Single-Server FCFS Queue

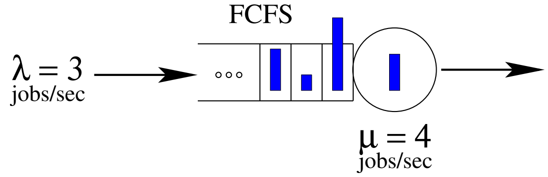

I: interarrival time

$\mathbf{E}[I] =$ ________________________

$\lambda =$ avg arrival rate $=$ ________________________

$S =$ job size $=$ # seconds a job takes on this server *when run alone*

$\mathbf{E}[S] =$ ________________________ on this server

$\mu =$ avg service rate $=$ ________________________

$T =$ response time $= t_{dept} - t_{arrive}$

$T_Q =$ queueing time (a.k.a., delay)

Question: How are T and T_Q related?

$N =$ number jobs in system

$N_Q =$ number jobs in queue

Single-Server FCFS Queue

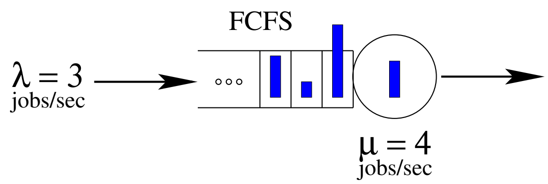

Assume throughout that $\lambda < \mu$:

Question: : As $\lambda \uparrow$, what happens to T? to N?

Question: As $\mu \uparrow$, what happens to T? to N?

Question: Suppose $I \sim$ Deterministic $\equiv \frac{1}{\lambda}$, and $S \sim$ Deterministic $\equiv \frac{1}{\mu}$:

(a) What is T? ______________________

(b) What is T_Q? ______________________

Defn: **Load** (denoted by ρ) is the fraction of time that the server is busy.

Question: How do you think ρ relates to λ and μ?

[We will formally derive ρ in Chpt 6]

Question: Is it possible to have $\rho = 0.1$ but T_Q very high?

Single-Server FCFS Queue

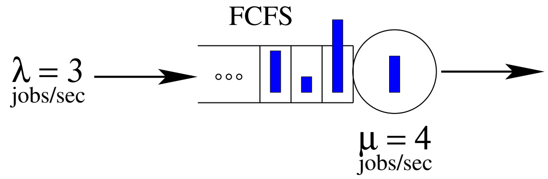

Question: So far we've assumed $\lambda < \mu$. What happens to N when $\lambda > \mu$?

Throughput

Story: Before I was at CMU, I was at MIT. There, I developed a scheduling algorithm for Web servers that dramatically reduced $\mathbf{E}[T]$.

Story Question: Given that my algorithm lowered $\mathbf{E}[T]$, does that imply that throughput went up?

[We will come back to this.]

Question: Which of the following systems has higher throughput?

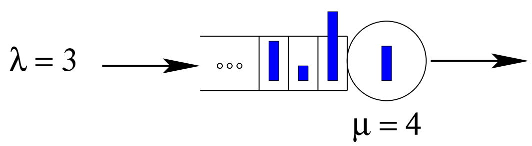

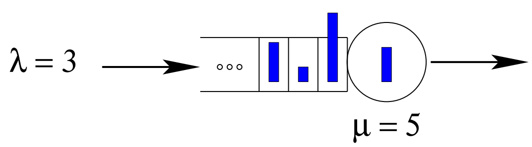

Defn: **Throughput** is the average rate of completions, i.e., jobs/sec that complete. It is traditionally denoted by X, even though it's not a r.v.

Question: What is throughput when $\lambda > \mu$?

Revisit story question ....

Throughput, examples

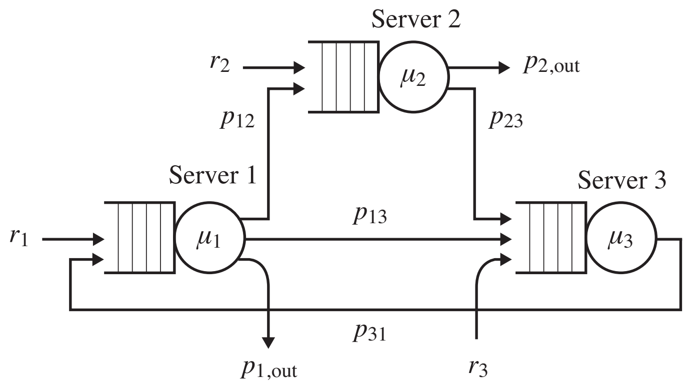

Question: What is X?

Question: What is X_i?

Question: What constrains r_i?

Throughput, examples

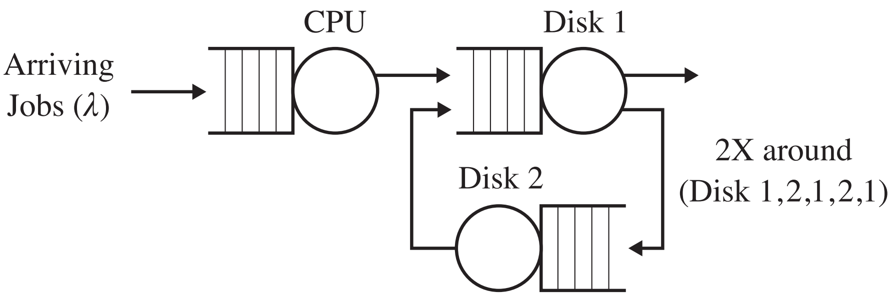

Question: What is X above?

