ECE560: Computer Systems Performance Evaluation

Lecture # 13 –

Queueing Systems (II)

Instructor: Dr. Liudong Xing

Administration Issues

- **Annotated Bibliography**
  - Due: **March 22, Friday**
  - Refer to Section 2.2 in the Project Description for the guidelines
- **If you are looking for jobs or internship:**
  - **Wednesday, March 27**: Job, Internship, and Graduate School Expo, 12-3:00 p.m., Tripp Athletic Center
- **Survey results (13 responses)**

1. How was the pace of the class?

   Too slow ____ Slow_____ Just right_13____ Fast_____ Too fast _____

2. How was the workload of ECE560 class, when compared to your other classes?

   Much less__1____ Less__2____ Similar__10_____ More_____ Much more__

3. How is the level of detail of material covered in the lectures?

   Not enough ____ Just right _11____ Too much _2____

4. How clear are the lectures/examples?

   Not clear _____ Okay __0.5___ Clear __1.5__ Very clear __11___

5. How was the midterm exam?

   Too easy ____ Easy _5___ Just right __8__ Hard ____ Too hard _____

6. How were the homework assignments?

   Too easy ____ Easy _1___ Just right _11___ Hard __1__ Too hard _____

Review of Lecture #11

- Overview of queueing systems (Kendall notation, Little's Law)
- Performance measures
  - Traffic intensity (α)
  - Server utilization (ρ)
  - π_n: steady-state probability that there are *n* customers in the queueing system
  - Throughput (γ): rate at which jobs successfully depart from the system
  - Average number of jobs in the system (L)
  - Average time in the system (W)
  - Average number of jobs in the queue (L_q)
  - Average time in the queue (W_q)
- D/D/1 queueing systems
- M/M/1 queueing systems

Review (M/M/1)

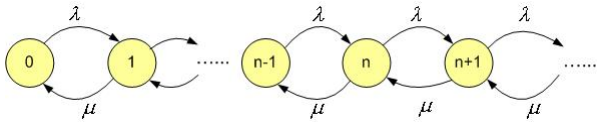

- Performance measures:

$$\alpha = \rho = \lambda / \mu$$

$$\pi_0 = 1 - \frac{\lambda}{\mu} = 1 - \rho,\ \pi_n = \left(\frac{\lambda}{\mu}\right)^n \pi_0 = \rho^n (1-\rho)$$

$$\gamma = \mu P[> 0 \text{ jobs in the system}]$$
$$= \mu(1 - P[0 \text{ jobs in the system}])$$
$$= \mu(1-\pi_0) = \mu(1-(1-\rho)) = \mu\rho = \lambda$$

$$L = \sum_{n=0}^{\infty} n\pi_n = (1-\rho)\sum_{n=0}^{\infty} n\rho^n$$

$$= (1-\rho)\rho\sum_{n=1}^{\infty} n\rho^{n-1} = \frac{(1-\rho)\rho}{(1-\rho)^2} = \frac{\rho}{1-\rho}$$

$$W = L/\lambda = \frac{\rho}{1-\rho}/\lambda = \frac{1}{\mu-\lambda}$$

$$L_q = L - (1 * P[\text{Server is not empty}]$$
$$= L - (1 - P[0 \text{ jobs in the system}])$$
$$= L - (1-\pi_0) = L - (1-(1-\rho))$$

$$= L - \rho = \frac{\rho}{1-\rho} - \rho = \frac{\rho^2}{1-\rho}$$

$$W_q = L_q/\lambda = \frac{\rho^2}{1-\rho}\frac{1}{\lambda} \text{ or } W_q = W - W_s = \frac{\rho}{(1-\rho)\lambda} - \frac{1}{\mu} = \frac{\rho^2}{1-\rho}\frac{1}{\lambda}$$

What if the system has finite capacity (e.g., 1 buffer in the queue)?

