

CSE/EE 461 – Lecture 16

Bandwidth Allocation

Janet Davis
jlnd@cs.washington.edu
February 11, 2004
Reading: Peterson 6.1 – 6.2

Midterm Feedback:

What's working well

- Fishnet projects
- Engaging lectures – thanks, I'm trying! ☺
 - Presenter/Tablet PC
 - Lecture notes
- End-of-class feedback

L4.2

Midterm Feedback:

Suggestions

- Homework
 - Too vague, too divorced from lecture, too time-consuming, too harshly graded
 - We'll try to grade HW2 more graciously
 - I'll try to do better at designing the next homework
 - I'll take this into account for final grades
 - Let's do more examples in quiz section
- Lecture is too fast, too much to write
 - I'm working on a more relaxed pace – Stop me!
- Jumping around text is confusing
 - Sorry, too late to avoid this! Any strategies?

L4.3

Last time...

- Medium Access Control
 - How can senders share access to a single network?
 - Token Ring networks
 - Taking turns is efficient but complex!
 - Wireless networks
 - Hidden and exposed terminals
 - Collision avoidance with RTS/CTS

L4.4

This time...

- Bandwidth Allocation
 - How can many senders share bandwidth across the Internet?
 - Issues: Congestion and Fairness
 - What's inside a router?

L4.5

HTTP



- How fast should the web server send packets?

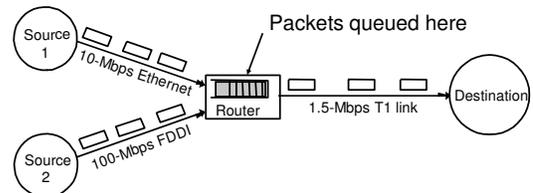
L4.6

Bandwidth Allocation

- Congestion:
 - sending too fast swamps the network
- Fairness:
 - Different users should get their fair share of the bandwidth
- Often treated together (e.g. TCP) but can be considered separately

L4.7

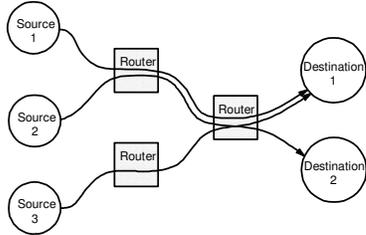
Congestion



- Buffer intended to absorb bursts when input rate > output
- But if sending rate is persistently > drain rate, queue builds
- Dropped packets represent wasted work

Chapter 6, Figure 1d/jw // CSE/EE 461, Autumn 2001

Fairness



- Each *flow* from a source to a destination should get an equal share of the *bottleneck* link

Chapter 6, Figure 2d/jw // CSE/EE 461, Autumn 2001

Bandwidth Allocation Taxonomy

- Router-centric vs Host-centric
 - Whose responsibility is it to learn and act on network conditions?
- Reservation-based vs Feedback-based
 - Are allocations guaranteed at the beginning of the flow, or adjusted based on feedback?
- Rate-based vs Window-based
 - Is allocation in terms of sending rate or buffer space in routers?

L4.10

Design Choices

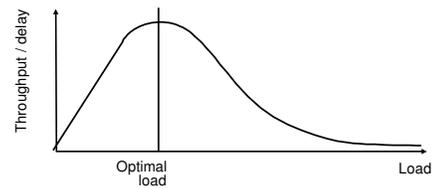
- Two positions: What are advantages and disadvantages of each?
 - Router-centric, reservation-based, rate-based

 - Host-centric, feedback-based, window-based

L4.11

Evaluating Congestion Management

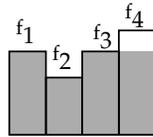
- How efficient is the bandwidth allocation?



Chapter 6, Figure 3d/jw // CSE/EE 461, Autumn 2001

Evaluating Fairness

- How do we compute the fairness of an allocation?
 - If all flows have an equal share on a link it's "fair"
 - e.g., min-max fairness



- But how unfair are unequal allocations?

L4.13

Jain's Fairness Index

- Jain's fairness index:
 - For n flows each receiving a fraction f_i of the bandwidth
 - Always between 0 and 1
 - What is fairness if all flows get the same bandwidth?
 - What if only k out of n flows get bandwidth?

$$fairness = \frac{(\sum f_i)^2}{n \sum f_i^2}$$

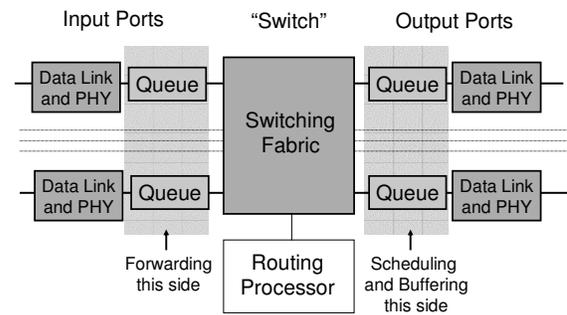
L4.14

What's in a Router?



L4.15

Model of a Router



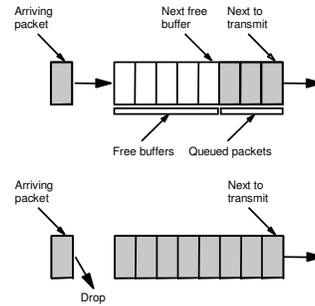
L4.16

Scheduling and Buffer Management

- Two different functions implemented at the output queue
 - Order in which to send packets
 - e.g.,
- A scheduling discipline
 - Order in which to send packets
 - e.g.,
- A buffer management policy
 - Which packets get dropped
 - e.g.,

L4.17

FIFO with Tail Drop



L4.18

Key Concepts

- Issues for bandwidth allocation: congestion and fairness
- Congestion occurs when buffers inside the network fill with excess traffic.
- Fairness means that competing traffic flows gain a “fair share” of the available bandwidth.
- Packets are queued at router outputs.

L4.19

Next time...

- Host-based congestion control in TCP

L4.20