EE 122: Networks Performance & Modeling

Ion Stoica

TAs: Junda Liu, DK Moon, David Zats

http://inst.eecs.berkeley.edu/~ee122/fa09

(Materials with thanks to Vern Paxson, Jennifer Rexford, and colleagues at UC Berkeley)

Outline

- Motivations
- Timing diagrams
- Metrics
- Little’s Theorem
- Evaluation techniques

Motivations

- Understanding network behavior
- Improving protocols
- Verifying correctness of implementation
- Detecting faults
- Monitor service level agreements
- Choosing providers
- Billing

Outline

- Motivations
- ➢ Timing diagrams
- Metrics
- Little’s theorem
- Evaluation techniques

Timing Diagrams

- Sending one packet
- Queueing
- Switching
  - Store and forward
  - Cut-through

Definitions

- **Link bandwidth (capacity):** maximum rate (in bps) at which the sender can send data along the link
- **Propagation delay:** time it takes the signal to travel from source to destination
- **Packet transmission time:** time it takes the sender to transmit all bits of the packet
- **Queuing delay:** time the packet need 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

Sending One Packet

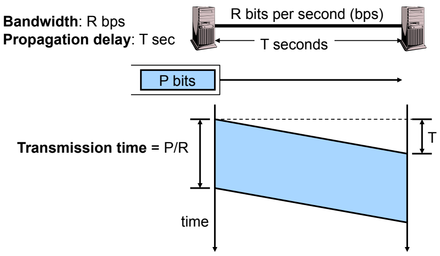

Propagation delay =T = Length/speed

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

7

Sending one Packet: Examples

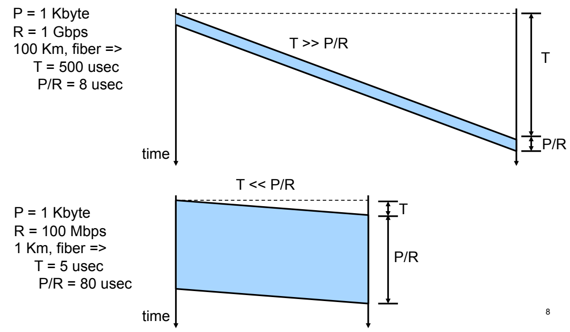

8

Queueing

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

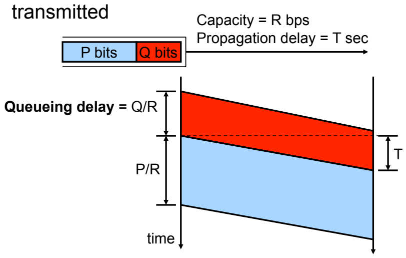

9

Queueing Example

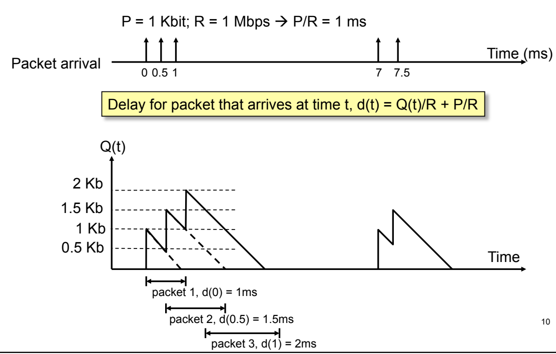

10

Switching: Store and Forward

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

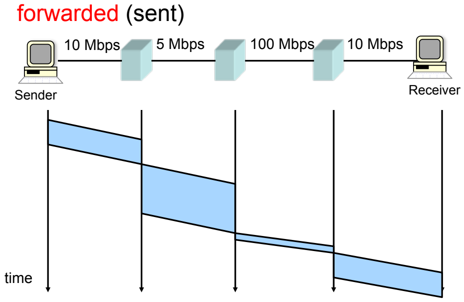

11

Store and Forward: Multiple Packet Example

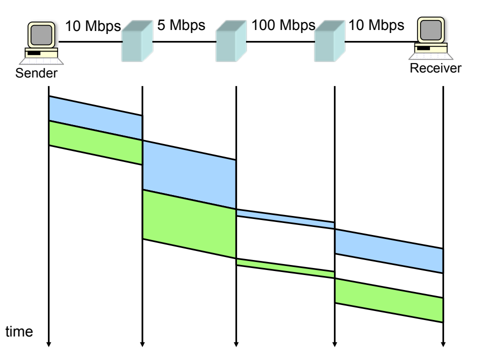