Topics

- More Birth-and-Death Queues (Cont'd)
  - **M/M/1/N Queues**
  - M/M/c Queues
  - M/M/ ∞ Queues
  - M/M/1/k/k Queues

M/M/1/N Queues

- The infinite buffer/queue assumption is unrealistic in practice!
- *N*: the system capacity/the maximum number of customers allowed in the system: 1 in service and (*N*-1) in the queue
- State transition diagram:

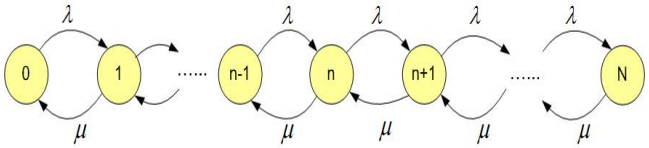

$$\lambda_n = \begin{cases} \lambda & n < N \\ 0 & n \geq N \end{cases}$$

$$\mu_n = \begin{cases} \mu & n \leq N \\ 0 & n > N \end{cases}$$

M/M/1/N: Steady-State Probabilities

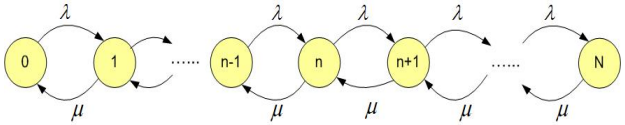

- The same Balance Equations as M/M/1, thus,

$$\lambda\pi_0 = \mu\pi_1$$
$$\lambda\pi_1 = \mu\pi_2$$
$$\cdots\cdots\cdots\cdots$$
$$\lambda\pi_{N-1} = \mu\pi_N$$

In general

$$\Rightarrow \quad \pi_n = \frac{\lambda}{\mu}\pi_{n-1} = \left(\frac{\lambda}{\mu}\right)^n \pi_0 = \alpha^n \pi_0$$

for $n = 0, 1, 2, \ldots, N$, $\quad \alpha = \lambda/\mu$

$$\sum_{n=0}^{N} \pi_n = \pi_0 \sum_{n=0}^{N} \alpha^n = 1 \Rightarrow \pi_0 = \frac{1}{\sum_{n=0}^{N} \alpha^n} = \frac{1}{\frac{1-\alpha^{N+1}}{1-\alpha}} = \frac{1-\alpha}{1-\alpha^{N+1}}$$

M/M/1/N: Performance Measures (I)

$$\alpha = \lambda / \mu$$

$$\pi_0 = \frac{1-\alpha}{1-\alpha^{N+1}},$$

$$\pi_n = \alpha^n \pi_0 = \frac{\alpha^n (1-\alpha)}{1-\alpha^{N+1}} \text{ for } n = 0,1,2,..., N$$

Blocking probability (P_B): probability an arriving job is turned away due to a full buffer

$$P_B = \pi_N = \alpha^N \pi_0 = \frac{\alpha^N (1-\alpha)}{1-\alpha^{N+1}}$$

For example: for α=0.5, if want a blocking probability of $< 10^{-6}$, what is the minimum system capacity?

M/M/1/N: Performance Measures (II)

The effective / average arrival rate of customers into the system:

$$\lambda_{eff} = \lambda(1 - P_B)$$

$$P_B = \pi_N = \alpha^N \pi_0 = \frac{\alpha^N (1-\alpha)}{1-\alpha^{N+1}}$$

The true server utilization rate ρ – probability that the server is busy:

$$\rho = 1 - \pi_0 = 1 - \frac{1-\alpha}{1-\alpha^{N+1}},$$ or

$$\rho = \lambda_{eff} W_s = \frac{\lambda_{eff}}{\mu} = \lambda (1 - P_B) / \mu = \alpha (1 - P_B)$$

$$\rho = \frac{\alpha(1-\alpha^N)}{1-\alpha^{N+1}}$$

M/M/1/N: Performance Measures (III)

Throughput (γ): rate at which jobs successfully depart from the system

$$\gamma = \mu P[>0 \text{ jobs in the system}] + 0P[0 \text{ jobs}]$$
$$= \mu(1 - P[0 \text{ jobs in the system}])$$
$$= \mu(1-\pi_0) = \mu(1-(\frac{1-\alpha}{1-\alpha^{N+1}}))$$
$$= \mu\frac{\alpha(1-\alpha^N)}{1-\alpha^{N+1}} = \frac{\lambda(1-\alpha^N)}{1-\alpha^{N+1}}$$

