

Internet Architecture II

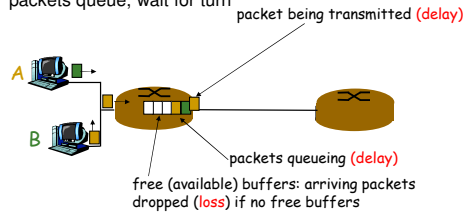
EECS 122: Lecture 3

Department of Electrical Engineering and Computer Sciences
University of California
Berkeley

How do loss and delay occur?

packets *queue* in router buffers

- packet arrival rate to link exceeds output link capacity
- this often happens because arrivals are random and bursty
- packets queue, wait for turn



January 19, 2006

EECS122 Lecture 2 (AKP)

4

Today's Outline

- Administrivia
- Last lecture we focused on protocols and layering
- Today, let's begin talking about performance
 - The internet as a network of Queues
 - The components of delay
 - Little's Law
 - Throughput
 - Measuring performance

January 19, 2006

EECS122 Lecture 2 (AKP)

2

Packet loss

- queue (aka buffer) preceding link in buffer has finite capacity
- when packet arrives to full queue, packet is dropped (aka lost)
- lost packet may be retransmitted by previous node, by source end system, or not retransmitted at all

January 19, 2006

EECS122 Lecture 2 (AKP)

5

Why care about performance

- The internet is a datagram network with only "best effort" service
- If "best effort" is lousy the network will be useless
 - Remember the World Wide Wait?
- Some applications such as video are quite performance sensitive
- Can't understand the behavior of a network without some understanding of how it performs


January 19, 2006

EECS122 Lecture 2 (AKP)

3

One packet, One Link

Packet Size: P bits
Bandwidth: R bps
Propagation delay: T sec



At what time is the sending buffer empty? P/R : Packet Xmission Delay
does the first bit get delivered?: T : Propagation Delay
does the last bit get delivered?: $T+P/R$: Arrival time at B

January 19, 2006

EECS122 Lecture 2 (AKP)

6

Propagation Delay, T

Packet Size: P bits
Link Speed/Bandwidth: R bps
Link Length: L meters

- Given P, R and L can you determine/approximate T?
 - No!
- S: speed of the underlying medium
L: length of the link
 $T = L/S$

1/speed = 3.3 usec in free space
4 usec in copper
5 usec in fiber


January 19, 2006

EECS122 Lecture 2 (AKP)

7

Switching: Store and Forward

- A packet is **stored** (enqueued) before being **forwarded** (sent)



January 19, 2006

EECS122 Lecture 2 (AKP)


10

One Packet, One Link

Packet Size: P bits
Link Speed/Bandwidth: R bps
Propagation delay: T sec

At what time is the sending buffer empty? P/R : Packet Xmission Time
does the first bit get delivered?: T : Propagation Delay
does the last bit get delivered?: T+P/R: Arrival time at B

P/R << T:



P/R >> T:




January 19, 2006

EECS122 Lecture 2 (AKP)


8

Queueing

Packet Size: P bits
Link Speed/Bandwidth: R bps
Propagation delay: T sec



- The queue has Q bits when packet arrives → packet has to wait for the queue to drain before being transmitted




January 19, 2006

EECS122 Lecture 2 (AKP)


11

Timing Diagram

P = 1 Kbyte
R = 1 Gbps
100 Km, fiber =>
T = 500 usec
P/R = 8 usec



P = 1 Kbyte
R = 100 Mbps
1 Km, fiber =>
T = 5 usec
P/R = 80 usec




January 19, 2006

EECS122 Lecture 2 (AKP)

9

Store and Forward: Multiple Packet Example




January 19, 2006

EECS122 Lecture 2 (AKP)

12

Processing Delay



- Make sure packet has not been corrupted
- Forward to right output link queue
- Schedule the next packet to be transmitted from queue

January 19, 2006

EECS122 Lecture 2 (AKP)

13

Example

- Packets of average length P arrive at an average rate of α , where the arrival process follows a particularly nasty distribution
- The network consists of many hops over copper, fiber, satellite, undersea cable etc.
- There are many other contending sessions
- The average delay of a packet is D

Question, how many packets are in the network on average?