12

Switching: Cut-Through

- A packet starts being forwarded (sent) as soon as its header is received

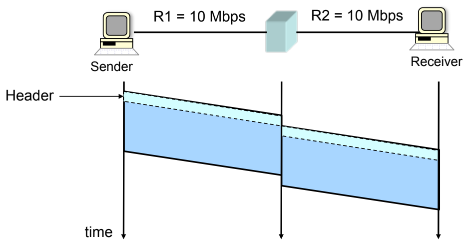

What happens if R2 > R1 ?

Outline

- Motivations
- Timing diagrams
- Metrics
  - Throughput
  - Delay
- Little's Theorem
- Evaluation techniques

Throughput

- Throughput of a connection or link = total number of bits successfully transmitted during some period [t, t + T) divided by T
- Link utilization = (throughput of the link)/(link rate)
- Bit rate units: 1Kbps = 10^3bps, 1Mbps = 10^6bps, 1 Gbps = 10^9bps [For memory: 1 Kbyte = 2^{10} bytes = 1024 bytes]
  - Some rates are expressed in packets per second (pps) → relevant for routers/switches where the bottleneck is the header processing

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/RTT = 64,000*8/0.2s = 2.6 Mbps

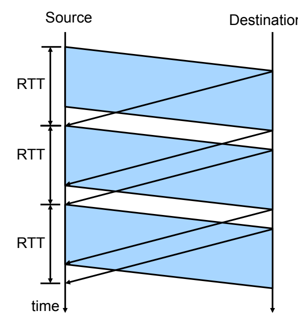

Throughput: Fluctuations

- Throughput may vary over time

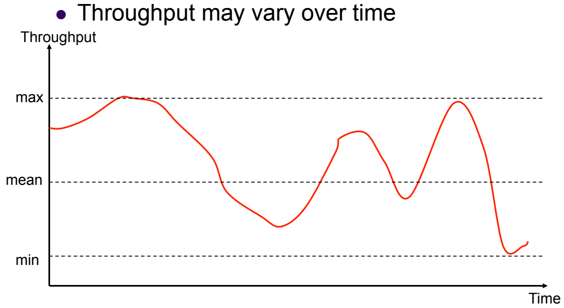

Delay Related Metrics

- Delay (Latency) of bit (packet, file) from A to B
  - The time required for bit (packet, file) to go from A to B
- Jitter
  - Variability in delay
- Round-Trip Time (RTT)
  - Two-way delay from sender to receiver and back
- Bandwidth-Delay product
  - Product of bandwidth and delay → "storage" capacity of network

Delay Illustration

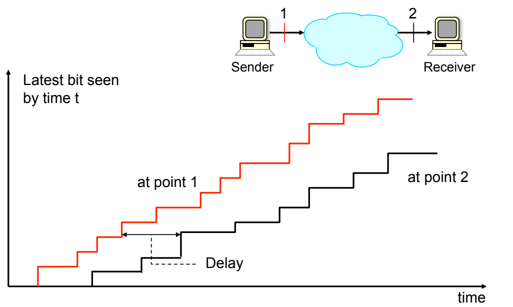

19

Outline

- Motivations
- Timing diagrams
- Metrics
- ➢ Little's Theorem
- Evaluation techniques

20

Little's Theorem

- Assume a system at which packets arrive at rate λ
- Let d be mean delay of packet, i.e., mean time a packet spends in the system
- Q: What is the mean (average) number of packets in the system (N) ?

