - Approximate answers (e.g. iSLIP, submaximal match)
- Recent result (Dai et al, 2000)
 - Maximal size matching + speedup of two guarantees 100% utilization for most traffic assumptions


June 3, 2003

CSE 222A – Lecture 16 – Polpourri

Thanks to Nick Weaver for some slides

Modern router

- **IQ + VoQ + OQ**
 - Speedup of 2
 - Central scheduler
 - Fixed-sized internal cells
- **Pro**
 - Can achieve utilization of ~ 1
 - Can scale to multiple Tb/s
- **Con**
 - Multiple congestion points
 - Complexity



June 3, 2003

CSE 222A – Lecture 16 – Potpourri

Thanks to Nick Weaver for some slides

Typical function breakdown

- **Input interface**
 - Forwarding
 - Virtual output queuing
- **Switch**
 - Scheduling input interface requests to output interfaces
 - Multicast scheduling
- **Output interface**
 - Queue packets for transmission
 - Classification
 - Buffer management (which pkt to drop)
 - Scheduling (which pkt to send from buffer)

June 3, 2003

CSE 222A – Lecture 16 – Potpourri

Thanks to Nick Weaver for some slides

Next bottlenecks

- **Buffering at high speed**
 - SRAM density too low for $BW \cdot D$ of 40Gbps link
 - DRAM too slow
 - SRAM memory management as cache for DRAM
- **Scheduler and arbiter overhead**
 - Limits size of switch and link BW
 - Two-state switch (Chang et al, 2000); no scheduler
- **High density (100's-1000's of line cards)**
 - Physical distance to support density; electrical links degrade
 - Optical links; optical cross connect (MEMs, tunable lasers)
- **Time to market, Power/Heat**

June 3, 2003

CSE 222A – Lecture 16 – Potpourri

Thanks to Nick Weaver for some slides

Next time...

- I will be in Princeton NJ (not my choice)
- **Geoff Voelker will give a lecture on distributed web caching (don't skip this, there may be a related question on the final)**
 - I will arrange to have his slides put on the Web
- I will see you folks at the review session or the final
- Good luck, its been a pleasure.

June 3, 2003

CSE 222A – Lecture 16 – Potpourri

Thanks to Nick Weaver for some slides