January 19, 2006

EECS122 Lecture 2 (AKP)

16

Components of Per Hop Delay


- **Propagation delay:** time it takes the signal to travel from source to destination. *Only random component*
- **Packet transmission time:** time it takes the sender to transmit all bits of the packet
- **Queuing delay:** time the packet needs to wait before being transmitted because the queue was not empty when it arrived
- **Processing Time:** time it takes a router/switch to process the packet header, manage memory, etc

January 19, 2006

EECS122 Lecture 2 (AKP)

14

Delays and Queues



January 19, 2006

EECS122 Lecture 2 (AKP)

17

Why Queuing Delay Confounds Performance Analysis

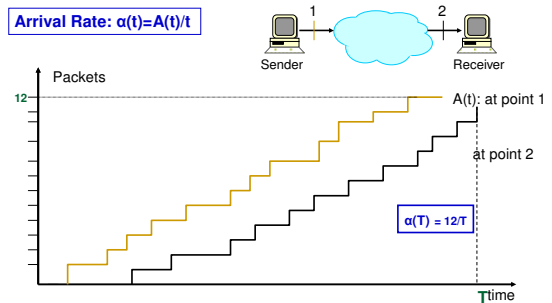
- Arrival processes are uncertain
- Modeled probabilistically
 - Analyzing a network of queues is very hard
 - Thousands of papers exist trying to calculate measures of performance assuming specific arrival processes
- Simulated based on real data
 - Traces collected from actual network traffic
 - Drawing general conclusions from traces is very hard
- Good news: enough is known to be able to get many useful insights

January 19, 2006

EECS122 Lecture 2 (AKP)

15

Delays and Queues



January 19, 2006

EECS122 Lecture 2 (AKP)

18

Delays and Queues

Arrival Rate: $\alpha = A(t)/t, t \rightarrow \infty$

Avg Packet Delay: $D = (D(1)+D(2)+\dots+D(n))/n$ as $n \rightarrow \infty$

Avg Occupancy = Area

January 19, 2006
EECS122 Lecture 2 (AKP)
19

Little's Law

- Does Little's Law hold if:
 - the packet sizes are variable?
 - Yes
 - if the packets are not received in the order they are sent?
 - Yes
 - If there is loss in the system?
 - Depends

January 19, 2006
EECS122 Lecture 2 (AKP)
22

Delays and Queues

Arrival Rate: $\alpha = A(t)/t, t \rightarrow \infty$

Avg Packet Delay: $D = (D(1)+D(2)+\dots+D(n))/n$ as $n \rightarrow \infty$

Avg Occupancy = Area

January 19, 2006
EECS122 Lecture 2 (AKP)
20

Application of Little's Law

- Suppose packets are of unequal length and the average length is 1000 bits
- The arrival rate of packets is 200 packets/sec
- $R = 1 \text{ Mb/s}$
- What fraction of the time is the link busy?
 - This is defined as Link Utilization

January 19, 2006
EECS122 Lecture 2 (AKP)
23

Little's Law

- Shaded Area up to time T is equal to both
 1. $D(1)+D(2)+\dots+D(A(T))$
 2. $\int_0^T Q(t) dt$
- Divide and multiple 1 by $A(T)$:
 1. $[D(1)+D(2)+\dots+D(A(T)) / A(T)] A(T)$
- Divide both (rewritten) 1 and 2 by T and take limits
 - $D \alpha = Q$

average occupancy = (average Delay) X (average arrival rate)

January 19, 2006
EECS122 Lecture 2 (AKP)
21


Application of Little's Law

- There can be 1 or 0 packets in the system
- α is still 200
- $D = P/R = 1000/1000000 = 1 \text{ millisecond}$
- By Little's Law utilization is $\alpha D = 0.2$
- Link is idle 80% of the time

January 19, 2006
EECS122 Lecture 2 (AKP)
24

Utilization for a “typical queue”

- $\alpha P/R \sim 0$: average queuing delay small
- $\alpha P/R \rightarrow 1$: delays become large
- $\alpha P/R > 1$: more “work” arriving than can be serviced, average delay infinite!
- (The text refers to $\alpha P/R$ as traffic intensity)



January 19, 2006

EECS122 Lecture 2 (AKP)

25

What about broadcast networks?

