Lecture 7

CS625: Advanced Computer Networks
Fall 2004

Tuesday, 12 August 2003

Bhaskaran Raman
CSE, IIT-Kanpur

http://www.cse.iitk.ac.in/users/braman/courses/cs625-fall2004/outline.html

Outline for Today

- BGP instability (from previous class)
- Packet switching, congestion, and queues
- Fair queuing
- *Scribe for today?*

Packet Switching and Congestion

- *Congestion:* when the network cannot carry datagrams fast enough
  - Buffer space in routers may get exhausted
  - Blocking can be implemented in circuit-switched systems
- *Congestion control* was considered to be a problem of avoiding buffer exhaustion
- [Nag87]: What happens with *infinite storage?*

Congestion and Infinite Storage

- Packets have a finite Time-to-Live (TTL)
  - Router decrements packet's TTL by amount of time spent by the packet in its buffer
  - Decrement at least by 1
- Under overload, *all* packets would be dropped

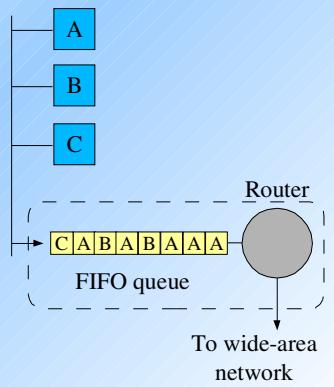

Congestion control/avoidance

- Transport protocol must react to congestion by reducing the sending rate
  - But, no motivation for a host to be *well-behaved*
- Possible solutions:
  - Co-operative, authoritarian (policing), market
  - Market: optimal strategy for each person is optimal for all
  - *Fair queuing*

Fair Queuing [Nag87]

- Separate queues for each source
- Round-robin processing of queues
- More aggressive source will suffer
- This does not prevent *malicious* users from eating up bandwidth

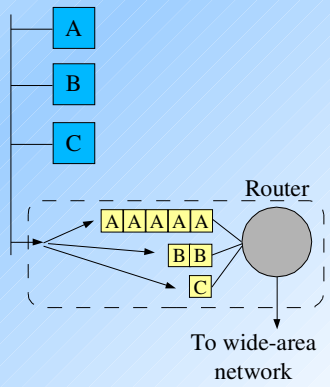

Congestion control and Fair Queuing [DKS89]

- Congestion control
  - Can be done at *source* or at *router*
  - Queuing is a router mechanism
- Queuing deals with:
  - Bandwidth allocation (*which* packets to transmit)
  - Promptness (*when* to transmit)
  - Buffer space (which packets to *discard*)

Notion of a *User*

- Four possibilities
  - Source
  - Destination
  - Source+Destination
  - Individual process
- Each has its draw-backs

Fair Queuing Details

- How to allocate buffer space *fairly*?
  - Drop packet, when necessary, from the longest queue
- How to deal with variable packet length?
  - Bit-wise Round-robin (BR) is ideal
  - Simulate it by estimating *finish time*
  - With and without packet preemption
  - These variants are equivalent asymptotically

Fair Queuing Details (continued)

- How to achieve promptness?
  - Built *credit* over a period of time
  - Flows which are inactive get served prompty
  - Useful for *telnet*

Fair Queuing: Other Considerations

- What happens when there are a network of gateways?
  - See simulations in paper [DKS89]
- Bottlenecks to scaling:
  - Per-packet overhead
    - Deficit-Round-Robin (DRR) to address this
  - Per-flow memory (state)
    - Core-Stateless-Fair-Queuing (CSFQ) to address this

Further topics this week

- TCP congestion control
- QoS: IntServ and DiffServ