M/G/1 Model with 4 Categories of Scheduling Policies

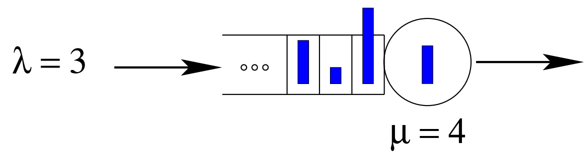

	Non-Preemptive	Preemptive
Non-Size-Based		
Size/Class-Based		

Table 1: *4 quadrants*

Question: Assume that S is continuous. Define these important quantities:

(a) $\mathbf{E}[S]$

(b) $\mathbf{E}[S^2]$

(c) ρ

Non-Preemptive, Non-Size-Based scheduling policies

Question: Name and define some non-preemptive, non-size-based policies

Question: What did HW 5 teach us on how these compare wrt $\mathbf{E}[T_Q]$? $\mathbf{E}[T]$?

Question: What is some intuition behind this result?

Question: What's the problem with non-preemptive, non-size-based policies?

Non-preemptive Priority Scheduling

Question: When does non-preemptive priority come up?

MODEL:

- n priority classes.
- Class 1 has ______________ priority. Class n has ______________ priority. Within a class serve jobs FCFS.
- Arrival rates: $\lambda_1, \lambda_2, \ldots, \lambda_n$
- Job size distributions: $S_1, S_2, \ldots, S_n$

Question: What is ρ_i, the load made up of jobs of class i?

Question: How do we denote the total load made up by classes 1 through k?

Non-preemptive Priority Scheduling, cont.

Let $T_Q(k)^{\text{NP-Prio}}$ denote the queueing time for a job of class k.
Let $T(k)^{\text{NP-Prio}}$ denote the response time for a job of class k.

$$\mathbf{E}\left[T_Q(k)\right]^{\text{NP-Prio}} = \frac{\rho}{(1-\sum_{i=1}^{k}\rho_i)(1-\sum_{i=1}^{k-1}\rho_i)} \cdot \frac{\mathbf{E}\left[S^2\right]}{2\mathbf{E}\left[S\right]}$$

Question: What is $\mathbf{E}\left[T_Q(1)\right]$? How does this compare with $\mathbf{E}\left[T_Q\right]^{\text{FCFS}}$?

Question: What is $\mathbf{E}\left[T_Q(n)\right]$? How does this compare with $\mathbf{E}\left[T_Q\right]^{\text{FCFS}}$?

Question: For smaller k, how does $\mathbf{E}\left[T_Q(k)\right]^{\text{NP-Prio}}$ compare with $\mathbf{E}\left[T_Q\right]^{\text{FCFS}}$?

Question: Why does $\mathbf{E}\left[T_Q(k)\right]$ still see full $\mathbf{E}\left[S^2\right]$?

Non-preemptive Shortest-Job-First (SJF)

Question: Define SJF

Question: Let ρ_x denote the load made up of jobs of size $\leq x$. What is ρ_x?

<u>Defn</u>: $T_Q(x)^{\text{SJF}}$ is the queueing time for a job of size x

$$\mathbf{E}\left[T_Q(x)\right]^{\text{SJF}} = \frac{\rho}{(1-\rho_x)^2} \cdot \frac{\mathbf{E}\left[S^2\right]}{2\mathbf{E}\left[S\right]}$$

Non-preemptive Shortest-Job-First (SJF), cont.

$$\mathbf{E}\left[T_Q(x)\right]^{\text{SJF}} = \frac{\rho}{(1-\rho_x)^2} \cdot \frac{\mathbf{E}\left[S^2\right]}{2\mathbf{E}\left[S\right]}$$

Question: How does $\mathbf{E}\left[T_Q(x)\right]^{\text{SJF}}$ compare with $\mathbf{E}\left[T_Q(x)\right]^{\text{FCFS}}$ for small x? How about large x?

Question: What is $\mathbf{E}\left[T_Q\right]^{\text{SJF}}$? How do you expect it compares to $\mathbf{E}\left[T_Q\right]^{\text{FCFS}}$?

Question: Why is $\mathbf{E}\left[S^2\right]$ still there?

Preemptive policies

Question: Why do preemptive policies have lower response times than non-preemptive policies? Explain for PS and P-LCFS.

Question: In HW 5 we investigated preemptive policies, PS and P-LCFS. How did these compare? What was their response time?

$$\mathbf{E}[T]^{\text{PS}} =$$

$$\mathbf{E}[T]^{\text{P-LCFS}} =$$

Question: Let's look under the hood ...

Let $T(x)$ denote the response time for a job of size x.

$$\mathbf{E}[T(x)]^{\text{PS}} =$$

$$\mathbf{E}[T(x)]^{\text{P-LCFS}} =$$

Preemptive, Size-Based Scheduling

Fill in Picture: Response time as a function of load

Question: Can we possibly improve upon PS? What is there left to do?

Three preemptive size/class-based policies:

1. Preemptive Priorities
2. PSJF
3. SRPT

Preemptive Priorities

MODEL:

- n priority classes.
- Class 1 has ______________ priority. Class n has ______________ priority.
- Arrival rates: $\lambda_1, \lambda_2, \ldots, \lambda_n$
- Job size distributions: $S_1, S_2, \ldots, S_n$
- $\rho_i =$______________.
- Preempt-resume.

