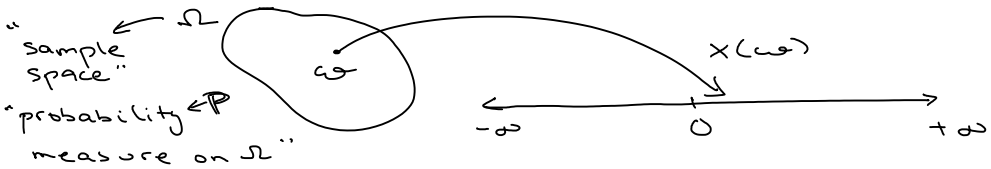
INTRODUCTION TO STOCHASTIC PROCESSES

Recall: An $\mathcal{S}$-valued random is a map from an underlying probability space to the set $\mathcal{S}$.

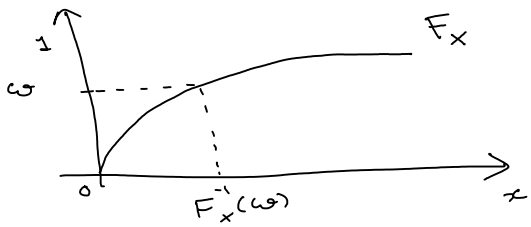
Example: a Real-valued R.V. X



Intuitively: Suppose we want to sample from the distribution $F_X(x) = \text{Prob}(X \le x)$ using MATLAB

Step 1: Generate $\omega = \text{unif}[0,1]$ using rand()

Step 2: Find $F_x^{-1}(\omega)$



$$\text{Prob}(X \le x) = \text{Prob}(\omega \le F_X(x)) = F_X(\omega)$$

Therefore in this case: $\Omega = [0,1]$

$\mathbb{P}$ = the uniform distribution (Lebesgue measure) on Ω

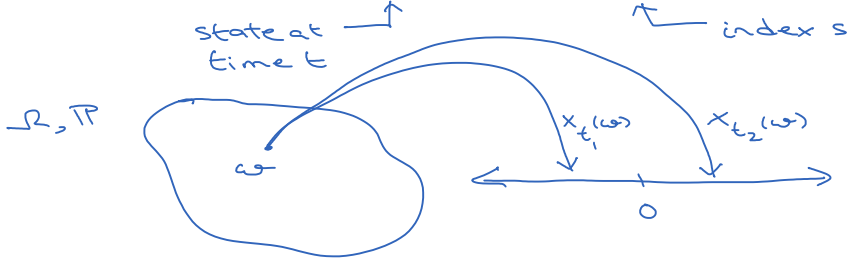
$X(\omega) = F_x^{-1}(\omega)$

Note: In reality, probability space has a third element: a σ-algebra $\mathcal{F}$ of subsets of Ω which define what events we can find the probabilities for. In this course we will ignore this level of rigor.

Defn: A stochastic process with state space $\mathcal{S}$ is a family of $\mathcal{S}$-valued random variables indexed by time $t \in T$

$$\{ X_t : t \in T \}$$

X_t: state at time t; T: index set



VIEW 1: X_t is an $\mathcal{S}$-valued random variable $\quad X_t : \Omega \to \mathcal{S}$

VIEW 2: $X(t,\omega)$ for a fixed ω is a "sample path" or a "realization"

$X(\cdot,\omega) : T \to \mathcal{S}$

4 cases for X_t:

$t \in T$ \ $X \in S$	S discrete	S - continuous
T - discrete	e.g. # jobs seen in system by n^{th} arrival	e.g. waiting time of n^{th} arrival
T - continuous	e.g. # jobs of class k at time t	e.g. unfinished workload at time t

Alternately: A stochastic process is a sequence-valued (discrete T) or a function-valued (continuous T) random variable.

The difference between a r.v. and a s.p. is in the questions

- Long-term averages
  - "Expected # jobs in system?"
  - "Average response time of jobs?"
- How will future depend on present
  - "Distribution of stock prices a month from now"
  - "How long until the queue empties"
- Likelihood of boundary / tail / rare events
  - "How often does buffer overflow"
  - "Probability of waiting more than 1 hour"

The Bernoulli Process

- Illustrates some basic classes of stoch. processes
- Motivates the Poisson process (Lect 2/3)

Defn: A Bernoulli process is a sequence of i.i.d. Bernoulli random variables $\{X_i\}$ with

$$Pr(X_i = 1) = p \quad \text{(arrival/success/Heads)}$$
$$Pr(X_i = 0) = 1-p$$

$\Rightarrow$ a discrete-time, discrete space stoch. process

a possible sample path: $X(\cdot, \omega) = (0, 0, 1, 0, 1, 1, 0, \ldots)$

Before we analyze the properties of Bernoulli Process, let us first think how may we construct the probability space $(\Omega, \mathbb{P})$ for it.

(For most of the course we will not mention the probability space and assume there is one in the background (except when we start comparing stoch. processes). Therefore this is a good place to build our understanding and then we will not bring it up again.)

Q: How could we simulate Bernoulli process sample paths in MATLAB?

Ans 1: Suppose we had access to a p-biased coin.

→ $\omega_i = 1$ if coin comes up Heads
0 otherwise

→ $X(i,\omega) = \omega_i$ (the identity map) $\omega = (0,1,0,0,..) \Rightarrow X(\cdot,\omega) = (0,1,0,0,..)$

So $\Omega = \{0,1\}^{\mathbb{N}}$ (sequences on 0,1)
$\mathbb{P}$ = the product measure of i.i.d. Bernoulli

Ans 2: Use the rand() function

→ ω_i = rand()

→ $X(i,\omega) = 1$ if $\omega_i \le p$
0 otherwise

(e.g.: $p = 0.5$
$\omega = (0.1, 0.33, 0.8, 0.2, ..) \Rightarrow X(\cdot,\omega) = (0,0,1,0,..)$)

So $\Omega = [0,1]^{\mathbb{N}}$ (sequences of $[0,1]$)
$\mathbb{P}$ = the product measure of i.i.d Unif$[0,1]$

Properties of Bernoulli Process

Q: What does it mean for an ω-sequence to be independent?

A: All finite subsets $(X_{i_1}, X_{i_2}, \dots, X_{i_k})$ are independent.

↳ These are called "Finite Dimensional Distributions" (f.d.d.)

Q: Let A_i be the time of i^{th} arrival. Distribution of A_1?

A: Geometric(p) : $\mathbb{P}(A_1 = n) = (1-p)^{n-1} p$

Q: Distribution of $N(t)$ = # of arrivals by time t

A: Binom(t,p) $\mathbb{P}(N(t) = n) = \binom{t}{n} p^n (1-p)^{t-n}$

Bernoulli process is an

(1) <u>independent process</u> : $P_r(X_{i_1} \le x_1, X_{i_2} \le x_2, \dots, X_{i_k} \le x_k)$
$= P_r(X_{i_1} \le x_1) P_r(X_{i_2} \le x_2) \dots P_r(X_{i_k} \le x_k)$

(2) <u>stationary process</u> :
$P_r(X_{i_1+t} \le x_1, X_{i_2+t} \le x_2, \dots) = P_r(X_{i_1} \le x_1, X_{i_2} \le x_2, \dots)$

(i.e. f.d.d. are invariant to shifting in time)

<u>Define</u>: $\tau_i = A_i - A_{i-1}$

τ_i: i^{th} interarrival time; A_i: i^{th} arrival time

<u>Q</u>: What is the distribution of the sequence $\{\tau_i\}$

<u>A</u>: i.i.d. Geom(p)

Therefore, Bernoulli processes define a <u>renewal process</u> with Geom(p) interarrival times.

<u>Defn</u>: A <u>Renewal process</u> $\{N(t), t \in T\}$ is a stochastic process with i.i.d. interarrival times X_i ; such that

$$N(t) = \max\left\{n : \sum_{i=1}^{n} X_i \le t\right\}$$

$N(t)$ counts the number of arrivals/renewals.

We have seen 3 classes of stoch. processes. Following is the list of other important stoch. processes we will encounter:

1. Independent processes
2. Stationary processes
3. Renewal processes
4. Markov processes (or Markov Chains when state space is discrete)
5. Semi-Markov processes
6. Random Walks

<u>Note</u>: These are not disjoint classes

<u>e.g.</u> All renewal processes are random walks

<u>e.g.</u> Markov processes can be stationary or non-stationary.

DISCRETE TIME MARKOV CHAINS (DTMCs)

the simplest and most important modeling tool

Defn: A discrete time Markov chain is a stochastic process $\{X_n : n \in \mathbb{Z}\}$ s.t.

$$P_r(X_n = i \mid X_{n-1} = j, X_{n-2} = x_{n-2}, \ldots, X_0 = x_0)$$
$$= P_r(X_n = i \mid X_{n-1} = j) \quad \text{(Markov property)}$$
$$\text{(usually)} \quad = P_{ij} \quad \text{(Stationarity / Time homogeneity)}$$

Defn. The matrix $P = [P_{ij}] = \begin{bmatrix} P_{11} & P_{12} & \cdots \\ P_{21} & P_{22} & \cdots \\ \vdots & & \end{bmatrix}$ is called the probability transition matrix.

Q: What is the n-step probability of $i \to j$ transition

$$P_{ij}^{(n)} = P_r(X_n = j \mid X_0 = i)$$

A: A few ways to do this depending on what time we condition on:

$$P_{ij}^{(n)} = \sum_k P_{ik} P_{kj}^{(n-1)} \quad \text{(Backward C-K eqn)}$$
$$= \sum_k P_{ik}^{(n-1)} P_{kj} \quad \text{(Forward C-K eqn.)}$$
$$= \sum_k P_{ik}^{(m)} P_{kj}^{(n-m)}$$

Chapman-Kolmogorov Eqns.

In matrix notation:

$$P^{(n)} = \left[P_{ij}^{(n)}\right] = P^n, \qquad \pi(n) = \pi(0) P^n$$

$\pi(n)$: distribution at $t=n$; $\pi(0)$: distribution at $t=0$

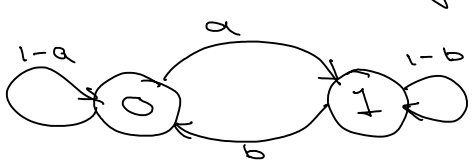
Example: A two-state DTMC

$X_n = 1$ if it rains on day n

$= 0$ does not rain on day n

$P_{01} = a$

$P_{10} = b$



$$P = \begin{bmatrix} 1-a & a \\ b & 1-b \end{bmatrix}$$

Easy to show via induction:

$$P^{(n)} = P^n = \begin{bmatrix} \frac{b + a(1-a-b)^n}{a+b} & \frac{a - a(1-a-b)^n}{a+b} \\ \frac{b - b(1-a-b)^n}{a+b} & \frac{a + b(1-a-b)^n}{a+b} \end{bmatrix}$$

In this course, we will primarily concern ourselves with the <u>limiting</u> or <u>long-term</u> behavior of systems:

$$\lim_{n \to \infty} P^{(n)}$$

For this example: if $|1-a-b| < 1$

$$\lim_{n \to \infty} P^{(n)} = \begin{bmatrix} b/a+b & a/a+b \\ b/a+b & a/a+b \end{bmatrix}$$ ← rows are identical

Q: What does this say?

A: ① There is a "limiting probability" of being in each state

② This limiting probability is independent of the starting state.

For the rest of this lecture we will be interested in answering

Q: What properties of P guarantee

(1) $\lim_{n \to \infty} P_{ij}^{(n)}$ exist

(2) $\lim_{n \to \infty} P_{ij}^{(n)} = \pi_j$ (i.e., independent of i)

(3) " is non-zero $\forall j$

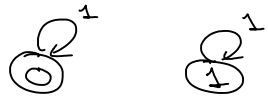
First we will consider finite state DTMCs

Then infinite state DTMCs

Limiting Behavior of Finite State DTMCs

Let us consider what can go wrong

Example 1



Q: Do $\lim_{n\to\infty} P_{ij}^{(n)}$ exist? Ans: Yes

Q Independent of start state Ans: No

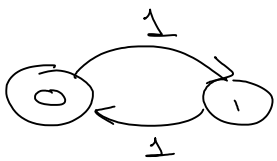
In this example, the DTMC is not irreducible.

Defn: State j is accessible from state i if $P_{ij}^{(n)} > 0$ for some n.

Defn: States i & j communicate ($i \leftrightarrow j$) if they are accessible from each other

Defn: A DTMC is irreducible if all states communicate with each other.

Example 2



Q Do $\lim_{n\to\infty} P_{ij}^{(n)}$ exist? Ans: No

In the above example, the DTMC is not aperiodic.

But example 2 is not as bad as example 1. Consider a weaker notion: stationary distribution

Defn: $\pi = (\pi_1, \pi_2 \ldots \pi_n)$ is said to be a stationary distribution for DTMC P if

$$\pi = \pi \cdot P \quad \text{and} \quad \sum_j \pi_j = 1$$

Equiv: $\pi_j = \sum_i \pi_i P_{ij}$

Q: Why "stationary"?

A: If we choose $X_0 \sim \pi$ and take one step acc. to P, then $X_1 \sim \pi$.

Another way to define stationary distrib is through the time-average or Cesàro-average

$$P_j = \lim_{t\to\infty} \frac{1}{t} \sum_{n=0}^{t} \mathbb{1}_{\{X_t = j\}}$$

For the DTMC in example 2, $[\frac{1}{2}\ \frac{1}{2}]$ is the unique stationary dist.

Thm: For a finite state DTMC

(1) irreducibility $\Rightarrow$ unique stationary distrib.

(2) irreducibility + aperiodicity $\Rightarrow$ limiting distrib exists & limiting dist. $\equiv$ stationary dist.

I omit a formal proof, please refer to any standard text.

(1) can be shown by noting that the stationary dist. is the left-eigenvector of P with eigenvalue 1.

$\Rightarrow$ need to show: irreducibility $\Rightarrow$ the multiplicity of 1 as an eigenvalue is 1

$\Rightarrow$ argue that $\begin{bmatrix}1\\ \vdots \\ 1\end{bmatrix}$ is the unique right eigenvector with eigenvalue 1.

(2) can be shown by first arguing that (aperiodicity + irreducib.) $\Rightarrow \exists n : P^{(n)}_{ij} > \varepsilon$ for some $\varepsilon > 0$ $\forall i,j$

Now argue that columns of $P^{(n)}$ will converge to scaled unit vectors (the right eigenvector $\begin{bmatrix}1\\ \vdots \\ 1\end{bmatrix}$) geometrically.

Defn: A DTMC is said to be in "steady state" if the initial state is chosen acc. to the stationary dist.

Note: Recall the equation for stationary distrib:

$$\pi_j = \sum_i \pi_i P_{ij}$$

This equation is valid for any process; not just DTMCs.

Let $N_j(t)$ = # transitions out of state j in t steps

$N_{ij}(t)$ = # $i \to j$ transitions in t steps

Q: How large can $\left| N_j(t) - \sum_i N_{ij}(t) \right|$ be?

Ans: At most 1. Can not exit j twice without an $i \to j$

So $$-1 \le N_j(t) - \sum_i N_{ij}(t) \le 1$$

$$\Rightarrow \quad -\frac{1}{t} \le \frac{N_j(t)}{t} - \sum_i \left(\frac{N_{ij}(t)}{N_i(t)}\right)\cdot\left(\frac{N_i(t)}{t}\right) \le \frac{1}{t}$$

If: $\lim_{t\to\infty} \frac{N_i(t)}{t} = \pi_i$ and $\lim_{t\to\infty} \frac{N_{ij}(t)}{N_i(t)} = P_{ij}$ a.s.

Then: $$\pi_j = \sum_i \pi_i P_{ij}$$

This is an example of a "Sample path Law" { holds regardless of probabilistic assumptions

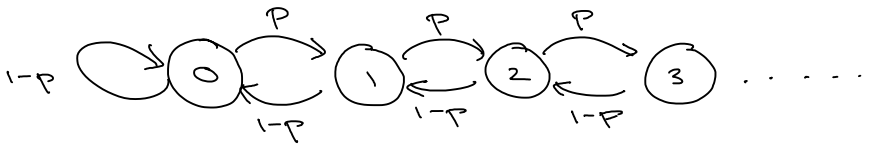
Remark: Instead of choosing state j, we can choose any partition A, A^c of the state space.

Then $$\left| \#_{\text{transitions}}\, A \to A^c - \#_{\text{transitions}}\, A^c \to A \right| \le 1.$$

gives $$\sum_{i\in A} \pi_i \sum_{j\in A^c} P_{ij} = \sum_{j\in A^c} \pi_j \sum_{i\in A} P_{ji}$$

Limiting Behavior of Infinite-State DTMCs

Q: Are aperiodicity and irreducibility enough for ∞-state?
Consider the following (aperiodic + irreducible) DTMC:



Do limiting probs exist if

1) $p > 1/2$ → NO, we "drift off" to $+\infty$
2) $p = 1/2$?
3) $p < 1/2$ → YES

We can solve for the stationary distrib.

$$\pi_0 \cdot p = \pi_1 \cdot (1-p) \quad \Rightarrow \pi_1 = \pi_0 \left(\tfrac{p}{1-p}\right)$$

$$\pi_1 \cdot p = \pi_2 (1-p) \quad \Rightarrow \pi_2 = \pi_1 \left(\tfrac{p}{1-p}\right) = \pi_0 \left(\tfrac{p}{1-p}\right)^2$$

$$\pi_i \cdot p = \pi_{i+1} (1-p) \quad \Rightarrow \pi_{i+1} = \pi_i \left(\tfrac{p}{1-p}\right) = \pi_0 \left(\tfrac{p}{1-p}\right)^{i+1}$$

Find π_0 using $\sum_i \pi_i = 1 \Rightarrow \pi_0 = \underbrace{\left[1 + \frac{p}{1-p} + \left(\frac{p}{1-p}\right)^2 + \dots\right]^{-1}}$

sum diverges if $p \geq 1/2$
$\Rightarrow \pi_j = 0 \ \forall j$

To understand ∞-state DTMCs, we need to understand recurrence probabilities and recurrence times.

Defn: Recurrence probability for state j

$$f_j = P_j(\text{eventually return to state } j \mid \text{start in } j)$$

Defn: If $f_j = 1$ then state j is recurrent
If $f_j < 1$ then state j is transient

Q: How do we test for recurrence / transience?

A: $\mathbb{E}[\#\text{visits to state } j] = \infty$ if j recurrent
$= \frac{1}{1-f_j} < \infty$ if j transient

But, also

$$\mathbb{E}[\#\text{visits to } j] = \mathbb{E}\left[\sum_n \mathbb{1}_{\{X_n=j\}} \mid X_0=i\right]$$
$$= \sum_n \mathbb{E}\left[\mathbb{1}_{\{X_n=j\}} \mid X_0=i\right]$$
$$= \sum_n p_{ij}^{(n)}$$

Thm: If state j is transient then $\sum_n p_{jj}^{(n)} < \infty$
If state j is recurrent then $\sum_n p_{jj}^{(n)} = \infty$

Defn: The mean recurrence time of state j is the expected time between consecutive visits to j.

Formally: if $f_{jj}^{(n)} = \Pr[\text{first visit to } j \text{ at } t=n \mid X_0 = j]$

then

"mean recurrence time" or "hitting time" $\rightarrow$ $h_{jj} = \sum_{n=1}^{\infty} n f_{jj}^{(n)}$

Let $N_j(t) = \#\text{visits to } j \text{ in } t \text{ steps}$

$N_j(t)$ is a renewal process { a process that counts events with i.i.d. interarrival times

THM (Elementary Renewal Theorem): Let $N(t)$ be a renewal process with i.i.d. interarrival times X_i with mean $\mathbb{E}[X]$.

Then

(1) $\lim_{t\to\infty} \frac{N(t)}{t} = \frac{1}{\mathbb{E}[X]}$ w.p. 1

(2) $\lim_{t\to\infty} \frac{\mathbb{E}[N(t)]}{t} = \frac{1}{\mathbb{E}[X]}$

Proof sketch: Assume $\mathbb{E}[X] < \infty$.

(i) $N(t) \to \infty$ as $t \to \infty$ w.p. 1

(ii) Define $S_n = \sum_{i=1}^{n} X_i$

$\lim_{n\to\infty} \frac{S_n}{n} = \mathbb{E}[X]$ w.p. 1 { $\Pr(\{\omega : \frac{S_n}{n} \to \frac{1}{\mathbb{E}[X]}\}) = 1$

(iii) $S(N(t)) \le t \le S(N(t)+1)$

epoch immed. preceding t ↳ epoch immed following t

$$\Rightarrow \underbrace{\frac{S(N(t))}{N(t)}}_{\to \mathbb{E}[X] \text{ by (ii)}} \le \frac{t}{N(t)} \le \underbrace{\frac{S(N(t)+1)}{N(t)+1}}_{\to \mathbb{E}[X] \text{ by (ii)}} \cdot \underbrace{\frac{N(t)+1}{N(t)}}_{\to 1 \text{ by (i)}}$$

$$\Rightarrow \frac{N(t)}{t} \to \frac{1}{\mathbb{E}[X]} \quad \text{w.p. 1}$$

- When $\mathbb{E}[X] = \infty$, the theorem can be proved by considering $X_i' = \min\{X_i, c\}$ & letting $c \to \infty$
- Proving (2): $\lim_{t\to\infty} \frac{\mathbb{E}[N(t)]}{t} = \frac{1}{\mathbb{E}[X]}$

needs a bit more work. It does not follow from (1). ▩

Returning to recurrence times, from the Elt. renewal thm.

$$\pi_j = \lim_{n\to\infty} P_{ij}^n = \lim_{t\to\infty} \frac{N_j(t)}{t} = \frac{1}{h_{jj}}$$

$$\Rightarrow \pi_j = \frac{1}{h_{jj}}$$

Defn: A recurrent state j is either:

positive recurrent if $h_{jj} < \infty$

null recurrent if $h_{jj} = \infty$

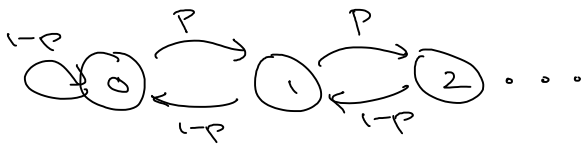
Thm: Transience / Pos. Recurrence / Null Recurrence are class properties.

That is, if $i \leftrightarrow j$ (i & j communicate)

either (1) both i, j are transient

(2) both i, j are positive recurrent

(3) both i, j are null recurrent

So 

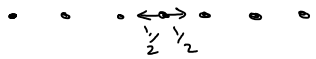
(1) $p > 1/2$

(2) $p < 1/2$

(3) $p = \frac{1}{2}$

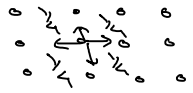
Interesting fact:

(1) the uniform 1-D infinite random walk



is null recurrent.

(2) the uniform 2-D infinite random walk



is also null recurrent

(3) the uniform 3-D infinite random walk is transient!

Returning to our original question:

THM: An irreducible, aperiodic DTMC is either

(1) Transient/null recurrent, and $\lim_{n\to\infty} P_{ij}^{(n)} = 0 \quad \forall j.$

(2) Positive recurrent, and $\lim_{n\to\infty} P_{ij}^{(n)} = \pi_j > 0$

& π is the unique stationary distribution

Defn: A Markov chain is ergodic if it is

- aperiodic,
- irreducible, and
- positive recurrent.

Remark: For a finite state DTMC, positive recurrence is guaranteed by irreducibility.

An important use of Ergodicity

Q: Suppose we want to estimate the stationary mean #jobs in our queing system. How do we do this?

Ans: solve $\pi = \pi P$ $\quad (\pi_j = \lim_{n\to\infty} P_{ij}^{(n)})$

$$\mathbb{E}[N] = \sum_n n\,\pi_n$$

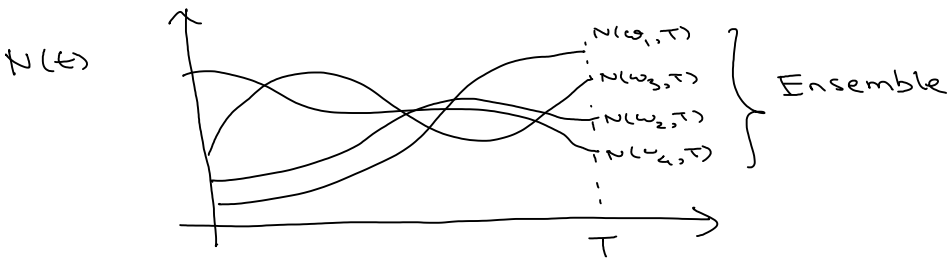
Q: What is we can't solve for the stationary dist?

Ans: Simulation. Let's explore 2 ways...

Method 1: Ensemble Averages

1. Pick an arbitrary initial state (say empty system)
2. Simulate for a "large enough" time T
3. Look at the final state $N(\omega, T)$ { ω: the seq. of coin tosses

This is just one sample. Repeat n times to get n independent samples $N(\omega_1, T), N(\omega_2, T) \cdots, N(\omega_n, T)$



As $n \to \infty$ $\mathbb{E}[N(T)] = \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} N(\omega_i, T)$ is called the ensemble average.

i.i.d. samples $\Rightarrow$ can generate confidence intervals.

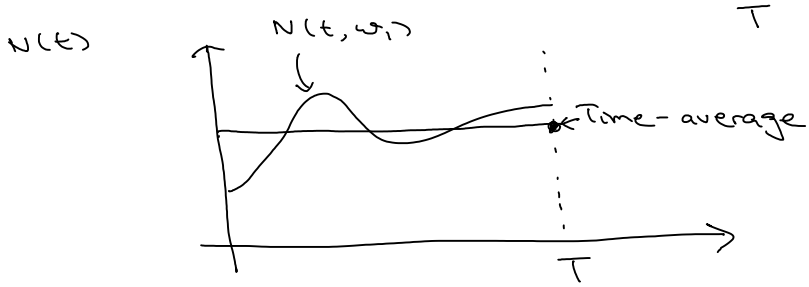
As $T \to \infty$ $\overline{N}^{\text{ensemble}} = \lim_{T\to\infty} \mathbb{E}[N(T)] = \sum_i i\pi_i$

Q: Suppose we only want to estimate $\overline{N}^{\text{ensemble}}$. Is there an easier method?

Method 2: Time-average

1. Pick an arbitrary initial state
2. Simulate for a "large enough" time T
3. Report

$$\overline{N}^{\text{time-avg}}(T, \omega_1) = \frac{1}{T} \int_0^T N(t, \omega_1)\, dt$$



We would like to say:

$$\lim_{T\to\infty} \overline{N}^{\text{time-average}}(T, \omega) = \mathbb{E}[N] \quad \text{w.p.1}$$

THM: For an ergodic system the ensemble average

$$\bar{N}^{\text{ensemble}} = \lim_{T \to \infty} \mathbb{E}[N(T)]$$

exists, and

$$\bar{N}^{\text{time-average}} \xrightarrow{\text{a.s.}} \bar{N}^{\text{ensemble}}.$$

That is, for almost all sample paths or realizations, the time average along that sample path converges to the ensemble average. □

Remark: The above is in fact how ergodicity is usually defined. For DTMCs, irreducibility + aperiodicity + pos. recurrence guarantee ergodicity.

Remark: We said "large enough" T. What we meant was T large enough so the system reaches steady state
⇒ $N(\omega, T)$ is a sample from stationary-distrib.

Q: How can we guarantee that T was chosen so we get a stationary sample

A: We can't. We can analyze the mixing time of the DTMC to ensure it is sufficiently close.

BUT: There is a fascinating field called PERFECT SIMULATION and a technique called COUPLING FROM THE PAST. Here we start multiple "coupled" simulations (roughly speaking) until at time T all the simulations reach the same state.

The coupling process then guarantees that this common state is a sample from the stationary distribution!

OPERATOR VIEW OF TRANSITION MATRIX P

We saw that P acts "from the right" to evolve probabilities:

$$\pi(0) P^n = \left[\sum_i \pi_i(0) P^{(n)}_{i1} \quad \sum_i \pi_i(0) P^{(n)}_{i2} \quad \sum_i \pi_i(0) P^{(n)}_{i3} \cdots \right]$$

$$\doteq \left[\pi_1(n) \quad \pi_2(n) \quad \pi_3(n) \cdots \right]$$

$$\doteq \pi(n)$$

There is also a view in which P "acts from the left" to evolve functions on $\mathcal{S}$.

Consider: $f : \mathcal{S} \to \mathbb{R}$ (function from states to $\mathbb{R}$)

i.e. $f = \begin{bmatrix} f(1) \\ f(2) \\ \vdots \end{bmatrix}$

Q: What is $P^{(n)} f$?

A: Look at the first entry: $\begin{bmatrix} x \\ ? \end{bmatrix} = \begin{bmatrix} P^{(n)}_{11} & P^{(n)}_{12} \cdots \\ & P^{(n)} \end{bmatrix} \begin{bmatrix} f(1) \\ f(2) \\ \vdots \end{bmatrix}_{f}$

$$x = \sum_i P^{(n)}_{1i} f(i)$$

$$= \mathbb{E}\left[f(X_n) \mid X_0 = 1 \right]$$

$\Rightarrow P^{(n)} f$ evolves the expected value of f given the start state

Also written as $\underbrace{(P^{(n)} f)(i)}_{\text{a map } \mathcal{S} \to \mathbb{R}} \doteq \mathbb{E}\left[f(X_n) \mid X_0 = i \right]$

THE EXPONENTIAL DISTRIBUTION

- the mainstay of queueing theory (model for job sizes, interarrival times)
- memoryless property of Exp dist allows modeling using discrete state space

If X is distributed as Exponential(λ) ← "rate"

density: $f_X(x) = \lambda e^{-\lambda x}$

dist. fn: $F_X(x) = \mathbb{P}_r(X \le x) = 1 - e^{-\lambda x}$

complementary cdf: $\bar{F}_X(x) = e^{-\lambda x}$

mean: $\mathbb{E}[X] = 1/\lambda$

variance: $Var[X] = 1/\lambda^2$

Defn: Squared coefficient of variation (SCV) for a r.v. X

$$C^2(X) = \frac{Var(X)}{(\mathbb{E}[X])^2}$$

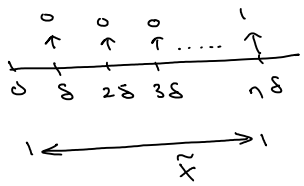
(a "scale-invariant" measure of variability)

For $X \sim Exp(\lambda)$: $C^2(X) = 1$

Intuition: Exponential dist is the continuous analog of the Geometric dist.

Consider: For $i = 1, 2, 3, \ldots$, we flip a coin at times $t = i\delta$ with success prob $p = (\lambda\delta)$

$\tilde{X}$ = time until first success



$$\mathbb{P}_r(\tilde{X} \ge t) = \mathbb{P}_r(\tilde{X} \ge \delta(t/\delta))$$
$$= \mathbb{P}_r(Geom(\lambda\delta) \ge t/\delta)$$
$$= (1-\lambda\delta)^{t/\delta} \xrightarrow{\delta \to 0} e^{-\lambda t}$$

It will be very helpful to think Exponential as a sequence of coin flips with δ infinitesimally small.

