New Topic: Scheduling Theory

This will be our topic for the remainder of the course. Today's focus will be on **performance metrics** for scheduling policies.

Question: Why should we care about scheduling?

1 Traditional Performance Metrics

$\mathbf{E}[T]$: mean response time
$\mathbf{E}[T_Q]$: mean waiting time
$\mathbf{E}[N]$: mean number in system
$\mathbf{E}[N_Q]$: mean number in queue

Question: Your goal is to improve mean response time. Someone tells you they have a super algorithm that improves mean waiting time by a factor of 100. Before you buy the algorithm, what question should you ask?

2 Commonly Used Performance Metrics

Work in system, W: remaining work left to do in the system

Utilization of device, ρ: fraction of time that the device is busy

Defn: **Arrival Sequence:** sequence of arrival times and job sizes.

Question: You are told that 2 scheduling policies, run on the same arrival sequence, result in the same work in system over all time and the same device utilization. Does that mean that the two policies also have the same response time?

Definition: A **work-conserving** scheduling policy is one that always performs work on some job when there is a job in the system. Also, the policy does not create new work.

Question: Do all work-conserving policies have the same mean response time?

3 Today's Trendy Metrics: Slowdown

Defn: The **slowdown** of a job is its response time divided by its size:

$$\text{Slowdown} = \frac{T}{S}$$

Question: Why is mean slowdown preferable to mean response time?

Question: What does mean slowdown tell us about max slowdown?

4 Today's Trendy Metrics: Tails

Defn: The **tail of response time** is $\mathbf{P}\{T > t\}$. HARD to derive!

Chebyshev's Bound: Goal: Bound $\mathbf{P}\{T \geq \mathbf{E}[T] + a\}$

Chebyshev Cantelli Bound: Goal: Bound $\mathbf{P}\{T \geq \mathbf{E}[T] + a\}$

5 Minimizing the Tail

Question: Given specific SLO: 0.5 sec. How to minimize $\mathbf{P}\{T > 0.5 \text{ sec}\}$?

The ubiquitous "minimize the 99%-tile" goal

Minimizing the tail is discussed in: [Mor Harchol-Balter. "Open problems in queueing theory inspired by datacenter computing." Queueing Systems , vol. 97, no. 1, 2021, pp. 3–37].

6 Starvation Metrics

People like to say that policies that bias towards short jobs **starve** large jobs. Example: SRPT is accused of starving long jobs.

Question: Does starvation really exist if $\rho < 1$?

Question: Let's instead ask about **fairness**. What performance metric might allow us to judge if large jobs are being treated unfairly as compared with small jobs?

Question: Is there a scheduling policy that provides equal expected slowdown to all job sizes?

Question: Fairness is deceptive. Intuition says "there's no free lunch," meaning that a scheduling policy that helps some jobs *must* hurt others. Suppose I tell you that switching from scheduling policy A to scheduling policy B resulted in strictly improving the response time of *almost all* jobs, where *no job* ended up with worse response time. Is this possible?

7 Commercial Break: Announcements

1. Zhouzi's office hours today: GHC 6003.

2. There's only 1 week left to schedule a 1-on-1 research meeting with me. If you don't schedule the meeting, you will forfeit 5% of your grade.

3. Midterm 2: Thursday, Dec 4th, 5 p.m. - 7 p.m. (but I will give you extra time if needed) in our current classroom. You will have a 5x7 inch index card as a cheat sheet. You can write on both sides. As usual, the midterm will have at least 50% of the questions taken directly from class or existing homework. The goal is for EVERYONE to pass.

8 Where we're going ... scheduling policies

For all remaining chapters, assume M/G/1, where G has finite mean and variance. Assume $\rho < 1$, and some work-conserving scheduling policy.

Defn: A **non-preemptive** service order is one that does not preempt a job once it starts service (i.e., each job is run to completion).

Defn: A **preemptive** service order is one where we are allowed to stop (preempt) a running job and continue running it later where it left off.

	Non–Preemptive	Preemptive
Don't Use Size		
Use Size/ Class		

9 Non-preemptive policies that don't use size

Here are 3 non-preemptive policies that don't make use of job size:

FCFS: When the server frees up, it always chooses the job at the head of the queue to be served and runs that job to completion.

LCFS: When the server frees up, it always chooses the last job to arrive and runs that job to completion.

RANDOM: When the server frees up, it always chooses a random job to run next and runs that job to completion.

Question: When would one use LCFS?

Question: Which of these has lowest mean response time?

10 Non-preemptive scheduling theorem

Question: Consider all non-preemptive service orders that do not make use of job sizes. Which of the following statement(s) is true?

(a) All these policies have the same $\mathbf{E}[N]$ and $\mathbf{E}[T]$?

(b) All these policies have the same distribution on N.

(c) All these policies have the same distribution on T.

(d) None of the above.

Question: Why was it important that the scheduling policies not use job size?

11 Higher moments of response time

Question: How does $\widehat{N}(z)$ differ for FCFS, LCFS, and Random?

Question: Can we use DLL to convert $\widehat{N}(z)$ to $\widetilde{T}(s)$?

Question: Do FCFS, Random, and LCFS have the same $\mathbf{Var}(T)$?

GOAL FOR REST OF LECTURE: We will now derive $\widetilde{T_Q}^{LCFS}(s)$.
But first, it will help to review Busy Periods!

12 Reviewing Busy Periods

Given M/G/1 queue with arrival rate λ and job sizes S.

B_x = Busy period started by job of size x.

$B = B_S$ = Busy period started by arbitrary job

B_W = Busy period started by W work

Question: Derive $\widetilde{B_x}(s)$:

Question: Derive $\widetilde{B}(s)$:

Question: Derive $\widetilde{B_W}(s)$:

13 Derive $\widetilde{T_Q}^{LCFS}(s)$

After lots of differentiating we can show from the transforms for LCFS and FCFS that:

$$\mathbf{E}\left[T_Q^2\right]^{LCFS} = \frac{\lambda \mathbf{E}\left[S^3\right]}{3(1-\rho)^2} + \frac{\left(\lambda \mathbf{E}\left[S^2\right]\right)^2}{2(1-\rho)^3}$$

By contrast:

$$\mathbf{E}\left[T_Q^2\right]^{FCFS} = \frac{\lambda \mathbf{E}\left[S^3\right]}{3(1-\rho)} + \frac{\left(\lambda \mathbf{E}\left[S^2\right]\right)^2}{2(1-\rho)^2}$$

Question: How do $\mathbf{E}\left[T_Q^2\right]^{LCFS}$ and $\mathbf{E}\left[T_Q^2\right]^{FCFS}$ compare?