ANALYSIS:
Waiting Time: Time until job first starts serving
Residence Time: Time from when job first receives service until it departs

$$\mathbf{E}\left[T(k)\right]^{\text{P-Prio}} = \frac{\sum_{i=1}^{k} \rho_i \frac{\mathbf{E}\left[S_i^2\right]}{2\mathbf{E}\left[S_i\right]}}{(1 - \sum_{i=1}^{k-1} \rho_i)(1 - \sum_{i=1}^{k} \rho_i)} + \frac{\mathbf{E}\left[S_k\right]}{1 - \sum_{i=1}^{k-1} \rho_i}$$

Question: How do waiting time and residence time compare under P-Prio and NP-Prio?

Preemptive Shortest Job First – PSJF

Question: Define PSJF. How does this differ from SJF?

Question: What is ρ_x? Why is it relevant? Who is visible to a job of size x?

ANALYSIS:
Waiting Time: Time until job first starts serving
Residence Time: Time from when job first receives service until it departs

$$\mathbf{E}\left[T(x)\right]^{\text{PSJF}} = \frac{\lambda \int_0^x t^2 f_S(t) dt}{2(1-\rho_x)^2} + \frac{x}{1-\rho_x}$$

Question: What is $\mathbf{E}\left[T\right]^{\text{PSJF}}$?

Question: Why should we expect that PSJF is superior to SJF?

SRPT

Question: Define SRPT. How does this differ from PSJF? How do we program SRPT?

ANALYSIS:
Waiting Time: Time until job first starts serving
Residence Time: Time from when job first receives service until it departs

$$\mathbf{E}\left[T(x)\right]^{\text{SRPT}} = \frac{\lambda \int_0^x t^2 f_S(t)dt + \lambda x^2 \left(\overline{F_S}(x)\right)}{2(1-\rho_x)^2} + \int_{t=0}^{x} \frac{dt}{1-\rho_t}$$

Question: Compare the formulas for $\mathbf{E}\left[T\right]^{\text{PSJF}}$ and $\mathbf{E}\left[T\right]^{\text{SRPT}}$.

$$\mathbf{E}\left[T(x)\right]^{\text{PSJF}} = \frac{\lambda \int_0^x t^2 f_S(t)dt}{2(1-\rho_x)^2} + \frac{x}{1-\rho_x}$$

Theorem [Schrage 1966]: SRPT produces optimal mean response time for any arrival sequence (any arrival times, job sizes).

Comparison of policies

Here's what you can expect to see when the job size distribution, S, has high C_S^2:

FILL IN GRAPH

Note: This shows $\mathbf{E}[T]$, but actually all moments of T are known (see textbook).

EXTRA: Scheduling and starvation

Question: What is starvation? Does it really exist?

Definition: $Slowdown(x) = \frac{T(x)}{x}$

Question: What is $\mathbf{E}\left[Slowdown(x)\right]^{\text{PS}}$?

Question: For policies which favor short jobs, what do you think $\mathbf{E}\left[Slowdown(x)\right]$ looks like as a function of x?

The All-Can-Win-Theorem: If $\rho \leq \frac{1}{2}$, then for **all** job sizes x,

$$\mathbf{E}\left[Slowdown(x)\right]^{\text{SRPT}} \leq \mathbf{E}\left[Slowdown(x)\right]^{\text{PS}} .$$

(Bansal and Harchol-Balter "Analysis of SRPT scheduling: Investigating unfairness" ACM SIGMETRICS 2001.)

EXTRA: Opt scheduling when we don't know job size

1. The FB policy – Optimal when job size distribution has DFR.

2. Shortest-Expected-Remaining-Processing-Time (SERPT) – Simple and almost always near optimal

   Scully, Harchol-Balter, Scheller-Wolf "Simple near-optimal scheduling for the M/G/1" SIGMETRICS 2020.

3. Gittins Index Policy – Always optimal, but complex.

   First exact analysis of both SERPT and Gittins: Scully, Harchol-Balter, Scheller-Wolf. "SOAP: One clean analysis of all age-based scheduling policies." SIGMETRICS 2018.

EXTRA: Opt scheduling when jobs have value and size

DRAW PICTURE

Question: What scheduling policy makes the most intuitive sense?

The $c\mu$ Rule

Real-world complications ...

Scully and Harchol-Balter. "The Gittins policy in the M/G/1 queue." 19th International Symposium on Modeling and Optimization in Mobile, Ad hoc, and Wireless Networks 2021.

EXTRA: Scheduling when jobs have deadlines

Set up: Each job has a deadline. Job is worthless after its deadline.

Question: What's wrong with Earliest Deadline First?

Better idea: Drop-If-Hopeless (DIH)

Even better idea: Drop-The-Largest-Hopeless (DIH)

(see Discussion in Section 7.2 of this paper: Harchol-Balter. "Open problems in queueing theory inspired by datacenter computing." Queueing Systems, vol. 97, no. 1, pp. 3-37, 2021.)