

CSE/EE 461 – Lecture 21

Network Support for QOS

David Wetherall
djw@cs.washington.edu

Last Time ...

- Congestion Avoidance
- Focus
 - How to we avoid congestion?
- Topics
 - Random Early Detection (RED) gateways
 - Explicit Congestion Notification (ECN)

Application
Presentation
Session
Transport
Network
Data Link
Physical

This Lecture


- Network support for QOS
- Focus
 - What network mechanisms provide which kinds of quality assurances?
- Topics
 - Scheduling and Buffer management
 - Fair Queuing
 - Intserv
 - Diffserv

Application
Presentation
Session
Transport
Network
Data Link
Physical

djw // CSE/EE 461, Autumn 2001

L21.3


Roadmap – Various Mechanisms

Simple to build, Weak assurances  Complex to build, Strong assurances	FIFO with Drop Tail	Classic Best Effort
	FIFO with RED	Congestion Avoidance
	Weighted Fair Queuing	Per Flow Fairness
	Differentiated Services	Aggregate Guarantees
	Integrated Services	Per Flow Guarantees

djw // CSE/EE 461, Autumn 2001

L21.4

What's in a Router?

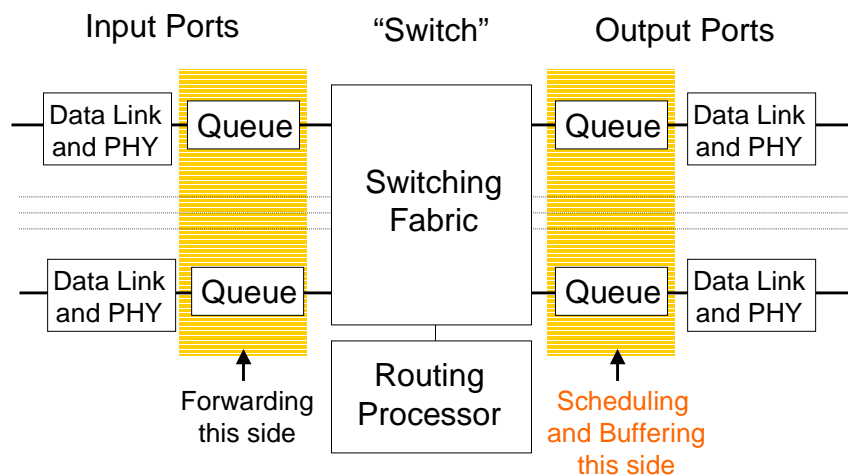


- By convention, draw input ports on left, output on right. (But in reality a single physical port handles both directions.)

djw // CSE/EE 461, Autumn 2001

L21.5

Model of a Router



djw // CSE/EE 461, Autumn 2001

L21.6

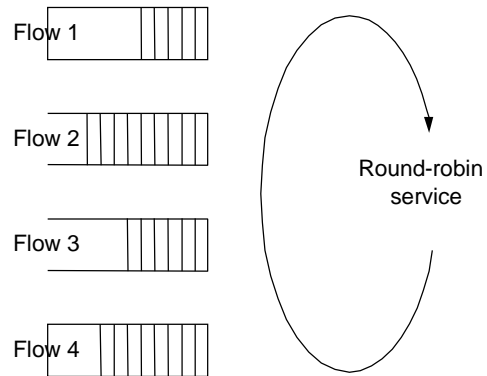
Scheduling and Buffer Management

- Two different functions implemented at the queue
- A scheduling discipline
 - This is the order in which we send queued packets
 - Examples: FIFO or priority-based
- A buffer management policy
 - This decides which packets get dropped or queued
 - Examples: Drop tail or random drop

Fair Queuing (FQ)

- FIFO is not guaranteed (or likely) to be fair
 - Flows jostle each other and hosts must play by the rules
 - Routers don't discriminate traffic from different sources
- Fair Queuing is an alternative scheduling algorithm
 - Maintain one queue per traffic source (flow) and send packets from each queue in turn
 - Actually, not quite, since packets are different sizes
 - Provides each flow with its "fair share" of the bandwidth

Fair Queuing




djw // CSE/EE 461, Autumn 2001

L21.9

Fair Queuing

- Want to share bandwidth
 - At the “bit” level, but in reality must send whole packets
- Approximate with finish times for each packet
 - finish (F) = arrive + length*rate; rate depends on # of flows
 - Send in order of finish times, except don't preempt (stop) transmission if a new packet arrives that should go first



- More generally, assign weights to queues (Weighted FQ, WFQ)

djw // CSE/EE 461, Autumn 2001

L21.10


Supporting QOS Guarantees

- Flowspecs. Formulate application needs
 - Need descriptor, e.g. token bucket, to ask for guarantee
- Admission Control. Decide whether to support a new guarantee
 - Network must be able to control load to provide guarantees
- Signaling. Reserve network resources at routers
 - Analogous to connection setup/teardown, but at routers
- Packet Scheduling. Use different scheduling and drop mechanisms to implement the guarantees
 - e.g., set up a new queue and weight with WFQ at routers

IETF Integrated Services

- Fine-grained (per flow) guarantees
 - Guaranteed service (bandwidth and bounded delay)
 - Controlled load (bandwidth but variable delay)
- RSVP used to reserve resources at routers
 - Receiver-based signaling that handles failures
- WFQ used to implement guarantees
 - Router classifies packets into a flow as they arrive
 - Packets are scheduled using the flow's resources

Resource Reservation Protocol (RSVP)



djw // CSE/EE 461, Autumn 2001

L21.13

RSVP Issues

- RSVP is receiver-based to support multicast apps
- Only want to reserve resources at a router if they are sufficient along the entire path
- What if there are link failures and the route changes?
- What if there are sender/receiver failures?

djw // CSE/EE 461, Autumn 2001

L21.14

IETF Differentiated Services


- A more coarse-grained approach to QOS
 - Packets are marked as belonging to a small set of services, e.g, premium or best-effort, using the TOS bits in the IP header
- This marking is policed at administrative boundaries
 - Your ISP marks 10Mbps (say) of your traffic as premium depending on your service level agreement (SLAs)
 - SLAs change infrequently; much less dynamic than Intserv
- Routers understand only the different service classes
 - Might separate classes with WFQ, but not separate flows

djw // CSE/EE 461, Autumn 2001

L21.15

Two-Tiered Architecture

Mark at Edge routers
(per flow state,
complex)



djw // CSE/EE 461, Autumn 2001

L21.16

Key Concepts

- Different scheduling and drop mechanisms can be used to support different QOS assurances
- Weighted Fair Queuing (WFQ) provides proportional fairness between different flows
- Integrated Services provides per-flow guarantees
 - Need admission control to make any absolute guarantees
- Differentiated Services provides coarse guarantees
 - But potentially simpler to implement