

2/11/2010

Queueing: FA

Ideally,

Role of routers in congestion control: many, but here are a few.

- Actively monitor onset of congestion
- Provide early indication of congestion
- Isolate misbehaving flows so that end-host congestion control is effective. — or unresponsive

"Active Queue Management (AQM)"

FA → does this really well.

Others that do subsets of these → discuss in the next class.

How to isolate: ideal → max-min fairness; efficient use of resources → also buffer and delay allocation.

~~FA goal: allocate bandwidth, buffer space and delay~~


~~in a fair fashion across multiple end-to-end flows. It is ineffective at all three~~

What we have today →

- ① - FCFS: ~~interturns all these goals and messes up the outcomes above~~
- ② - Router backplane switching does this using starvation-free fair scheduling. Simple, efficient, but...
 - Too granular: focuses on input ports as opposed to flows. parking lot problem where flows that traverse multiple routers get exponentially lesser share.

Question: → How to build a fair mechanism that is easy to implement/use?

Nagle's idea: Instead of links, allocate bandwidth by isolating flows.



- Bad
- ① # Queues = # flows ever seen (provision ^{many} queues for worst case)
 - ② Unfairness = relative proportions of n/w shares received

$$= \frac{\text{Max Pkt}}{\text{Min Pkt}} \rightarrow \infty$$
- Good
- ③ Work complexity = $O(1)$ ~~per flow~~
 - O(1) enqueueing
 - O(1) dequeueing

FG proposal of DKS:

- Fixes the unfairness problems
 - Buffer allocation and incentives / disincentives
 - But somewhat complex to implement
 - Other salient features.
- Simulate bit-by-bit RR.

Upon packet arrival, calculate finish time and insert into a sorted queue based on finish times. → not just round number

- per flow state

~~Finish time = size + max (current time, finish time of prev. packet in flow)~~

Finish Round = $\frac{\text{size}}{\text{in bits}} + \max(\text{current round}, \text{finish round of prev. pkt in flow})$

$\frac{dR(t)}{dt}$ → can change over time as the number of active flows grows and shrinks.

also $R(t)$ is fractional. → packets arriving almost at the same time will see different $R(t)$

① Insert $O(\log m)$
 $\hookrightarrow m$ is the number of packets in the queue.

② For each flow, track finish time of last packet in flow.
 $O(\text{num active flows})$

FG Buffer management: Buffer has a certain maximum length.

Two options: "Buffer stealing"

remove packet from flow with maximum byte occupancy } can be done in $O(\log n)$ if done naively }
 } $O(1)$ time.

But don't update

What to do when buffer is full.

Delay allocation with a δ

- Salient properties:
- ① Fair
 - ② Low delay for telnet-flows
 - ③ Good congestion control still important to make good efficient use of n/w.
 - ④ Incentives to adopt good congestion control
 - ⑤ Punishment to misbehaving flows.

Can we achieve the same fairness as FA but at $O(1)$ cost per packet, and lower space requirement, and low amount of state tracking.

RR of Nagle $\rightarrow O(1)$ but unfairness.

idea:

RR + deficit

Additional variables

Quantum: Q

Deficit per-queue: D_i

~~At round t , send $Q_i + D_i$ bytes from queue i~~

Round Robin: when servicing Queue i

check if size of packet at head, $H_i \leq Q + D_i$

yes \rightarrow send $B_i = |H_i|$ bytes

no \rightarrow send $B_i = 0$

$D_i \leftarrow B_i + D_i$

Thm: Difference between FA and DR.


FA \rightarrow maximum fine-grained difference across flows MAX

DR \rightarrow " " " $\approx 3MAX$

~~Another difference~~

Complexity: Enqueue: finding the queue, append to tail of queue

Dequeue: $Q \geq MAX \rightarrow$ guaranteed to send on packet every ~~good~~ time a queue is visited



$(N-1)$ MAX delay.

Unfairness: $3MAX$

Problem: queues are way too many

How to bound the number of Queues: SFR.

If the Hash index is larger than the number of active flows at any given point then fairness guarantees become static.

From the last lecture:

- Buffer sizing
- TCP Throughput
- TFR

QoS