- There are n nodes sharing the broadcast channel
- When a node transmits a packet it collide with one from another node which has an overlapping transmission time
- Collisions can be detected by the sender
- What should a sender do upon detecting a collision?
- Throughput and delay depend heavily on scheme
 - Back off for a fixed time: very low throughput
 - Random back off time: better throughput
- More detail later...


January 19, 2006

EECS122 Lecture 2 (AKP)

28

Throughput

- The average number of bits successfully delivered per second.
 - Throughput of a lossless link: $R X$ (Utilization)



January 19, 2006

EECS122 Lecture 2 (AKP)

26

Practical Matters

- Many mathematical queuing models are very difficult to analyze
 - Analysis often makes simplifying assumptions on $A(t)$
 - Analysis is hairy and incomprehensible to most
 - Analysis only works for toy problems
 - Mostly useful to generate insights
- A realistic network is very difficult to simulate
 - How do you know your improvement to Bittorrent is actually better?
- Measurements
 - gather data from a real network
 - realistic, specific
 - Victim of that particular network's design
- Simulations: run a program that pretends to be a real network
 - e.g., NS network simulator, Nachos OS simulator
- Usually use combination of methods


January 19, 2006

EECS122 Lecture 2 (AKP)

29

Example: Windows Based Flow Control

- Connection:
 - Send W bits (window size)
 - Wait for ACKs
 - Repeat
- Assume the round-trip-time is RTT seconds
- Throughput = W/RTT bps
- Numerical example:
 - $W = 64$ Kbytes
 - $RTT = 200$ ms
 - Throughput = $W/T = 2.6$ Mbps



January 19, 2006

EECS122 Lecture 2 (AKP)

27

“Real” Internet delays and routes

- What do “real” Internet delay & loss look like?
- **Traceroute program**: provides delay measurement from source to router along end-end Internet path towards destination. For all i :
 - sends three packets that will reach router i on path towards destination
 - router i will return packets to sender
 - sender times interval between transmission and reply.



January 19, 2006

EECS122 Lecture 2 (AKP)

30

“Real” Internet delays and routes

traceroute: gaia.cs.umass.edu to www.eurecom.fr

Three delay measurements from
gaia.cs.umass.edu to cs-gw.umass.edu

```
1 cs-gw (128.119.240.254) 1 ms 1 ms 2 ms
2 border1-rt-fa5-1-0.gw.umass.edu (128.119.3.145) 1 ms 1 ms 2 ms
3 cft-vbns.gw.umass.edu (128.119.3.130) 6 ms 5 ms 5 ms
4 jn1-at1-0-0-19.wor.vbns.net (204.147.132.129) 16 ms 11 ms 13 ms
5 jn1-so7-0-0-0.wae.vbns.net (204.147.136.136) 21 ms 18 ms 18 ms
6 abilene-vbns.abilene.ucaid.edu (198.32.11.9) 22 ms 18 ms 22 ms
7 nycm-wash.abilene.ucaid.edu (198.32.8.46) 22 ms 22 ms 22 ms
8 62.40.103.253 (62.40.103.253) 104 ms 109 ms 106 ms
9 de2-1.de1.de.geant.net (62.40.96.129) 109 ms 102 ms 104 ms
10 de.fr1.fr.geant.net (62.40.96.50) 113 ms 121 ms 114 ms
11 renater-gw.fr1.fr.geant.net (62.40.103.54) 112 ms 114 ms 112 ms
12 nio-n2.cssi.renater.fr (193.51.206.13) 111 ms 114 ms 116 ms
13 nice.cssi.renater.fr (195.220.96.102) 123 ms 125 ms 124 ms
14 r3i2-nice.cssi.renater.fr (195.220.98.110) 126 ms 126 ms 124 ms
15 eurecom-valbonne.r3i2.ft.net (193.48.50.54) 135 ms 128 ms 133 ms
16 194.214.211.25 (194.214.211.25) 126 ms 128 ms 126 ms
17 ..
18 ..
19 fantasia.eurecom.fr (193.55.113.142) 132 ms 128 ms 136 ms
```

trans-oceanic link

means no response (probe lost, router not replying)

January 19, 2006

EECS122 Lecture 2 (AKP)

31

What we learned

- How to think of the internet as a network of queues
- The components of delay
- Little's Law
- Utilization
- Simple examples of throughput
 - Link, TCP connection, broadcast network
- Measuring delays

January 19, 2006

EECS122 Lecture 2 (AKP)

32