Properties of Exp dist

1. Memorylessness

$$Pr[X \geq s+t \mid X \geq t] = \frac{e^{-\lambda(s+t)}}{e^{-\lambda t}} = e^{-\lambda s} = Pr[X \geq s]$$

Meaning: If a job is distributed $Exp(\lambda)$ and has already been in service for 5 sec; its <u>remaining size</u> is still $\sim Exp(\lambda)$

2. Minimum of independent exponential r.v.s.

Let $X_1 \sim Exp(\lambda_1)$

$X_2 \sim Exp(\lambda_2)$

$X_1 \perp X_2$ { X_1, X_2 are independent

$Y \doteq \min(X_1, X_2)$

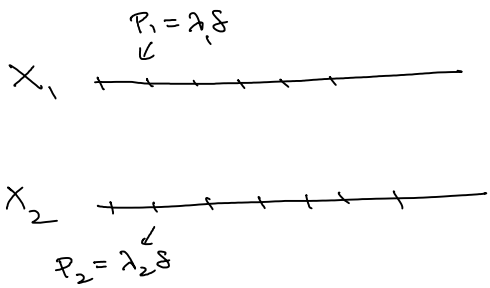
(i) $Y \sim Exp(\lambda_1 + \lambda_2)$

(ii) $Pr(Y = X_1) = \frac{\lambda_1}{\lambda_1+\lambda_2}$

<u>Pf:</u> $Pr[Y \geq t] = Pr[X_1 \geq t, X_2 \geq t] = Pr[X_1 \geq t]\, Pr[X_2 \geq t]$

$= e^{-\lambda_1 t}\, e^{-\lambda_2 t}$

$= e^{-(\lambda_1+\lambda_2)t}$



For (ii) think of coin flipping

for $\delta << 1$;

$Pr(\text{both } X_1, X_2 \text{ flip H}) = \lambda_1\lambda_2\delta^2$

$Pr(X_1 \text{ flip H} \mid \text{either } X_1 \text{ or } X_2 \text{ flip H})$

$$= \frac{\lambda_1\delta}{\lambda_1\delta + \lambda_2\delta - \lambda_1\lambda_2\delta^2} \xrightarrow{\delta \to 0} \frac{\lambda_1}{\lambda_1+\lambda_2}$$

LAPLACE and *z*- TRANSFORMS

- Transforms capture information about the moment sequence of distributions in a succinct form
- Often in our analysis, we will need to add r.v.s

$$C = A + B \quad \Rightarrow \quad F_C(x) = \int_{y=0}^{x} F_A(x-y)\, dF_B(y)$$

Convolution: $f_C = f_A * f_B$

This is messy in "distribution world" but in "transform world" this becomes pointwise multiplication.

- Preferred strategy: convert Chapman-Kolmogorov or stochastic recurrences into transforms $\Rightarrow$ solve for transform of # in system, or waiting time
$\Rightarrow$ recognize the distrib from transform / numerically invert.

Laplace Transform (or, Laplace-Stieltjes transform)

For a non-negative r.v. X, & for $s > 0$:

$$\mathcal{L}_X(s) = \mathbb{E}[e^{-sX}] = \int_0^\infty e^{-sx}\, dF_X(x)$$

Example: $X \sim \exp(\lambda)$

$$\mathcal{L}_X(s) = \int_0^\infty e^{-sx} \lambda e^{-\lambda x}\, dx = \lambda/(\lambda+s)$$

Properties of Laplace transform:

(i) $\mathcal{L}_X(0) = 1$

- This is a very useful sanity check for calculations.
- Also used to solve for unknown constants.

(ii) Finding moments.

Consider:

$$\frac{d}{ds}\mathcal{L}_X(s) = \int_0^\infty -x\, e^{-sx}\, dF_X(x)$$

$$\overset{s=0}{=} -\mathbb{E}[X]$$

More generally: $\left.\frac{d^n}{ds^n}\mathcal{L}_X(s)\right|_{s\to 0^+} = (-1)^n\, \mathbb{E}[X^n]$

(iii) Let $Z = X + Y$ & $X \perp Y$

$$\mathcal{L}_Z(s) = \mathbb{E}[e^{-sZ}] = \mathbb{E}[e^{-sX} e^{-sY}] = \mathbb{E}[e^{-sX}]\mathbb{E}[e^{-sY}]$$
$$= \mathcal{L}_X(s)\,\mathcal{L}_Y(s)$$

(iv) Let $Y \sim \text{Bernoulli}(p)$

$$Z = X_1 \mathbb{1}_{\{Y=1\}} + X_2 \mathbb{1}_{\{Y=0\}}$$

$$= X_1 \cdot Y + X_2(1-Y)$$

(X_1, X_2 need not be independent)

Then $\mathcal{L}_Z(s) = p\,\mathcal{L}_{X_1}(s) + (1-p)\,\mathcal{L}_{X_2}(s)$

Pf: $\mathcal{L}_Z(s) = \mathbb{E}_{X_1,X_2,Y}\left[e^{-sX_1 Y} e^{-sX_2(1-Y)}\right]$

$$= \mathbb{E}_{X_1,X_2}\left[\mathbb{E}_Y\left[e^{-sX_1 Y} \cdot e^{-sX_2(1-Y)}\right]\right]$$

$$= \mathbb{E}_{X_1,X_2}\left[p\,e^{-sX_1} + (1-p)\,e^{-sX_2}\right]$$ ↙ Linearity of expectation

$$= p\,\mathcal{L}_{X_1}(s) + (1-p)\,\mathcal{L}_{X_2}(s)$$

The z-transform

- for discrete r.v. conventional to use z-transform instead of Laplace(-Stieltjes) transform

For non-negative discrete r.v. X & $z \in \mathbb{C}$ (↖ complex plane), $|z| \le 1$

$$\hat{G}_X(z) = \mathbb{E}[z^X] = \sum_{n=0}^{\infty} p_n z^n$$

Example: $X \sim \text{Geom}(p)$

$$\hat{G}_X(z) = p \sum_{n=1}^{\infty} (1-p)^{n-1} z^n = \frac{pz}{1-(1-p)z}$$

Properties: * $\hat{G}_X(1) = 1$ (analog of $\mathcal{L}_X(0) = 1$)

* $\hat{G}_X(0) = P_r(X=0)$ (Laplace analog: $\mathcal{L}_X(\infty) = P_r(X=0)$)

* Moments from z-transform

Consider $\frac{d}{dz} G_X(z) = \sum_{n=0}^{\infty} p_n (n z^{n-1})$

(assuming we can differentiate inside summation)

$$\Rightarrow \quad G'_X(z)\Big|_{z=1} = \mathbb{E}[X]$$

More generally: $G_X^{(k)}(z) = \sum_{n=0}^{\infty} p_n \big(n(n-1)\cdots(n-k+1)\big) z^{n-k}$

$$\Rightarrow \quad G_X^{(k)}(z)\Big|_{z=1} = \mathbb{E}\big[X(X-1)(X-2)\cdots(X-k+1)\big]$$

Example: $X \sim \text{Geom}(p)$

$$G_X(z) = \sum_{n=1}^{\infty} p(1-p)^{n-1} z^n = pz\left(1 + (1-p)z + (1-p)^2 z^2 + \cdots\right)$$

$$= \frac{pz}{1-(1-p)z}$$

THE POISSON PROCESS

- most results in queueing theory rely on Poisson process as model of arrival process of jobs (there are some recent work on self-similar or long-range dependent arrivals)
- One justification is the

Palm-Khintchine Thm: If we merge n i.i.d. renewal processes with rate $\frac{\lambda}{n}$; then as $n \to \infty$ the aggregate traffic is Poisson (λ).

(the actual result is more general and only needs the renewal processes to be independent.)

Some definitions:

A stochastic process $\{N(t), t \in T\}$

(i) is a counting process if N makes jumps of size 1

(ii) has independent increments if $(N(t+s) - N(s))$ is indepen. of $(N(u), u \le s)$

(iii) has stationary increments if $(N(t+s) - N(s))$ is invariant to s.

Defn 1: A Poisson process with rate λ $(PP(\lambda))$ is a counting process with

(i) $N(0) = 0$

(ii) stationary increments

(iii) independent increments

(iv) # arrivals in an interval of length follows the Poisson(λt) distrb:

$$Pr\left(N(t) = n\right) = e^{-\lambda t}(\lambda t)^n / n!$$

Note: $\mathbb{E}[N(t)] = \lambda t$; hence λ is the "rate"

Definition 1 is useful for proving a stoch process is PP or for generalizing to other arrival processes (as we will see later)

Definition 2 below is useful to build intuition, and for simulating a $PP(\lambda)$ arrival stream.

Defn 2: A Poisson process with rate λ is a counting process with i.i.d. $Exp(\lambda)$ interarrival times.

Claim: Definitions 1 & 2 are equivalent

Proof: Defn 1 $\Rightarrow$ Defn 2.

Let Z_i be the i^{th} interarrival time

$P_\lambda(Z_1 > t) = P_\lambda(N(t)=0) = e^{-\lambda t} \Rightarrow Z_1 \sim Exp(\lambda)$

$$P_\lambda(Z_2 > s+t \mid Z_1 \in [s, s+h]) \overset{h<<1}{\sim} P_\lambda(N(s+t)=1 \mid N(s+h)=1)$$
$$= P_\lambda(N(t+h)=0)$$
$$\overset{h<<1}{\sim} e^{-\lambda t} \Rightarrow Z_2 \perp Z_1 ; Z_2 \sim Exp(\lambda)$$

Defn 2 $\Rightarrow$ Defn 1

We will use the coin flipping analogy for $Exp(\lambda)$

Divide $[0,t]$ into t/δ disjoint, exhaustive intervals of length δ each.



$P_\lambda(\geq 2 \text{ arrivals in a } \delta\text{-interval}) = o(\delta)$

$P_\lambda(\geq 2 \text{ arrivals in any } \delta \text{ interval}) \leq \frac{t}{\delta} \cdot o(\delta) = o(1)$

a function $f(\delta)$ is "small oh" of δ: $f(\delta) \in o(\delta)$ if $\lim_{\delta \to 0} \frac{f(\delta)}{\delta} = 0$

So via defn 2: with Probability $1 - o(1)$: $N(t) \sim Bin(t/\delta, \lambda\delta)$

Now use Stirling's approximation: $\lim_{t \to \infty} \frac{t!}{\left(\frac{t}{e}\right)^t \sqrt{2\pi t}} = 1$

to show $Bin(t/\delta, \lambda t) \xrightarrow{D} Poisson(\lambda t)$

$\hookrightarrow$ convergence in distribution

□

There is also a third equivalent, although seemingly much weaker, definition:

Defn 3: A Poisson process with rate λ is a counting process with:

i) $N(0) = 0$

ii) independent increments

(iii) stationary increments

(iv) $P_\lambda(N(\delta)=1) = \lambda\delta + o(\delta)$

$P_\lambda(N(\delta)\geq 2) = o(\delta)$

{ "orderliness" or arrivals one at a time.

Properties of Poisson processes

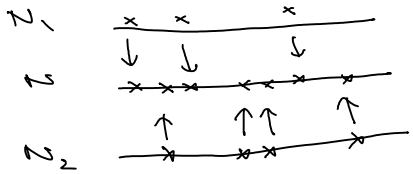
1. POISSON MERGING

$N_1(t) \sim PP(\lambda_1)$, $N_2(t) \sim PP(\lambda_2)$, $N_1 \perp N_2$

then the merged process

$N(t) = N_1(t) + N_2(t)$

is a Poisson process with rate $(\lambda_1+\lambda_2)$.



Pf 1: $Z_1 \sim Exp(\lambda_1)$, $Z_2 \sim Exp(\lambda_2)$

Z for merged process $\sim \min\{Z_1, Z_2\}$

$\sim Exp(\lambda_1+\lambda_2)$

and after every event of N, both N_1, N_2 "restart"

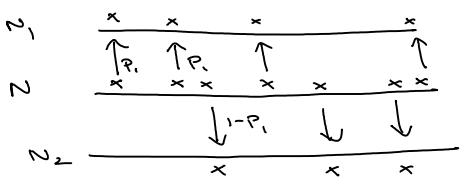
Pf 2: $N_1(t) \sim Poisson(\lambda_1 t)$; $N_2(t) = Poisson(\lambda_2 t)$

& $Poisson(\lambda_1 t) + Poisson(\lambda_2 t) \sim Poisson((\lambda_1+\lambda_2)t)$ □

2. POISSON SPLITTING

Thm: Given $N(t) \sim PP(\lambda)$; if (independently) each event is marked type 1 with prob. p_1 & type 2 with $p_2 = 1-p_1$

then $N_1 \sim PP(\lambda p_1)$; $N_2 \sim PP(\lambda p_2)$ & $N_1 \perp N_2$.



Pf: Intuitively with coin flipping interpretation; in a δ-interval N comes heads with prob $(\lambda\delta)$ & then marked as class 1 with prob p_1. ⇒ success prob. for N_1 is $(\lambda p_1 \delta)$

OR: interarrival times for N_1 ~ sum of Geom(p) many i.i.d. Exp(λ) r.v.s

$= Exp(\lambda p_1)$ ← homework

to verify independence, we revert to definition 1.

$$P_{\mathbb{Q}}\left(N_1(t)=n_1, N_2(t)=n_2\right) = \sum_n P_{\mathbb{Q}}\left(N_1(t)=n_1, N_2(t)=n_2 \mid N(t)=n\right) P_{\mathbb{Q}}(N(t)=n)$$

$$= \binom{n_1+n_2}{n_1} p_1^{n_1} p_2^{n_2} \; e^{-\lambda t} \frac{(\lambda t)^{n_1+n_2}}{(n_1+n_2)!}$$

$$= \left(e^{-\lambda p_1 t} \frac{(\lambda p_1 t)^{n_1}}{n_1!}\right)\left(e^{-\lambda p_2 t} \frac{(\lambda p_2 t)^{n_2}}{n_2!}\right)$$

$\Rightarrow$ independent Poisson distributions $\square$

3. Uniformity

Thm: $0 < A_1 < A_2 < \cdots < A_k \le t$ be the random variables denoting the (ordered) arrival times for a PP conditioned on $\{N(t)=k\}$

Let $\tilde{A}_1, \tilde{A}_2, \ldots, \tilde{A}_k$ be a permutation of $A_1, \ldots, A_k$ chosen uniformly at random

Then $\{\tilde{A}_i\}$ are i.i.d. Unif$[0,t]$ random variables.

Corollary: The joint distribution of $A_1, A_2, \ldots A_k$ has density

$f(a_1, a_2, \ldots, a_k) = \frac{k!}{t^k}$ in the k-dimensional "quadrant"

$0 \le a_1 \le a_2 \le \ldots \le a_k \le t$

Pf: The corollary is easier to prove:

$$\text{Prob}\left(A_1 \in (a_1, a_1+h), A_2 \in (a_2, a_2+h), \ldots, A_k \in (a_k, a_k+h) \mid N(t)=k\right)$$

$$= \Big[P_{\mathbb{Q}}(N(a_1)=0)\; P_{\mathbb{Q}}(N(a_1+h)-N(a_1)=1) \cdot P_{\mathbb{Q}}(N(a_2)-N(a_1+h)=0) \cdots\Big] \Big/ P_{\mathbb{Q}}(N(t)=k)$$

$$= \frac{e^{-\lambda a_1}\,(\lambda h)\; e^{-\lambda(a_2-a_1)}\,(\lambda h) \cdots e^{-\lambda(a_k-a_{k-1})}\,(\lambda h)\, e^{-\lambda(t-a_k)}}{e^{-\lambda t}(\lambda t)^k / k!}$$

$$= \frac{e^{-\lambda t}(\lambda h)^k}{e^{-\lambda t}(\lambda t)^k / k!} = \frac{k!}{t^k}(h)^k.$$

$\square$

Some Generalizations of PP

1. Compound Poisson Process
   - Relax the counting process constraint
   - Replace by i.i.d. $\{V_i\}$ jumps

Compound PP ⇒ arrivals follow PP(λ) but each arrival is a batch of size V_i

(2) Non-stationary PP

Relax stationarity to : $N(t+s) - N(s)$ is Poisson distributed with mean $\int_s^{s+t} \lambda(u)\,du$.

$\lambda(u)$: arrival intensity at time u.

Q: How can we simulate a non-stationary PP ?

A: Let $\Lambda \geq \sup_t \{\lambda(t)\}$

- Simulate a stationary PP(Λ).
- Accept an arrival at time A_i with prob $\frac{\lambda(A_i)}{\Lambda}$.

(3) Brownian motion

- Relax counting process requirement
- $N(t) \sim N(\lambda t, \sigma^2 t)$ ← Normal distrib with mean λt, variance $\sigma^2 t$

⇒ Brownian motion has continuous sample paths.

CONTINUOUS TIME MARKOV CHAINS (CTMCs)

We will extend what we know about DTMCs to develop analysis of CTMCs.

Recall:

Defn: DTMC is a discrete-time discrete-space stoch. process $\{X_n : n \in \mathbb{N}\}$ satisfying

$$\Pr(X_{n+1}=j \mid X_n=i, X_{n-1}=x_{n-1}, \dots, X_0=x_0)$$
$$= \Pr(X_{n+1}=j \mid X_n=i) \quad \text{(Markov property)}$$
$$= P_{ij} \quad \text{(Stationarity)}$$

DTMC	vs.	CTMC
(i) state transitions at $\{1,2,\dots\}$		(i) state transitions can happen at any $t \geq 0$
(ii) on every visit to state i we stay for $\text{Geom}(1-P_{ii})$ time * for stationary DTMC		(ii) on every visit to state i we stay for an exponentially dist time * for stationary CTMC
(iii) P_{ij} Markovian		(iii) P_{ij} Markovian

First the formal defintion, then we will see how to intuitively think about CTMCs

Defn: A CTMC is a continuous time discrete space stoch process $\{X_t : t \geq 0\}$ satisfying:

$$\Pr(X_{t+s}=j \mid X_s=i ; X_u=x_u, u \leq s)$$
$$= \Pr(X_{t+s}=j \mid X_s=i) \quad \text{(Markov Prop.)}$$
$$\equiv P_{ij}^{(s,s+t)} \quad \text{(Notation)}$$
$$= P_{ij}^{(t)} \quad \text{(Stationarity)}$$

We will only concern ourselves with stationary CTMCs, but I will point out which of the various "views" allow non-stationary generalization.

Let τ_i = time until CTMC leaves state i given currently in i

Claim: continuous time + Markov property + stationary $\Rightarrow$ $\tau_i \sim$ Exponentially distributed

Pf:

$$P_i(\tau_i > s+t \mid \tau_i > s) \equiv P_i(X(u)=i,\ s\le u\le s+t \mid X(u)=i,\ 0\le u\le s)$$

$$= P_i(X(u)=i,\ s\le u\le s+t \mid X(s)=i) \quad [\text{Markov prop.}]$$

$$= P_i(X(u)=i,\ 0\le u\le t \mid X(0)=i) \quad [\text{stationarity}]$$

$$\equiv P_i(\tau_i > t)$$

$\Rightarrow \tau_i$ is memoryless

$\Rightarrow$ either Exp ✓ or ~~Geom~~ (Discrete time)

□

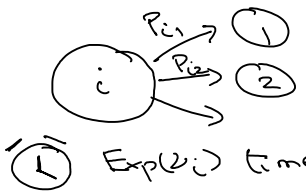
→ A few different ways to see CTMCs (some good for intuition, some good for ability to generalize)

CTMC-VIEW 1 (stationary): Every time we enter state i

(1) Spend $\text{Exp}(\nu_i)$ time before making a transition

(2) On leaving state i, transition to state j with probability P_{ij} (w.l.o.g. $P_{ii}=0$)

That is:



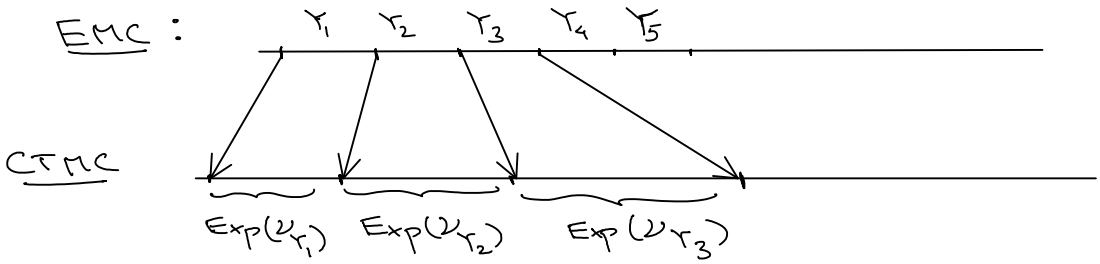
Defn: Let $\{Y_n\}$ be the discrete time process of the sequence of states visited by a CTMC.
$\{Y_n\}$ is a DTMC with transition probs $[P_{ij}]$ and is called the Embedded Markov Chain (EMC) corresponding to CTMC.

VIEW 1 gives us the stationary distrib of a CTMC based on what we already know:

Step1: We know how to get stationary dist for the embedded MC $\Leftarrow$ call it $\tilde{\pi}$

$$\tilde{\pi} = \tilde{\pi} P \iff \tilde{\pi}_j = \sum_{i \neq j} \tilde{\pi}_i P_{ij} \qquad (*)$$

Step2: EMC:



1 time step of EMC $\Longrightarrow$ stretched by $Exp(\nu_{Y_n})$
therefore stationary distrib of CTMC is "length-biased" or weighted version of $\tilde{\pi}$

Let π be the CTMC stationary distribution

$$\text{Length biasing} \Rightarrow \pi_j \propto \frac{\tilde{\pi}_j}{\nu_j}$$

$$\Rightarrow \tilde{\pi}_j \propto \pi_j \nu_j \qquad (**)$$

Substitute $(**)$ into $(*)$

$$\underbrace{\pi_j \nu_j}_{-q_{jj}} = \sum_{i \neq j} \pi_i \underbrace{\nu_i P_{ij}}_{q_{ij}} \qquad \Rightarrow \qquad \sum_i \pi_i q_{ij} = 0$$

Lets write $\sum_i \pi_i q_{ij}$ in matrix form

$$[\pi_1 \quad \pi_2 \quad \dots \quad \pi_n] \begin{bmatrix} q_{1j} \\ q_{2j} \\ \vdots \\ q_{jj} = -\nu_j \\ \vdots \\ q_{nj} \end{bmatrix} = 0$$

$Q \equiv$ "infinitesimal generator" OR "rate matrix"

Stationary distrib for a CTMC

$$\boxed{\pi Q = 0 \quad , \quad \sum_i \pi_i = 1}$$

Q: What are the row sums for Q?

A: $$\sum_j q_{ij} = \sum_{j \neq i} q_{ij} + q_{ii}$$

$$= \nu_i \sum_{j \neq i} P_{ij} - \nu_i = \nu_i - \nu_i = 0$$

Therefore Q is not full rank

$\Rightarrow$ left null space is not empty

$\Rightarrow$ a π satisfying $\pi Q = 0$ exists.

Remark: So far view 1 has only given us the stationary distribution, not the time dependent behavior. We will turn to this shortly.

Remark: In solving for π we only used the fact

$$\mathbb{E}[\tau_i] = \frac{1}{\nu_i}$$

and not that $\tau_i \sim \text{Exp}(\nu_i)$

If we allow τ_i to be any general distribution, we get a Semi-Markov process (SMP)

SMPs are continuous time processes that are not Markov.

Aside: Markov renewal process is a discrete time Markov process representation of an SMP: $X_n \equiv (Y_n, Z_{Y_{n-1}})$

- X_n: state of Markov Renewal process
- Y_n: state of EMC
- $Z_{Y_{n-1}}$: the generally distributed sojourn time spent in Y_{n-1}

(Note: this allows $Z_{Y_{n-1}}$ to depend on Y_{n-1} & Y_n!)

□

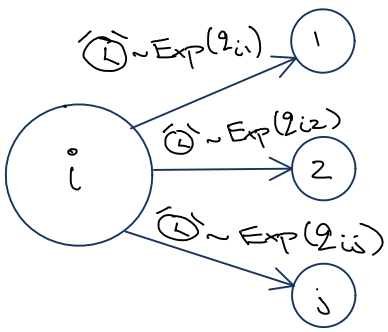
Now we will see another view that will be more useful when thinking about modeling systems as CTMCs.

CTMC-VIEW 2 (Stationary): A stoch. process s.t. every time we enter state i

(1) all states generate an independent Exp timer
ie. state k generates $X_k \sim Exp(\underbrace{\nu_i P_{ik}}_{q_{ik}})$

(2) stay in i for $\tau_i = \min_k \{X_k\}$

& then transition to $j = \underset{k}{\operatorname{argmin}} \{X_k\}$



What good is view 2?

A lot. Each $i \to j$ transition is a possible event in our system.

↳ arrival of job
↳ service completion
↳ abandonment from queue

If $\left\{\begin{array}{l}\text{interarrival times}\\ \text{service requirements}\\ \text{patience times}\end{array}\right\}$ are independent Exp distributed, then

these are the times in view-2.

Transition from i happens acc. to the event that happens first

Claim: VIEW 1 $\iff$ VIEW 2

Heuristic Proof

Pf 1: VIEW 2 $\Rightarrow$ VIEW 1

recall $X_k \sim Exp(q_{ik}) = Exp(\nu_i P_{ik})$

$\Rightarrow \tau_i = \min_k \{X_k\}$ has distrib $Exp(\underbrace{\sum_k q_{ik}}_{=\nu_i})$

$$Pr\left(\operatorname*{argmin}_k \{X_k\} = j\right) = \frac{q_{ij}}{\sum_k q_{ik}} = P_{ij}$$

Pf 2: Imagine the coin flipping world

VIEW 1



$Pr(\text{success in } \delta\text{-flip}) = \nu_i \delta$

$Pr(\text{transition to } j \mid \text{success}) = (P_{ij})$

$\Rightarrow Pr(\text{transition to } j) = (P_{ij}\nu_i)\delta$

$\equiv q_{ij}\delta$

VIEW 2



flip independent coins for each state we can go to

$Pr(\text{success for } j) = (\nu_i P_{ij})\,\delta \equiv (q_{ij})\delta$

$Pr(\text{transition to } j) = Pr(\text{success for } j, \text{ failures for rest})$

$= (q_{ij}\delta) \prod_{k \neq i,j} (1 - q_{ik}\delta)$

$= q_{ij}\delta + o(\delta)$

As $\delta \to 0$ the stochastic processes converge in distribution.

□

The coin-flipping view gives us another heuristic approach to solving the time dependent behavior of CTMC.

⇒ Approximate CTMC by a DTMC where events happen every δ-time units acc. to the success probs $(q_{ij}\delta)$.

Let $\hat{\pi}(t) = [\hat{\pi}_1(t) \;\; \hat{\pi}_2(t) \;\; \dots \;\; \hat{\pi}_n(t)]$

be distrib at time t in the DTMC

Let $\hat{P} = [\hat{P}_{ij}]$ where $\begin{cases} \hat{P}_{ij} = (q_{ij}\delta) & i \neq j \\ \hat{P}_{ii} = (1 - \nu_i\delta) \\ \quad\;\; = (1 + q_{ii}\delta) \end{cases}$

Using DTMC results:

$$\hat{\pi}(t+\delta) = \hat{\pi}(t)\hat{P} \iff \hat{\pi}_j(t+\delta) = \sum_{i \neq j} \hat{\pi}_i(t) q_{ij}\delta + \hat{\pi}_j(t)(1 - \nu_i\delta)$$

$$\Rightarrow \frac{\hat{\pi}_j(t+\delta) - \hat{\pi}(t)}{\delta} = \sum_{i \neq j} \hat{\pi}_i(t) q_{ij} - \hat{\pi}_j(t)\nu_i$$

Now, letting $\delta \to 0$, $\hat{\pi} \to \pi$

$$\boxed{\frac{d}{dt}\pi(t) = \pi(t)Q}$$

{ replace Q by $Q(t)$ for non-stationary

at stationarity, $\frac{d}{dt}\pi(t) = 0 \Rightarrow \pi Q = 0$.

This also explains Q as the

* infinitesimal generator : it gives the infinitesimal change in $\pi(t)$
* rate matrix : the entries of Q are the rates of transitions

So much for $\pi Q = 0$. Following is a more intuitive way to write this equation, which suggests a generalization:

Start with: $\sum_i \pi_i q_{ij} = 0$

$\Leftrightarrow \quad \sum_{i \neq j} \pi_i q_{ij} - \pi_j \nu_j = 0$

$\Leftrightarrow \quad \pi_j \sum_i q_{ji} = \sum_{i \neq j} \pi_i q_{ij}$

($\pi_j \sum q_{ji}$: rate leave j; $\sum \pi_i q_{ij}$: rate enter j)

$q_{jj} \equiv -\nu_j$

$\nu_j = \nu_j \sum_i P_{ji}$

$= \sum_i q_{ji}$

Flow balance equations: $\pi_j \sum q_{ji} = \sum_i \pi_i q_{ij}$

Generally: A, A^c a partition of state space

$\Rightarrow \quad \sum_{i \in A} \pi_i \sum_{j \in A^c} q_{ij} = \sum_{j \in A^c} \pi_j \sum_{i \in A} q_{ji}$

(rate $A \to A^c$; rate $A^c \to A$)

Remarks: Similar to what we did with DTMCs, we can show that flow balance equations are infact a sample path law

$N_{ij}(t) = \#\, i \to j$ transitions in $[0,t]$

$T_i(t) =$ cumulative time spent in state i in $[0,t]$

$$\sum_i N_{ji}(t) = \sum_i N_{ij}(t) \pm 1$$

$$\Rightarrow \quad \sum_i \frac{N_{ji}(t)}{t} = \sum_i \frac{N_{ij}(t)}{t} \pm \frac{1}{t}$$

$$\Rightarrow \quad \frac{T_j(t)}{t} \sum_i \frac{N_{ji}(t)}{T_j(t)} = \sum_i \frac{T_i(t)}{t} \cdot \frac{N_{ij}(t)}{T_i(t)} \pm \frac{1}{t}$$

$t \to \infty \longrightarrow \quad \pi_j \sum_i q_{ji} = \sum_i \pi_i \cdot q_{ij}$

When does a limiting π exist?

THEOREM: Given an irreducible CTMC,

if $\exists\ \pi_i$ s.t. $\pi_i > 0\ \forall i$ and

$$\pi_j \nu_j = \sum_{i \neq j} \pi_i q_{ij} \quad , \quad \sum_i \pi_i = 1$$

then the π_i's are the limiting probs for the CTMC.

Finally we turn to the <u>third view</u> of CTMCs. This is the most technical & rigorous approach:

<u>CTMC-view 3</u>: A stoch. process defined via the transition semi-group of probability transition matrices $\{P^{(s,t)}\}$ s.t.

1. $P^{(s,s)}(i,j) = I$
2. For any $s \leq u \leq t$: $P^{(s,t)} = P^{(s,u)} P^{(u,t)}$

If $P^{(s,t)} = P^{(0,t-s)}\ \forall s \leq t$, the transition semigroup is time-homogeneous.

The infinitesimal generator $\mathcal{Q}$ is defined as follows:

$$\lim_{h \downarrow 0} \frac{P^{(t,t+h)} v - P^{(t,t)} v}{h} = \mathcal{Q} v$$

(for all v for which the limit exists)

<u>Remark</u>: The above view is just a little away from defining Markov processes via time-evolution operators.

What operates on what?

* $\{P\}$ operate on distributions from right to evolve them forward

$$\pi(s)\, P^{(s,t)} = \pi(t)$$

* $\{P\}$ form linear operators on $|\mathcal{S}|$-dimensional normed vector space

$$v = \begin{bmatrix} v(1) \\ v(2) \\ \vdots \end{bmatrix}$$

$$\left(P^{(s,t)} v\right)(i) = \sum_j P^{(s,t)}_{ij} v_j = \mathbb{E}\left[v(X_t) \mid X_s = i\right]$$

□

Kolmogorov Backward / Forward equations

Using view 3, let

$$P_{ij}^{(t)} = P_r(X_{t+s}=j \mid X_s=i)$$

and then: $P_{ij}^{(t+s)} = \sum_k P_{ik}^{(s)} P_{kj}^{(t)}$ [semi-group property]

$$\Rightarrow \quad P_{ij}^{(t+s)} - P_{ij}^{(t)} = \sum_{k\neq i} P_{ik}^{(s)} P_{kj}^{(t)} - (1-P_{ii}^{(s)}) P_{ij}^{(t)}$$

dividing by s; letting $s \to 0$

$$\lim_{s\to 0} \frac{P_{ij}^{(t+s)} - P_{ij}^{(t)}}{s} = \lim_{s\to 0} \left(\sum_{k\neq i} \frac{P_{ik}^{(s)}}{s} P_{kj}^{(t)} \right) - \lim_{s\to 0} \left(\frac{1-P_{ii}^{(s)}}{s} \right) P_{ij}^{(t)}$$

q_{ik} = rate of $i \to k$; ν_i = total rate leave i

LOF (leap of faith) for exchanging lim and Σ:

$$\dot{P}_{ij}^{(t)} = \sum_{k\neq i} q_{ik} P_{kj}^{(t)} - \nu_i P_{ij}^{(t)}$$

Kolmogorov Backward Eqns.