Question: What is X_{disk1} above?

Question: What is X for the finite buffer queue below?

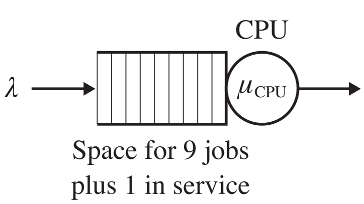

Commercial Break: Announcements

1. You must turn in Probability Assessment to stay in the class. The TAs and I will grade this and get back to you.

2. Remember to bookmark class website:
   www.cs.cmu.edu/~harchol/Perfclass/class.html
   - Go to **Announcements-and-Homeworks link** to get HW 1
   - Check frequently for updates and hints
   - New homework is posted every Friday (often earlier).
   - Homework is due every Friday at the START of class.
   - No late homeworks. Get help early!
     - Zhouzi – zhouzil@andrew.cmu.edu – GHC 6003 – M 3:30 - 5 p.m.
     - Mor – harchol@cs.cmu.edu – GHC 7207 – W 5:30 - 7 p.m.
     - Keerthana – kgurusha@andrew.cmu.edu – GHC 7004 – Th 5:30 - 7 p.m.

3. This class meets M, W, F. You can get some advanced notice of what is being covered in each class looking at the syllabus on the Announcements-and-Homework link!

4. Please come to Mor's office hours TODAY: 5:30 p.m. in GHC 7207.

5. Please keep up with the reading for this class. Spend 3 hours after every class reading the relevant chapter. Book is a sequence of Questions & Answers. Read with a pen in hand. Cover up the answers and think!

6. Friday's class will cover generating random variables for simulation. This is Chpt 4 of the book. For those in PnC, you have seen this, but might still want a review. Up to you, but there are a couple simulation problems in HW 1.

7. No class Monday – Labor Day.

Closed Systems

So far: **Open Systems**

- exogenous arrival process
- almost all of theory lives here

Now: **Closed Systems**

- arrivals triggered by job completions
- many computer systems live here
- two types of closed systems

Type 1: Interactive Closed System:

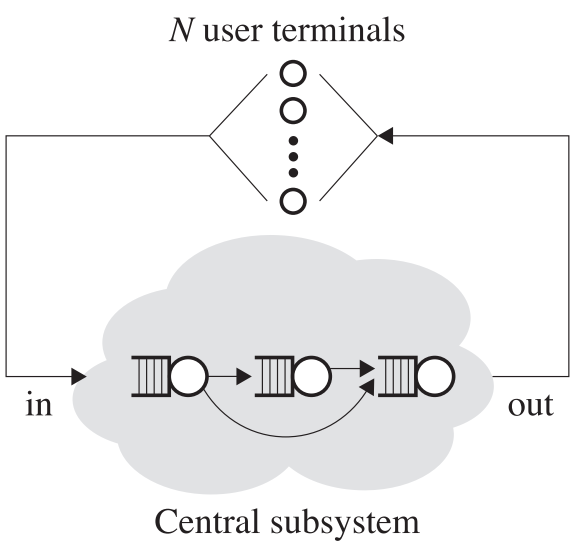

Type 2: Batch Closed System:

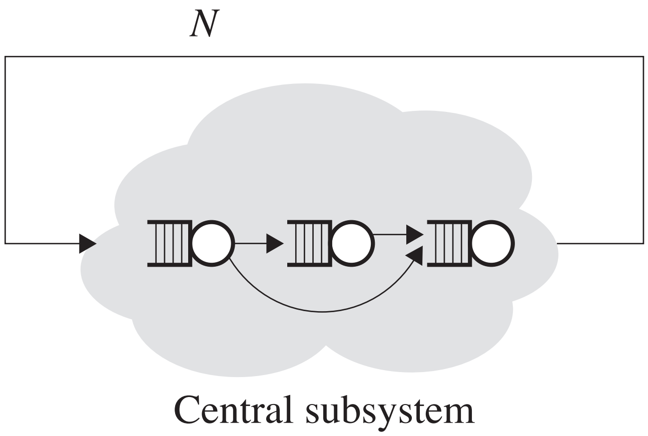

Comparing Closed vs. Open Systems

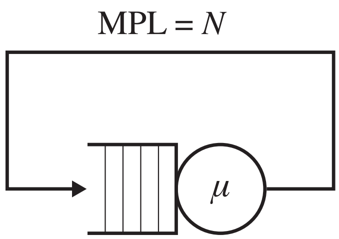

Question: What is X?

Question: What is $\mathbf{E}[R]$?

Question: What happens to X and $\mathbf{E}[R]$ if we change μ to 2μ?

Question: How does this behavior differ from an open system?

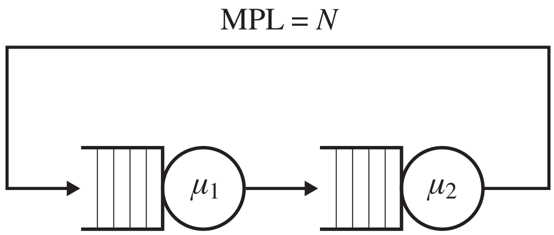

Question: What is X?

If time: Some puzzles ...