

## Fair Queueing: Demers, Keshav, Shenker [Demers89a]

CSci551: Computer Networks  
SP2006 Thursday Section  
John Heidemann

## Key ideas

- want to have smart routers to control congestion
- routers can enforce fair allocation of bandwidth
- why?
  - can prevent malicious users from hurting others
  - can keep bulk traffic sources from unnecessarily delaying non-bulk traffic

## Places to Implement Congestion Control

- applications (admission control)
  - ex: rate-adaptive applications (not admission control), telephone network (“all circuits are busy”), sometimes web sites show “busy signal” (503 unavailable?)
  - advantages: pre-emptive, maybe easier (just once per connection, not ongoing)
- source and destination (transport, end-to-end)
  - ex: TCP, DEChit
  - advantages: dumb (cheaper), faster? routers, easy to change end points—maybe more flexible
- routers (“gateways”)
  - via adaptive routing (ex: load-sensitive routing)
    - advantages: xxx
    - disads: xxx
  - via queueing policy (ex: NOT FCFC (FIFO, drop tail), FQ Fair Queueing)
    - advantages: rtr has more information, and info about all the flows (if network is shared), can enforce fairness, deal with malicious users

## Queueing Policies

- Many policies have been considered:
  - FIFO (“drop tail”) [and also drop head]
  - round robin (per flow)
  - weighted round robin
  - fair queueing [this one]
  - token bucket
  - virtual clock
  - class-based queueing (per class of traffic)
  - stochastic fair queueing (statistical)

## Alternatives

- how quickly do you provide feedback
- what kind of fairness do you provide
  - fair to whom: flows, src-dest pairs
    - be careful: if you have more flows, you get more bandwidth
    - although in principle you could change the definition of who
  - how fair: min-max fair, guarantee
- how efficient you are (router go idle?)
  - ex. circuit switched networks doing TMDA let the router go idle
  - but FQ wants to keep rtr busy
- how much state you must keep
  - information per flow
  - *per flow state* is a big concern in big routers
- how do you signal congestion
  - dropping vs. explicit feedback (DEChit, ECN)

## Fair Queueing

- fairness
  - to whom: *conversations* (now called *flows*: src-dest addr-port pairs)
  - quality: always (not just statistical)
  - definition: max-min
- state
  - $R(t)$ —cumulative number of rounds serviced from time 0 to time  $t$
  - $S_i$  and  $F_i$ —start and finish rounds for flow  $i$

## Max-Min Fairness

- definition:
  - no user gets more than request
  - no better allocation
  - “condition 2 is recursively true if we remove the minimal user and reduce the total resource accordingly”
- another view:
  - $\mu_i = \text{MIN}(\mu_{\text{fair}}, \rho_i)$
- considering allocation
  - find the minimal allocation
  - give everyone that much if undercapacity
  - if over capacity, divide what's left equally
  - if capacity left over, repeat with anyone left who wants more

## Examples of Min-Max Fairness

- offered load  $\Rightarrow$  accepted load (dropped)
- $1/3, 1/3, 1/3 \Rightarrow 1/3, 1/3, 1/3 (0, 0, 0)$
- $1, 1, 1 \Rightarrow 1/3, 1/3, 1/3 (2/3, 2/3, 2/3)$
- $1/3, 1/6, 1/6 \Rightarrow 1/3, 1/6, 1/6 (0, 0, 0)$
- $2/3, 1/6, 1/6 \Rightarrow 2/3, 1/6, 1/6 (0, 0, 0)$
- $1, 1/6, 1/6 \Rightarrow 2/3, 1/6, 1/6 (1/3, 0, 0)$
- $1, 1, 1/6 \Rightarrow 5/12, 5/12, 1/6 (xxx)$
- $1, 2/3, 1/6 \Rightarrow 5/12, 5/12, 1/6 (xxx)$
- what about bad users? always request 1, always get  $1/N$ 
  - why bother to under request? hopefully people will ask for only what they need, OR/AND you could do billing AND this FQ paper talks about giving time bonuses to people who don't use their full share


## Developing Fair Queueing

- start with *bit-wise fair queueing*
  - not realizable in practice, but allows strong definitions of fairness
- generalize to *packet-level fair queueing*
  - same arguments for correctness apply, but can now be realized

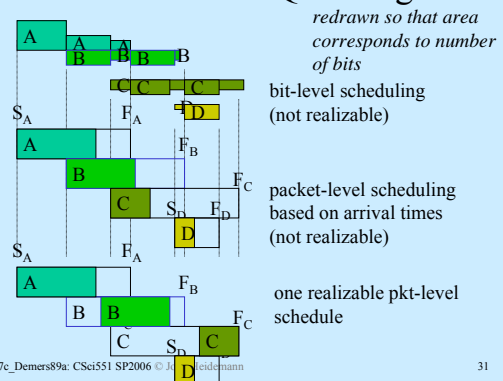
## Term Definitions

- $R(t)$ : cumulative number of rounds through  $t$ 
  - a round is sending one bit of each flow
- $dR(t)/dt$ : rate you service rounds
- $S_i, F_i$ : start and end times of a sender pkt  $i$
- to allow *promptness*
  - $B_i$ : bid (request for bonus if you haven't sent much)

## Bit-wise Fair Queueing



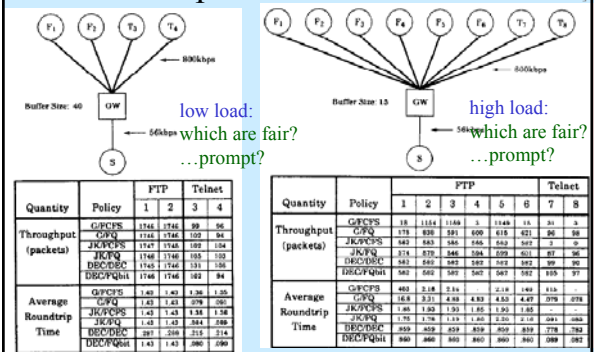
## Bit-wise Fair Queueing



## Modeling Fair Queueing

- Analytic modeling (section 2.3)
  - evaluate mean queueing delay
- Simulation modeling
  - compare a variety of “flow control” alternatives: window, TCP, DECbit

## Sample of Sim Results [Demers89a, scenarios 1 & 2]



## Limitations Fair Queueing

- problem of per-flow state
  - Internet backbone routers have many, many flows per second => can't keep per flow state
- recent approaches
  - police at *edges* of network (*core* is dumb but fast)
    - “core stateless” fair queueing [Stoica et al]
  - or can we keep *less than per-flow state*
    - maybe only state about “big” flows
    - active area of research

## Other questions/observations?

- XXX