If we start with $P_{ij}^{(t+s)} = \sum_k P_{ik}^{(t)} P_{kj}^{(s)}$; we will get

$$\dot{P}_{ij}^{(t)} = \sum_{k\neq j} P_{ik}^{(t)} q_{kj} - P_{ij}^{(t)} \nu_j$$

Kolmogorov Forward Eqns (KFE)

Assuming $\lim_{t\to\infty} P_{ij}^{(t)} \to \pi_j \quad \forall i,j$

$$\text{KFE} \Rightarrow \quad 0 = \sum_{k\neq j} \pi_k q_{kj} - \pi_j \nu_j \iff \boxed{\pi_j \nu_j = \sum_{k\neq j} \pi_k q_{kj}}$$

Using Laplace transforms to solve transient behavior

For stationary CTMC : $\frac{d}{dt}\pi(t) = \pi(t)Q$

But this is a little cumbersome to deal with for solving $\pi(t)$.

but $P^{(t+s)} = P^{(t)}P^{(s)} \Rightarrow P^{(t+s)} - P^{(t)} = P^{(t)}(P^{(s)} - I)$

which we would like to think is : $dP^{(t)} = P^{(t)}Q$

or $P^{(t)} = e^{tQ}$ $\quad \begin{cases} e^{tQ} = I + tQ + \frac{t^2Q^2}{2!} + \cdots \\ \text{when converges.} \end{cases}$

this is approximately correct (at least for finite state)

$$\Rightarrow \quad \pi(t) = \pi(0)e^{tQ}$$

Now let $\tilde{\pi}(s) = \int_0^\infty \pi(t)e^{-st}\,dt$ [Laplace transform of $\pi(t)$

$$\Rightarrow \quad \tilde{\pi}(s) = \int_0^\infty \pi(t)e^{-st}\,dt = \int_0^\infty \pi(0)e^{tQ}e^{-st}\,dt$$

$$= \int_0^\infty \pi(0)e^{t(Q-sI)}\,dt = \pi(0)[Q-sI]^{-1}$$

So inverting $Q-sI$ gives $\pi(t)$ though

$$\int_0^\infty \pi(t)e^{-st} \equiv \tilde{\pi}(s) = \pi(0)[Q-sI]^{-1}$$

"Resolvent" of the CTMC

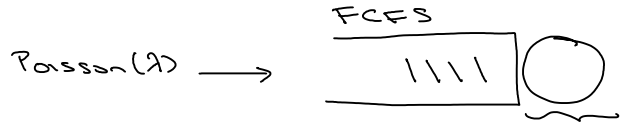
In practice, we will first start with z-transform:

$$\hat{N}(t,z) = \sum_i \pi_i(t)z^i = \pi(0)P^{(t)}\begin{bmatrix} 1 \\ z \\ z^2 \\ \vdots \end{bmatrix}$$

Then define $\tilde{N}(z,s) = \int_0^\infty \hat{N}(t,z)e^{-st}\,dt$

and try to get this into a nice form! □

Our first Q-ing example



Q: What is a (minimal) Markovian state space?

A: Here is one possibility:

Option 1: state(t) = { remaining size of all jobs in system (in order of arrival), remaining time until next arrival }

Clearly this is Markov; future evolution is completely specified by this "detailed" state descriptor

⇒ Above we have taken the view:

* size of a job is decided on arrival (i.e. he arrives with a size)
* the next arrival time is decided at current arrival time

Option 2: Decide "on the fly"

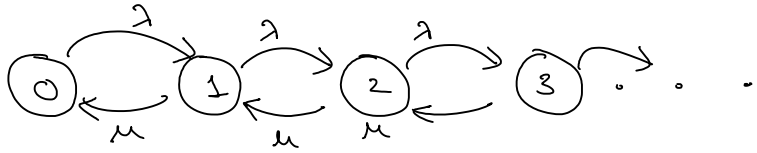
Here we exploit the property of Exponential distribution:

look at next δ-interval → w.p. $(\mu\delta)$ the job at server finishes
→ w.p. $(\lambda\delta)$ an interarrival time expires
→ w.p. $(1-(\lambda+\mu)\delta)$ no event happens

Therefore, sufficient to maintain

state(t) = N(t) = # jobs in system at time t

Graphical representation of CTMC:



Note: labels on arcs are the rates, not probabilities.

Q: What is the Q-matrix?

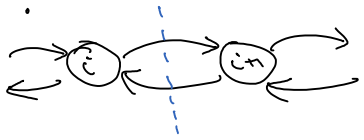
A: $$Q = \begin{bmatrix} -\lambda & \lambda & 0 & 0 & 0 & \cdots \\ \mu & -(\lambda+\mu) & \lambda & 0 & 0 & \\ & \mu & -(\lambda+\mu) & \lambda & 0 & \\ & & \mu & -(\lambda+\mu) & \lambda & \\ & & \vdots & & \ddots & \end{bmatrix}$$

Solving for π:

Method 1: Use $\pi Q = 0 \Rightarrow$

$$\lambda \pi(0) = \mu \pi(1)$$
$$(\lambda+\mu)\pi(1) = \lambda \pi(0) + \mu \pi(2)$$
$$(\lambda+\mu)\pi(2) = \lambda \pi(1) + \mu \pi(3)$$
$$\vdots$$

Method 2: use flow balance across these cuts



$$\Rightarrow \quad \lambda \pi(0) = \mu \pi(1) \quad \Rightarrow \pi(1) = \frac{\lambda}{\mu}\pi(0)$$
$$\lambda \pi(1) = \mu \pi(2) \quad \Rightarrow \pi(2) = \frac{\lambda}{\mu}\pi(1) = \left(\frac{\lambda}{\mu}\right)^2 \pi(0)$$
$$\vdots$$
$$\lambda \pi(i) = \mu \pi(i-1) \quad \Rightarrow \pi(i) = \frac{\lambda}{\mu}\pi(i-1) = \left(\frac{\lambda}{\mu}\right)^i \pi(0)$$

$$\sum_i \pi(i) = 1 \Rightarrow \pi(0)\left[1 + \frac{\lambda}{\mu} + \left(\frac{\lambda}{\mu}\right)^2 + \cdots\right] = 1$$

$$\Rightarrow \pi(0) = \left(1 - \frac{\lambda}{\mu}\right)$$

Finally: $$\pi(i) = \left(1 - \frac{\lambda}{\mu}\right)\left(\frac{\lambda}{\mu}\right)^i$$

Q: Expected steady-state # jobs in system

A: $$\mathbb{E}[N] = \sum_{n=0}^{\infty} n\,\pi(n)$$

$$= \left(1 - \frac{\lambda}{\mu}\right)\left[1 \cdot \frac{\lambda}{\mu} + 2\cdot\left(\frac{\lambda}{\mu}\right)^2 + 3\cdot\left(\frac{\lambda}{\mu}\right)^3 + \cdots\right] = \frac{\lambda/\mu}{1-\lambda/\mu}$$

Handy trick: $S = a + 2a^2 + 3a^3 + \cdots$

$$aS = \quad a^2 + 2a^3 + \cdots$$
$$S(1-a) = a + a^2 + a^3 + \cdots = a/(1-a)$$
$$\Rightarrow \quad S = a/(1-a)^2$$

INTRODUCTION TO QUEUEING

We saw one queueing system in the last lecture:



This is written as an M/M/1 queue in Kendall's notation.

- M/M/1: arrival process / service distrib./job size distrib / # servers

Other options for arrival process / service distrib:

- M : memoryless / Markovian (⇒ Exp interarrival / job sizes)
- D : Deterministic
- G : General { also written as GI for general interarrival to emphasize i.i.d.
- Ph : Phase type { used to approximate GI : more later
- Er_k : k-stage Erlang
- H_k : Hyperexponential { mixture of k Exp(μ_i); special case: H^* → p → Exp(μ), 1−p → 0
- M_t : non-stationary Markovian

The full Kendall's notation: A/B/m/K/P

- A: arrival process
- B: service distrib
- m: #servers
- K: buffer { optional, K ≥ m }
- P: population size { for closed systems }

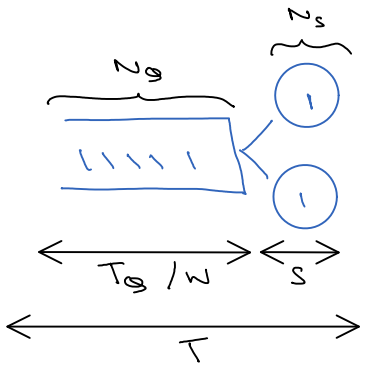
Other optional parameters

* scheduling rule (if not FCFS) : e.g. PS (Processor Sharing)
  RR (Round-Robin)
  LCFS (Last Come First Serve)
* abandonment distrib/ patience times : written as "+ M" or "+ GI"

Example : $M/G/2/LCFS + M$

⇒ Poisson arrival
Generally distrib. job sizes
2 servers
Last-Come First Served scheduling (non-preemptive)
i.i.d Exp patience times

Queueing terminology / Performance metrics



1. Time average metrics
   * N_Q = # jobs in queue
   * N_S = # jobs in service
   * $N(t) = N_Q(t) + N_S(t)$

   } $\mathbb{E}[\ ]$, $\mathrm{var}[\ ]$, $P_\infty[\ \geq n]$

2. Customer-average metrics
   * System time / response time / sojourn time / flow time
     $T \equiv t_{depart} - t_{arrive}$

   * Time in queue
     $T_Q \equiv T - S$ (job size)

   * Waiting time / delay
     $W \equiv t_{enter\text{-}service} - t_{arrive}$

($T_Q = W$ for non-preemptive service)

$\mathbb{E}[\ ]$, $\mathrm{var}[\ ]$, $P_\infty[\ \geq t]$

3. Fairness metrics

* Slowdown $= \mathbb{E}\left[\frac{T(x)}{x}\right]$ { $T(x)$ = expected response time of size x jobs

4. Utilization = fraction of time server busy

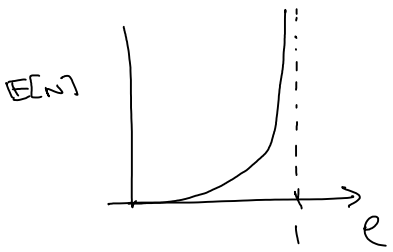
5. Throughput (X) = completion rate of jobs

Back to M/M/1:

(1) utilization $= 1 - \pi_0 = \frac{\lambda}{\mu} \equiv \rho$

λ/μ is also called the offered load

(2) $\mathbb{E}[N] = \sum_{n=0}^{\infty} n(1-\rho)\rho^n = \frac{\rho}{1-\rho}$



$var[N] = \frac{\rho}{(1-\rho)^2}$

$Pr[N \geq n] = \rho^n$

(3) Throughput $= \lambda$

method 1: for open systems with no loss: throughput = arrival rate

method 2: jobs complete at rate μ when server busy

$\Rightarrow X = \mu\left[1 - \pi(0)\right] = \mu \cdot \frac{\lambda}{\mu} = \lambda$

method 2 is more general → closed system; → finite buffer abandonments

Q: What about $\mathbb{E}[T]$; $Pr[T \geq t]$?

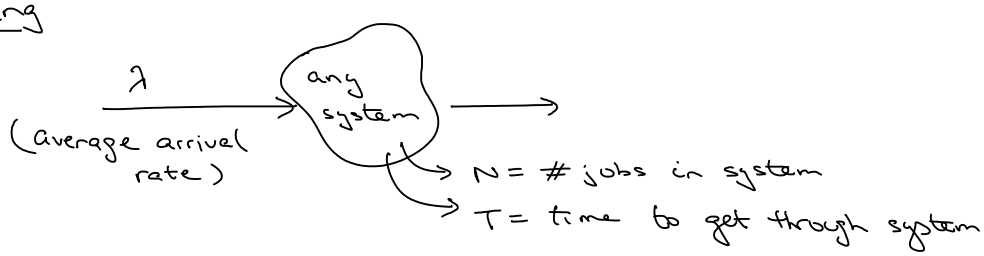
These are "customer-centric" metrics ⇒ depends on state observed by customers

Additional tools to go from N to T { 1. Little's Law; 2. PASTA

LITTLE'S LAW (1961 : 50 years after Erlang's work!)

- One of the most fundamental results in queueing-theory

Setting

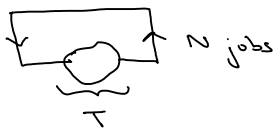


We will use the following version of L.L.

> **Little's Law (steady-state version):** For any ergodic system
>
> $$\mathbb{E}[N] = \lambda\, \mathbb{E}[T]$$

Remark: For a closed system

$$N = \lambda\, \mathbb{E}[T]$$



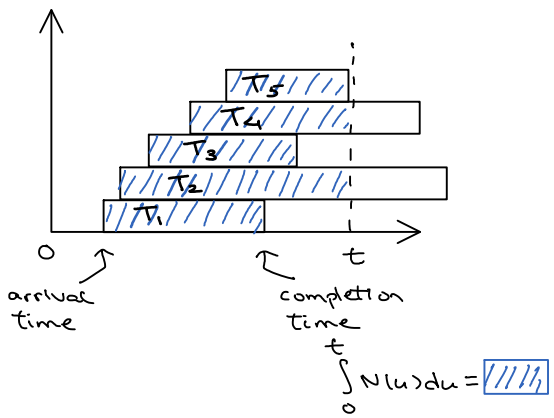
We will prove the "sample-path" version of Little's Law.

(sample path LL + regenerative process theory ⇒ steady-state LL)

Setup: $A(t)$ = # arrivals by time t

$C(t)$ = # completions by time t

T_j = system time of j^{th} arrival



Assume:

$$\lim_{t\to\infty} \frac{A(t)}{t} = \lambda$$

$$\lim_{n\to\infty} \frac{\sum_{j=1}^{n} T_j}{n} = T \quad \left\{ \Rightarrow \lim_{n\to\infty} \frac{T_n}{n} \to 0 \right.$$

We have:

$$\sum_{j=1}^{C(t)} T_j \le \int_0^t N(u)du \le \sum_{j=1}^{A(t)} T_j$$

$$\Rightarrow \left(\frac{A(t)}{t}\right)\left(\frac{C(t)}{t}\right)\left(\frac{\sum_{j=1}^{C(t)} T_j}{C(t)}\right) \le \frac{\int_0^t N(u)du}{t} \le \left(\frac{A(t)}{t}\right)\left(\frac{\sum_{j=1}^{A(t)} T_j}{A(t)}\right)$$

by assumptions: $A(t) \to \infty$ as $t \to \infty$, $\lim_{n\to\infty} \frac{T_n}{n} = 0$ $\Big\}$ $\lim_{t\to\infty} \frac{C(t)}{A(t)} \to 1$

$\Longrightarrow$ $\lambda T (1-\varepsilon) \leq N \leq \lambda T$ where $\varepsilon > 0$ is arbitrary

$\Rightarrow$ $\boxed{N = \lambda T}$

Remark: Since the above is for a sample path ω satisfying the assumptions., we really mean $N(\omega) = \lambda(\omega) T(\omega)$.

Remark: We made almost no assumptions on our "system"

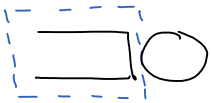
(i) arbitrary arrival/service distributions

(ii) no i.i.d. assumptions

(iii) no assumption of FCFS service

Q: L.L. applied to the buffer

A: $\mathbb{E}[N_Q] = \lambda \mathbb{E}[T_Q]$



Q: L.L. applied to the server

A: $\mathbb{E}[N_S] = \lambda \mathbb{E}[S]$

"utilization law"



Q: Say we arbitrarily mark jobs as red and blue

A: $\mathbb{E}[N_{red}] = \lambda_{red} \mathbb{E}[T_{red}]$

Two ways to remember Little's Law:

(1) an accounting trick:

λT = reward rate when jobs pay T_j on arrival

N = reward rate when jobs pay 1 per unit time while in system

(2) $\lambda T \approx$ # jobs that arrive during my sojourn $\approx N$

Generalizations of Little's Law:

(i) If jobs depart in the order of arrival { e.g. FCFS with single server } and Poisson arrivals:

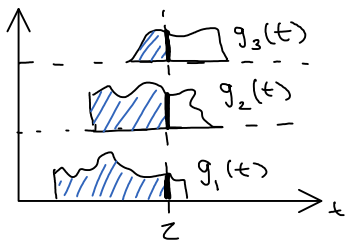
$$\mathbb{E}[N(N-1)\cdots(N-i+1)] = \lambda^i \mathbb{E}[T^i]$$

(ii) $H = \lambda G$.

Setup: we can view $N = \lambda T$ as a special case where each job contributes $\int_0^\infty g_j(t)\,dt$ with $g_j(t) = \mathbb{1}\{j \text{ in system at } t\}$

Generalize: arbitrary $g_j(t)$ but with support on j's system time

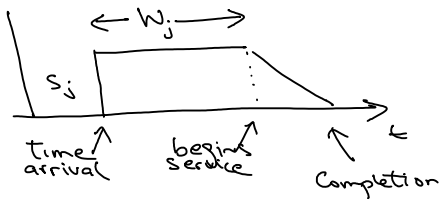
$$G_j = \int_0^\infty g_j(t)\,dt$$
$$H(t) = \sum_{j=1}^{\infty} g_j(t)$$
$$\Rightarrow \boxed{H = \lambda G}$$



Example: $g_j(t)$ = remaining service time with non-premptive service

$$G_j = S_j W_j + \frac{1}{2} S_j^2$$

$H(t)$ = unfinished workload $\equiv V(t)$



$$H = \lambda G = \mathbb{E}[V] = \lambda\left[\mathbb{E}[SW] + \frac{1}{2}\mathbb{E}[S^2]\right]$$

for size-independent policy $S_j \perp W_j$

$$\Rightarrow \boxed{\mathbb{E}[V] = \lambda\left(\mathbb{E}[S]\,\mathbb{E}[W] + \frac{1}{2}\mathbb{E}[S^2]\right)}$$

Brumelle's formula

(We will return to this for M/G/1 & M/G/k)

PASTA (Poisson Arrivals See Time Averages)

for M/M/1 ; Little's law $\Rightarrow$ $\mathbb{E}[T] = \frac{1}{\lambda}$ $\mathbb{E}[N] = \frac{1}{\lambda} \cdot \frac{\lambda/\mu}{1-\lambda/\mu} = \frac{1}{\mu-\lambda}$

Q: var(T) ? $P_r(T \geq t)$?

A: We could use the higher moment version of L.L. But a slicker method is using PASTA.

Let

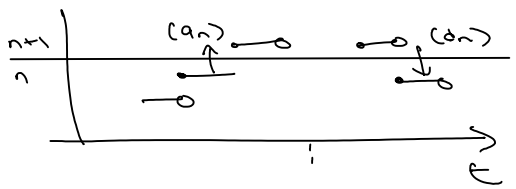
a_n = fraction of jobs that see n jobs on arrival (excluding themself)

P_n = long run fraction of time there are n jobs

d_n = fraction of jobs that leave behind n jobs on departure

Q: How are a_n & d_n related

A: For an ergodic system $a_n = d_n$



$\left(\#\ (n \to n+1) \text{ arrivals} = \#\ (n+1 \to n) \text{ departures}\right) \pm 1$

Q: Are a_n & P_n related?

A: Not nec. eg. D/D/1 with $\lambda = 1$ job/sec, $S = \frac{1}{2}$ sec $\begin{cases} a_0 = 1 \\ P_0 = P_1 = \frac{1}{2} \end{cases}$

PASTA says a_n & P_n are related when arrival process is Poisson.

Actually, not quite ! What we need is

Lack of Anticipation Assumption (LAA)

$A(t)$ = # arrivals by time t

$X(t)$ = process we are observing

LAA: for every $t \geq 0$, $\{A(t+u) - A(t), u > 0\} \perp \{X(s) : 0 \leq s \leq t\}$

Pf of PASTA:

Define: $p_B = \lim_{t\to\infty} Pr(X(t) \in B)$

$$a_B = \lim_{t\to\infty} \lim_{\delta\to 0} Pr\left(X(t)\in B \mid \Lambda(t+\delta) - \Lambda(t) = 1\right)$$

$$= \lim_{t\to\infty} \lim_{\delta\to 0} \frac{Pr\left(X(t)\in B,\ \Lambda(t+\delta)-\Lambda(t)=1\right)}{Pr\left(\Lambda(t+\delta)-\Lambda(t)=1\right)}$$

$$= \lim_{t\to\infty} \lim_{\delta\to 0} \frac{Pr\left(\Lambda(t+\delta)-\Lambda(t)=1 \,\middle|\, X(t)\in B\right) Pr(X(t)\in B)}{Pr\left(\Lambda(t+\delta)-\Lambda(t)=1\right)}$$

$$\overset{(LAA)}{=} \lim_{t\to\infty} \lim_{\delta\to 0} \frac{Pr\left(\Lambda(t+\delta)-\Lambda(t)=1\right)}{Pr\left(\Lambda(t+\delta)-\Lambda(t)=1\right)} \cdot Pr(X(t)\in B)$$

$$= p_B$$

Q: Is Poisson arrival process sufficient?

A: No. E.g. n^{th} job size = $\left((n+1)^{st}$ interarrival time$\right)/2$

$\Rightarrow a_0 = 1$

But $p_0 \neq 1$

Remark: PASTA is useful in simulations because you only need to track statistics at times of job arrivals

Extension: i.i.d. Retrials see Time averages

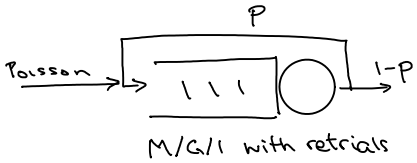
Setup: departures cycle back to queue with probab. p (independent of state)

retrials will see d_n; but $d_n = a_n$

LAA for external arrivals $\Rightarrow a_n = p_n$

$\Rightarrow d_n = p_n$.

$\Rightarrow$ retrials see p_n.

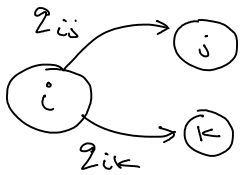


SINGLE STATION MARKOVIAN QUEUEING SYSTEMS

* single station : customer queues form at a single location
* queueing network : a collection of queueing stations with some routing policy for the jobs

Before proceeding, let's recap the basic tools / laws :

1. CTMCs



q_{ij} = rate of event which causes $i \to j$ transition
= rate of Exp distribution
e.g. : λ : arrival
μ : service completion
ν : abandonment

stationary dist : $\pi Q = 0$

transient $\frac{d}{dt}\pi(t) = \pi(t) Q$

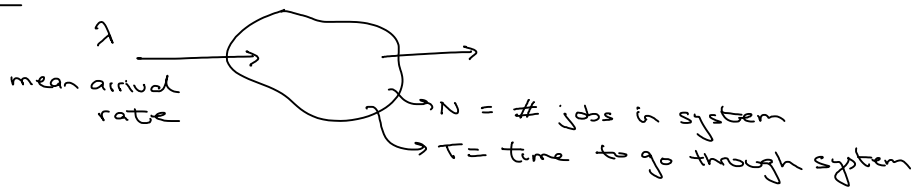
$$Q = \begin{bmatrix} -\nu_1 & q_{12} & q_{13} & \cdots \\ q_{21} & -\nu_2 & q_{23} & \cdots \\ q_{31} & q_{32} & -\nu_3 & \\ \vdots & & & \end{bmatrix} \qquad \nu_i = \sum_{j \neq i} q_{ij}$$

Flow-balance : A, A^c a partition of state space

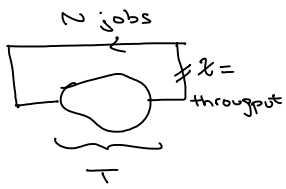
$$\sum_{i \in A} \pi(i) \sum_{j \in A^c} q_{ij} = \sum_{j \in A^c} \pi(j) \sum_{i \in A} q_{ji}$$

2. Little's Law

(a)



	Open System	Closed System
Sample path :	$N = \lambda T$	$N = \chi T$
Steady state :	$\mathbb{E}[N] = \lambda \mathbb{E}[T]$	$N = \chi \mathbb{E}[T]$



(b) Generalization $H = \lambda G$

3. Poisson Arrivals See Time Averages (PASTA)

a_n, P_n, d_n } fraction of { arrivals, time, departures } that { See, see, leave behind } n jobs

$a_n = d_n$ { when arr / dep one-at-a-time

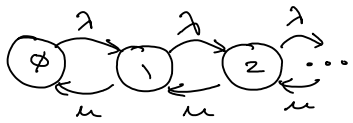
$a_n = P_n$ (PASTA) $\in$ LAA

We have already seen our first single station system

(I) M/M/1 (or M/M/1/FCFS)
↳ usually suppressed



Graphical representation: of CTMC with state = N(t)



stationary dist : $\pi(i) = (1-\rho)\rho^i$

utilization : $1-\pi(0) = \rho = \frac{\lambda}{\mu}$ ≡ "offered load"

$$\mathbb{E}[N] = \frac{\rho}{1-\rho} \quad ; \quad \mathbb{E}[N_Q] = \mathbb{E}[N] - \mathbb{E}[N_s] = \frac{\rho}{1-\rho} - \rho = \frac{\rho^2}{1-\rho}$$

Little's Law $\Rightarrow$ $\mathbb{E}[T] = \frac{1}{\lambda}\cdot\mathbb{E}[N] = 1/(\mu-\lambda)$

Remark: By combining Little's and PASTA, we can get $\mathbb{E}[T]$ without solving for π.

(i) Little $\Rightarrow$ $\mathbb{E}[N] = \lambda\,\mathbb{E}[T]$ ⊛

(ii) $\mathbb{E}[T \mid \text{see } n \text{ jobs on arrival}] = (n+1)\cdot 1/\mu$

$$\Rightarrow \mathbb{E}[T] = \sum_{n=0}^{\infty} a_n \cdot (n+1)\frac{1}{\mu} \overset{\text{(PASTA)}}{=} \sum_{n=0}^{\infty} \pi(n)(n+1)\frac{1}{\mu} = \frac{1}{\mu}\left(\mathbb{E}[N]+1\right) \quad \text{⊛⊛}$$

Eliminate $\mathbb{E}[N]$ from ⊛ & ⊛⊛ $\Rightarrow \mathbb{E}[T] = 1/(\mu-\lambda)$.

The above is an example of the Mean Value Analysis (MVA) technique

Step 1: Write T in terms of state (N) on arrival
– then take expectation using linearity of expectations

Step 2: Use Little's law to obtain a second relation b/w $\mathbb{E}[T]$ & $\mathbb{E}[N]$

Q. If $\lambda \to 2\lambda$; then what μ_{new} keeps $\mathbb{E}[T]$ the same

A: $\mathbb{E}[T]$ depends on $(\mu-\lambda)$ $\Rightarrow \Delta\mu = \Delta\lambda$ $\Rightarrow$ $\mu_{new} = \mu + \lambda$

Distribution of $T^{M/M/1}$:

* $(T \mid \text{see } n \text{ jobs on arrival}) \sim \sum_{i=1}^{n+1} X_i \qquad X_i \sim Exp(\mu)$

* PASTA $\Rightarrow (N+1) \sim Geom(1-\rho)$

$$\Rightarrow T \overset{d}{=} \sum_{i=1}^{N+1} X_i \overset{(H.W.1)}{=} Exp(\mu(1-\rho)) = Exp(\mu-\lambda)$$

Here is the proof for completeness :

Laplace transform of T : $\mathcal{L}_T(s) = \mathbb{E}[e^{-sT}]$

$$= \mathbb{E}_N\left[\mathbb{E}[e^{-sT} \mid N]\right]$$

$$= \mathbb{E}_N\left[\left(\frac{\mu}{\mu+s}\right)^{N+1}\right]$$

$$= \hat{G}_{(1-\rho)}\left(\frac{\mu}{\mu+s}\right) \qquad \begin{cases} \hat{G}_p(z) = \dfrac{pz}{1-(1-p)z} \\ \text{z-transform of } Geom(p) \end{cases}$$

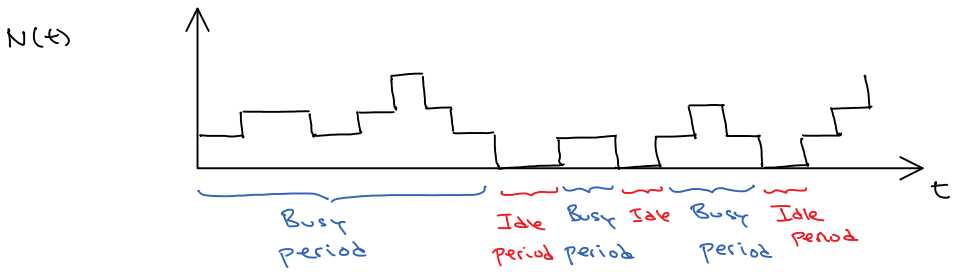
$$= \frac{(\mu-\lambda)}{(\mu-\lambda)+s} \qquad \{ \text{Laplace transform of } Exp(\mu-\lambda)$$

□

Busy / Idle periods

Busy period : the period between $0 \to 1$ transition until next $1 \to 0$ transition

Idle period : period b/w $1 \to 0$ transition until next $0 \to 1$ transition.



Q : What is the distribution of Idle period duration?

A : $I \sim Exp(\lambda)$

Q : Expectation of Busy period duration?

A: Let $\mathbb{E}[B]$ = Expected busy period duration

B = B. period duration started by first arrival

S_1 = size of first arrival

$N_A(S_1)$ = # jobs that arrive while S_1 is at server

Then:

$$B = S_1 + \sum_{i=1}^{N_A(S_1)} B_i$$

↑ each arrival starts its own B. period $\stackrel{d}{=}$ i.i.d. B

"Busy period recurrence"

$$\Rightarrow \mathbb{E}[B \mid S_1, N_A] = S_1 + \sum_{i=1}^{N_A(S_1)} \mathbb{E}[B_i]$$

$$= S_1 + N_A(S_1)\,\mathbb{E}[B]$$

$$\Rightarrow \mathbb{E}[B \mid S_1] = S_1 + \mathbb{E}[N_A(S_1)]\,\mathbb{E}[B]$$

$$= S_1 + \lambda S_1 \mathbb{E}[B]$$

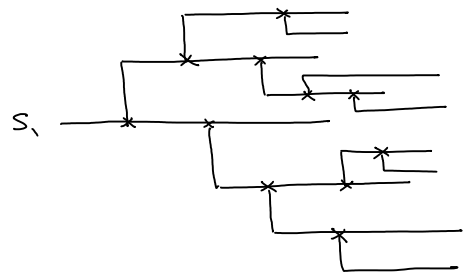
$$\Rightarrow \mathbb{E}[B] = \mathbb{E}[S_1] + \lambda \mathbb{E}[S_1]\,\mathbb{E}[B]$$

$$= \frac{1}{\mu} + \frac{\lambda}{\mu}\mathbb{E}[B]$$

$$\Rightarrow \boxed{\mathbb{E}[B] = \frac{\mathbb{E}[S]}{1-\rho} = \frac{1}{\mu-\lambda}}$$

□

The branching process view.



horizontal segment = a single service time

: arrivals during this segment

: branch = B. period started by arrival

B. period started by S_1 = sum of all horizontal segments

How about that! $\mathbb{E}[B] = \mathbb{E}[T]$.

Do you think $B \sim Exp(\mu-\lambda)$ as well?

In HW2 you will find the transform of M/M/1 busy period

* M/M/1 Busy period distribution is a mess { Bessel functions

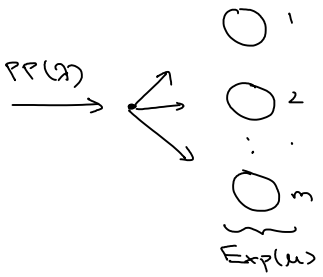
* but we will see we can write it as a mixture of infinitely many Exp (ν)

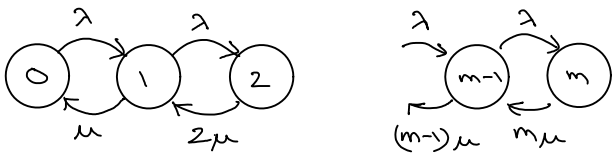
$$\Rightarrow \overline{F}_B(b) = \int_{\nu=\underline{\nu}}^{\bar{\nu}} e^{-\nu b}\, d\phi(\nu)$$

↳ "Spectral decomposition"

$\underline{\nu}$ = the "slowest" component gives information about convergence rate to steady state. More on this later!