Alternative way to compute γ (by looking at the input side)

Everything that arrives and is not blocked must eventually depart

$$\gamma = \lambda(1-P_B) = \lambda(1-\alpha^N\frac{1-\alpha}{1-\alpha^{N+1}}) = \frac{\lambda(1-\alpha^N)}{1-\alpha^{N+1}}$$

M/M/1/N: Performance Measures (IV)

Average number of customers in the system ***L***:

$$L = \sum_{n=0}^{N} n\pi_n = \frac{1-\alpha}{1-\alpha^{N+1}}\sum_{n=0}^{N} n\alpha^n$$
$$= (1-\alpha)\alpha \Big/ 1-\alpha^{N+1}\sum_{n=1}^{N} n\alpha^{n-1}$$
$$= \frac{(1-\alpha)\alpha}{1-\alpha^{N+1}}\left(\sum_{n=0}^{N}\alpha^n\right)' = \frac{(1-\alpha)\alpha}{1-\alpha^{N+1}}\left(\frac{1-\alpha^{N+1}}{1-\alpha}\right)'$$
$$= \frac{\alpha}{1-\alpha} - (N+1)\frac{\alpha^{N+1}}{1-\alpha^{N+1}}$$

Average response time ***W***:

$$W = L/\lambda_{eff} = \left[\frac{\alpha}{1-\alpha} - \frac{(N+1)\alpha^{N+1}}{1-\alpha^{N+1}}\right]/\lambda_{eff}$$

Note: effective arrival rate λ_{eff} used in Little's Law!

$$\lambda_{eff} = \lambda(1-P_B)$$

M/M/1/N: Performance Measures (V)

Average queue length L_q:

$$\begin{aligned} L_q &= L - (1 * P[\text{Server is not empty}] \\ &= L - (1 - P[0 \text{ jobs in the system}]) \\ &= L - (1 - \pi_0) \end{aligned}$$

Average waiting time Wq:

$$W_q = L_q / \lambda_{eff}$$

or

$$W_q = W - W_s = W - \frac{1}{\mu}$$

M/M/1/N: Performance Measures (Summary)

$$\alpha = \lambda / \mu$$

$$\pi_0 = \frac{1-\alpha}{1-\alpha^{N+1}},$$

$$\pi_n = \alpha^n \pi_0 = \frac{\alpha^n (1-\alpha)}{1-\alpha^{N+1}} \text{ for } n = 0,1,2,..., N$$

$$P_B = \pi_N = \alpha^N \pi_0 = \frac{\alpha^N (1-\alpha)}{1-\alpha^{N+1}}$$

$$\lambda_{eff} = \lambda(1 - P_B)$$

$$\rho = 1 - \pi_0 = 1 - \frac{1-\alpha}{1-\alpha^{N+1}},$$

$$\gamma = \lambda(1 - P_B) = \lambda(1 - \alpha^N \frac{1-\alpha}{1-\alpha^{N+1}}) = \frac{\lambda(1-\alpha^N)}{1-\alpha^{N+1}}$$

$$L = \frac{\alpha}{1-\alpha} - (N+1)\frac{\alpha^{N+1}}{1-\alpha^{N+1}}$$

$$W = L / \lambda_{eff} = \left[\frac{\alpha}{1-\alpha} - \frac{(N+1)\alpha^{N+1}}{1-\alpha^{N+1}}\right] / \lambda_{eff}$$

$$L_q = L - (1 - \pi_0)$$

$$W_q = L_q / \lambda_{eff} \text{ or } W_q = W - W_s = W - \frac{1}{\mu}$$

Approximation of a Finite-Buffer System by the Infinite Buffer Model

- For M/M/1 with an infinite buffer:

$$\pi_0 = 1 - \frac{\lambda}{\mu} = 1 - \rho = 1 - \alpha, \text{where} \quad \rho = \alpha = \frac{\lambda}{\mu}$$

$$\pi_n = \left(\frac{\lambda}{\mu}\right)^n \pi_0 = \rho^n (1-\rho) = \alpha^n (1-\alpha)$$