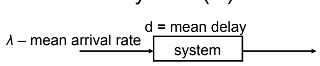

21

Example

- $\lambda = 1$
- $d = 5$

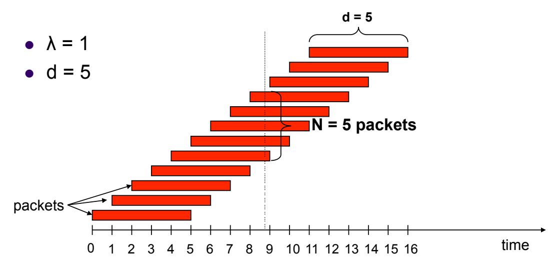

- A: $N = \lambda \times d$
  - E.g., $N = \lambda \times d = 5$

22

Little's Theorem: Proof Sketch

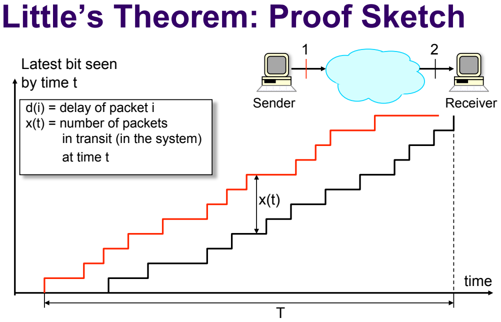

What is the system occupancy, i.e., average number of packets in transit between 1 and 2 ?

23

Little's Theorem: Proof Sketch

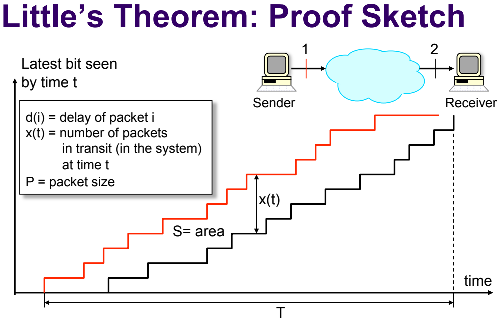

Average occupancy = S/T

24

Little's Theorem: Proof Sketch

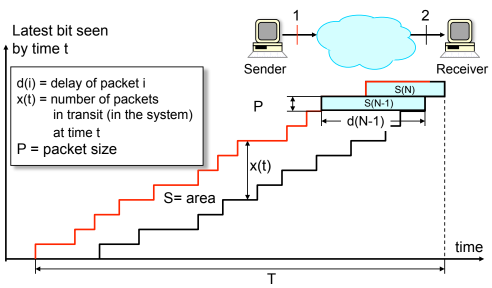

S = S(1) + S(2) + … + S(N) = P*(d(1) + d(2) + … + d(N))

25

Little's Theorem: Proof Sketch

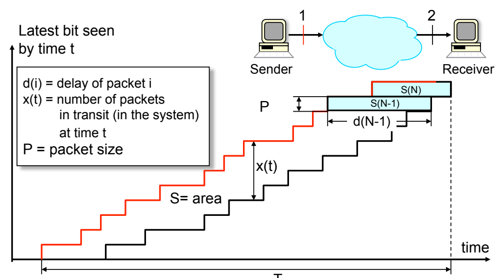

Average occupancy — S/T = (P*(d(1) + d(2) + … + d(N)))/T
= ((P*N)/T) * ((d(1) + d(2) + … + d(N))/N)

Average arrival rate — Average delay

26

Little's Theorem: Proof Sketch

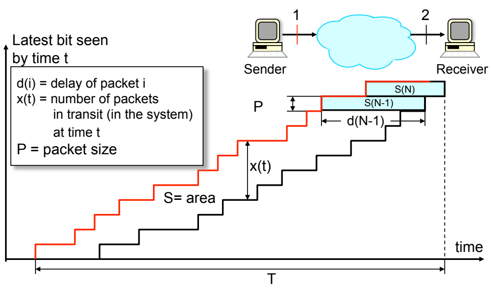

Average occupancy = (average arrival rate) x (average delay)

27

Outline

- Motivations
- Timing diagrams
- Metrics
- Little's Theorem
- ➢ Evaluation techniques

28

Evaluation Techniques

- Measurements
  - gather data from a real network
  - e.g., ping www.berkeley.edu
  - realistic, specific
- Simulations: run a program that pretends to be a real network
  - e.g., NS network simulator, Nachos OS simulator
- Models, analysis
  - write some equations from which we can derive conclusions
  - general, may not be realistic
- Usually use combination of methods