(II) M/M/m/m (a.k.a. Erlang-B loss system)





arrival served if an idle server, otherwise lost.

Q: Do we need $\lambda < \mu$?

A: No. finite buffer systems are always stable.

stable $\Leftrightarrow$ ergodic from now on

The key quantity of interest in finite buffer systems is:

"Blocking probability" $\equiv$ Pr{arrival finds all slots full & is lost}

Flow-balance:

$$\pi(0)\lambda = \pi(1)\mu \Rightarrow \pi(1) = \pi(0)\frac{\lambda}{\mu}$$

$$\pi(1)\lambda = \pi(2)2\mu \Rightarrow \pi(2) = \pi(0)\frac{1}{2!}\left(\frac{\lambda}{\mu}\right)^2$$

$$\pi(i-1)\lambda = \pi(i)i\mu \Rightarrow \pi(i) = \pi(0)\frac{1}{i!}\left(\frac{\lambda}{\mu}\right)^i$$

$$Pr(\text{Blocking}) = \pi(m) = \frac{\frac{1}{m!}\left(\frac{\lambda}{\mu}\right)^m}{\sum_{i=0}^{m}\frac{1}{i!}\left(\frac{\lambda}{\mu}\right)^i}$$

{ Erlang-B formula

Q. What would Little's Law for $\mathbb{E}[N]$ be?

A: $\mathbb{E}[N] = \chi\,\mathbb{E}[T]$

$\chi = \lambda(1-\pi(m))$

$\mathbb{E}[T] = \mathbb{E}[S]$ for arrivals which enter system

$$\Rightarrow \mathbb{E}[N] = \frac{\lambda}{\mu}\cdot\left(1 - Pr(\text{Blocking})\right)$$

(II B) M/M/∞

Let $m \to \infty$ in M/M/m/m:

$$\pi(n) = \pi(0) \frac{1}{n!} \left(\frac{\lambda}{\mu}\right)^n$$

$$\pi(0) = \left[\sum_{n=0}^{\infty} \frac{1}{n!} \left(\frac{\lambda}{\mu}\right)^n\right]^{-1} = e^{-\frac{\lambda}{\mu}}$$

$$\Rightarrow \pi(n) = e^{-\lambda/\mu} \frac{(\lambda/\mu)^n}{n!}$$

$$\Rightarrow N^{M/M/\infty} \sim \text{Poisson}\left(\frac{\lambda}{\mu}\right)$$

We will see later that:

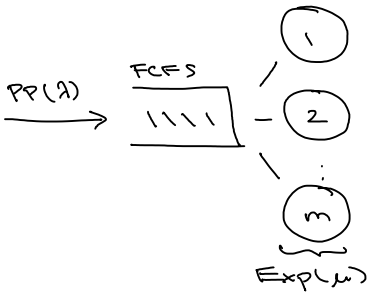
$$N^{M/G/m/m} \overset{d}{=} N^{M/M/m/m}$$

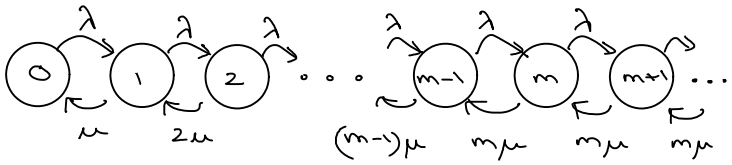
Therefore, distribution of # jobs (hence $\mathbb{E}[N]$, $\mathbb{E}[T]$) is "insensitive" to service distribution (only $\mathbb{E}[S]$ matters)

- A reason why for a long time queueing theory research focused on Markovian assumptions believing them to be harmless.

Remark: Insensitivity does not hold
* for M/G/m/m and M/G/∞ in transience, or
* if arrival process is not Poisson.

(III) M/M/m (a.k.a. Erlang-C system)





Q: Stability conditions?

A: $\lambda < m\mu$

$$\pi(n) = \begin{cases} \pi(0) \frac{1}{n!} \left(\frac{\lambda}{\mu}\right)^n & n \le m \\ \pi(m) \left(\frac{\lambda}{m\mu}\right)^{(n-m)} & n > m \end{cases}$$

Q: How should we define utilization?

A1: "offered load" $\rho = \lambda/\mu = \mathbb{E}[\#\text{ busy servers}]$ $\quad \{\in [0, m)$

A2: fraction of capacity used $\equiv \frac{\lambda}{m\mu}$ $\quad \{\in [0,1)$

I prefer $\rho = \lambda/\mu$ since $(m - \lceil \rho \rceil)$ = "# spare servers" critically effects how response time depends on job-size distribution

Another quantity of interest:

"Delay probability" (P_d) $= P_A\{\text{arrival finds all servers busy}\}$

$= P_A\{W > 0\}$

$= \sum_{n=m}^{\infty} \pi(n)$

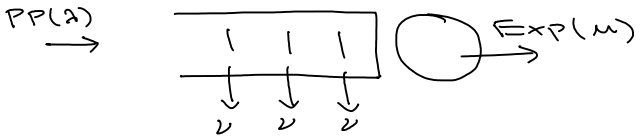
Q: Distribution of W?

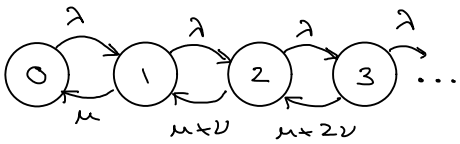
A. W $\begin{cases} \xrightarrow{1-P_d} 0 \\ \xrightarrow{P_d} \text{Exp}(m\mu - \lambda) \end{cases}$

$\left\{ (N+1-m \mid N \ge m) \stackrel{d}{=} \underbrace{N}_{\text{M/M/1}} + 1 \right.$

M/M/1 with service rate $m\mu$

(IV) M/M/1 + M





Each job in queue abandons at rate ν.

Q: Stability conditions

A: System with abandonments are stable irrespective of λ.

$$\pi(n) = \pi(0) \prod_{i=1}^{n} \frac{\lambda}{\mu + (i-1)\nu}$$

Q: Fraction of jobs that abandon?

A: $\lambda\,(1 - Pr(\text{abandon})) = \text{throughput} = \mu \cdot Pr(\text{busy})$

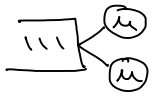
$$\Rightarrow \quad 1 - \pi(0) = \rho \cdot (1 - Pr(\text{abandon}))$$

Next lecture:

* stochastic orderings + coupling
  (a formalism to compare random variables)
* Which is better: M/M/1 or M/M/2 ?



"one fast server"



"many slow servers"

STOCHASTIC ORDERINGS & STOCHASTIC COUPLING - I

Random variables : mappings that allow us to compare different experiments and outcomes numerically.

Easy to compare numbers : $3 < 4$

For random variables : "$X < Y$" may mean one of many things

- $\mathbb{E}[X] < \mathbb{E}[Y]$
- $\mathbb{E}[X^i] < \mathbb{E}[Y^i]$
- $Pr[X > a] < Pr[Y > a]$
- $\mathbb{E}[X \mid X > a] < \mathbb{E}[Y \mid Y > a]$
- $X(\omega) < Y(\omega)$ for all sample paths ω

or, more broadly 'X better than Y'

- $\mathbb{E}[X] < \mathbb{E}[Y]$ & $Var[X] < Var[Y]$
  (e.g. queueing delay)
- $\mathbb{E}[X] > \mathbb{E}[Y]$ & $Var[X] < Var[Y]$
  (e.g. portfolio return)

Depending on application & need, any of these or other notion can be used. We will discuss some of the more common notions and where they show up in queueing systems:

1. Almost sure ordering $X \leq_{as} Y$
2. Stochastic ordering $X \leq_{st} Y$
3. Likelihood ratio ordering $X \leq_{lr} Y$
4. Hazard ratio ordering $X \leq_{hr} Y$
5. Increasing convex ordering $X \leq_{icx} Y$
   Concave $X \leq_{icv} Y$
6. Convex ordering $X \leq_{cx} Y$

1. Almost-Sure Ordering

- strongest ordering
- closest to comparing random variables like ordinary numbers.

Defn: X is almost surely smaller than Y, or $X \leq_{as} Y$ iff $\mathbb{P}(\{\omega : X(\omega) \leq Y(\omega)\}) = 1$.

Note: I wrote $\mathbb{P}$ to emphasize that X and Y must be defined on a common probability space (say $(\Omega, \mathbb{P})$) for this ordering to make sense

Example: $N(t) \equiv$ Poisson counting process

$N(5) \leq_{as} N(10)$

Example: $X \sim N(0,1)$ $\quad Y \sim N(10,1)$



Q: Is $X \leq_{as} Y$?

A: Depends on joint distribution of X, Y

If $Y = X + 10$, then yes

If $Y \perp X$, then no

Application: a.s. ordering of GI/GI/1/FCFS waiting time

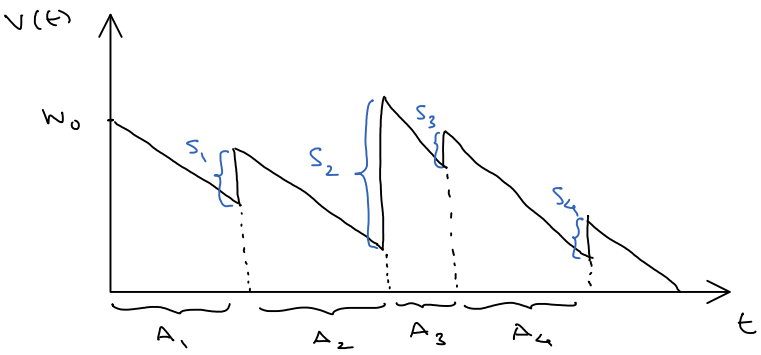
Start by looking at a "sample path view" of FCFS systems

W_0 = work at time 0

$\{A_n\}$ = interarrival time of n^{th} job ($T_n - T_{n-1}$, $n = 1, 2, \ldots$)

$\{S_n\}$ = size of n^{th} job

$V(t)$ = total unfinished work at time t



W_n = waiting time of n^{th} job

Q: $W_1 = ?$

A: $W_1 = (W_0 - A_1)^+$

W_0: work at $t=0$

A_1: this much drains by when job #1 arrives

$(\ )^+$ because can't have negative waiting time

Q: $W_2 = ?$

A $W_2 = (W_1 + S_1 - A_2)^+$

W_1: work job #1 saw in front of him

S_1: how much job #1 added

A_2: how much drained by when job #2 showed up

In general,

$$W_{n+1} = (W_n + S_n - A_{n+1})^+$$

"Lindley's Equation" or "Lindley's Recursion"

Now consider two G/G/1/FCFS systems:

$$W_0^{(1)} \leq_{as} W_0^{(2)}$$
$$A_n^{(1)} \geq_{as} A_n^{(2)}$$
$$S_n^{(1)} \leq_{as} S_n^{(2)}$$

Q: What can we say about $W_n^{(1)}$ vs. $W_n^{(2)}$?

$$A_1^{(1)} \geq_{as} A_2^{(1)} \Rightarrow -A_1^{(1)} \leq_{as} -A_1^{(2)}$$
$$W_0^{(1)} \leq_{as} W_0^{(2)}$$

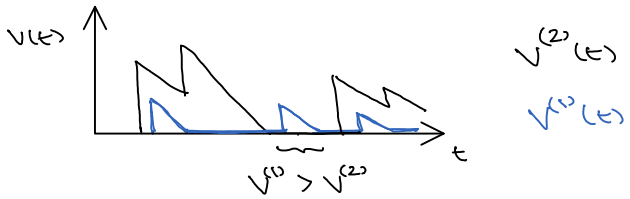
$$= W_0^{(1)} - A_1^{(1)} \leq_{as} W_0^{(2)} - A_1^{(2)}$$
$$\Rightarrow (W_0^{(1)} - A_1^{(1)})^+ \leq_{as} (W_0^{(2)} - A_1^{(2)})^+$$
$$\Rightarrow W_1^{(1)} \leq_{as} W_1^{(2)}$$

By induction : $W_n^{(1)} \leq_{as} W_n^{(2)}$

So: almost-sure ordering of $\{A_n\}$ and $\{S_n\}$ implies almost-sure ordering of waiting times $\{W_n\}$

Q: Is $V^{(1)}(t) \leq_{as} V^{(2)}(t)$ for all t?

A: No.



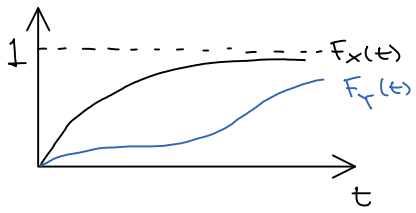
While almost sure ordering seems too restrictive, it is the vehicle through which we will prove weaker notions of ordering.

---

2. Stochastic ordering

Defn: X is said to be stochastically smaller than Y, $X \leq_{st} Y$ iff $\Pr[X \leq a] \geq \Pr[Y \leq a]$

Pictorially:



Remark: Unlike $\leq_{as}$, $\leq_{st}$ is a property of distribution functions. The random variables need not be defined on the same prob. space

Example:

$X \sim N(0,1)$ $Y \sim N(1,1)$ $\Rightarrow$ $X \leq_{st} Y$

$X \sim Exp(\mu_X)$ $Y \sim Exp(\mu_Y)$ $\Rightarrow$ $X \leq_{st} Y \Leftrightarrow \mu_X \geq \mu_Y$

$X \sim Bernoulli(p)$ $Y \sim Bernoulli(q)$ $\Rightarrow$ $X \leq_{st} Y \Leftrightarrow p \leq q$

Properties:

(1) $X \leq_{st} Y$ iff $\mathbb{E}[g(X)] \leq \mathbb{E}[g(Y)]$ $\forall$ increasing functions $g(\cdot)$

(Proof later)

(2) Stochastic ordering is preserved under

- convolution $\left(\begin{matrix} X_1 \leq_{st} Y_1 \\ X_2 \leq_{st} Y_2 \end{matrix} \Rightarrow X_1 + X_2 \leq_{st} Y_1 + Y_2 \right)$
- increasing functions $\left(g \text{ inc.}, X \leq_{st} Y \Rightarrow g(X) \leq_{st} g(Y) \right)$
- weak limits $X_n \leq_{st} Y_n$, $X_n \Rightarrow X$, $Y_n \Rightarrow Y$ $\Rightarrow X \leq_{st} Y$

Pf: $X_n \sim F_n(\cdot)$; $Y_n \sim G_n$

$X_n \leq_{st} Y_n \Rightarrow F_n(t) \geq G_n(t)$ $\forall n, t$

$\Rightarrow F(t) \geq G(t)$ at continuity pts. of F, G which are dense in $\mathbb{R}$.

(3) This property is so important, we state it as a theorem:

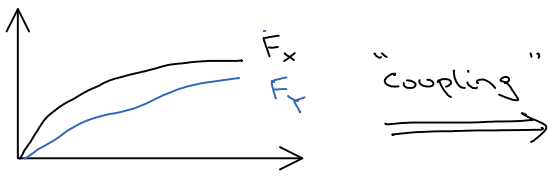
Thm: $X \leq_{st} Y$ iff there exist random variables $\tilde{X}, \tilde{Y}$ s.t.

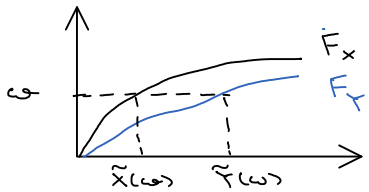
(i) $\tilde{X} \overset{d}{=} X$ } Marginals agree

(ii) $\tilde{Y} \overset{d}{=} Y$ }

(iii) $\tilde{X} \leq_{as} \tilde{Y}$ } $\tilde{X}, \tilde{Y}$ defined on a common probability space

Explanation and proof by picture:





(often called "quantile coupling")

$X \leq_{st} Y$ but not defined on a common prob. space

$\tilde{X}, \tilde{Y}$ defined on $(\Omega, \mathbb{P})$

$\Omega = [0,1]$, $\mathbb{P} = \mathcal{U}[0,1]$

Note: $\tilde{X} \leq_{as} \tilde{Y}$

$\tilde{X}$ & $\tilde{Y}$ $\equiv$ "coupled versions" of X & Y

Remark: there are multiple ways to couple distributions.

Example $X \sim Exp(\mu_X)$ $Y \sim Exp(\mu_Y)$ $(\mu_X \geq \mu_Y)$

(1) sample $\tilde{X} \sim Exp(\mu_X)$

$$\tilde{Y} = \tilde{X} * \left(\frac{\mu_X}{\mu_Y}\right)$$

(2) sample $\tilde{X} \sim Exp(\mu_X)$

$$\tilde{Y} = \begin{cases} \tilde{X} & \text{with prob. } \mu_Y/\mu_X \\ \tilde{X} + Exp(\mu_Y) & \text{otherwise} \end{cases}$$

Using either method $\tilde{X} \leq_{as} \tilde{Y}$

Proof of property (1): $X \leq_{st} Y$ iff $\mathbb{E}[g(X)] \leq \mathbb{E}[g(Y)]$ $\forall$ inc. g

$\Leftarrow$ consider $g_a(t) = \mathbb{1}_{\{t>a\}}$

$$\mathbb{E}[g_a(X)] \leq \mathbb{E}[g_a(Y)] \Rightarrow P_X[X>a] \leq P_Y[Y>a]$$

$\Rightarrow$ consider coupled versions $\tilde{X}, \tilde{Y}$,

$X \leq_{st} Y \xrightarrow[\text{coupling}]{} \tilde{X} \leq_{as} \tilde{Y} \Rightarrow g(\tilde{X}) \leq_{as} g(\tilde{Y})$ ↙ "decoupling"

$\Rightarrow g(X) \leq_{st} g(Y)$

$\Rightarrow \mathbb{E}[g(X)] \leq \mathbb{E}[g(Y)]$

□

Remark: If $\mathbb{E}[X] = \mathbb{E}[Y]$ then $X \leq_{st} Y$ iff $F_X = F_Y$.

Proving stochastic orderings for queueing system

The following example illustrates the general approach

Example 1: Consider the G/G/1/FCFS example from earlier this lecture, but now:

$$W_0^{(1)} \leq_{st} W_0^{(2)} \; ; \; A_n^{(1)} \geq_{st} A_n^{(2)} \; ; \; S_n^{(1)} \leq_{st} S_n^{(2)}$$

Prove: $W_n^{(1)} \leq_{st} W_n^{(2)}$, $n = 1, 2, 3, \ldots$

Step 1: inputs with "st" ordering $\xrightarrow{\text{coupling}}$ inputs with "as" ordering

$$\left.\begin{array}{l} W_0^{(1)} \leq_{st} W_0^{(2)} \\ A_n^{(1)} \geq_{st} A_n^{(2)} \\ S_n^{(1)} \leq_{st} S_n^{(2)} \end{array}\right\} \Rightarrow \left\{\begin{array}{l} \text{prob space } (\Omega, P) \text{ \& } \tilde{W}, \tilde{A}_n, \tilde{S}_n \\ \text{s.t.} \quad \tilde{W}_0^{(1)} \leq_{as} \tilde{W}_0^{(2)} \\ \qquad \tilde{A}_n^{(1)} \geq_{as} \tilde{A}_n^{(2)} \\ \qquad \tilde{S}_n^{(1)} \leq_{as} \tilde{S}_n^{(2)} \end{array}\right.$$

Step 2: carry out a sample path "as" analysis in the coupled prob. space

$\Rightarrow$ From earlier in lecture: $\tilde{W}_n^{(1)} \leq_{as} \tilde{W}_n^{(2)}$

Step 3: Infer "st" orderings for original system

$$\left.\begin{array}{l} \tilde{W}_n^{(1)} \stackrel{d}{=} W_n^{(1)} \\ \tilde{W}_n^{(2)} \stackrel{d}{=} W_n^{(2)} \\ \tilde{W}_n^{(1)} \leq_{as} \tilde{W}_n^{(2)} \end{array}\right\} \Rightarrow W_n^{(1)} \leq_{st} W_n^{(2)}$$

Note: The coupling here was obvious, but sometimes it takes a little ingenuity to find a convenient coupling

Example 2: π = stationary distribution of # jobs in M/M/1

Consider an M/M/1 system with #jobs at $t=0$ as $N(0)$

Claim: $N(0) \geq_{st} \pi \Rightarrow N(t) \geq_{st} \pi \quad \forall\, t \geq 0$

Proof:

Step 1: coupling

First we need to create a second system that will be our proxy for the steady state of M/M/1

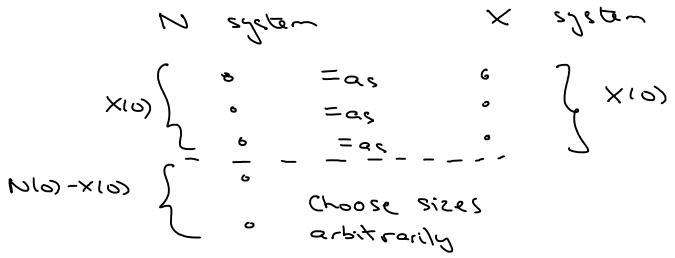
→ call this $X(t)$ with $X(0) \sim \pi$

(i.e. initial # jobs sampled acc. to π)

$N(0) \geq_{st} \pi \Rightarrow \exists$ coupled versions $N(0) \geq_{as} X(0)$

Coupling service times and arrivals

method 1 :- initial state



- Couple arrival times so arrivals happen together
- and corresponding arrivals to X & N systems have equal sizes

This is getting hairy really quickly !!

* Exploit the memoryless property of Exp()

method 2 : * Couple $X(0)$, $N(0)$ so that $N(0) \geq_{as} X(0)$

* define a Poisson(λ) process ← drives arrivals to both

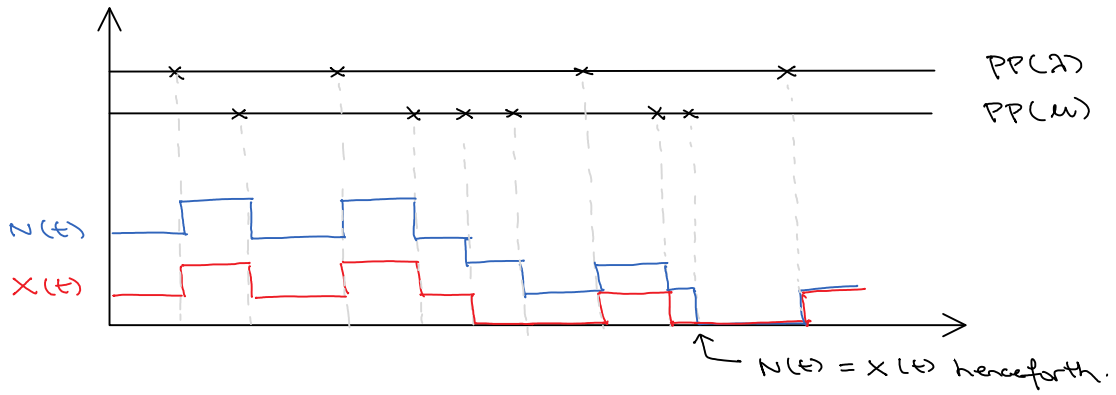
* define a Poisson(μ) process ← drives departures to both when busy

Q How do the sizes of n^{th} job in two systems compare?

A: We don't care! All we need is that the individual processes behave like legitimate M/M/1

Step 2:

(Picture proof)

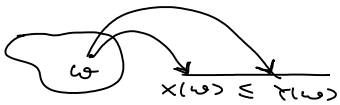


$\Rightarrow N(t) \geq_{as} X(t)$

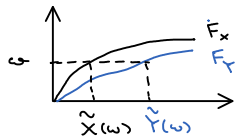
Step 3: $X(0) \sim \pi \Rightarrow X(t) \sim \pi, \forall t \Rightarrow N(t) \geq_{st} \pi, \forall t \geq 0$ □

STOCHASTIC ORDERINGS & STOCHASTIC COUPLING - II

Recap:

* $X \leq_{as} Y$ 

for ω with prob. 1
* $X \leq_{st} Y \iff Pr(X \leq a) \geq Pr(Y \leq a) \quad \forall a$
* almost-sure $\rightleftharpoons$ (coupling) stochastic ordering



* examples of inferring almost-sure & stochastic ordering for Q ing syst.

This lecture:

* M/M/1 vs. M/M/2
* ↳ more orderings.

Example 3: Which of the following is better?

M/M/1 vs. M/M/2



one fast server or two slow servers

- One of the fundamental system design questions.

Claim:

(i) $N^{M/M/1} \leq_{st} N^{M/M/2} \leq_{st} N^{M/M/1} + 1$

(ii) $T_Q^{M/M/1} \geq_{st} T_Q^{M/M/2}$

(iii) $T^{M/M/1} \leq_{icx} T^{M/M/2}$

icx = increasing convex (we will get to this shortly)

⇒ "better" depends on your choice of metric

* # jobs is stoch. less in M/M/1
* Waiting time (until service begins) is stoch smaller in M/M/2
* Time in system is "icx smaller" in M/M/1

Proof: 1. Create coupled versions $X(t)$ & $Y(t)$
$\hookrightarrow N^{M/M/1}(t)$ $\hookrightarrow N^{M/M/2}(t)$

2. prove $X(t) \leq_{as} Y(t) \leq_{as} X(t)+1 \quad \forall t \geq 0$

3. Since $\lim_{t\to\infty} X(t) \stackrel{d}{=} N^{M/M/1}$ & $\lim_{t\to\infty} Y(t) \stackrel{d}{=} N^{M/M/2}$

by closure of $\leq_{st}$ under weak convergence

$N^{M/M/1} \leq_{st} N^{M/M/2} \leq_{st} N^{M/M/1}+1$

4. PASTA + ordering on $N^{M/M/x} \Rightarrow$ ordering on T_Q, T.

Step 1:

Q: What should we set $X(0)$ & $Y(0)$?

A: Choice does not affect $\lim_{t\to\infty} X(\infty)$, $\lim_{t\to\infty} Y(\infty)$. So we choose any convenient condition which satisfies the invariant we want to prove. Eg $X(0) = Y(0) = 0$

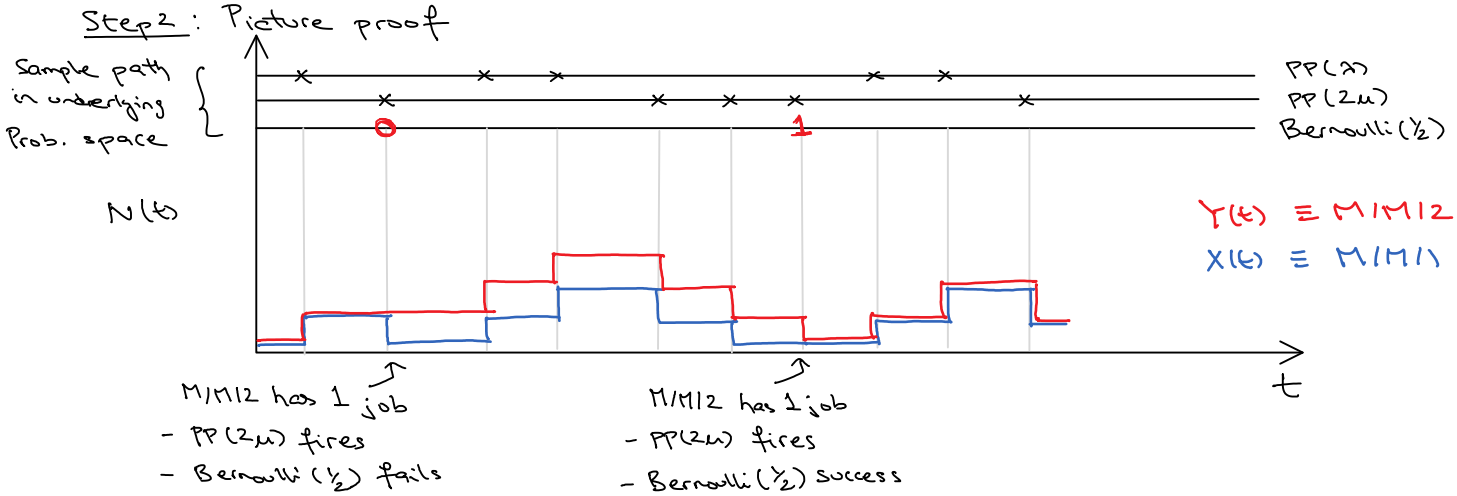
Q: Coupling of arrival and service processes?

A: Common Poisson(λ) stream couples arrivals.

Poisson(2μ) stream $\Rightarrow$ M/M/1 departures
M/M/2 departures when both servers busy

Bernoulli($\frac{1}{2}$) stream $\Rightarrow$ to "thin" Poisson(2μ) when M/M/2 has one job

Step 2: Picture proof



Step 4: want to show $T_Q^{M/M/1} \geq_{st} T_Q^{M/M/2}$

$T_Q^{M/M/1} = \sum_{i=1}^{N^{M/M/1}} S_i$ $\left(S_i \text{ i.i.d. } Exp(2\mu)\right)$ ⋮ $T_Q^{M/M/2} = \sum_{i=1}^{\left(N^{M/M/2}-1\right)^+} S_i$ $\left(S_i \text{ i.i.d. } Exp(2\mu)\right)$

also $N^{M/M/1} \geq_{st} \left(N^{M/M/2}-1\right)^+$

$\Rightarrow$ Couple $\tilde{N}^{M/M/1} \geq_{a.s.} \left(\tilde{N}^{M/M/2}-1\right)^+$

Couple $\tilde{S}_i =_{a.s.} \tilde{S}_i$

$\Rightarrow$ $\tilde{T}_Q^{M/M/1} \geq_{a.s.} \tilde{T}_Q^{M/M/2}$

$\Rightarrow$ $T_Q^{M/M/1} \geq_{st} T_Q^{M/M/2}$

Q: Is $T^{M/M/1} \leq_{st} T^{M/M/2}$?

A: Let us try the approach above

$T^{M/M/1} = \sum_{i=1}^{\left(N^{M/M/1}+1\right)} S_i$ ⋮ $T^{M/M/2} = S' + \sum_{i=1}^{\left(N^{M/M/2}-1\right)^+} S_i$

$S_i \sim Exp(2\mu)$ ⋮ $S_i \sim Exp(2\mu)$; $S' \sim Exp(\mu)$

know $N^{M/M/1} \leq_{st} N^{M/M/2}$

$\Rightarrow$ Couple $\tilde{N}^{M/M/1} \leq_{as} \tilde{N}^{M/M/2}$

in the worst case $n_1 = n_2 \ (\equiv n)$

$\Rightarrow T^{M/M/1} = \left(\sum_{i=1}^{n-1} S_i\right) + S_n + S_{n+1}$ ⋮ $T^{M/M/2} = \left(\sum_{i=1}^{n-1} S_i\right) + S'$

So, $S_n + S_{n+1} \leq_{st} S' \Rightarrow T^{M/M/1} \leq_{st} T^{M/M/2}$

but $\mathbb{E}[S_n + S_{n+1}] = \frac{1}{\mu} = \mathbb{E}[S']$

Can not claim $T^{M/M/1} \leq_{st} T^{M/M/2}$!!

(We will return to this question towards the end of the lecture) □

3. Likelihood ratio ordering

Sometimes we want to claim something stronger than $X \leq_{st} Y$:

e.g. $[X \mid X \in A] \leq_{st} [Y \mid Y \in A]$

Likelihood ratio order allows precisely this stronger statement.

Defn: X is smaller than Y in likelihood ratio sense, $X \leq_{lr} Y$

if

* $\frac{f_X(t)}{f_Y(t)}$ is non-increasing (continuous X, Y)
* $\frac{P(X=t)}{P(Y=t)}$ is non-decreasing (discrete X, Y)

Why the name likelihood ratio?