- For M/M/1/N with a finite buffer:

$$\pi_0 = \frac{1-\alpha}{1-\alpha^{N+1}}, \text{ where } \alpha = \frac{\lambda}{\mu}$$

$$\pi_n = \alpha^n \pi_0 = \frac{\alpha^n (1-\alpha)}{1-\alpha^{N+1}} \text{ for } n = 0,1,2,\ldots, N$$

For α =0.8 and N=16, these probabilities differ by less than 2.3%

For α =0.8 and N=32, the difference is only 0.06%

Therefore, the infinite buffer model M/M/1 is a very good approximation of a finite buffer system M/M/1/N, even for moderate buffer sizes.

Hands-On Problem

- Consider a computer system with one processor and a queue with 2 buffers. The job requests arrive to the processor at the rate of 16 requests per second with Poisson pattern. The time to service a job request at the processor is exponentially distributed with a mean of 50 milliseconds. Answer the following questions:
  - What is the probability of the processor being busy?
  - What is the effective arrival rate of job requests into the system?
  - Assume a job request in the queue and not being serviced can depart without service; this behavior is called "defect". Assume the defect process is also exponential with the constant rate of δ = 2 requests/second.
    - Draw the complete state-transition diagram for the queueing computer system with the above described defect behavior.
    - Pick any one state and write down the balance equation for that state.

Agenda

- More Birth-and-Death Queues
  - M/M/1/N Queues
  - **M/M/c Queues**
  - M/M/ ∞ Queues
  - M/M/1/k/k Queues

What if the system has 3 servers that share a common job queue with infinite capacity?

M/M/c Queues (I)

- Have c identical servers

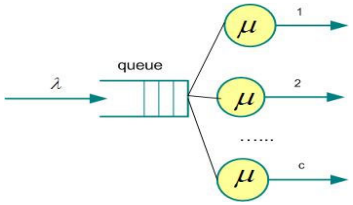

- A birth-and-death process with coefficients:

$$\lambda_n = \lambda, n = 0,1,2,3,.....$$
$$\mu_n = \begin{cases} n\mu, n = 1,2,..., c \\ c\mu, n \geq c \end{cases}$$

- Traffic intensity and server utilization:

$$\alpha = \lambda / \mu$$
$$\rho = \alpha / c = \lambda / c\mu$$

M/M/c Queues (II)

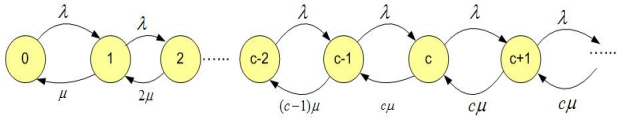

State	Rate out=Rate in
0	$\lambda\pi_0 = \mu\pi_1$
1	$(\lambda+\mu)\pi_1 = \lambda\pi_0 + 2\mu\pi_2$
2	$(\lambda+2\mu)\pi_2 = \lambda\pi_1 + 3\mu\pi_3$
…	……
c	$(\lambda+c\mu)\pi_c = \lambda\pi_{c-1} + c\mu\pi_{c+1}$
c+1	$(\lambda+c\mu)\pi_{c+1} = \lambda\pi_c + c\mu\pi_{c+2}$
…	……

By adding each two consecutive equations:

$$\lambda\pi_0 = \mu\pi_1$$
$$\lambda\pi_1 = 2\mu\pi_2$$
$$.....$$
$$\lambda\pi_{c-1} = c\mu\pi_c$$
$$\lambda\pi_c = c\mu\pi_{c+1}$$
$$\lambda\pi_{c+1} = c\mu\pi_{c+2}$$
$$.....$$

i.e.,

$$\pi_1 = \frac{\lambda}{\mu}\pi_0$$
$$\pi_2 = \frac{\lambda}{2\mu}\pi_1$$
$$..... ..$$
$$\pi_c = \frac{\lambda}{c\mu}\pi_{c-1}$$
$$\pi_{c+1} = \frac{\lambda}{c\mu}\pi_c$$
$$\pi_{c+2} = \frac{\lambda}{c\mu}\pi_{c+1}$$
$$..... ..$$

