EECS 122
Supplementary Notes on WFQ

Kevin Fall, kfall@cs.berkeley.edu

Weighted Fair Queuing (WFQ)

- Textbook is not quite precise on its description of WFQ, so these notes will attempt to clarify what is going on
- Recall WFQ is equivalent to PGPS (packet-by-packet GPS), but was discovered separately
- Both are packet-oriented approximations of GPS, and thus approximate max-min fairness

WFQ Outline

- Compute the time at which a packet would complete service using GPS
- Service packets in the order of these times (but not necessarily at these times)
- Scheduler maintains two variables:
  - the current round number
  - the highest per-queue finish number

The Round Number

- The round number is the number of rounds of service completed by a bit-by-bit round-robin scheduler at some time
- The round number may not be an integer. Fractions represent partially-complete rounds.
  - round number 4.5 would be half way through round 5

Active Connections

- A connection is active if the largest finish number in its queue is beyond (greater than) the current round number
- So, the length of a round is proportional to the number of active connections
  - the round rate is inversely proportional to the number of active connections
  - the round rate is proportional to the link rate

Computing Finish Numbers

- By knowing the current round number, calculate the finish number of a packet as its size (in bits) plus the greater of:
  - the finish number of the previous packet in the same queue (active connections)
  - the current round number (inactive connections)
- For the moment, assume all weights are 1.0; we will extend this later…

Computing Finish Numbers

- So, if $F_i(k,t)$ is the finish number of packet k on connection i at time t:

$$F_i(k,t) = \max\{F_i(k-1,t), R(t)\} + P_i(k,t)$$

- Update the round number $R(t)$ on each packet arrival or departure. $P_i(k,t)$ is the size of the kth packet arriving on connection i at time t

An Example [Keshav97]

- Assume a WFQ scheduler with all weights 1.0 using three queues (connections)
- Packets of size 1, 2, and 2 units arrive at time 0 on connections A, B, and C. A second packet on connection A arrives at time 4.
- The link rate is 1.0 unit/s
- Initialize the system with R(t)=0, and per-connection finish numbers as zero

WFQ Example [1]

			Queue A		Queue B		Queue C	
	Number of		Finish #	Remaining	Finish #	Remaining	Finish #	Remaining
Time	Connections	Round						
t=0	3	0	1	1	2	2	2	2

- The finish numbers for connections A, B, and C are set as 1, 2, and 2
- With three connections, the round rate is 1/3 rounds/sec

WFQ Example [2]

			Queue A		Queue B		Queue C	
	Number of		Finish #	Remaining	Finish #	Remaining	Finish #	Remaining
Time	Connections	Round						
t=0	3	0	1	1	2	2	2	2
t=3	2	1	1	0*	2	1	2	1

- After 3 units of time, each connection has received 3*0.33=1.0 units of service
- That is enough service for the first packet to depart, but only half enough for the packets on connections B and C

WFQ Example [3]

			Queue A		Queue B		Queue C	
	Number of		Finish #	Remaining	Finish #	Remaining	Finish #	Remaining
Time	Connections	Round						
t=0	3	0	1	1	2	2	2	2
t=3	2	1	1	0*	2	1	2	1
t=4	3	1.5	3.5	2	2	0.5	2	0.5

- During the time interval [3,4], only 2 connections, so the round rate increases to 0.5 r/s
- A packet of size 2 arrives on connection A at time 4, so rate returns to 0.33 r/s

WFQ Example [3]

			Queue A		Queue B		Queue C	
	Number of		Finish #	Remaining	Finish #	Remaining	Finish #	Remaining
Time	Connections	Round						
t=0	3	0	1	1	2	2	2	2
t=3	2	1	1	0*	2	1	2	1
t=4	3	1.5	3.5	2	2	0.5	2	0.5
t=5.5	1	2	3.5	1.5	2	0*	2	0*

- During the [4,5.5] time interval, connections B and C receive 1.5*0.33=0.5 units of service, allowing them to depart
- This leaves only 1 connection (rate 1.0)

WFQ Example [3]

			Queue A		Queue B		Queue C	
	Number of		Finish #	Remaining	Finish #	Remaining	Finish #	Remaining
Time	Connections	Round						
t=0	3	0	1	1	2	2	2	2
t=3	2	1	1	0*	2	1	2	1
t=4	3	1.5	3.5	2	2	0.5	2	0.5
t=5.5	1	2	3.5	1.5	2	0*	2	0*
t=7	0	3.5	3.5	0*	2	0	2	0

- During the [5.5,7] time interval, the round rate is 1.0, so connection A receives 1.5 units of service (all it needs).
- The scheduler goes idle at time 7

WFQ Example [4]

- The actual departure order is based on the finish number, so the order is either:
  - A, B, C, A
    - or
  - A, C, B, A
- B and C had same finish number, so break the tie randomly (or some other way)

WFQ Finish Numbers

- So, if $F_i(k,t)$ is the finish number of packet k on connection i at time t:

$$F_i(k,t) =$$
$$\max\{F_i(k-1,t), R(t)\} + P_i(k,t) / f(i)$$

- $f(i)$ is the weight on connection i
- In this case, the round increase rate is

$$1/\sum f(i)$$

Reference

- [Keshav97] - S. Keshav, *An Engineering Approach to Computer Networking*, 1997