Consider the following experiment: $Z \begin{cases} \to X & \text{with prob } p \\ \to Y & \text{with prob } (1-p) \end{cases}$

$$\text{Q:} \quad \frac{P_r[Z = X \mid Z = t]}{P_r[Z = Y \mid Z = t]} = \left(\frac{p}{1-p}\right) \frac{P_r(X=t)}{P_r(Y=t)}$$

} Likelihood ratio

So, $X \leq_{lr} Y$ says higher values are less and less likely to come from X.

Example:

$$X = \begin{cases} 1 & 0.5 \\ 2 & 0.1 \\ 3 & 0.2 \\ 4 & 0.2 \end{cases} \qquad Y = \begin{cases} 2 & 0.5 \\ 4 & 0.5 \end{cases}$$

(probs)

$X \leq_{st} Y$ but $[X \mid X \text{ even}] \geq_{st} [Y \mid Y \text{ even}]$.

Properties:

(1) $X \leq_{lr} Y$ implies $X \leq_{st} Y$ (choose $A = \Omega$)

(2) $X \leq_{lr} Y$ implies $g(X) \leq_{lr} g(Y)$ for all increasing $g(\cdot)$.

(3) $\leq_{lr}$ is not closed under convolution

(4) $\leq_{lr}$ is closed under weak convergence

4. Hazard rate ordering

- Weaker than $\leq_{lr}$, stronger than $\leq_{st}$
- When making scheduling decisions, we don't care about $[X \mid X \in A]$ for all events. We only care about how much longer we will have to wait given already waited for t

  ie. want

$$[X \mid X > t] \leq_{st} [Y \mid Y > t] \quad \forall t$$

This is what hazard rate ordering guarantees.

Defn: X is smaller than Y in hazard ratio sense, $X \leq_{hr} Y$ iff

$$\frac{\bar{F}_X(t)}{\bar{F}_Y(t)} \quad \text{is non-increasing in } t.$$

Alternately: if X & Y have densities f_X & f_Y, then $X \leq_{hr} Y$ iff

$$\frac{f_X(t)}{\bar{F}_X(t)} \geq \frac{f_Y(t)}{\bar{F}_Y(t)} \quad \forall t$$

The ratio $h_X(t) \triangleq \dfrac{f_X(t)}{\bar{F}_X(t)}$ is called the "hazard rate" or "failure rate"

Q: what does $h_X(t)$ mean intuitively?

A:

$$h_X(t)\,\delta t \simeq \frac{P_X[X \in (t, t+\delta t)]}{P_X[X > t]} = \text{Probability job finishes in the next } \delta t \text{ time given not finished until } t.$$

(The word "failure" comes from reliability theory where one looks at lifetimes of machine parts)

Properties: (1) $X \leq_{hr} Y$ implies $g(X) \leq_{hr} g(Y)$ for increasing g

(2) not closed under convolution

(3) closed under weak convergence.

5. Increasing Convex Ordering

Sometimes even $[X|X>t] \leq_{st} [Y|Y>t]$ is too strong for our purpose, and we may only want:

$$\mathbb{E}[(X-t)^+] \leq \mathbb{E}[(Y-t)^+]$$

This is the increasing convex ordering, so called because of the following definition:

Defn: X is smaller than Y in increasing convex order sense, $X \leq_{icx} Y$, iff

$$\mathbb{E}[g(X)] \leq \mathbb{E}[g(Y)] \quad \forall \text{ increasing convex } g(\cdot)$$

So if $X \leq_{icx} Y$

* $\mathbb{E}[X] \leq \mathbb{E}[Y]$
* $\mathbb{E}[X^n] \leq \mathbb{E}[Y^n]$ for non-negative X, Y, $n = 2, 3, \ldots$

Facts:

* $X \leq_{st} Y$ implies $X \leq_{icx} Y$ { $\leq_{st}$ needs $\mathbb{E}[g(X)] \leq \mathbb{E}[g(Y)]$ $\forall$ inc. $g(\cdot)$

Claim: $X \leq_{icx} Y$ iff $\mathbb{E}[(X-t)^+] \leq \mathbb{E}[(Y-t)^+]$ $\forall\, t$

Pf: $\Rightarrow$ $g(x) = (x-t)^+$ is an increasing convex function

$\Leftarrow$ let $h_t(x) = (x-t)^+$

then $g(x) = \int_{-\infty}^{\infty} h_t(x)\, g''(t)\, dt + (\text{constant})$

□

Intuitively $X \leq_{icx} Y$ means X is "less variable" than Y.

6. Convex Ordering

Def: X is smaller than Y in convex ordering sense, $X \le_{cx} Y$ iff

$$\mathbb{E}[g(X)] \le \mathbb{E}[g(Y)] \quad \forall \text{ convex fns. } g(\cdot)$$

Thm: $X \le_{cx} Y$ iff $X \le_{icx} Y$ and $\mathbb{E}[X] = \mathbb{E}[Y]$.

Pf: (1a) $X \le_{cx} Y \Rightarrow X \le_{icx} Y$

all increasing convex fns. are by definition convex

(1b) $X \le_{cx} Y \Rightarrow \mathbb{E}[X] = \mathbb{E}[Y]$

$g(x) = x$ & $g(x) = -x$ are both convex

(2) $X \le_{icx} Y$ & $\mathbb{E}[X] = \mathbb{E}[Y] \Rightarrow X \le_{cx} Y$

Let $g(\cdot)$ be a convex function

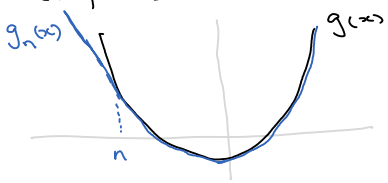
Case: $g(x) + \alpha x$ is increasing [& convex] for some α

$\Rightarrow \mathbb{E}[g(X) + \alpha X] \le \mathbb{E}[g(Y) + \alpha Y]$ $(X \le_{icx} Y)$

$\Rightarrow \mathbb{E}[g(X)] \le \mathbb{E}[g(Y)]$ $(\mathbb{E}[X] = \mathbb{E}[Y])$

Case: if not, approximate $g(\cdot)$ monotonically as

$$g_n(x) = \begin{cases} g(x) & x \ge -n \\ g(-n) + \underbrace{g'_+(-n)}_{\text{right derivative at } x=-n}(x+n) & \text{o.w.} \end{cases}$$



□

Thm: $X \le_{icx} Y \iff \exists Z : X \le_{st} Z \le_{cx} Y$; **Thm:** $X \le_{icx} Y \iff \exists Z : X \le_{cx} Z \le_{st} Y$.

$X \le_{cx} Y$ implies

* $\mathbb{E}[X^n] \le \mathbb{E}[Y^n]$ for all n if X, Y non-negative
* $\mathbb{E}[(X - \mathbb{E}[X])^n] \le \mathbb{E}[(Y - \mathbb{E}[Y])^n]$ for all even n
* $X_e \le_{st} Y_e$ (Homework 2)
  ↳ stationary excesses / equilibrium dist.

$\le_{cx}$, $\le_{icx}$ are also called "variability orderings" because of the above consequences.

Properties

- Closed under convolution
- Not closed under weak convergence (in general)

  but if $X_n \to X$, $Y_n \to Y$ in distribution

  And $\mathbb{E}[(X_n)^+] \to \mathbb{E}[X^+]$, $\mathbb{E}[(Y_n)^+] \to \mathbb{E}[Y^+]$

  then $X \leq_{cx} Y$

Similar to $\leq_{st}$, we have a coupling type result for $\leq_{cx}$

> **Thm:** $X \leq_{cx} Y$ if and only if there exist $\tilde{X}, \tilde{Y}$ defined on a common prob. space satisfying
>
> (i) $\tilde{X} =_{st} X$
>
> (ii) $\tilde{Y} =_{st} Y$
>
> (iii) $\mathbb{E}[\tilde{Y} \mid \tilde{X}] = \tilde{X}$ a.s.
>
> That is $\{\tilde{X}, \tilde{Y}\}$ is a martingale.

Pf: $\Rightarrow$ is non-trivial

$\Leftarrow$ Let $g(\cdot)$ be convex

$$\mathbb{E}[g(X)] = \mathbb{E}[g(\tilde{X})] = \mathbb{E}\left[\mathbb{E}[g(\tilde{Y}) \mid \tilde{X}]\right]$$

$$\geq \mathbb{E}[g(\mathbb{E}[\tilde{Y} \mid \tilde{X}])]$$

$$= \mathbb{E}[g(\tilde{X})]$$

$$= \mathbb{E}[g(X)]$$

□

Corollary: Let $X = \sum_{i=1}^{m} A_i$ $\quad A_i \sim Exp(m\mu)$

$Y = \sum_{i=1}^{n} B_i$ $\quad B_i \sim Exp(n\mu)$

then $m \geq n \Rightarrow X \leq_{cx} Y$

Exercise: Use the above to show $T^{M/M/1} \leq_{icx} T^{M/M/2}$

□

STOCHASTIC ORDERINGS AND STOCHASTIC COUPLING - III

Uses of stochastic coupling

(1) comparing queueing systems (e.g. M/M/1 vs M/M/2)

(2) proving existence of stationary distrib. (we may see an example later)

(3) finding rate of convergence to steady-state for CTMCs

Convergence rate to steady-state distribution for M/M/1

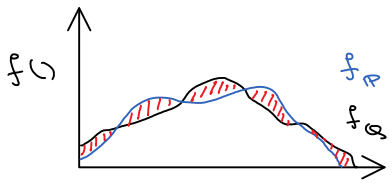
Goal: Given we start an M/M/1 at $t=0$ with $X(0)=x$, how far is $X(t)$ from the steady-state distribution?

Measuring distance b/w distributions: the total variation norm

Defn: Let $\mathbb{P}$ & $\mathbb{Q}$ be measures on space $\mathcal{S}$.
The total variation norm b/w $\mathbb{P}$ & $\mathbb{Q}$ is defined as

$$\| \mathbb{P} - \mathbb{Q} \|_{tv} = 2 \sup_{A \subseteq \mathcal{S}} \left| \mathbb{P}(A) - \mathbb{Q}(A) \right|$$

Pictorially:



$\| \mathbb{P} - \mathbb{Q} \|_{tv}$ = red area

If the "2" seems mysterious, it shows up because of the following alternate definition:

$$\| \mathbb{P} - \mathbb{Q} \|_{tv} = \sup_{\|f\|_\infty \le 1} \left| \mathbb{E}_{\mathbb{P}}[f(x)] - \mathbb{E}_{\mathbb{Q}}[f(x)] \right|$$

Example: $X \sim$ Bernoulli(p) $\quad$ $Y \sim$ Bernoulli(q)

$X \to 1$ w.p. p, 0 w.p. $1-p$ $\quad$ $Y \to 1$ w.p. q, 0 w.p. $(1-q)$

$\mathcal{S} = \{0,1\}$, So $\|X - Y\|_{tv} = 2\max\{|p-q|, |(1-p)-(1-q)|\} = 2|p-q|$

How does stochastic coupling enter the picture?

As follows:

Thm: $\| \mathbb{P} - \mathbb{Q} \|_{tv} = 2 \inf \Pr[\tilde{X} \neq \tilde{Y}]$

where $\tilde{X}$ & $\tilde{Y}$ are coupled random variables with distributions $\mathbb{P}$ & $\mathbb{Q}$ respect. ($\Pr$ denotes the measure induced by the common prob space).

Pf: we will only prove $\| \mathbb{P} - \mathbb{Q} \|_{tv} \leq 2 \Pr_{\tilde{X},\tilde{Y}}[\tilde{X} \neq \tilde{Y}]$ for any coupling

for any $A \subseteq S$

$$\begin{aligned}
\mathbb{P}(X \in A) - \mathbb{Q}(Y \in A) &= \Pr(\tilde{X} \in A, \tilde{X} = \tilde{Y}) + \Pr(\tilde{X} \in A, \tilde{X} \neq \tilde{Y}) \\
&\quad - \Pr(\tilde{Y} \in A, \tilde{X} = \tilde{Y}) - \Pr(\tilde{Y} \in A, \tilde{X} \neq \tilde{Y}) \\
&= \Pr(\tilde{X} \in A, \tilde{X} \neq \tilde{Y}) - \Pr(\tilde{Y} \in A, \tilde{X} \neq \tilde{Y})
\end{aligned}$$

$$\begin{aligned}
\Rightarrow \| \mathbb{P} - \mathbb{Q} \|_{tv} &= 2 \sup_A \left| \mathbb{P}(A) - \mathbb{Q}(A) \right| \\
&= 2 \sup_A \left| \mathbb{P}(X \in A) - \mathbb{Q}(Y \in A) \right| \\
&\leq 2 \sup_A \Pr(\tilde{X} \in A, \tilde{X} \neq \tilde{Y}) \\
&= 2 \Pr(\tilde{X} \neq \tilde{Y})
\end{aligned}$$

□

The coupling under which $\| \mathbb{P} - \mathbb{Q} \|_{tv} = 2 \Pr(\tilde{X} \neq \tilde{Y})$ is called the "maximal coupling". For our purposes, any reasonable coupling would suffice

Back to M/M/1

Let π denote the stationary distrib for # jobs.

- We want to find $\| X(t) - \pi \|_{tv}$
- and use coupling to do so

⇒ have to define $\{X(t)\}_{t \geq 0}$ & a stationary M/M/1 on a common space

Here is how we will do this

(1) Start an M/M/1 with $X(0) = x$ jobs

(2) Start another M/M/1 with $N(0)$ jobs, $N(0)$ sampled from π

(3) Let $X(t)$ & $N(t)$ evolve independently, until the first time they coincide:

$$\tau = \min_{t \geq 0} \{ X(t) = N(t) \}$$

τ is called the coupling time (a stopping time)

(4) for $t \geq \tau$; define $X(t) = N(t)$

we say "$X(t)$ & $N(t)$ couple at τ"

So how does this help us?

- Note $N(t) \stackrel{d}{=} \pi \quad \forall t > 0$
- Therefore by the theorem above

$$\| X(t) - \pi \|_{tv} \leq 2 P_r (X(t) \neq N(t))$$
$$\leq 2 P_r (\tau > t)$$

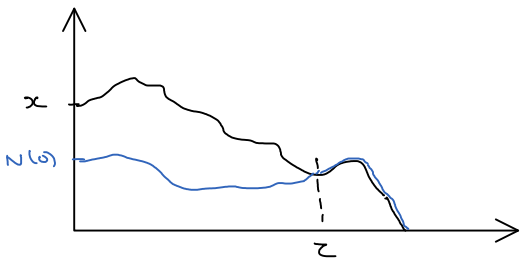
$\{ N(t), X(t)$ are coupled$\}$

$\{ \{X(t) \neq N(t)\} \subseteq \{\tau > t\} \}$

So we have reduced the $\| \ \|_{tv}$ problem to analyzing the coupling time.

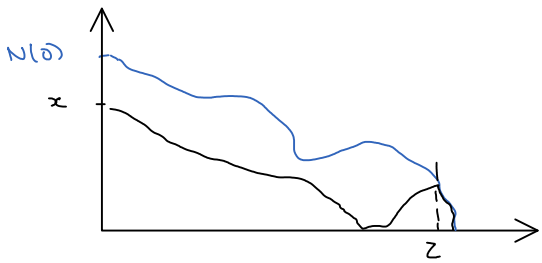
We can now have two cases:

Case 1: $N(0) \leq X(0) = x$



$\tau \leq$ time $X(t)$ first hits 0

Case 2: $N(0) > X(0) = x$



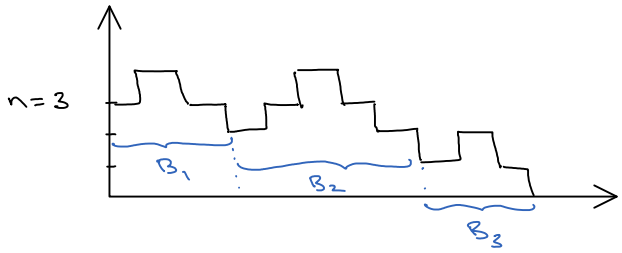
$\tau \leq$ time $N(t)$ first hits 0

Define:

H_n = Hitting to 0 given start with n jobs

Q: What is the distribution of H_n?

A: $\sum_{i=1}^{n} B_i$ where $B_i \sim$ i.i.d. M|M|1 busy periods



So we can write:

$$Pr(Z > t) \leq \sum_{n=0}^{x-1} P_x(H_x > t)\left[(1-\epsilon)\epsilon^n\right] + \sum_{n=x}^{\infty} P_x(H_n > t)\left[(1-\epsilon)\epsilon^n\right] \quad (*)$$

⇒ Reduced coupling time problem to busy period problem.

But we understands Busy periods very well through their transform

$$P_x(H_1 > t) = P_x(B > t) = P_x\left(e^{(\sqrt{\mu}-\sqrt{\lambda})^2 B} > e^{(\sqrt{\mu}-\sqrt{\lambda})^2 t}\right)$$ {A common trick to get tail bounds}

$$\leq \mathbb{E}\left[e^{(\sqrt{\mu}-\sqrt{\lambda})^2 B}\right] e^{-(\sqrt{\mu}-\sqrt{\lambda})^2 t}$$ {Markov's ineq}

$$= \mathcal{L}_B\left(-(\sqrt{\mu}-\sqrt{\lambda})^2\right) e^{-(\sqrt{\mu}-\sqrt{\lambda})^2 t}$$

$$\leq \sqrt{\frac{\mu}{\lambda}}\; e^{-(\sqrt{\mu}-\sqrt{\lambda})^2 t} \qquad \text{(Homework 2)}$$

Similarly: $\mathcal{L}_{H_n}(s) = \left(\mathcal{L}_B(s)\right)^n \Rightarrow$

$$P_x(H_n > t) \leq \left(\sqrt{\frac{\mu}{\lambda}}\right)^n e^{-(\sqrt{\mu}-\sqrt{\lambda})^2 t}$$

Substitute in ⊛

$$P_x(\tau > t) \le \left[\sum_{n=0}^{x-1} \sqrt{\frac{\mu}{\lambda}}^{x} (1-\rho)\rho^n + \sum_{n=x}^{\infty} \sqrt{\frac{\mu}{\lambda}}^{n} (1-\rho)\rho^n \right] e^{-(\sqrt{\mu}-\sqrt{\lambda})^2 t}$$

$$\le \left(\sqrt{\frac{\mu}{\lambda}}^{x} + 1 \right) e^{-(\sqrt{\mu}-\sqrt{\lambda})^2 t}$$

□

Notes :- the rate $e^{-(\sqrt{\mu}-\sqrt{\lambda})^2 t}$ is infact tight, i.e. ∃ matching lower bound

- $\frac{1}{(\sqrt{\mu}-\sqrt{\lambda})^2}$ is also called the "relaxation time" ≡ time until dist to equilibrium decreases by a constant factor

- Exponential Ergodicity : Processes for which

$$\| P_x(t) - \pi \| \le V(x) e^{-r(t)}$$

for strictly positive fns $V(\cdot)$ & $r(\cdot)$ are said to be exponentially or geometrically ergodic

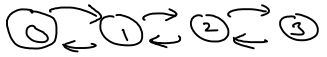
---

QUEUEING NETWORKS - I : REVERSIBILITY, BURKE'S THEOREM

We are moving from single station queueing systems to network of queues. E.g.

- supply chain networks
- assembly of parts inside factory
- Emergency depts as networks of doctors/nurses/CT-Xray/beds
- Computers as a network of CPU, memory, internet bandwidth

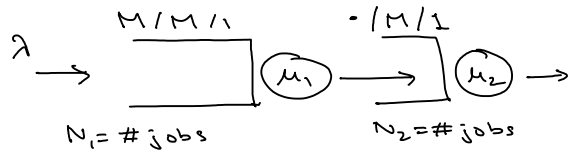
We could solve the CTMCs easily until now

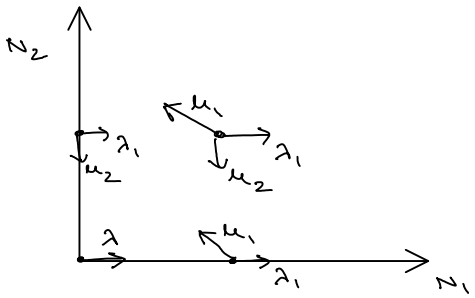


{ write $\pi(i)$ in terms of $\pi(i-1)$ ..
{ find $\pi(0)$ using normalization

But with networks, our state space will become multi-dimensional

E.g. a simple tandem network





How can we solve these CTMCs ?

- if all buffers are finite ⇒ finite state space
  ⇒ can still solve $\pi Q = 0$
- infinite buffers?

Luckily there is a large class of networks for which we can instantly compute the stationary distribution.

"Networks of reversible queueing systems"

Today we will lay the groundwork to understand the intuition behind the most general result, which we will see next week.

Reversible Markov Chains

Consider an ergodic (stationary) CTMC in steady state. $\begin{cases} Q = [q_{ij}] \\ \{\pi_s\} = \text{stat. dist} \end{cases}$

$$\cdots \rightarrow 3 \rightarrow 5 \rightarrow 1 \rightarrow 2 \rightarrow \cdots$$

Now, imagine playing a movie of the above CTMC in reverse

$$\cdots \leftarrow 3 \leftarrow 5 \leftarrow 1 \leftarrow 2 \leftarrow \cdots$$

This is another valid stochastic process. In fact:

Claim: The reverse process is also a CTMC with transition rates

$$q^*_{ij} = \frac{\pi_j q_{ji}}{\pi_i}$$

Proof: (Reverse process is a CTMC)

"View 1 of CTMC": spend time $\text{Exp}(\nu_i)$ on visiting state i and then transition according to P_{ij}

- the reverse process also spends $\text{Exp}(\nu_i)$ in state i
- $P^*_{ij} = \Pr[\text{reverse process goes to } j \mid \text{currently in } i]$
  $= \Pr[\text{forward process came from } j \mid \text{currently in } i]$

(all we need to know is P^*_{ij} is a legitimate prob. dist $\Rightarrow \sum_j P^*_{ij} = 1$ and that transition probs are independent of time spent in states i and j)

(Transition rates)

Consider a long time period T

$$\pi_j q_{ji} = \lim_{T \to \infty} \frac{\#(j \to i) \text{ transitions in } [0,T] \text{ for forward CTMC}}{T}$$

$$= \lim_{T \to \infty} \frac{\#(i \to j) \text{ transitions for reversed CTMC}}{T} = \pi_i^* q^*_{ij}$$

So $\pi_i^* q_{ij}^* = \pi_j q_{ji}$

but $\pi_i^* = \pi_i$ [it is the same movie played in reverse]

$$\Rightarrow \quad q_{ij}^* = \frac{\pi_j q_{ji}}{\pi_i}$$

□

Intuitively, what should reversibility mean?

That the reverse process is indistinguishable from forward process

That is, $\forall i,j$ $\quad q_{ij} = q_{ij}^* = \frac{\pi_j q_{ji}}{\pi_i}$

$$\Rightarrow \quad \pi_i q_{ij} = \pi_j q_{ji}$$

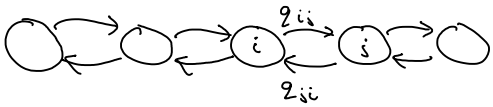
That is all !!

> **Defn:** A CTMC is time-reversible if
> $\pi_i q_{ij} = \pi_j q_{ji} \quad \forall i,j$

> **Thm:** If a CTMC is time-reversible, then the forward and reverse chains are stochastically identical ($q_{ij} = q_{ij}^*$)

Remarks:

1. All CTMCs of the following type are time-reversible



$\pi_i q_{ij} = \pi_j q_{ji}$ are the flow-balance eqns.

2. When presented with a hairy CTMC, one should first check if a π satisfying

$$\pi_i q_{ij} = \pi_j q_{ij}$$ "detailed balance equations"

exists.

3. Metropolis-Hastings: a popular sampling algorithm based on reversibility

↳ want to sample from π : know π_i/π_j but not the normalization const.

⇒ design q_{ij} satisfying detailed balance ⇒ do a random walk to sample from π

Burke's theorem for M/M/1

We already remarked that the M/M/1 CTMC is time reversible

⇒ if we play the tape of an M/M/1 in reverse

- arrivals of forward process → departures of reverse
- departures of forward process → arrivals of reverse process

it still looks like an M/M/1.

This leads to the following incredible result:

> **Burke's theorem:** For an M/M/1 with arrival rate λ, in steady-state
>
> (i) the departure process is Poisson(λ)
>
> (ii) At all times t, number of jobs $N(t)$ is independent of departure process prior to t.

These are non-intuitive

(i) the inter-departure times are sometimes Exp(μ), sometimes Exp(μ) + Exp(λ).

In Homework 1, we proved a single interarrival time was Exp(λ)

Burke's says, they are all independent!!

(ii) One would imagine that knowing a lot of departures happened recently would mean $N(t)$ was small.

Burke: Not so!

Reversibility is not intuitive! The properties it leads to are not intuitive!

But this is precisely why CTMC is a strong tool to have.

Sample path view (e.g. Lindley recursion) does not tell us about these phenomena.

Proof of Burke's theorem:

(i) Forward departure process = Reverse arrival process (always)

$=_{st}$ Forward arrival process (reversibility)

= Poisson(λ)

(ii) For the forward process: $N(t) \perp$ future arrivals



(forward process)

But: future arrivals of forward = prior departures of reverse



(reverse process)

Reverse process is a legit M/M/1

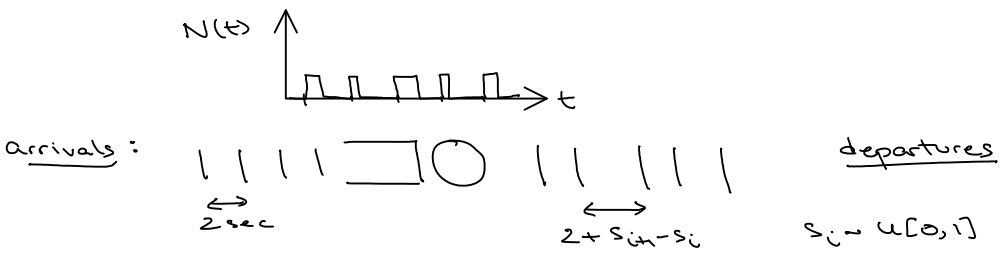
$\Rightarrow$ for an M/M/1 : $N(t) \perp$ (departures prior to t)

□

Q: Would reversibility hold for all single server systems?

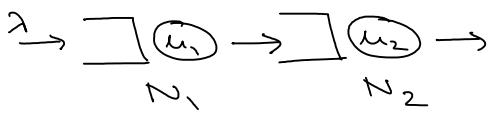
A: No, here is a counter example

D/U/1 – arrivals every 2sec
– Sizes U[0,1]



M/M/1 → ·/M/1 tandem system

Let's look again at



$N_2(t)$ only depends on departures from queue 1 prior to t, and acc to Burke this is independent of $N_1(t)$.

Therefore,

$$Pr[N_1 = n_1, N_2 = n_2] = Pr[N_1 = n_1]\, Pr[N_2 = n_2]$$

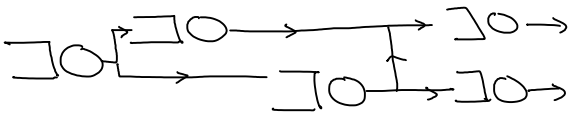
Queue 1 is an M/M/1. By Burke's arrivals to queue 2 are also Poisson(λ)

$$\pi(n_1, n_2) = \underbrace{(1-\rho_1)\rho_1^{n_1}\ (1-\rho_2)\rho_2^{n_2}}_{\text{"Product form distribution"}} \qquad \begin{cases} \rho_1 = \lambda/\mu_1 \\ \rho_2 = \lambda/\mu_2 \end{cases}$$

* you can verify that the above $\pi(n_1, n_2)$ satisfy the balance equations.

Remark: We can extend this line of analysis to general "acyclic" or "feed-forward" networks.

e.g.



Q: Is Burke's theorem true for cyclic networks?

A: No. eg.

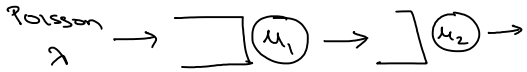


For $\mu \gg \lambda$, an arrival immediately cycles back w.p. p. ⇒ arrivals to queue occur in geometrically distributed batches.

Remarkably: product form still holds for cyclic networks! (Next time)

QUEUEING NETWORKS - II : SINGLE-CLASS (OPEN/CLOSED) JACKSON NETWORKS

Last lecture: M/M/1 → ·/M/1



We proved: $Pr(N_1=n_1, N_2=n_2) = Pr(N_1=n_1)\,Pr(N_2=n_2)$ { "product form distribution"

$= (1-\rho_1)\rho_1^{n_1}\,(1-\rho_2)\rho_2^{n_2}$

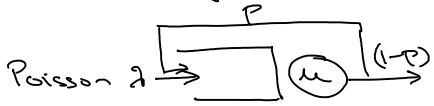
Proof Outline: reversibility of server 1 M/M/1

⇒ (i) departures from Q1 = arrivals to Q_2 are Poisson(λ)

(ii) $N_2(t)$ = function of Q1 departures prior to t ⊥ $N_1(t)$

Extends to acyclic networks, but not for cyclic.

Example: simplest cyclic network

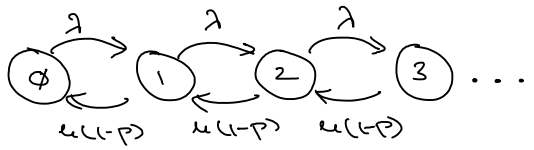


aggregate arrivals to buffer are not Poisson (let $\mu >> \lambda$)

GOAL for today: a more general methodology that allows us to analyze a large class of cyclic networks.

Note that the M/M/1 with retrials is still "nice":

CTMC:



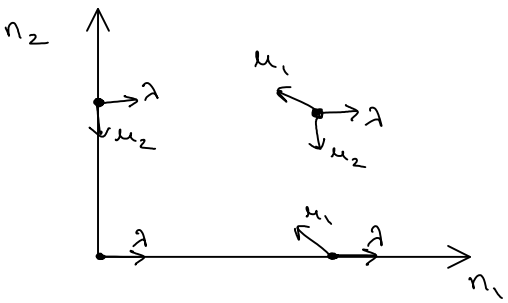
This CTMC is reversible! So, aggregate departures are Poisson even though aggregate arrivals are not!

May be we can say that cyclic network are time reversible?

Let's check if tandem $M/M/1 \to \cdot/M/1$ is time reversible?

time reversibility $\Leftrightarrow$ $\forall$ states i,j $\quad \pi_i q_{ij} = \pi_j q_{ji}$

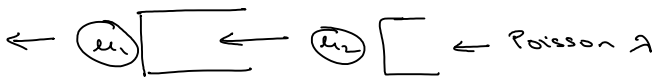
Transition diagram:



Trouble: if $q_{ij} > 0$, then $q_{ji} = 0$ No matching transitions!!

CTMC is not time reversible.

But reversed process is a valid CTMC. Maybe we can think of it as another queng network? (Goal: find q^*_{ij})



Guess: reverse process is another tandem $M/M/1 \to \cdot/M/1$

(i) Poisson λ arrivals to queue 2

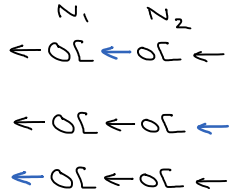
(ii) Queue 2 departures = Queue 1 arrivals

Q: q^*_{ij} for this system?

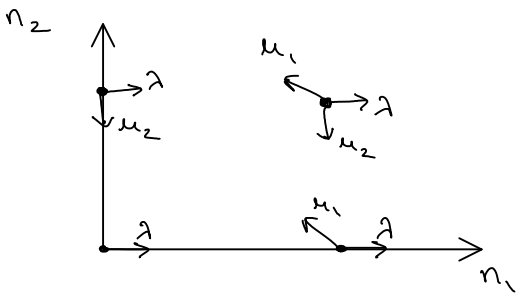
A: $q^*((n_1,n_2),(n_1+1,n_2-1)) = \mu_2$

$q^*((n_1,n_2),(n_1,n_2+1)) = \lambda$

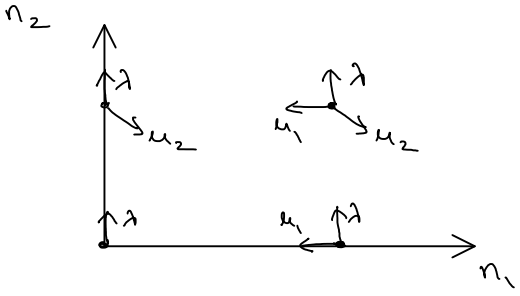
$q^*((n_1,n_2),(n_1-1,n_2)) = \mu_1$



Transition diagrams:



FORWARD PROCESS



REVERSE PROCESS.