M/M/c (III): Solution to Steady-State Probabilities

- Thus:

$$\pi_n = \begin{cases} \frac{(\lambda/\mu)^n}{n!}\pi_0 = \frac{\alpha^n}{n!}\pi_0, & n = 0,1,2,...,c \\ \frac{(\lambda/\mu)^n}{c!\,c^{n-c}}\pi_0 = \frac{\alpha^n}{c!\,c^{n-c}}\pi_0, & n \geq c \end{cases}$$

- All probabilities have to sum up to one:

$$\sum_{n=0}^{\infty}\pi_n = 1$$

- Therefore:

$$\pi_0\left[\sum_{n=0}^{c-1}\frac{\alpha^n}{n!} + \sum_{n=c}^{\infty}\frac{\alpha^n}{c!\,c^{n-c}}\right] = 1 \Rightarrow \pi_0 = \left[\sum_{n=0}^{c-1}\frac{\alpha^n}{n!} + \sum_{n=c}^{\infty}\frac{\alpha^n}{c!\,c^{n-c}}\right]^{-1}$$

$$\therefore \pi_0 = \left[\sum_{n=0}^{c-1}\frac{\alpha^n}{n!} + \frac{\alpha^c}{c!(1-\alpha/c)}\right]^{-1}$$

M/M/c (IV): Performance Measures (I)

Average number of customers in the system:

$$L = \sum_{n=0}^{\infty} n\,\pi_n$$

Average response time:

$$W = L/\lambda$$

Average waiting time:

$$W_q = W - W_s$$

Average queue length:

$$L_q = \lambda W_q$$

Probability that at least one of the c servers is busy:

$$1 - \pi_0$$

The probability that an arriving customer has to wait for service

P[wait] ?

M/M/c(V): Performance Measures (II)

- A customer must queue for service *iff* there are c or more customers already in the system:

$$P[wait] = P[n \geq c] = \sum_{n=c}^{\infty} \pi_n = \pi_0 \left[\sum_{n=c}^{\infty} \frac{\alpha^n}{c!\, c^{n-c}} \right]$$

$$= \pi_0 \frac{\alpha^c}{c!(1-\alpha/c)} = \frac{\pi_c}{1-\alpha/c} = \frac{\pi_c}{1-\rho}$$

$$\therefore P[wait] = \frac{\alpha^c/c!}{(1-\rho)\left(\sum_{n=0}^{c-1}\frac{\alpha^n}{n!} + \frac{\alpha^c}{c!(1-\alpha/c)}\right)} = \frac{\alpha^c/c!}{(1-\rho)\sum_{n=0}^{c-1}\frac{\alpha^n}{n!} + \frac{\alpha^c}{c!}} = C[c,\alpha]$$

-- Erlang's C formula *or* Erlang's delay formula

- It can be shown that

$$W_q(t) = P[\text{time in queue q} \leq t] = 1 - C[c,\alpha]e^{-\mu(c-\alpha)}$$

and

$$W_q = E[q] = \frac{C[c,\alpha]W_s}{c(1-\rho)}, \quad L_q = \lambda W_q = \frac{C[c,\alpha]\rho}{(1-\rho)}$$

Hands-On Problem

- A storage system consists of two disk drives sharing a common queue with infinite capacity. The I/O requests arrive to the storage system at the rate of 40 requests per second with *Poisson* pattern. The time to service an I/O request at each disk drive is exponentially distributed with a mean of 45 milliseconds.
  - Draw the state transition diagram (show at least the first 4 states).
  - What is the probability that each disk drive is busy (i.e., average disk drive utilization)?
  - What is the probability that the entire storage system is idle?
  - What is the probability that both disk drives are busy and exactly two I/O requests are waiting in the queue.