29

Simulation

- Model of traffic
- Model of routers, links
- Simulation:
  - Time driven:
    X(t) = state at time t
    X(t+1) = f(X(t), event at time t)
  - Event driven:
    E(n) = n-th event
    Y(n) = state after event n
    T(n) = time when even n occurs
    [Y(n+1), T(n+1)] = g(Y(n), T(n), E(n))
- Output analysis: estimates, confidence intervals

30

Simulation Example

- Use trivial time-driven simulation to illustrate statistical multiplexing
- Probabilistically generate the bandwidth of a flow, e.g.,
  - With probability 0.2, bandwidth is 6
  - With probability 0.8, bandwidth is 1
- Average bandwidth, avg=0.2*6 + 0.8*1 = 2
- peak/avg = 6/2 = 3

One Flow

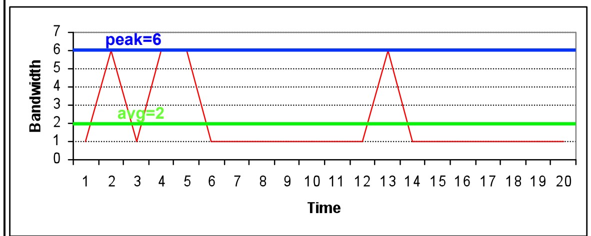

peak / avg = 3

Two Flows

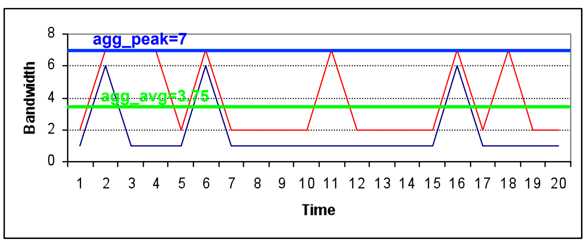

agg_peak / agg_avg = 7/3.75 = 1.86
(agg_avg = average of aggregate bandwidth)
(agg_peak=maximum value of aggregate bandwidth)

50 Flows

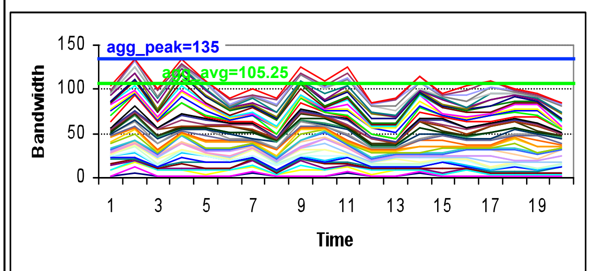

agg_peak / agg_avg = 7/3.75 = 135/105.25 = 1.28

Statistical Multiplexing

- As number of flows increases, agg_peak/agg_avg decreases
  - For 1000 flows, peak/avg = 2125/2009=1.057
- Q: What does this mean?
- A: Multiplexing a large enough number of flows "eliminates" burstiness
  - *Use average bandwidth to provision capacity, instead of peak bandwidth*
  - E.g., For 1000 flows
    - Average of aggregate bandwidth = 2,000
    - Sum of bandwidth peaks = 6,000

Evaluation: Putting Everything Together

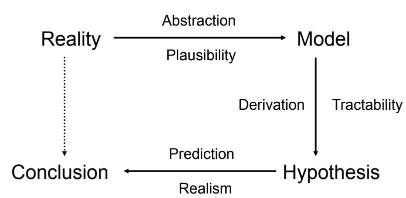

- Usually favor plausibility, tractability over realism
  - Better to have a few realistic conclusions than none (could not derive) or many conclusions that no one believes (not plausible)

Next Lecture

- Architecture, Layering, and the “End-to-End Principle”
- Read 1.4 & 1.5 of Kurose/Ross
- Pick up class computer account forms, if you haven’t done it already

- Project 1 (tiny world or warcrafts) out today
  - First part (client) due Oct 7 @ 11:59:59pm
  - Second part (server) due Oct 26 @ 11:59:59pm

37