Notation: $n \equiv (n_1, n_2, \ldots)$

$n + e_i \equiv (n_1, n_2, \ldots, n_i + 1, n_{i+1} \ldots)$

$n + e_i - e_j \equiv (n_1, \ldots, n_i + 1, \ldots, n_j - 1, \ldots)$

Assuming guess is correct: $\pi_i q_{ij} = \pi_j q^*_{ji}$ (Always)

let $i \equiv n = (n_1, n_2)$, $j = (n - e_1 + e_2)$

$\Rightarrow \quad \pi(n) \underbrace{q_{ij}}_{\mu_1} \stackrel{?}{=} \pi(n - e_1 + e_2) \underbrace{q^*_{ji}}_{\mu_2}$

Substitute $\pi(n) = (1-\rho_1)(1-\rho_2)\rho_1^{n_1}\rho_2^{n_2}$

$$\Leftrightarrow \quad (1-\rho_1)(1-\rho_2)\rho_1^{n_1}\rho_2^{n_2} \cdot \mu_1 \stackrel{?}{=} (1-\rho_1)(1-\rho_2)\rho_1^{n_1-1}\rho_2^{n_2+1} \cdot \mu_2$$

$$\Leftrightarrow \quad \rho_1 \mu_1 \stackrel{?}{=} \rho_2 \mu_2$$

$$\Leftrightarrow \quad \lambda/\mu_1 \cdot \mu_1 \stackrel{?}{=} \lambda/\mu_2 \cdot \mu_2 \qquad \checkmark$$

So the guess seems to check out.

How do we know our guess for q^*_{ij} are correct?

Thm: For a CTMC with generator $Q = [q_{ij}]$, if $\exists\ \pi_i$ $(\sum_i \pi_i = 1)$, q^*_{ij} satisfying

(i) $\pi_i q^*_{ij} = \pi_j q_{ji} \quad \forall i, j$

(ii) $\sum_j q^*_{ij} = \sum_j q_{ij} \quad \forall i$

then π are the stationary probs, and q^*_{ij} are the rates for reverse chain.

Proof: Sum (i) over j

$$\sum_j (\pi_i q^*_{ij}) = \sum_j \pi_j q_{ji}$$

$$\Rightarrow \quad \pi_i \sum_j q^*_{ij} = \sum_j \pi_j q_{ji} \qquad \text{using (ii)}$$

$$\Rightarrow \quad \pi_i \sum_j q_{ij} = \sum_j \pi_j q_{ji}$$

which are the balance equations for forward chain. □

Steps for solving for π for queueing networks:

(i) guess the queueing process corresp. to reverse CTMC
(ii) use the guess to obtain q^*_{ij}
(iii) verify $\sum_j q^*_{ij} = \sum_i q_{ij}$
(iv) verify $\{\pi_i q^*_{ij} = \pi_j q_{ji}\}$ yield a valid distribution (i.e. true for all i,j pairs)

Re-visiting $M/M/1 \to \cdot/M/1$

The $\pi_i q_{ij} = \pi_j q^*_{ji}$ in this case are:

	Forward	Reverse
(i) $\pi(n_1,n_2)\mu_1 = \pi(n_1-1,n_2+1)\mu_2$	Completion at 1	Completion at 2
(ii) $\pi(n_1,n_2)\mu_2 = \pi(n_1,n_2-1)\lambda$	Completion at 2	Arrival at 2
(iii) $\pi(n_1,n_2)\lambda = \pi(n_1+1,n_2)\mu_1$	Arrival at 1	Completion at 1

massage these so all RHS are $\pi(n_1,n_2)$:

Rate enter (n_1,n_2) = Rate leave (n_1,n_2)	
(i) $\pi(n_1+1,n_2-1)\mu_1 = \pi(n_1,n_2)\mu_2$	due to arrival to server 2 = due to departure from server 2
(ii) $\pi(n_1,n_2+1)\mu_2 = \pi(n_1,n_2)\lambda$	due to external departure = due to external arrival
(iii) $\pi(n_1-1,n_2)\lambda = \pi(n_1,n_2)\mu_1$	due to arrival to server 1 = due to departure from server 1

"Local-balance equations"

Remark: The following weaker phrasing of $\pi_i q_{ij} = \pi_j q^*_{ji}$ is called "local balance" or "detailed balance"

(A) rate entering state n due to arrival at server i = rate of leaving state n due to completion at server i

(B) rate entering state n due to departure to outside = rate of leaving state n due to arrival from outside

Local balance holds for Jackson networks, but not in general.

SINGLE-CLASS OPEN JACKSON NETWORKS

* Network of J servers
* at server i
  - arrivals from outside : Poisson with rate α_i
  - service rate depending on n_i : $\mu_i(n_i)$
  - routing probabilities : $P_{i0} = Pr\left(\text{job leaves system after completion at server } i\right)$

    $P_{ij} = Pr\left(\text{job routed to server } j \text{ after completion at server } i\right)$

Assumption: $P = [P_{ij}]$ $(i,j = 1,2,\ldots,J)$ has spectral radius less than 1.

$P^{(n)}_{ij} = Pr[\text{external arrival to server } i \text{ is at server } j \text{ after } n \text{ service completions}]$

spectral radius$(P) < 1 \Rightarrow P^n \to 0 \Rightarrow$ all jobs leave system eventually

Q: What is the total arrival rate to server i?

A: If λ_i = total arrival rate to server i

$$\lambda_i = \alpha_i + \sum_j \lambda_j P_{ji}$$

or, in matrix notation: $\lambda = \begin{bmatrix} \lambda_1 \\ \vdots \\ \lambda_J \end{bmatrix}$, $\alpha = \begin{bmatrix} \alpha_1 \\ \vdots \\ \alpha_J \end{bmatrix}$

$$\lambda = \alpha + P^T \lambda \quad \Rightarrow \quad \lambda = \underbrace{(I - P^T)^{-1}}_{\text{invertible when sp.radius}(P) < 1} \alpha$$

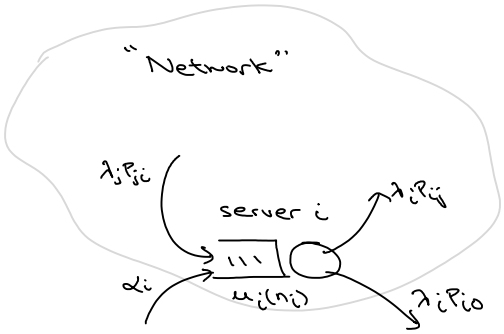
Remark: $\mu_i(n_i)$ allows us to model many single-station queues

(i) M/M/1 : $\mu_i(n_i) = \mu_i$ for $n_i > 0$

(ii) M/M/m_i : $\mu_i(n_i) = \mu_i \min(m_i, n_i)$

(iii) M/M/∞ : $\mu_i(n_i) = \mu_i n_i$

Summarizing:



Step 1: Guess stoch. process corresp. to reverse chain

$\Rightarrow$ Jackson network with parameters α_i^*, P_{ij}^*

Q: What should α_i^* be?

A: External arrival rate to server i for reverse process

= departure rate to outside from server i for forward process

= $\lambda_i P_{io}$

Q: What should P_{io}^* be?

A: departure rate to outside from server i in reverse = $\lambda_i P_{io}^*$

= arrival rate from outside to server i in forward = α_i

$$\Rightarrow \quad P_{io}^* = \frac{\alpha_i}{\lambda_i}$$

Q: P_{ij}^* ?

A: Total $i \to j$ rate in reverse = $\lambda_i P_{ij}^*$ $\quad \Rightarrow P_{ij}^* = \dfrac{\lambda_j P_{ji}}{\lambda_i}$

= Total $j \to i$ rate for forward = $\lambda_j P_{ji}$

Thm: For an open Jackson network with parameters $\{\alpha_i\}, \{P_{ij}\}, \{\mu_i(\cdot)\}$

(i) the reverse process is another open Jackson network with

- state-dependent service rates $\mu_i(n_i)$
- external arrival rates $\alpha_i^* = \lambda_i P_{i0}$
- routing probabilities $P_{i0}^* = \alpha_i / \lambda_i$

$$P_{ij}^* = \lambda_j P_{ji} / \lambda_i$$

(ii) the stationary distribution $\{\pi(n)\}$ is:

$$\pi(n_1, n_2, \cdots, n_J) = \pi_1(n_1)\,\pi_2(n_2) \cdots \pi_J(n_J)$$

where $$\pi_i(n_i) = \left(\prod_{k=1}^{n_i} \frac{\lambda_i}{\mu_i(k)} \right) \underbrace{\left[1 + \sum_{n=1}^{\infty} \prod_{k=1}^{n} \frac{\lambda_i}{\mu_i(k)} \right]^{-1}}_{\text{normalization constant}}$$

Remark: (i) The steady state distribution has product form

$\Rightarrow N_1, N_2, \ldots N_J$ are independent and are distributed as M/M/1 with arrival rate λ_i & state-dependent service $\mu_i(n_i)$

(ii) The theorem implies that the streams of external departures are independent Poisson, even though the internal flows are not Poisson!

(iii) for the simple case of $\mu_i(n_i) = \mu_i \quad [n_i > 0]$

$$\pi(n_1, n_2, \ldots, n_J) = (1-\rho_1)\rho_1^{n_1} \cdot (1-\rho_2)\rho_2^{n_2} \cdot \cdots \cdot (1-\rho_J)\rho_J^{n_J}$$

$$(\rho_i = \lambda_i / \mu_i)$$

Proof: You should verify that $\sum_j q_{ij}^* = \sum_j q_{ij}$ holds

We will see how to use $\pi_i q_{ij}^* = \pi_j q_{ji}$ to guess the form for $\pi(n)$.

- without loss of generality, assume $\alpha_1 > 0$

Consider $i = (n_1, 0, 0, \cdots 0)$ $\quad q_{ji} = \alpha_1$

$j = (n_1 - 1, 0, 0, \cdots 0)$ $\quad q_{ij}^* = \mu_1(n_1) P_{i0}^* = \mu_1(n_1) \frac{\alpha_1}{\lambda_1}$

Now $\pi_i q_{ij}^* = \pi_j q_{ji}$

$$\Rightarrow \pi(n_1, 0, \cdots 0)\, \mu_1(n_1) \frac{\alpha_1}{\lambda_1} = \pi(n_1 - 1, 0, \cdots 0)\, \alpha_1$$

$$\Rightarrow \pi(n_1, 0, 0 \cdots, 0) = \pi(n_1 - 1, 0, 0, \cdots) \frac{\lambda_1}{\mu_1(n_1)}$$

$$\vdots$$

$$= \pi(0, 0, \cdots 0) \left(\prod_{k=1}^{n_1} \frac{\lambda_1}{\mu_1(k)} \right)$$

- without loss of generality assume $P_{12} > 0$

Consider $i = (n_1, n_2, 0, \cdots 0)$ $\quad q_{ji} = \mu_1(n_1 + 1) P_{12}$

$j = (n_1 + 1, n_2 - 1, 0, \cdots 0)$ $\quad q_{ij}^* = \mu_2(n_2) P_{21}^* = \mu_2(n_2) \frac{\lambda_1 P_{12}}{\lambda_2}$

Now $\pi_i q_{ij}^* = \pi_j q_{ji}$

$$\Rightarrow \pi(n_1, n_2, 0, \cdots 0)\, \mu_2(n_2) \frac{\lambda_1 P_{12}}{\lambda_2} = \pi(n_1 + 1, n_2 - 1, \cdots 0)\, \mu_1(n_1 + 1) P_{12}$$

$$\Rightarrow \pi(n_1, n_2, 0, \cdots 0) = \pi(n_1 + 1, n_2 - 1, 0, \cdots 0) \frac{(\lambda_2 / \mu_2(n_2))}{(\lambda_1 / \mu_1(n_1 + 1))}$$

$$\vdots$$

$$= \pi(n_1 + n_2, 0, \cdots 0) \frac{\left(\prod_{k=1}^{n_2} \frac{\lambda_2}{\mu_2(k)} \right)}{\left(\prod_{k=n_1+1}^{n_1+n_2} \frac{\lambda_1}{\mu_1(k)} \right)}$$

$$= \pi(0, 0, \cdots 0) \left(\prod_{k=1}^{n_1+n_2} \frac{\lambda_1}{\mu_1(k)} \right) \left(\prod_{k=1}^{n_2} \frac{\lambda_2}{\mu_2(k)} \right) \Bigg/ \left(\prod_{k=n_1+1}^{n_1+n_2} \frac{\lambda_1}{\mu_1(k)} \right)$$

$$= \pi(0, 0, 0 \cdots) \left(\prod_{k=1}^{n_1} \frac{\lambda_1}{\mu_1(k)} \right) \left(\prod_{k=1}^{n_2} \frac{\lambda_2}{\mu_2(k)} \right)$$

You get the general idea ...

Now verify that $\pi_i q_{ij}^* = \pi_j q_{ji}$ hold for all transitions

(i) $\pi(n)\, \underbrace{q(n, n+e_i)}_{\text{external arrival to } i} = \pi(n+e_i)\, \underbrace{q^*(n+e_i, n)}_{\text{external departure from } i}$

$= \alpha_i \qquad = \mu_i(n_i+1)\frac{\alpha_i}{\lambda_i}$

(ii) $\pi(n)\, \underbrace{q(n, n-e_i)}_{\text{ext. departure from } i} = \pi(n-e_i)\, \underbrace{q^*(n-e_i, n)}_{\text{ext. arrival to } i}$

$= \mu_i(n_i) P_{i0} \qquad = \alpha_i^* = \lambda_i P_{i0}$

(iii) $\pi(n)\, \underbrace{q(n, n-e_i+e_j)}_{\text{job shuffle } i\to j} = \pi(n-e_i+e_j)\, \underbrace{q^*(n-e_i+e_j, n)}_{\text{job shuffle } j\to i}$

$= \mu_i(n_i) P_{ij} \qquad = \mu_j(n_j+1) P_{ji}^*$

$= \mu_j(n_j+1)\frac{\lambda_i P_{ij}}{\lambda_j}$

□

SINGLE-CLASS CLOSED JACKSON NETWORKS

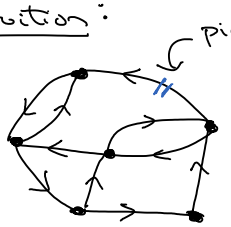
Closed — no external arrivals ($\alpha_i = 0$)
— no external departures ($P_{i0} = 0$)
} a fixed population of N jobs in the network

Assumption: $P = [P_{ij}]$ is irreducible

$\lambda_i = \alpha_i + \sum_j \lambda_j P_{ji}$ becomes $\theta_i = \sum \theta_j P_{ji}$

"visit ratios"

Intuition:



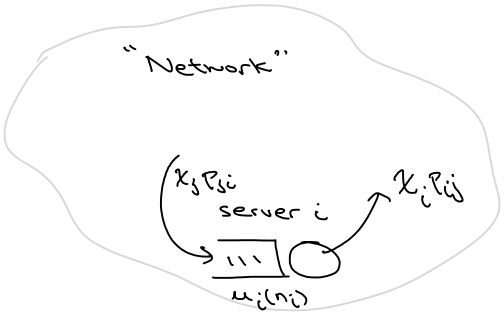
$\chi \equiv$ rate at which jobs cross this link "throughput"

$\theta_i \equiv \mathbb{E}\left[\#\text{ visits to server } i \text{ b/w two consecutive traversals of the link}\right]$

Remark: (1) different choice of links changes χ & v_i
but $\frac{v_i}{v_j}$ is invariant (hence "visit ratio")

(2) rate of completion of jobs at server $i \equiv \chi_i = \chi v_i$
↳ $\mathbb{E}$[time b/w traversals of links] $= 1/\chi$
within ↗ period we visit i an expected v_i times
$\Rightarrow \mathbb{E}$[time b/w visiting server i] $= \frac{1}{\chi v_i}$
$\Rightarrow$ throughput at server $i = \chi v_i$

Summarizing:



Step 1: Guess stoch. process corresponding to reverse chain
$\Rightarrow$ closed Jackson network with routing P^*_{ij}

Q: What should P^*_{ij} be?
A: Total rate of jobs transferred $i \to j$ in reverse $= \chi_i P^*_{ij}$
$= \quad " \quad " \quad " \quad " \quad "$ $j \to i$ in forward $= \chi_j P_{ji}$

$$\Rightarrow P^*_{ij} = \frac{\chi_j P_{ji}}{\chi_i} = \frac{v_j}{v_i} P_{ji}$$

Q: Do we expect $\pi(n_1, n_2, \ldots n_J) = \pi_1(n_1) \ldots \pi_J(n_J)$?
A: No. $\pi_i(n_i) \neq$ for $n_i < N$
but $\pi(n_1, n_2, \ldots n_J) = 0$ if $n_1 + n_2 + \cdots + n_J \neq N$

Thm: For a closed single-class Jackson network with parameters $\{P_{ij}\}$, $\{\mu_i(\cdot)\}$

(i) the reverse process is another open Jackson network with
- state-dependent service rates $\mu_i(n_i)$
- routing probabilities $p^*_{ij} = \frac{\vartheta_j}{\vartheta_i} P_{ji}$

(ii) the stationary distribution $\{\pi(n)\}$ is:

$$\pi(n_1, n_2, \ldots, n_J) = \frac{1}{G(J,N)} \prod_{i=1}^{J} g_i(n_i) \quad \text{for} \quad \sum n_i = N$$

where,

$$g_i(n_i) = \prod_{k=1}^{n_i} \frac{\vartheta_i}{\mu_i(k)}$$

$$G(j,x) = \sum_{|n|_1 = x} \prod_{i=1}^{j} g_i(n_i)$$

Proof: (Exercise) □

Remarks:

(1) $G(J,N)$ is the normalizing constant and is called the "partition function"

(2) The throughput χ can be expressed in terms of $G(\cdot)$ as

$$\chi = \frac{G(J, N-1)}{G(J,N)}$$

(3) $$G(j,n) = \sum_{l=0}^{n} G(j-1, n-l)\, g_j(l)$$

which gives the following recursive algorithm

The Convolution algorithm:

1. Initialize: $G(1,n) = g_1(n) \quad n = 0,1,\ldots,N$; $G(j,0) = 1 \quad j = 1,2,\ldots J$
2. For $j = 2, \ldots, J$

   For $n = 1, 2, \ldots, N$

   $$G(j,n) = \sum_{l=0}^{n} G(j-1, n-l)\, g_j(l)$$

□

Q: Consider $\mu_i(n_i) = \mu_i$ [vanilla $\cdot/M/1$ servers]

What happens as $N \to \infty$

A: Assume $\frac{\nu_1}{\mu_1} > \frac{\nu_2}{\mu_2} > \dots > \frac{\nu_J}{\mu_J}$

Note that: $\pi(n_1, n_2, \dots n_J) \propto \left(\frac{\nu_1}{\mu_1}\right)^{n_1} \left(\frac{\nu_2}{\mu_2}\right)^{n_2} \cdots \left(\frac{\nu_J}{\mu_J}\right)^{n_J}$

$\Rightarrow$ as $N \to \infty$, states where $n_i = \omega(1)$ for $i \neq 1$ are exponentially rare

$\Rightarrow$ all but $\Theta(1)$ jobs accumulate at server 1

server 1 becomes "bottleneck server" or "rate limiting server"

$\Rightarrow$ $\chi_i = \mu_1 \frac{\nu_i}{\nu_1}$ $\forall i$

$N_1 \to \infty$ but we can find the marginal distribution $N_2, N_3, \dots N_J$

$$\pi(n_2, n_3, \dots, n_J) = \pi_2(n_2)\,\pi_3(n_3) \cdots \pi_J(n_J)$$

where: $$\left[\pi_i(n_i) = \left(1 - \frac{\mu_1/\nu_1}{\mu_i/\nu_i}\right)\left(\frac{\mu_1/\nu_1}{\mu_i/\nu_i}\right)^{n_i}\right]$$

□

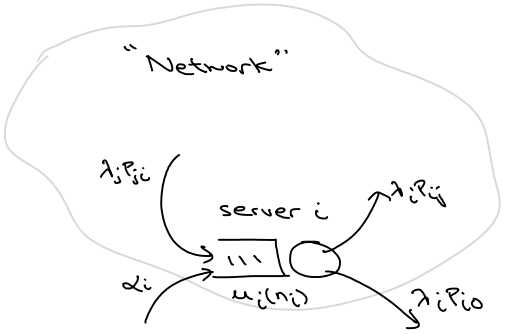
Next time: Symmetric policies

Classed Jackson networks

BCMP theorem.

QUEUEING NETWORKS - III : SYMMETRIC POLICIES, MULTICLASS (OPEN) JACKSON NETWORKS

Last lecture : Jackson network of M/M/1 type queues exhibits product form stationary distribution.:



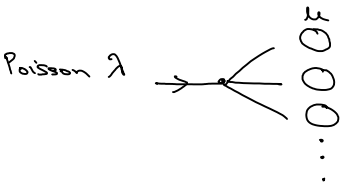
$$\pi(n_1, n_2, \ldots n_J) = \pi_1(n_1)\,\pi_2(n_2)\ldots\ldots\ \pi_J(n_J)$$

where: $\pi_i(n_i) = P_r$ [server i has n_i jobs in isolation under Poisson λ_i arrivals]

Goals for today : Multiclass Jackson networks
- ↳ routing probs. depend on class of jobs
- ↳ allow generally distributed sized under special scheduling policies (product form compatible/ symmetric)

We will look at two simple examples of symmetric queues before the most general result.

Example 1 : M/G/∞ (infinite server system)

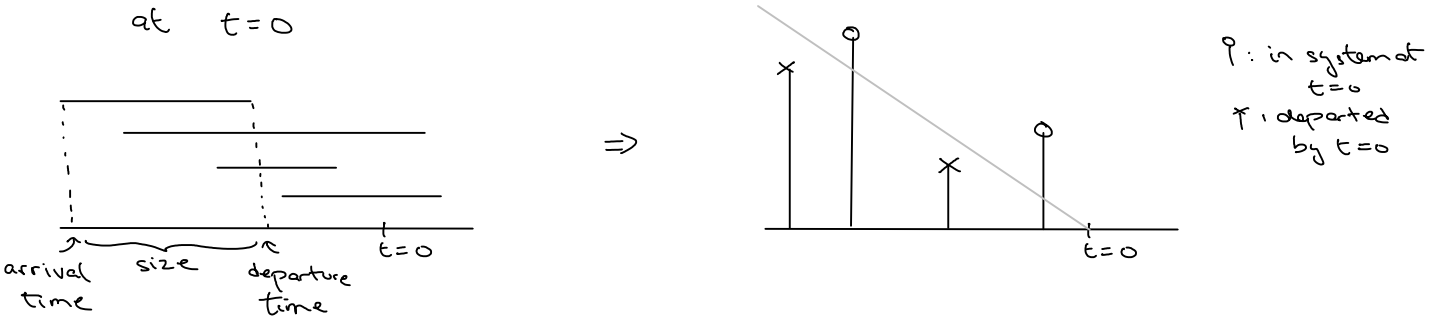


Recall M/M/∞ : $P_\pi[N^{M/M/\infty} = n] = e^{-\rho}\,\frac{\rho^n}{n!}$

where $\rho = \lambda/\mu$.

Thm: For the M/G/∞ system with general service distribution with mean $\mathbb{E}[S]$ under steady-state
(i) $N \sim$ Poisson$(\lambda\,\mathbb{E}[S])$
(ii) departure process is Poisson(λ)

Proof: Imagine we started the system at $t=-\infty$ and observe it at $t=0$



$$\Pr[\text{an arrival at } t=-s \text{ is around at } t=0] = \Pr[\text{size} > s] = \overline{F}(s)$$

Therefore: (# jobs at $t=0$) = (# arrivals of a non-stationary Poisson process in $(-\infty, 0]$ with intensity $\lambda(t) = \lambda \overline{F}(-t)$)

$$\sim \text{Poisson with mean } \underbrace{\int_{-\infty}^{0} \lambda \overline{F}(-t)\, dt}_{\lambda \mathbb{E}[S] = \rho}$$

For second part: assume we are only watching jobs with size $(x, x+\delta)$

arrival process $\equiv$ Poisson with rate $\lambda f(x)\delta$

departure process $\equiv$ arrival process shifted forward by x time units

$\equiv$ Poisson with rate $\lambda f(x)\delta$

Now: arrivals = merge of infinitely many independent Poisson streams

departures = merge of infinitely many independent Poisson streams shifted by different amounts

$\equiv$ still Poisson(λ)

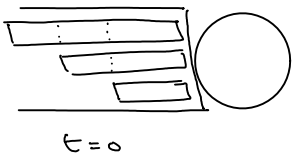
□

Example 2: $M/H_2/1/PS$ (Processor Sharing)

Processor sharing: n jobs in system $\Rightarrow$ each job gets $1/n$ fraction of server's capacity

Example:

$S_3 = 3$
$S_2 = 2$
$S_1 = 1$



$t = 0$

Assuming no arrivals:

S_1 departs at $t=3$
S_2 departs at $t=5$
S_3 departs at $t=6$

We will get to the full analysis of M/G/1/PS later in the lecture.

For warm up: H_2 job sizes (2-phase hyperexponential)

S → Exp(μ_1) w.p. p
S → Exp(μ_2) w.p. $(1-p)$

$\mathbb{E}[S] = \frac{p}{\mu_1} + \frac{1-p}{\mu_2}$; $\rho_1 \equiv \frac{\lambda p}{\mu_1}$ ("load due to phase 1")

$\rho = \lambda \mathbb{E}[S]$; $\rho_2 \equiv \frac{\lambda(1-p)}{\mu_2}$ ("load due to phase 2")

State $\equiv (n_1, n_2)$

n_1: # phase 1 jobs in system
n_2: # phase 2 jobs in system

Transitions:

$(n_1, n_2) \to (n_1+1, n_2)$: $q(\,,\,) = \lambda p$

$(n_1, n_2) \to (n_1, n_2+1)$: $q(\,,\,) = \lambda(1-p)$

$(n_1, n_2) \to (n_1-1, n_2)$: $q(\,,\,) = n_1\mu_1/(n_1+n_2)$

$(n_1, n_2) \to (n_1, n_2-1)$: $q(\,,\,) = n_2\mu_2/(n_1+n_2)$

Each of (n_1+n_2) jobs gets $\frac{1}{n_1+n_2}$ fraction of capacity

$\Rightarrow$ departure rate of a phase 1 job slows down to $\mu_1/(n_1+n_2)$

$\Rightarrow$ total departure rate of phase 1 jobs $= \mu_1 n_1/(n_1+n_2)$

Claim: M/H_2/1/PS CTMC is time reversible.

Proof: want to show $\forall$ state pairs i, j : $\pi_i q_{ij} = \pi_j q_{ji}$

$(n_1, n_2) \leftrightarrow (n_1-1, n_2)$

$$\Rightarrow \quad \pi(n_1,n_2)\, \mu_1 \frac{n_1}{n_1+n_2} = \pi(n_1-1,n_2)\, \lambda p$$

$$\Rightarrow \quad \pi(n_1,n_2) = \pi(n_1-1,n_2) \left(\frac{\lambda p}{\mu_1}\right) \frac{n_1+n_2}{n_1} = \pi(n_1-1,n_2)\, \rho_1 \frac{n_1+n_2}{n_1}$$

repeating:

$$= \pi(n_1-2, n_2) (\rho_1)^2 \frac{(n_1+n_2)(n_1+n_2-1)}{n_1(n_1-1)}$$

$$\vdots$$

$$= \pi(0, n_2) (\rho_1)^{n_1} \frac{(n_1+n_2)\cdots(n_2+1)}{n_1!} \qquad (*)$$

$(n_1, n_2) \leftrightarrow (n_1, n_2-1)$

$$\Rightarrow \quad \pi(n_1,n_2) \frac{\mu_2 n_2}{n_1+n_2} = \pi(n_1,n_2-1)\, \lambda(1-p)$$

$$\Rightarrow \quad \pi(n_1,n_2) = \pi(n_1,n_2-1) \left(\frac{\lambda(1-p)}{\mu_2}\right) \frac{n_1+n_2}{n_2} = \pi(n_1,n_2-1)\, \rho_2 \frac{n_1+n_2}{n_2} \qquad (**)$$

Substituting $(**)$ into $(*)$

$$\pi(n_1,n_2) = \pi(0)\, \rho_1^{n_1} \rho_2^{n_2} \frac{(n_1+n_2)!}{n_1!\, n_2!}$$

$$= \pi(0)\, \rho^n \left(\frac{\rho_1}{\rho}\right)^{n_1} \left(\frac{\rho_2}{\rho}\right)^{n_2} \frac{(n_1+n_2)!}{n_1!\, n_2!}$$

$$* \quad P_\pi[N=n] = \sum_{n_1+n_2=n} \pi(n_1,n_2) = \pi(0)\, \rho^n \sum_{n_1+n_2=n} \left(\frac{\rho_1}{\rho}\right)^{n_1} \left(\frac{\rho_2}{\rho}\right)^{n_2} \frac{(n_1+n_2)!}{n_1!\, n_2!}$$

$$= \pi(0)\, \rho^n \left(\frac{\rho_1}{\rho} + \frac{\rho_2}{\rho}\right)^n = \pi(0)\, \rho^n$$

$$\Rightarrow \quad \pi(n) = (1-\rho)\rho^n \qquad \text{(same as an M/M/1 !)}$$

$$* \quad P_\pi[N_1=n_1, N_2=n_2 \mid N=n] = \left(\frac{\rho_1}{\rho}\right)^{n_1} \left(\frac{\rho_2}{\rho}\right)^{n_2} \binom{n_1+n_2}{n_1}$$

Summary: (i) The stat. distrib. of # jobs in M/H$_2$/1/PS only depends on $\mathbb{E}[S]$

(ii) Conditioned on $\{N=n\}$, each job is independently Exp(μ_1) w.p. $\left(\frac{\rho_1}{\rho}\right)$ & Exp(μ_2) otherwise

(iii) time reversibility $\Rightarrow$ departures of phase 1 & phase 2 jobs form independent Poisson streams.

Symmetric scheduling policies

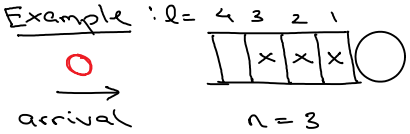
We now turn to the most general single station queueing system for which we can claim a reversibility / insensitivity type result.

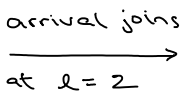
Consider a multiclass queueing system

- C job classes
  - arrivals of class are Poisson with rate λ_c
  - generally distributed job sizes (distribution F_c, mean $\mathbb{E}[S_c]$, $\rho_c = \lambda_c \mathbb{E}[S_c]$)
- server : when n jobs in system
  - service rate $\phi(n)$
  - arrival joins position ℓ with prob. $\gamma(n+1,\ell)$, $\ell = 1,2,..,n+1$.
  - job at position ℓ get $\delta(n,\ell)$ of server's capacity , $\ell = 1,2,..n$.

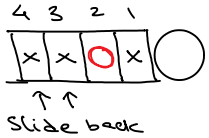
"Symmetric" $\equiv$ $\delta(n,\ell) = \gamma(n,\ell)$ $\forall n,\ell$

Example : $\ell =$ 4 3 2 1





$\ell =$ 4 3 2 1



Processor sharing (PS) : $\phi(n) = 1$ $\quad \gamma(n,\ell) = \delta(n,\ell) = 1/n$

Infinite server (IS) : $\phi(n) = n$ $\quad \gamma(n,\ell) = \delta(n,\ell) = 1/n$

Pre-emptive Last Come First Served (PLCFS) : $\phi(n) = 1$ $\quad \gamma(n,1) = 1$, $\delta(n,1) = 1$

Q: state space ?

A: need to keep track of
- # jobs
- classes of all jobs
- position in queue
- information about size.