Agenda

- More Birth-and-Death Queues
  - M/M/1/N Queues
  - M/M/c Queues
  - M/M/ ∞ Queues
  - M/M/1/k/k Queues

M/M/∞ Queues

- Have an infinite number of servers
  - A server is immediately provided for each arriving customer
- State transition diagram:

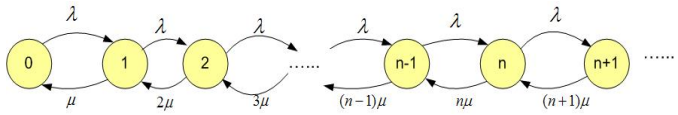

- State probabilities:

$$\pi_n = \frac{\left(\lambda/\mu\right)^n}{n!}\pi_0 = \frac{\alpha^n}{n!}\pi_0, n = 0,1,2,3,.....$$

$$\sum_{n=0}^{\infty}\pi_n = 1 \Rightarrow \pi_0\sum_{n=0}^{\infty}\frac{\alpha^n}{n!} = \pi_0 e^{\alpha} = 1$$

$$\therefore \pi_0 = e^{-\alpha}$$

$$\pi_n = e^{-\alpha}\frac{\alpha^n}{n!} \Rightarrow L = \alpha$$

M/M/∞ Queues: Example

- Calls in a telephone arrive randomly (Poisson) at an exchange at the rate of 140 per hour. If there is a very large number of lines available to handle the calls that last an average of 3 minutes, what is the average number of lines in use?

Agenda

- More Birth-and-Death Queues
  - M/M/1/N Queues
  - M/M/c Queues
  - M/M/ ∞ Queues
  - **M/M/1/k/k Queues**

Kendall Notation (review)

Standard notation for queueing systems:

A/B/c/K/m/Z

- **A**: arrival process or inter-arrival time distribution
  - 'M' = Poisson arrival process
  - 'D' = Deterministic (constant) arrival rate
  - 'G' = General arrival process
- **B**: service process or service time dist.
  - 'M' = Exponential service time dist.
  - 'D' = Deterministic (constant) service time
  - 'G' = General service time
- **c**: number of servers
- **K**: the capacity of the system (queue+server(s)) (default: ∞)
- **m**: total job/customer population (default: ∞)
- **Z**: scheduling discipline (default: FIFO)

M/M/1/k/k Queues (I)

- Machine repair model

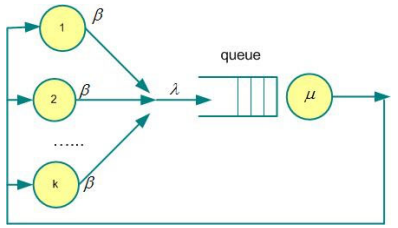

- There are only k customers (devices / machines)
- The machines break at a rate of β
  - The operating time between breakdowns: τ

$$\Pr[\tau \le t] = 1 - e^{-\beta t}$$

  - Mean operating time between breakdowns is *1/β*
- One server (repairman)
- State transition diagram?

M/M/1/k/k Queues (II)

- State transition diagram

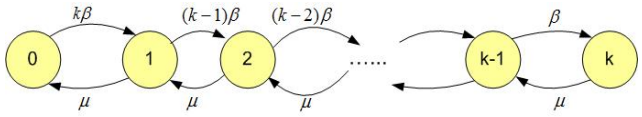

State	Rate out=Rate in
0	$k\beta\pi_0 = \mu\pi_1$
1	$[(k-1)\beta+\mu]\pi_1 = k\beta\pi_0 + \mu\pi_2$
…	……
k-1	$(\beta+\mu)\pi_{k-1} = 2\beta\pi_{k-2} + \mu\pi_k$
k	$\mu\pi_k = \beta\pi_{k-1}$

$$\Rightarrow \begin{cases} \pi_1 = \dfrac{k\beta}{\mu}\pi_0 \\ \pi_2 = \dfrac{(k-1)\beta}{\mu}\pi_1 \\ \cdots\cdots\cdots \\ \pi_{k-1} = \dfrac{2\beta}{\mu}\pi_{k-2} \\ \pi_k = \dfrac{\beta}{\mu}\pi_{k-1} \end{cases}$$

M/M/1/k/k Queues (III)

