Lecture 12

CS625: Advanced Computer Networks
Fall 2004

Tuesday, 26 August 2003

Bhaskaran Raman
CSE, IIT-Kanpur

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

Outline for Today

- Queuing: clarifications
- Integrated Services
- Announcements:
  - RED: what you need to know
  - Project web-page up – please check
  - I will give project feedback by this weekend
  - Short project status report due in 2 weeks (Friday, 5 Sep 2003)
- *Scribe for today?*

Queuing Clarifications

- Recall router queuing mechanisms:
  - First-Come-First-Served
  - Fair Queuing (FQ)
- FQ requires O(logN) time per packet
- Stochastic FQ (SFQ): use small number of queues, and hash flows onto queues
- Deficit-Round-Robin (DRR): carry over deficit to next round; O(1) scheduling
- Core-Stateless-Fair-Queuing (CSFQ): edge-routers insert label; no per-flow state in core-routers

Integrated Services [CSZ92]

- Data and Voice service in same network
- Architecture:
  - Service commitments
  - Scheduling algorithm at router
  - Service interface
  - Admission control
  - Signaling protocol

Classification of Applications

- Best-effort
- Real-time, with *play-back point*
- Properties:
  - Low delay desired
  - Absolute or statistical bound
  - Delay within play-back point is acceptable
  - Some packet loss can be tolerated
- Video can be *bursty* – take advantage of statistical multiplexing

Further Classification

- *Rigid* play-back versus *Adaptive*
- *Intolerant* to delay over play-back point versus *tolerant*
- Two classes:
  - Rigid and intolerant
  - Adaptive and tolerant

Service Commitments

- Guaranteed
- Predicted service
  - Network has predictable behaviour: past is guide to future
  - Network tries to minimize delay
- Best-effort

Traffic characterization: Token Bucket

- Token rate *r*, bucket size *b*
  - One token required for each bit transmission
  - Rate *r* represents the average flow rate
  - Bucket size *b* represents the max. token *credit*
- Define:
  - $n(0) = b$,
  - $n(i) = \min[b, n(i-1) - p(i) + r \times \{t(i) - t(i-1)\}]$
  - $n(i) >= 0$ implies flow conforms to (r, b)
- Parameter *b* represents "burstiness"
  - Can specify b(r) for any r, for a given flow

Scheduling for Guaranteed Service

- Weighted Fair Queuing (WFQ)
- Parekh and Gallager:
  - If, flow gets same clock-rate at all routers, and
  - If sum of rates of all flows at a router < capacity
  - Then, delay in network is bounded by b(r)/r for each flow
- Guaranteed service: delay does not depend on behaviour of other flows

Scheduling for Predicted Service

- Experiment: consider a single source bursty at some instant
  - WFQ: only that flow will see all excess delay
  - Some packets may miss their deadline
  - Earliest Deadline First (EDF) is more appropriate
- If (1) all flows have similar service desire, and (2) all packets have same deadline offset from arrival time, then EDF == FIFO
- FIFO good for *sharing*, WFQ for *isolation*

More on Scheduling

- Can implement priority classes
- Multi-hop ==> FIFO+
  - Keep track of *excess* delay at each router
- Unified scheduling:
  - One WFQ queue for each each guaranteed service flow
  - One for all predicted service flows, and best-effort
  - Multiple classes based on delay bounds

Also need...

- Admission control
- Signaling protocol
- Further in the week:
  - RSVP: Resource reSerVation Protocol
  - QoS: DiffServ – a simpler approach to multiple kinds of services