Let:

$$x = (c_1, c_2, \dots c_i, \dots c_n, x_1, x_2, \dots x_i, \dots x_n)$$

c_i: class of job in position i; x_i: attained service (age) of job in position i

Remark: Continuous state space $\Rightarrow$ Markov process

Claim: The reversed process is also a symmetric queue with the same parameters but with the residual sizes in the state space description.

Proof: Possible transitions

(1) departure of job in position i

$$x \to x - e_i \equiv (c_1, \dots, c_{i-1}, c_{i+1}, \dots, c_n, x_1, \dots, x_{i-1}, x_{i+1}, \dots x_n)$$

rate for forward process $= \phi(n) \cdot \gamma(n,i) \lim_{\varepsilon \to 0} \frac{\Pr[\text{size of } i^{th} \text{ job} \in (x_i, x_i+\varepsilon)]}{\Pr[\text{size of } i^{th} \text{ job} > x_i]}$

$\phi(n)$: total service capacity; $\gamma(n,i)$: fraction of capacity job i gets; the limit: departure rate = hazard rate if received service capacity was 1

$$= \phi(n)\, \gamma(n,i)\, \frac{f_{c_i}(x_i)}{\bar{F}_{c_i}(x_i)}$$

rate for $(x - e_i) \to x$ for reverse process $= \lambda_{c_i}\, \delta(n,i)\, f_{c_i}(x_i)$

λ_{c_i}: arrival of class c_i; $\delta(n,i)$: joins position i; $f_{c_i}(x_i)$: arrives with size x_i

$$\pi(x)\, q(x, x-e_i) = \pi(x-e_i)\, q^*(x-e_i, x)$$

$$\Rightarrow \quad \pi(x)\, \phi(n)\, \gamma(n,i)\, \frac{f_{c_i}(x_i)}{\bar{F}_{c_i}(x_i)} = \pi(x-e_i)\, \lambda_{c_i}\, \delta(n,i)\, f_{c_i}(x_i)$$

$$\Rightarrow \quad \pi(x) = \pi(x-e_i)\, \frac{\lambda_{c_i}}{\phi(n)} \cdot \bar{F}_{c_i}(x_i) \qquad \left(\gamma(n,i) = \phi(n,i)\right)$$

$$= \pi(x-e_i)\, \frac{\lambda_{c_i} \mathbb{E}[S_{c_i}]}{\phi(n)}\, \frac{\bar{F}_{c_i}(x_i)}{\mathbb{E}[S_{c_i}]}$$

$$= \pi(x-e_i)\, \frac{\rho_{c_i}}{\phi(n)}\, f_{e,c_i}(x_i)$$

[$f_{e_{c_i}}$ = density of stat. excess for F_{c_i}]

repeating:

$$\pi(x) = \pi(x-e_i)\ \frac{\rho_{c_i}}{\phi(n)}\ f_{e,c_i}(x_i)$$

$$\vdots$$

$$= \pi(0)\ \frac{\rho_{c_1}\rho_{c_2}\cdots\rho_{c_n}}{\phi(1)\,\phi(2)\cdots\phi(n)}\ f_{e,c_1}(x_1)\cdots f_{e,c_n}(x_n)$$

To verify our conjecture; check

(2) arrival of class c job in position i

$$x \to x + e_i^c = (c_1, c_2, \cdots c_{i-1}, c, c_i, \cdots, c_n, x_1, \cdots x_{i-1}, 0, x_i, \cdots x_n)$$

$$\Rightarrow \quad \pi(x)\ \lambda_c\ \delta(n+1,i) \overset{?}{=} \pi(x+e_i^c)\ \phi(n+1)\ \gamma(n+1,i)\cdot 1$$

↳ departure prob given residual size = 0

$$\Rightarrow \quad \pi(x+e_i^c) \overset{?}{=} \pi(x)\ \frac{\lambda_c\,\mathbb{E}[S_c]}{\phi(n+1)}\cdot\frac{1}{\mathbb{E}[S_c]}$$

$$= \pi(x)\ \frac{\rho_c}{\phi(n+1)}\ f_{e,c}(0) \qquad \checkmark$$

Therefore our conjecture is verified. □

$$Pr[N=n] = \frac{\pi(0)}{\phi(1)\cdots\phi(n)} \sum_{c_1,\cdots c_n} \rho_{c_1}\rho_{c_2}\cdots\rho_{c_n}\left(\int_{x_1} f_{e,c_1}(x)\,dx\right)\left(\int_{x_2} f_{e,c_2}(x)\,dx\right)\cdots$$

$$= \frac{\pi(0)}{\phi(1)\cdots\phi(n)} \sum_{c_1,\cdots c_n} \rho_{c_1}\rho_{c_2}\cdots\rho_{c_n}$$

$$= \frac{\pi(0)}{\phi(1)\cdots\phi(n)} \left(\sum_{c_1}\rho_{c_1}\right)\left(\sum_{c_2}\rho_{c_2}\right)\cdots\left(\sum_{c_n}\rho_{c_n}\right)$$

$$\boxed{Pr[N=n] = \pi(0)\cdot\frac{\rho^n}{\phi(1)\,\phi(2)\cdots\phi(n)}}$$

⇐ only depends on $\rho = \rho_1 + \rho_2 + \cdots + \rho_c$

$$\boxed{Pr[x \mid N=n] = \prod_{i=1}^{n}\left(\frac{\rho_{c_i}}{\rho}\ f_{e,c_i}(x_i)\right)}$$

Summary:

(i) Stationary distribution of number of jobs only depends on $\phi(\cdot)$ and $\rho = \sum_c \lambda_c \mathbb{E}[S_c]$

(ii) Conditioned on $\{N=n\}$, the job in each position independently belongs to class c_i with probability ρ_{c_i}/ρ and its age (or residual size) is distributed according to the stationary excess of F_{c_i}

(iii) reversibility $\Rightarrow$ departure process of class c_i jobs is Poisson (λ_{c_i})

Example: M/G/1/PS

$\phi(n) = 1$, $\rho = \lambda \mathbb{E}[S]$

$\Pr[N=n] = (1-\rho)\rho^n$

$f[x=(x_1,x_2,\ldots,x_n) \mid N=n] = \prod_{i=1}^{n} f_e(x_i)$

Therefore: $N^{M/G/1/PS} \stackrel{d}{=} N^{M/M/1/FCFS}$

residual sizes of jobs are i.i.d. samples from F_e

Multiclass Open Jackson Network (a.k.a. BCMP network)

J service stations, C job classes

Server i: Either FCFS with i.i.d. Exp(μ_i) size for all classes

OR Symmetric queue with generally distributed job sizes

external arrivals of class c: Poisson with rate α_i^c

routing probs: $P_{ij}^{cc'} = \text{Prob}\left[\begin{array}{c}\text{a class } c \text{ departure from server } i \text{ joins} \\ \text{server } j \text{ as a class } c' \text{ job}\end{array}\right]$

$P_{i0}^{c} = 1 - \sum_{c',j} P_{ij}^{cc'} = \text{Prob}\left[\begin{array}{c}\text{class } c \text{ departure from } i \\ \text{leaves the network}\end{array}\right]$

Note: By appropriately defining classes, we can make the routing prob. depend on the path taken through the network.

Net arrival rate of class c to node i

$$\lambda_i^c = \alpha_i^c + \sum_{j,c'} \lambda_j^{c'} p_{ji}^{c'c}$$

Thm: Let the state of the network be described by

$$\underline{x} = (x^{(1)}, x^{(2)}, \ldots, x^{(j)}, \ldots, x^{(J)})$$

($x^{(j)}$: state of node j)

$$x^{(j)} = (c_1^{(j)}, c_2^{(j)} \ldots, c_{n_j}^{(j)}, x_1^{(j)} \ldots, x_{n_j}^{(j)})$$

Let $\pi(\cdot)$ denote the density fn for steady state.

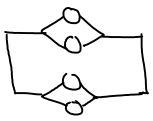
Then $$\pi(\underline{x}) = \prod_{j=1}^{J} \pi_j(x^{(j)})$$

where $\pi_j(\cdot)$ is the density fn. corresponding to node j in isolation with Poisson inputs with rates $\{\lambda_j^c\}$.

(We skip the proof as it is not very insightful)

Q: Does this product form extend to closed systems?

A: No. Example:



$N=2$, deterministic size of 1 unit at each station.

Remark: Same proof as for symmetric policies shows that M/G/m/m also exhibits insensitivity.

Coxian Distributions, Matrix Analytic (Matrix Geometric) Method

The story so far:

We have been looking at very simple queueing systems

(i) M/M/1 type systems.

i.e., CTMC looks like



- M/M/1
- M/M/m/m
- M/M/m
- M/M/1+M

also called "Birth-Death processes": transitions to a state above or below

(ii) symmetric scheduling policies (recall lect 10: $\gamma(n,i) = \delta(n,i)$)

e.g. Processor Sharing
Infinite Servers
Preemptive Last Come First Serve

- allows Generally distributed job sizes (still need Poisson arrivals)
- behave like M/M/1 systems (stationary number of jobs)

(iii) classed Jackson (BCMP) networks of (i) & (ii)

We have avoided contact with general distributions so far.

Today: a very useful tool that allows numerical approximation of many queueing systems with general distributions (and non-iid input; load balancing...) via CTMCs

---

Approximating general distributions using Exp()

To incorporate general service distributions, or general interarrival times but still be able to use the machinery of CTMCs, we first would like to approximate general distributions somehow using Exp() as building blocks.

We already saw one method last week (in analyzing $M/H_2/1/PS$)

1. H_r : r-phase hyperexponential

$$S \sim \begin{cases} \text{Exp}(\mu_1) & \text{w.prob. } p_1 \\ \vdots & \\ \text{Exp}(\mu_r) & \text{w. prob. } p_r \end{cases}$$

Two ways to think about it

(i) $F_S(x) = \sum_{i=1}^{r} p_i (1 - e^{-\mu_i x})$ (good for fitting distribs)

(ii) Suppose S represents a job size.
- ↳ when job gets to server; first rolls a biased r-faced dice
- ↳ samples from $\text{Exp}(\mu_i)$ corresp. to dice roll

(more useful while modeling through CTMCs)

Q: How does H_r compare with Exp() ?

A: Recall "squared coefficient of variation"

$$c_s^2 = \frac{\text{var}(S)}{\mathbb{E}[S]^2}$$

for $\text{Exp}(\mu)$: $c_s^2 = \frac{1/\mu^2}{(1/\mu)^2} = 1$

for H_r : $c_s^2 = \frac{\mathbb{E}[S^2]}{\mathbb{E}[S]^2} - 1 = \frac{2 \sum_{i=1}^{r} \frac{p_i}{\mu_i^2}}{\left(\sum_{i=1}^{r} \frac{p_i}{\mu_i}\right)^2} - 1$

$$= 2 \frac{\left(\sum_{i=1}^{r} \frac{p_i}{\mu_i^2}\right)\left(\sum_{i=1}^{r} p_i\right)}{\left(\sum_{i=1}^{r} \sqrt{\frac{p_i}{\mu_i^2}} \sqrt{p_i}\right)^2} - 1$$

$\left(\sum a_n b_n\right)^2 \le \left(\sum a_n^2\right)\left(\sum b_n^2\right)$

Cauchy-Schwarz

$$\geq 2 - 1$$

$$= 1$$

Key: H_r can only be used to approximate distributions "more variable" than Exp()

Fact: H_r are dense in the class of distributions with "completely monotone" densities (includes: Pareto, Weibull)

(See: Fitting mixtures of exponentials to long-tail distributions to analyze network performance models. By A. Feldman, W. Whitt.)

Q: How can we model distributions with $C^2 < 1$?

A: Convolution of Exp()

2. E_K : K-stage Erlang

$$S = \sum_{i=1}^{K} X_i \qquad \text{where } X_i \sim Exp(k\mu) \quad \text{i.i.d.}$$

$$var(S) = \sum_i var(X_i) = K \cdot \frac{1}{(K\mu)^2} = \frac{1}{K\mu^2}$$

$$\Rightarrow \quad C_s^2 = \frac{1}{K} < 1$$

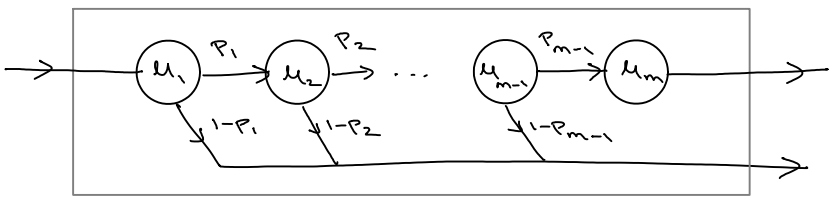
Remark: as $K \to \infty$, E_K approximates deterministic arbitrarily closely.

So : $H_r \to C^2 > 1$

$E_K \to C^2 < 1$

The next distribution unifies the two into a single framework.

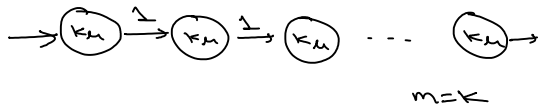
3. C_m : m-phase Coxian



alternately, $S \sim$

$$\begin{cases} Exp(\mu_1) & \text{w.p. } (1-p_1) \\ Exp(\mu_1) + Exp(\mu_2) & \text{w.p. } p_1(1-p_2) \\ Exp(\mu_1) + Exp(\mu_2) + Exp(\mu_3) & \text{w.p. } p_1 p_2 (1-p_3) \\ \vdots \end{cases}$$

Example: E_k using C_m

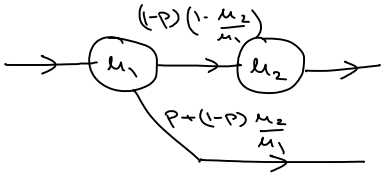


$m = k$

Example: H_2 using C_m

$$S \sim \begin{cases} \text{Exp}(\mu_1) & \text{w.p. } p \\ \text{Exp}(\mu_2) & 1-p \end{cases}$$

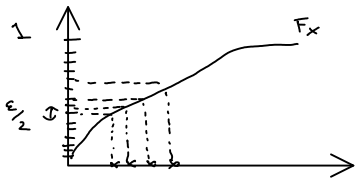
$(\mu_1 > \mu_2)$



Thm: Given an arbitrary distribution with Laplace transform $\mathcal{L}_X(s)$ & $\varepsilon > 0$, $\exists$ a Coxian distribution with Laplace transform $\mathcal{L}_Y(s)$ s.t.

$$|\mathcal{L}_X(s) - \mathcal{L}_Y(s)| \leq \varepsilon \quad \text{for all } s > 0$$

Proof idea:



First approximate X by mixture of deterministic

Then approximate deterministic by Erlang. □

In practice: when "fitting" an H_r or C_m to your data, match the first 2- or 3 moments

→ first moment: get ρ correct

→ second moment: drives $\mathbb{E}[T]$ in many queueing sys. (esp. in heavy traffic)

→ third moment: security buffer

For completeness, there is also a fourth class of distributions

4. Ph: Phase-type distributions.

Defined as distributions which represent time until absorption of some CTMC

Note: all Coxians are Phase-type.

Usually: a Ph distribution will "fit" data using fewer Exp() blocks

⇒ fewer states in CTMC (but trickier to fit Ph distributions)

Quasi-Birth-Death processes

Plan:

* We begin with a few examples to see what CTMCs we get when we use H_r job sizes / interarrivals.
* We will then look at these type of CTMCs more abstractly and see two more examples where they arise (QBD with repeating str.)
* Given their prevalence, we will look at a numerical algorithm for solving for stationary distribution of these CTMCs.

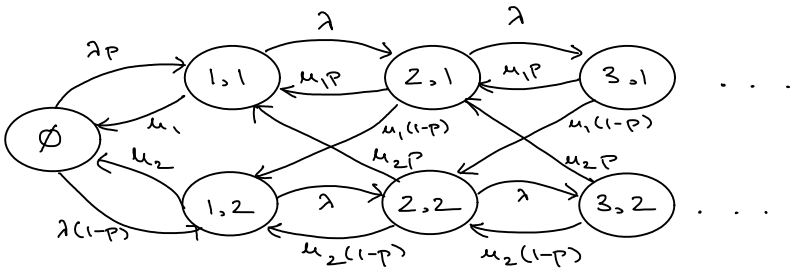
Example 1: $M/H_2/1/FCFS$

job sizes $S \sim \begin{cases} Exp(\mu_1) & \text{w.p. } p \\ Exp(\mu_2) & \text{w.p. } (1-p) \end{cases}$

Q: State space?

A: (# jobs in system, phase of job at server)

Graphically:



Key observation: the transition rates repeat!

Generator matrix (Q) =

	∅	(1,1)	(1,2)	(2,1)	(2,2)	(3,1)	(3,2)	
∅	$-\lambda$	λp	$\lambda(1-p)$					
(1,1)	μ_1	$-(\lambda+\mu_1)$		λ				
(1,2)	μ_2		$-(\lambda+\mu_2)$		λ			
(2,1)		$\mu_1 p$	$\mu_1(1-p)$	$-(\lambda+\mu_1)$		λ		...
(2,2)		$\mu_2 p$	$\mu_2(1-p)$		$-(\lambda+\mu_2)$		λ	
(3,1)				$\mu_1 p$	$\mu_1(1-p)$	$-(\lambda+\mu_1)$		
(3,2)				$\mu_2 p$	$\mu_2(1-p)$		$-(\lambda+\mu_2)$	

Or, more succinctly:

$$Q = \begin{bmatrix} L_0 & F_0 & & & \\ B_0 & L & F & & \\ & B & L & F & \\ & & B & L & F \\ & & & & \ddots \end{bmatrix}$$

where: L (local) $= \begin{bmatrix} -(\lambda+\mu_1) & \\ & -(\lambda+\mu_2) \end{bmatrix}$

F (forward) $= \begin{bmatrix} \lambda & \\ & \lambda \end{bmatrix}$

B (backward) $= \begin{bmatrix} \mu_1 p & \mu_1(1-p) \\ \mu_2 p & \mu_2(1-p) \end{bmatrix}$

Key: (i) block repeating structure

(ii) transitions only to adjacent "levels" (skip-free)

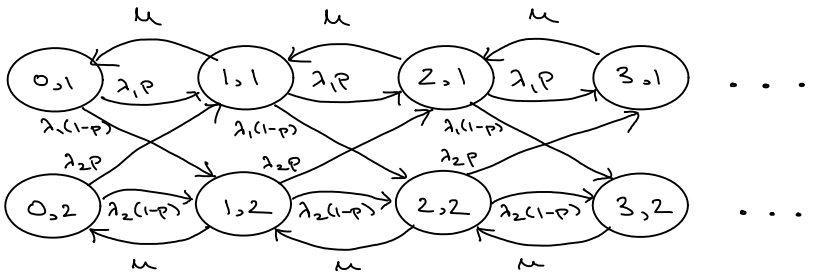
Note: rank(B) = 1 (we will come back to this)

Example 2: $H_2/M/1$

Interarrival times: $A \sim \begin{cases} \text{Exp}(\lambda_1) & \text{w.p. } p \\ \text{Exp}(\lambda_2) & \text{w.p. } (1-p) \end{cases}$

State space = (# jobs in system, phase of current interarrival time)

Graphically:



$$Q = \begin{bmatrix} 0 & 1 & 2 & 3 & \\ L_0 & F & & & \\ B & L & F & & \\ & B & L & F & \\ & & & & \ddots \end{bmatrix}$$

← repeating block structure

where,

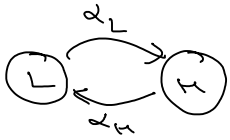
$$L = \begin{bmatrix} -(\lambda_1+\mu) & \\ & -(\lambda_2+\mu) \end{bmatrix} \quad ; \quad B = \begin{bmatrix} \mu & \\ & \mu \end{bmatrix} \quad ; \quad F = \begin{bmatrix} \lambda_1 p & \lambda_1(1-p) \\ \lambda_2 p & \lambda_2(1-p) \end{bmatrix} \quad ; \quad L_0 = \begin{bmatrix} -\lambda_1 & \\ & -\lambda_2 \end{bmatrix}$$

Now we will see two more examples of Markov chains with repeating block structure.

Example 3 : MMPP/M/1 queue

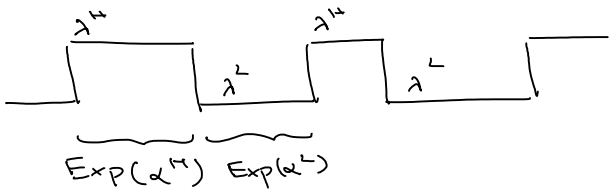
MMPP : Markov Modulated Poisson Process

Markov modulated : a background CTMC



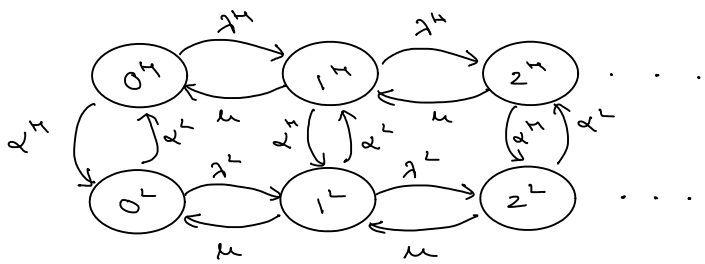
Poisson process : arrival process is $\text{Poisson}(\lambda^L)$ when in state L
$\text{Poisson}(\lambda^H)$ when in state H

That is, system behavior keeps switching b/w two different M/M/1



Arrivals times are not i.i.d. but positively correlated.

State space = (# jobs in system, state of background CTMC)
↳ the modulator process

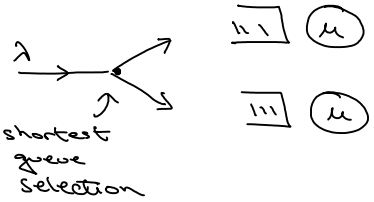


Generator matrix $Q = \begin{bmatrix} L_0 & F & & \\ B & L & F & \\ & B & L & F \\ & & & \ddots \end{bmatrix}$ (columns 0, 1, 2, 3 ...)

where,

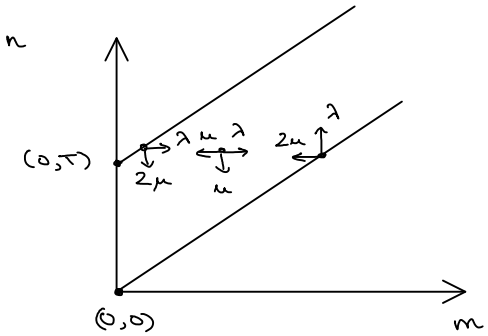
$$L = \begin{bmatrix} -(\lambda^H+\mu+\alpha^H) & \alpha^H \\ \alpha^L & -(\lambda^L+\mu+\alpha^L) \end{bmatrix} ; B = \begin{bmatrix} \mu & \\ & \mu \end{bmatrix} ; F = \begin{bmatrix} \lambda^H & \\ & \lambda^L \end{bmatrix} ; L_0 = \begin{bmatrix} -(\lambda^H+\alpha^H) & \alpha^H \\ \alpha^L & -(\lambda^L+\alpha^L) \end{bmatrix}$$

Example 4: Shortest queue with threshold jockeying



If difference between longest and shortest queue $> T$, then one job jumps (jockeys) from the longest to shortest queue

State $\equiv (m,n)$ where m = length of shortest queue
n = length of longest queue



Imagine slicing the state space horizontally:

level $n = \{(n-T,n), (n-T+1,n), \ldots, (n,n)\}$

states (m,n) with $n<T$ in level 0 (boundary states)

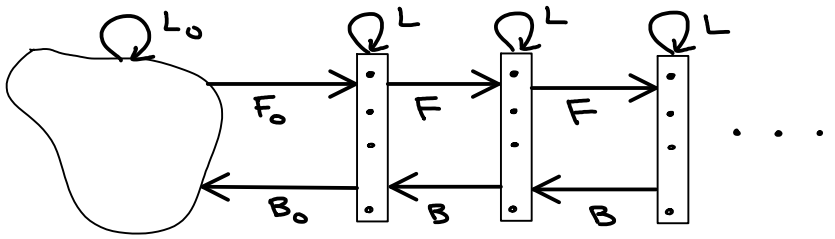
⇒ Gives repeating structure with

$$L = \begin{bmatrix} -(\lambda+2\mu) & \lambda & & \\ \mu & -(\lambda+2\mu) & \lambda & \\ & \ddots & \ddots & \ddots \\ & & 2\mu & -(\lambda+2\mu) \end{bmatrix}; \quad B = \begin{bmatrix} 0 & 2\mu & 0 & 0 & \cdots & 0 \\ & & \mu & & & \\ & & & \mu & & \\ & & & & \ddots & \\ 0 & 0 & 0 & 0 & \cdots & 0 \end{bmatrix}$$

$$F = \begin{bmatrix} 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & \cdots & 0 & 0 \\ & & & & \\ 0 & 0 & \cdots & \lambda & 0 \end{bmatrix}$$

Note: rank $(F) = 1$
(we will revisit this)

SUMMARY: CTMCs with the following structure are very useful



States within level i ($i \geq 1$) called "phases"

All levels $i \geq 1$ have m phases ($\Rightarrow$ F, B, L are $m \times m$ matrices)

These CTMCs are a special class of "Quasi Birth Death" chains
↳ transitions b/w adjacent levels.

GOAL: a numerical algorithm for solving for stationary distribution of QBDs with repeating block structure

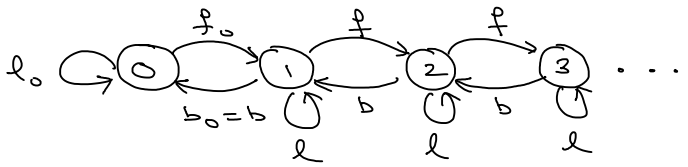
ROADMAP:

(i) we will develop the algorithm first using DTMCs
↳ first with $m=1$ (regular Birth-Death chain)
↳ then general m

(ii) finally we will use a simple trick to translate to CTMCs

The algorithm we will develop is called the Matrix Analytic Method or the Matrix Geometric Method.

Consider: a DTMC with one phase per level



$f + b + l = 1$

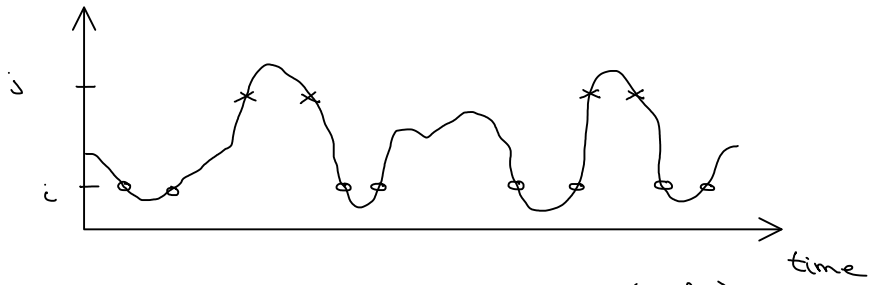
(these are all probabilities)

Obviously we can solve this toy DTMC more directly, but here we will develop an algorithm that we would be able to extend to $m > 1$.

Let $\{\pi_i\}$ = stationary distribution

Claim: $\frac{\pi_j}{\pi_i} = \mathbb{E}\left[\#\text{ visits to state } j \text{ between two consecutive visits to } i\right]$

example:



∘ : visits to i

× : visits to j

so $\frac{\pi_j}{\pi_i} = \lim_{t\to\infty} \frac{(\#\text{ of } \times)}{(\#\text{ of } \circ)} = \lim_{t\to\infty} \frac{1}{(\#\text{ of } \circ)} \sum_{k=1}^{(\#\text{ of } \circ)} \left(\#\text{ of } \times \text{ during } k^{th} \; i\to i \text{ excursion}\right)$

$= \mathbb{E}\left[\#\text{ visits to } j \text{ during an } i\to i \text{ excursion}\right]$

Intuitively: if $\pi(j) = 2\pi(i)$ then on average I visit state j 2 times for every visit to i

□

Define: $r_i^{(k)} = \mathbb{E}\left[\#\text{ visits to state } (i+k) \text{ before first return to } i, \text{ given we start in } i\right]$

Claim: $r_i^{(k)}$ is independent of i for $i \geq 1$.

Proof: Consider $i_1 \neq i_2$

↳ either first step is to $(i-1)$ in which case we do not visit $(i+k)$ ⇒ Note: $\Pr[\text{first step is backwards}]$ is same for i_1, i_2

↳ or first step is to $(i,i+1)$; in which case due to repeating structure the excursions are similar for i_1 & i_2.

Note: $\Pr$(first step is forward) is also same for i_1, i_2.

□

Note: the only properties we need for the last claim to hold is

(i) no jumps bigger than $+1$ (can have any negative jump)
≡ skip-free to the right

(ii) prob. of $i \to i+1$ is the same for all states. $= f$

So we conclude: $r_i^{(k)} = r^{(k)}$; Let $r \doteq r^{(1)}$

Therefore: $\pi_{i+1} = \pi_i r^{(1)} = \pi_{i-1} r^{(1)} r^{(1)} = \pi_{i-1} \left(r^{(1)}\right)^2$

but also $\pi_{i+1} = \pi_{i-1} r^{(2)}$

$\Rightarrow r^{(k)} = \left(r^{(1)}\right)^k = r^k$

$\Rightarrow \pi_{i+1} = \pi_1 r^i$

Q: Expression for r?

Recall: $r \doteq \mathbb{E}[$visits to $(i+1)$ before first return to i given start $i]$

Possible ways to visit $(i+1)$

- from i : only one opportunity for $i \to i+1$, take it with prob. f
- from $i+1$: in expectation r opportunities for $(i+1) \to (i+1)$, take each iid w.prob l
- from $i+2$: in expectation r^2 opportunities for $(i+2) \to (i+1)$, take each iid w.prob b

Combining:

$$\boxed{r = f + rl + r^2 b}$$

Claim: r is the smallest non-negative solution to the above equation.

Proof: Let x be the smallest non-negative solution

Find x using iteration

$$x(0) = 0 \quad ; \quad x(n+1) = f + x(n)\,l + x(n)^2\,b$$

Then $x(n) \uparrow$ and $x(n) \le r \quad \forall n$.

so $x(\infty) := x$ exists & $x \le r$

Also: $x \ge r$

define $r^{(k)}(n) = \mathbb{E}\Big[$ # visits to $(i+k)$ before first return to i over n time steps, given start in i $\Big]$

so $r^{(k)}(n) \uparrow r^{(k)}$ as $n \to \infty$

$$r^{(k)}(n) = \sum_{m=0}^{n} \underbrace{P^{(m)}_{i,i+k-1}}_{\substack{m\text{-step transition} \\ \text{prob}}} \cdot r^{(1)}(n-m) \le \sum_{m=0}^{n} P^{(m)}_{i,i+k-1}\, r^{(1)}(n) = r^{(k-1)}(n)\, r^{(1)}(n)$$

$$\Rightarrow \quad r^{(k)}(n) \le \left(r^{(1)}(n)\right)^k$$

subclaim: $r^{(1)}(n) \le x(n)$

Proof: by induction

$$r^{(1)}(0) = 0 = x(0)$$

Induction step:

$$\begin{aligned} r^{(1)}(n+1) &= f + r^{(1)}(n)\,l + r^{(2)}(n)\,b \\ &\le f + r^{(1)}(n)\,l + \left(r^{(1)}(n)\right)^2 b \\ &\le f + x(n)\,l + (x(n))^2\,b \\ &= x(n+1) \end{aligned}$$

□

since $r^{(1)}(n) \uparrow r^{(1)} := r$, $r \le x$

Thus $r = x$.