- State probabilities

$$\pi_n = \frac{(k-n+1)\beta}{\mu}\pi_{n-1} = \frac{(k-n+1)(k-n+2)\ldots k\beta^n}{\mu^n}\pi_0$$

$$= \frac{k!}{(k-n)!}\left(\frac{\beta}{\mu}\right)^n \pi_0 \qquad \text{n} = 0,1,2,\ldots\text{k}$$

$$\sum_{n=0}^{k}\pi_n = 1 \Rightarrow \pi_0\sum_{n=0}^{k}\frac{k!}{(k-n)!}\left(\frac{\beta}{\mu}\right)^n = 1$$

$$\therefore \pi_0 = \left[\sum_{n=0}^{k}\frac{k!}{(k-n)!}\left(\frac{\beta}{\mu}\right)^n\right]^{-1} = B[k,z]$$

called Erlang's B formula, where $z = \mu/\beta$.

- The server (repairman) utilization:

$$\rho = 1-\pi_0 = 1-B[k,z]$$

M/M/1/k/k Queues (IV): Performance Measures

- The average rate of customers to the queueing system λ

$$\rho = \alpha = \lambda / \mu = \lambda W_s \Rightarrow$$
$$\lambda = \rho / W_s = (1 - B[k, z]) / W_s$$

M/M/1/k/k Queues (V): Performance Measures

- Average time in the queue W_q:
  - Each customer (machine) has a complete cycle consisting of

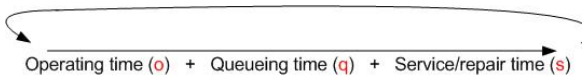

  - Average time to complete the cycle:

$$E[o] + E[q] + E[s] = 1/\beta + W_q + W_s$$

  - Overall "arrival" rate from k independent identical customer (average rate at which k machines break down and enter the queueing system) λ is:

$$\lambda = k / (E[o] + E[q] + E[s]) = k / (1/\beta + W_q + W_s)$$
$$\therefore W_q = k/\lambda - 1/\beta - W_s = k/\lambda - 1/\beta - 1/\mu$$

M/M/1/k/k Queues (VI): Performance Measures

- Average queue length L_q:

$$\begin{aligned} L_q &= \lambda W_q \\ &= \lambda(k/\lambda - 1/\beta - 1/\mu) \\ &= k - \lambda/\beta - \lambda/\mu \end{aligned}$$

- Average time in the system W:

$$\begin{aligned} W &= W_q + W_s \\ &= k/\lambda - 1/\beta - W_s + W_s \\ &= k/\lambda - 1/\beta \end{aligned}$$

- Average number of customers in the system L:

$$\begin{aligned} L &= \lambda W \\ &= \lambda(k/\lambda - 1/\beta) \\ &= k - \lambda/\beta \end{aligned}$$

M/M/1/k/k Queues (VI): Example

- Users of a time-sharing system have exponentially distributed think times with an average value of 20 seconds.
- Service time are exponentially distributed with a mean of 1.5 seconds.
- ***Q***: how many terminals can the system support if we want the average stretch factor is no more than 5?

Example Solution (Hint)

- Use M/M/1/k/k model

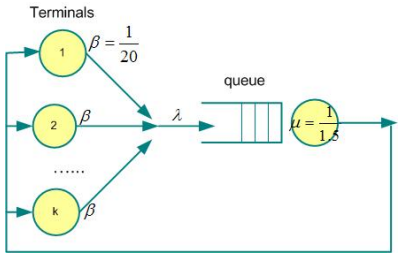

β = arrival rate per user = 1/20 per second

μ = service rate = 1/1.5 = 2/3 per second

Find maximum *k* satisfying W/Ws <= 5?

More in the textbook…

- M/M/c/c Queues
  - Also called the M/M/c loss systems because customers who arrive when all the c servers are busy are not allowed to wait for service and thus are lost to the system
  - Ch 5.2.4
- M/M/c/k/k Queues
  - Machine repair, multiple repairmen
  - Ch 5.2.7

Next Topics

- Embedded Markov Chain Queueing Systems

Things to Do

- Homework
- Keep working on the project