□

Iterative scheme for computing r:

$$r(0) = 0 \quad , \quad r(n+1) = f + r(n)\,l + r^2(n)\,b \quad ; \quad n = 0, 1, \ldots$$

Remark: The claim $r_i^{(k)} = r^{(k)} = r^k$

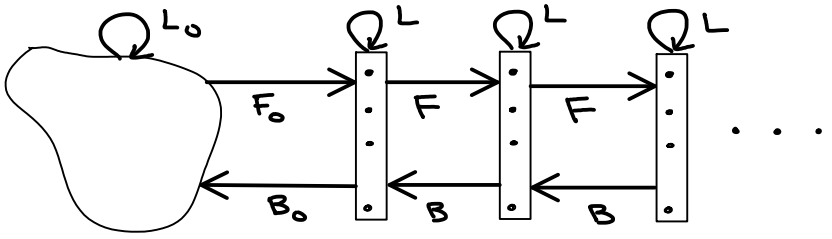
also works for

$$P = \begin{bmatrix} b_0 & f_0 & & & \\ b_1 & \ell & f & & \\ b_2 & b_1 & \ell & f & \\ b_3 & b_2 & b_1 & \ell & f \\ b_4 & b_3 & b_2 & b_1 & \ell \quad f \ddots \end{bmatrix}$$

in which case r is the smallest non-negative soln of

$$r = f + r\ell + \sum_{i=1}^{\infty} r^{i+1} b_i$$

Next consider each level ≡ m states/phases (still DTMC)



states: $\{ \underbrace{(0,1),(0,2)\ldots,(0,m)}_{\text{level 0}}, \underbrace{(1,1)\ldots(1,m)}_{\text{level 1}}, \underbrace{(2,1),\ldots(2,m)}_{\text{level 2}}, \ldots \}$

$$P = \begin{bmatrix} L_0 & F_0 & & \\ B_0 & L & F & \\ & B & L & F \ddots \end{bmatrix}$$

$\pi(i,j)$ ≡ stationary probability of state (i,j) (level ↓, phase within level ↘)

Define: $R_i^{(k)}(j,\ell) = \mathbb{E}\left[\begin{array}{l}\#\text{ visits to state } (i+k,\ell) \text{ before first return} \\ \text{to level } i \text{, given start in } (i,j)\end{array}\right]$

↳ any phase!

Claim: $\pi(i+k,l) = \sum_{j=1}^{m} \pi(i,j) R_i^{(k)}(j,l)$

Proof: Partition sample path into visits to level i

$$\#\text{ visits to } (i+k,l) = \sum_{j=1}^{m} \left(\#\text{ visits to } (i,j)\right)\left(\begin{array}{c}\text{avg. visits to } (i+k,l)\\ \text{during } (i,j)\to(i,j')\\ \text{excursion}\end{array}\right)$$

$$\underset{\pi(i+k)}{\underset{\text{phase } l}{\underset{\downarrow}{\left[\ \bullet\ \right]}}} = \underset{\pi(i)}{\left[\ \ \right]} \begin{bmatrix} R_i^{(k)}(1,l) \\ R_i^{(k)}(2,l) \\ \vdots \\ R_i^{(k)}(m,l) \end{bmatrix}$$

$R_i^{(k)}$

□

or, $\pi(i+k) = \pi(i) R_i^{(k)}$

As before: - $R_i^{(k)} = R^{(k)}$ (independent of i)

- $R^{(k)} = R^k$ where $R \doteq R^{(1)}$

$\Rightarrow \quad \pi(i+1) = \pi(1) R^i$

Q: Expression for R?

As before, condition on last step before visiting $(i+1,l)$

recall $R(j,l) = \mathbb{E}[\#\text{ visits to } (i+1,l) \text{ before first return to level } i, \text{ given start in } (i,j)]$

$$R(j,l) = \underbrace{1 \cdot F(j,l)}_{(i,j)\to(i+1,l)} + \sum_{k=1}^{m} \underbrace{R(j,k) L(k,l)}_{(i+1,k)\to(i+1,l)} + \sum_{k=1}^{m} \underbrace{(R^2)(j,k) B(k,l)}_{(i+2,k)\to(i+1,l)}$$

in matrix form:

$$\boxed{R = F + RL + R^2 B}$$

R is the minimal non-negative solution to the above equation.

Note: since $\pi(i) \to 0$ as $i \to \infty$, $\sigma(R) < 1$
↳ spectral radius

Computing R

Method 1: $R(0) = 0$

$$R(n+1) = F + R(n)L + R^2(n)B$$

Method 2: rewrite $R(I-L) = F + R^2 B$

so: $R(0) = 0$

$$R(n+1) = (F + R^2(n)B)\underbrace{(I-L)^{-1}}_{(I+L+L^2+\cdots) \text{ exists}}$$

Finishing the solution:

$$\pi(i+1) = \pi(1) R^i \quad \text{for } i \geq 1$$

$\Rightarrow$ $\pi(0)$ & $\pi(1)$ only unknowns.

Balance equations:

$$\pi(0) = \pi(0) L_0 + \pi(1) B_0$$

$$\pi(1) = \pi(0) F_0 + \pi(1) L + \pi(2) B$$

$$= \pi(0) F_0 + \pi(1)(L + RB)$$

Normalization:

$$\pi(0) \cdot \begin{bmatrix} \vdots \\ 1 \end{bmatrix} + \pi(1)(I + R + R^2 + \cdots) \begin{bmatrix} \vdots \\ 1 \end{bmatrix} = 1$$

$$\Rightarrow \pi(0) e_{m'} + \pi(1)(I-R)^{-1} e_m = 1$$

equivalently:

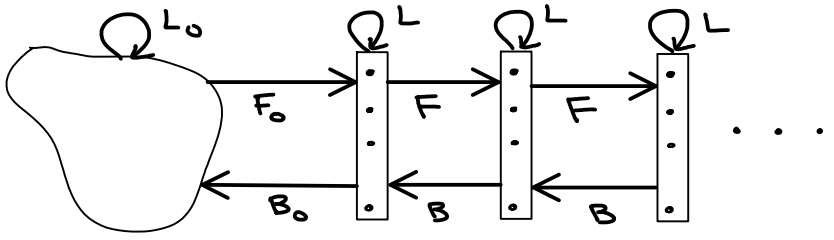
$$[\pi(0) \quad \pi(1)] = [\pi(0) \quad \pi(1)] \begin{bmatrix} L_0 & F_0 \\ B_0 & L+RB \end{bmatrix}$$

& replace last entry by 1 (here) and last column by $\begin{bmatrix} e_{m'} \\ (I-R)^{-1} e_m \end{bmatrix}$ (here)

Note: Stability conditions from F, L, B: $F + L + B = A$ is a probability transition matrix on $\{1, \ldots, m\}$ ← phases. Let $pA = p$ (stat. dist. for A)

The DTMC is stable if $\underbrace{pFe_m}_{\text{upward drift}} < \underbrace{pBe_m}_{\text{downward drift}}$ □

Finally consider: a CTMC with m phases per level



states: $\{\underbrace{(0,1),(0,2)\ldots,(0,m)}_{\text{level 0}},\underbrace{(1,1)\ldots(1,m)}_{\text{level 1}},\underbrace{(2,1),\ldots(2,m)}_{\text{level 2}},\ldots\}$

$$Q = \begin{bmatrix} L_0 & F_0 & & \\ B_0 & L & F & \\ & B & L & F \\ & & \ddots & \ddots \end{bmatrix}$$

(Generator for CTMC)

$\pi(i,j)$ ≡ stationary probability of state (i,j) (level i, phase within level j)

Trick: consider Δ small enough so that

$$P = I + \Delta Q$$

has all non-negative entries.

Claim: (i) P is a prob. transition matrix

(ii) π is also the stationary distribution for P

Proof: (i) by construction P is non-negative.

since Q is a CTMC generator, its rows sum to zero.

$\Rightarrow P = I + \Delta Q$ has all row sums equal to 1

(ii) $$\pi P = \pi(I + \Delta Q) = \pi + \Delta(\underbrace{\pi Q}_{=0}) = \pi$$

□

Therefore: $\pi(i+1) = \pi(i) R$

where R is given by

$$R = (\Delta F) + R(I + \Delta L) + R^2(\Delta B)$$

or

$$\boxed{F + RL + R^2 B = 0}$$

Iterative method to compute R:

$$R(0) = 0$$

$$R(n+1) = \left(F + R^2(n) B\right)(-L)^{-1}$$

To complete the solution:

$$[\pi(0) \quad \pi(1)] \begin{bmatrix} L_0 & F_0 & e_{m-} \\ B_0 & L+RB & (I-R)^{-1} e_m \end{bmatrix} = [0\ 0\ 0\ 0 \cdots 0\ 1]$$

Balance eqns. normalization

Tail of stationary distribution

Let $|\eta_1| \geq |\eta_2| \geq \cdots \geq |\eta_m|$ be the eigenvalues of R with (left) eigenvectors $v_1, v_2, \cdots v_m$

Since v_i are linearly independent, we can decompose

$$\pi_1 = \sum_{k=1}^{m} a_k v_k \qquad \text{for constants } a_1, \cdots a_m$$

$$\Rightarrow \pi_{i+1} = \pi_1 R^i = \sum_{k=1}^{m} a_k v_k R = \sum_{k=1}^{m} \eta_k^i a_k v_k$$

$\Rightarrow$ tail of π_i is geometrically decreasing with rate η_1.

[Except if $\eta_1 = \eta_2 = \cdots = \eta_b > \eta_{b+1} \geq \cdots$ in which case
$\pi_{i+1} \sim \eta_1^i \times$ (degree $b-1$ polynomial of i)]

Explicit solution for R for special cases

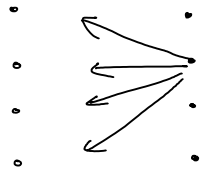
When F or B matrix have rank 1, it turns out we can say much more about the R matrix

Case: $B = v \cdot \alpha$ (e.g. $M/H_2/1/FCFS$)

That is $B = \begin{bmatrix} v_1 \\ \vdots \\ v_m \end{bmatrix} [\alpha_1 \dots \alpha_m]$

scale v so that $\alpha \cdot e_m = 1$ ($\Rightarrow \alpha$ is a prob. distribution)

Intuition:



level i level $(i+1)$

on $(i+1) \to i$ transition we reach (i,j) with prob. α_j irrespective of our phase in level $(i+1)$ (i.e., entrance distribution is α)

Now, balancing total flow b/w levels i & $(i+1)$:

$$\pi_i F e_m = \pi_{i+1} B e_m = \pi_{i+1} v \alpha e_m = \pi_{i+1} v$$

Flow balance (for CTMC):

$$\pi_{i-1} F + \pi_i L + \underbrace{\pi_{i+1} B}_{= \pi_{i+1} v \alpha = \pi_i F e_m \alpha} = 0$$

$$\Rightarrow \quad \pi_{i-1} F + \pi_i \left(L + F e_m \alpha \right) = 0$$

$$\text{or} \quad \pi_i = \pi_{i-1} \underbrace{\left[-F \left(L + F e_m \alpha \right)^{-1} \right]}_{R}$$

$$\Rightarrow \quad \boxed{R = -F \left(L + F e_m \alpha \right)^{-1}}$$

Case: $F = \omega \cdot \beta$ e.g. $H_2/M/1$, Shortest queue with jockeying

That is $F = \begin{bmatrix} \omega_1 \\ \vdots \\ \omega_m \end{bmatrix} [\beta_1 \dots \beta_m]$

Scale ω s.t. $\beta e_m = 1$ $\Rightarrow \beta$ is a prob. distribution.

Intuitively: when we have a level $i \to$ level $(i+1)$ transition, we enter $(i+1, l)$ with probability β_l irrespective of phase at level i.

Claim: When $F = \omega \cdot \beta$, then

$R = \omega \cdot a$ for some vector $a = [a_1 a_2 \dots a_m]$

Proof: (i) DTMC view

$$R(j,l) = (1-\omega_j)\cdot 0 + \omega_j \underbrace{\sum_{k=1}^{m} \beta_k S(k,l)}_{\text{only depends on } l, \text{ not } j}$$

$S(k,l) = \mathbb{E}[\#\text{visits to } (i+1,l) \text{ before first visit to level } i \text{ given start in } (i+1,k)]$

(ii) using iteration scheme

$\to$ for CTMC $R(0) = 0 = \omega \cdot 0$

$$\begin{aligned} R(n+1) &= -\left(F + R^2(-B)\right) L^{-1} \\ &= -(\omega\cdot\beta + \omega a_n \omega a_n B) L^{-1} \\ &= -\omega(\beta + a_n \omega a_n B) L^{-1} \\ &= \omega a_{n+1} \end{aligned}$$

□

Therefore: $R^i = (\omega a)(\omega a) \cdots (\omega a)$

$$\begin{aligned} &= \omega (a\omega)(a\omega) \cdots (a\omega) a \\ &= \eta^{i-1} \omega a \\ &= \eta^{i-1} R \end{aligned}$$ η: a scalar

$$\Rightarrow \quad \pi(i+1) = \pi(1) R^{i} = \pi(1) \eta^{i-1} R$$
$$= \eta^{i-1} \pi(2)$$

η is the unique root in $(0,1)$ of:

* $\det(F + L\eta + B\eta^2) = 0$ (for CTMC)

* $\det(F + (L-I)\eta + B\eta^2) = 0$ (for DTMC)

(or determine via $a_{n+1} = -(\beta + (a_n \omega) a_n B) L^{-1}$ for CTMC)

Matrix Analytic Method Cheat Sheet

We motivated the concept of the R matrix through its physical interpretation which also gave us the equation

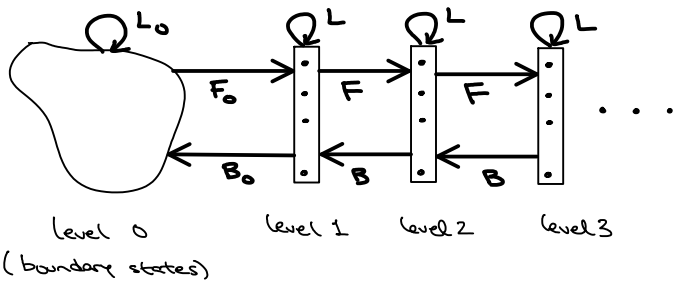
$$F + RL + R^2B = 0$$

This is useful when you want to extend these results for research etc.

If you only want to use Matrix Analytic Method in its simplest form, here is a less rigorous derivation:

(1) start with balance eqns.

$$\pi_0 L_0 + \pi_1 B_0 = 0$$
$$\pi_0 F_0 + \pi_1 L + \pi_2 B = 0$$
$$\pi_1 F + \pi_2 L + \pi_3 B = 0$$
$$\pi_2 F + \pi_3 L + \pi_4 B = 0$$
$$\vdots$$



(2) "guess" $\pi_i = \pi_1 R^{(i-1)}$

(3) $\pi_i F + \pi_{i+1} L + \pi_{i+2} B = 0 \iff \pi_i (F + RL + R^2 B) = 0$

(4) since this is true for all $i \geq 1$, we get

$$F + RL + R^2 B = 0$$

(5) Solve for R iteratively: $R(0) = 0$

$$R(n+1) = -(F + R^2(n) B) L^{-1}$$

(6) Solve for π_0, π_1

$$[\pi_0 \ \pi_1] \begin{bmatrix} L_0 & F_0 & e_{m'} \\ B_0 & (L+RB) & (I-R)^{-1} e_m \end{bmatrix} = [\overbrace{0 \cdots 0}^{m'} \ \overbrace{0 \cdots 0}^{m} \ 1]$$

Acknowledgement: These lecture notes are based on a tutorial delivered by Ivo Adan at Carnegie Mellon University on 24 April, 2007.

Mean Value Analysis of M/GI/1 ; Renewal Reward Theory

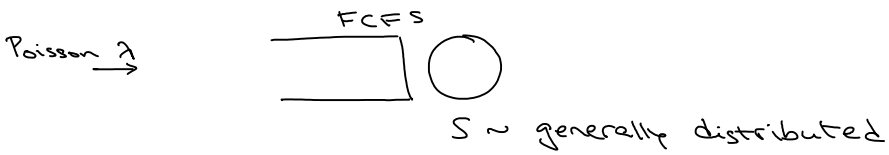
Matrix analytic method allows modeling general distributions using Exp() as building blocks

→ versatile but numeric

don't get much insight about what is happening, what parameters matter

Today: a direct closed-form analysis of M/GI/1 queue

The M/GI/1/FCFS queue



$\rho \doteq \lambda \mathbb{E}[S]$

Metrics: T_Q : time spent in buffer (before commencing service)

T : total system time (response / sojourn time)

N_Q, N } # jobs in { buffer, system = buffer + server

Recall: Mean Value Analysis (MVA) of $T_Q^{M/M/1/FCFS}$

$$\mathbb{E}[T_Q] = \mathbb{E}[\text{work seen on arrival ahead of me}]$$

$$= \mathbb{E}\left[\sum_{i=1}^{N_Q^{arrival}} S_i\right] + \mathbb{E}[\text{residual size of job at server} \mid \text{server busy}] \cdot \Pr[\text{server busy}]$$

$$\overset{(PASTA)}{=} \mathbb{E}[N_Q] \cdot \frac{1}{\mu} + \rho \cdot \underbrace{\frac{1}{\mu}}_{\text{memoryless property}}$$

$$\overset{(Little's\ law)}{=} \lambda \mathbb{E}[T_Q] \cdot \frac{1}{\mu} + \rho/\mu$$

$$\boxed{\mathbb{E}[T_Q^{M/M/1}] = \left(\frac{\rho}{1-\rho}\right)\mathbb{E}[S]}$$

Q: What breaks down for M/G/1 ?

A: Can't claim

$$\mathbb{E}[\text{residual size of job at server} \mid \text{busy}] = \mathbb{E}[S]$$

Let: $\mathcal{E}_s \doteq$ residual size of job at server seen by a Poisson arrival (i.e. time average) given server is busy

In terms of $\mathcal{E}_s$:

$$\boxed{\mathbb{E}\left[T_Q^{M/G/1}\right] = \left(\frac{\rho}{1-\rho}\right)\mathbb{E}[\mathcal{E}_s]}$$

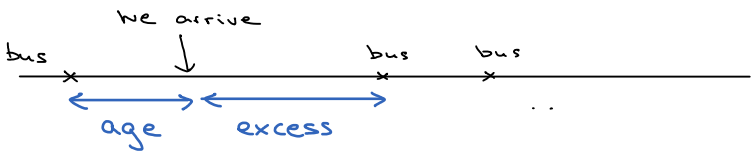
To learn more about $\mathcal{E}_s$, we will take a brief detour through renewal theory : (i) inspection paradox

(ii) renewal reward theory

Inspection Paradox

Suppose : we are waiting at a bus stop

buses arrive every 10 minutes on average



X_i : i.i.d. interarrival times

Q: How long should we expect to wait if $X_i \sim$ deterministic ?

A: Waiting time should be Unif[0,10] ; so on average 5 minutes.

Q: If $X_i \sim Exp(1/10)$?

A: By memoryless property, waiting time is $Exp(1/10)$

$\Rightarrow$ expected waiting time = 10 minutes.

This waiting time has a special name

Excess (of a renewal process) = time between arrival of a random observer and the subsequent event (renewal epoch)

Age (of a renewal process) = time between arrival of a random observer and the preceding event (renewal epoch)

By symmetry, when $X_i \sim Exp(1/10)$ then the age is $Exp(1/10)$ as well.

By linearity of expectation, the expected duration of the interarrival time we land in is 20 minutes, even though $\mathbb{E}[X] = 10$!

This is the inspection paradox: since we are more likely to land during a long period, we observe a "length-biased" interarrival time distribution.

Next we will see an informal derivation of density of excess: f_ε
For a rigorous proof, we will wait until renewal reward.

An informal derivation of f_ε

Let X_i be distributed according to $F_X(\cdot)$.

First: condition on landing in an interarrival time $= x$

$$f_\varepsilon(y \mid x)\, dy = Prob\left[\varepsilon \in (y, y+dy) \mid \text{land in } X_n = x\right]$$

$$= \begin{cases} \frac{dy}{x} & x > y \qquad (\varepsilon \sim Unif[0,x]) \\ 0 & x \le y \end{cases}$$

Now uncondition, and based on our argument,

$$\mathbb{P}_{\mathcal{E}}[\text{we land in } X_n = x] \propto x\, f_x(x)$$

$$\Rightarrow \quad f_{\mathcal{E}}(y)\,dy \propto \int_{x=0}^{\infty} \left(f_{\mathcal{E}}(y|x)\,dy\right)\left(x\, f_x(x)\,dx\right)$$

$$\propto \int_{x=y}^{\infty} \frac{dy}{x} \cdot x\, f_x(x)\,dx$$

$$\propto \bar{F}_x(y)\,dy$$

since $\int_0^{\infty} \bar{F}_x(y)\,dy = \mathbb{E}[x]$

$$\boxed{f_{\mathcal{E}}(y) = \frac{\bar{F}_x(y)}{\mathbb{E}[x]}}$$

Also called the stationary excess or equilibrium density corresp. to $F(\cdot)$. (recall HW 1 & 2!)

Renewal Reward Theory

Recall:

Defn.: A renewal process $N(t)$ is a counting process with i.i.d. interarrival times

Let $X_1, X_2, \ldots$ be the interarrival times

$$S_n = \sum_{i=1}^{n} X_i \quad \equiv \text{ time of } n^{th} \text{ renewal epoch}$$

$$N(t) = \sum_{n=1}^{\infty} 1_{\{S_n \le t\}}$$

$(0, S_1), (S_1, S_2), \ldots$ are called renewal cycles.

We also proved the following result :

Thm (Elementary Renewal Theorem)

$$\lim_{t \to \infty} \frac{N(t)}{t} = \frac{1}{\mathbb{E}[X]} \quad \text{w.p. 1}$$

The following are some basic facts and concepts of renewal theory (we will not need these, but these are good to be aware of)

Defn : The renewal function $H(t)$ is defined as the expected number of renewals by time t :

$$H(t) = \mathbb{E}[N(t)]$$

The renewal density $h(t)$ is defined as

$$h(t) = \frac{d H(t)}{dt}$$

Fact: $h(t)$ obeys the following integral equation

$$h(t) = f_x(t) + \int_0^t h(t-u) f_x(u) \, du$$

which gives the following relationship between Laplace transforms of $h(\cdot)$ & $f_x(\cdot)$

$$\mathcal{L}_H(s) = \frac{\mathcal{L}_F(s)}{1 - \mathcal{L}_F(s)} .$$

□

Now we add the feature of rewards to a renewal process

Define $r_n(t)$ = the instantaneous rate of earning reward during the n^{th} renewal cycle at time $(t+S_{n-1})$.

Let: $R_n = \int_{t=0}^{X_n} r_n(t)\,dt$ = total reward earned during n^{th} renewal cycle

$r(t) = r_{N(t)+1}\left(t - S_{N(t)}\right)$ = rate of earning reward at time t

$R(t) = \int_0^t r(u)\,du$ = total reward earned until time t

$$\mathbb{E}[R] = \mathbb{E}[R_n]$$
$$\mathbb{E}[X] = \mathbb{E}[X_n]$$

Assumption: r_n can depend on the n^{th} renewal cycle but is independent across renewal cycles.

That is $(r_n(), X_n)$ are an i.i.d. sequence but r_n & X_n can be correlated

Theorem (Renewal Reward): If $\mathbb{E}[R] < \infty$ and $\mathbb{E}[X] < \infty$

$$\lim_{t\to\infty} \frac{R(t)}{t} = \frac{\mathbb{E}[R]}{\mathbb{E}[X]} \quad \text{w.p. 1}$$

That is: avg rate at which reward is earned is equal to the expected reward earned in a renewal cycle divided by the expected length of renewal cycle.

This is extremely useful!

Suppose we are analyzing a process $Y(t)$ and want the stationary distribution of $g(Y)$

(1) Identify regeneration epochs ≡ event at which process $Y()$ probabilistically restarts
(e.g. an M/G/k process becoming idle)

(2) Define reward at time t as

* $r(t) = \mathbb{1}_{\{g(Y(t)) \le x\}} \Rightarrow \lim_{t\to\infty} \frac{R(t)}{t} = Pr[g(Y) \le x]$

* $r(t) = e^{-s\,g(Y(t))} \Rightarrow \lim_{t\to\infty} \frac{R(t)}{t} = \mathbb{E}[e^{-s\,g(Y)}]$

(Laplace transform of $f(Y)$)

(3) Only need to analyze a single renewal cycle!
⇒ no need to analyze stationary distributions of Markov chains etc.

A couple of examples will make this clearer, and then we will prove the theorem.

Example 1: Consider an M/G/1 queue
Find $Pr[\text{server is busy}]$

Q: What are renewal cycles?

A: One (idle + busy) period: due to memoryless property of Poisson process, system restarts on emptying

So $r(t) = \begin{cases} 1 & \text{if busy at } t \\ 0 & \text{if idle} \end{cases}$

$\lim_{t\to\infty} \frac{R(t)}{t} = Pr[\text{server busy}]$

R_n = total reward earned during n^{th} renewal

= n^{th} busy period

$\Rightarrow \mathbb{E}[R] = \mathbb{E}[B]$

X_n = total length of n^{th} renewal cycle

= (idle + busy) period

$\Rightarrow \mathbb{E}[X] = \mathbb{E}[I] + \mathbb{E}[B]$

$$\Rightarrow \Pr(\text{server busy}) = \frac{\mathbb{E}[B]}{\mathbb{E}[B] + \mathbb{E}[I]}$$

since we know $\mathbb{E}[I] = 1/\lambda$, $\Pr(\text{busy}) = \lambda \mathbb{E}[S] = \rho$

this gives

$$\boxed{\mathbb{E}[B] = \frac{\mathbb{E}[S]}{1-\rho}}$$

Example 2 : $\mathbb{E}[\mathcal{E}_s]$ for M|G|1|1|FCFS

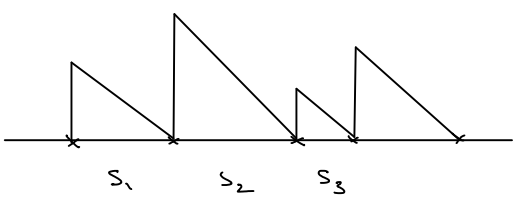
Intuitively want $r(t)$ = excess at time t (and server to be busy)

Consider,



↳ i.i.d job sizes (imagine server is continuously busy)

So what does $r(t) = \mathcal{E}(t)$ look like?



$$R_n = \frac{S_n^2}{2} \Rightarrow \mathbb{E}[R] = \frac{\mathbb{E}[S]^2}{2}$$

$$X_n = S_n \Rightarrow \mathbb{E}[X] = \mathbb{E}[S]$$

$$\mathbb{E}[\mathcal{E}_s] = \lim_{t\to\infty} \frac{R(t)}{t} = \frac{\mathbb{E}[S^2]}{2\mathbb{E}[S]}$$

So $$\boxed{\mathbb{E}[\mathcal{E}_s] = \frac{\mathbb{E}[S^2]}{2\mathbb{E}[S]}}$$

Going back to inspection paradox and bus arrivals:

* when $X_i \sim$ deterministic

$$\mathbb{E}[\text{excess}] = \frac{\mathbb{E}[X^2]}{2\mathbb{E}[X]} = \frac{\mathbb{E}[X]^2}{2\mathbb{E}[X]} = \frac{\mathbb{E}[X]}{2}$$

* when $X_i \sim Exp()$

$$\mathbb{E}[\text{excess}] = \frac{\mathbb{E}[X^2]}{2\mathbb{E}[X]} = \frac{2\mathbb{E}[X]^2}{2\mathbb{E}[X]} = \mathbb{E}[X]$$

Returning to MVA of $\mathbb{E}[T_Q^{M/G/1}]$

$$\mathbb{E}[T_Q^{M/G/1}] = \left(\frac{\rho}{1-\rho}\right)\mathbb{E}[\mathcal{E}_s]$$

$$= \frac{\rho}{1-\rho}\,\frac{\mathbb{E}[S^2]}{2\mathbb{E}[S]}$$

$$= \left(\frac{\rho\,\mathbb{E}[S]}{1-\rho}\right)\left(\frac{\mathbb{E}[S^2]}{2\,\mathbb{E}[S]^2}\right)$$

$$\boxed{\mathbb{E}[T_Q^{M/G/1}] = \mathbb{E}[T_Q^{M/M/1}]\left(\frac{C_s^2+1}{2}\right)}$$

Pollaczek-Khinchin formula (1930)

Observations:

* $\mathbb{E}[T_Q]$ (hence also $\mathbb{E}[T]$) grows linearly in variance of service distribution

  Intuition: large jobs block the server, more jobs are likely to arrive when a large job at server

* $\mathbb{E}[T_Q]$, $\mathbb{E}[T]$ increase in load as $\frac{1}{1-\rho}$.

Alternate derivation using Brumelle's formula

Recall the following generalization of Little's Law

for job j define : $q_j(t)$ (supported on its sojourn time)

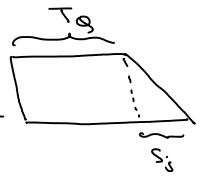
$$G_j = \int_0^\infty q_j(t)\,dt \qquad \therefore \; G = \mathbb{E}[G_j] \quad \text{(customer average)}$$

$$H(t) = \sum_{j=1}^{\infty} q_j(t) \qquad \therefore \; H = \mathbb{E}[H] \quad \text{(time average)}$$

then $H = \lambda G$.

Let $q_j(t)$ = residual size of job j at time t

For non-preemptive scheduling $q_j(t) = s_j$



$H(t) = V(t)$ ≡ unfinished work

For size independent scheduling T_Q independent of S_j

$$\Rightarrow \quad H = \lambda G \quad \Rightarrow \quad \mathbb{E}[V] = \lambda\left(\mathbb{E}[S T_Q] + \frac{\mathbb{E}[S^2]}{2}\right)$$

$$= \lambda\left(\mathbb{E}[S]\,\mathbb{E}[T_Q] + \frac{\mathbb{E}[S^2]}{2}\right)$$

For FCFS scheduling

$$\mathbb{E}[T_Q] = \mathbb{E}[\text{unfinished work seen by arrival}]$$

using PASTA

$$= \mathbb{E}[\text{time average unfinished work}]$$

$$= \mathbb{E}[V]$$

$$\Rightarrow \quad \mathbb{E}[T_Q] = \lambda\left(\mathbb{E}[S]\,\mathbb{E}[T_Q] + \frac{\mathbb{E}[S^2]}{2}\right)$$

$$\boxed{\mathbb{E}[T_Q] = \frac{\lambda\,\mathbb{E}[S^2]}{2(1-\lambda\,\mathbb{E}[S])}}$$

M/GI/1 using Transforms; FCFS vs. LCFS

M/GI/1/FCFS intuition from Mean value analysis:

- variability in job size distribution hurts
- jobs more likely to see server busy with a large job causing long queues

Goals for today:

- analysis of distribution of T, T_Q, N, N_Q for M/GI/1/FCFS
- beyond FCFS
  - LCFS : Last-Come-First Served
  - Preemptive LCFS

M/GI/1/FCFS via transforms

Recall analysis for M/M/1

1. CTMC for N (# in system)
2. Use PASTA to get distribution of T from N for a "tagged" arrival

Q Is N(t) a CTMC for M/GI/1 FCFS?

A: No, need to keep track of age or residual size of job at server.

age : rate of departure = hazard rate

residual size : departure w.p. 1 when residual size = 0

It is possible to analyze the process $(N(t), A(t))$

$N(t)$: # jobs in sys; $A(t)$: age of job at server

But tedious

(Kleinrock : method of supplementary variables)

Perhaps if we only inspect state at certain special times, these might form a Markov chain?

Idea: look at system only at instants of departures

state = # jobs left behind by departure

This is called the "embedded Markov chain".

Q: Let's call this random variable N_d. How does N_d helps us?

A: We showed earlier that when arrivals and departures occur one at a time

$$N_d \stackrel{d}{=} N_a \qquad (\text{\# seen by an arrival})$$

and using PASTA

$$N_d \stackrel{d}{=} N_a = N.$$

Define:

A_S = # arrivals of Poisson Process (λ) during service time S

$a_k = Pr[A_S = k]$

$$G_{A_S}(z) = \mathbb{E}[z^{A_S}] = \sum_{k=0}^{\infty} a_k z^k \qquad (z\text{-transform of } A_S)$$

$$= \mathbb{E}_S\left[\mathbb{E}[z^{A_S}]\right]$$

$$= \mathbb{E}_S\left[\sum_{k=0}^{\infty} e^{-\lambda s}\frac{(\lambda s)^k}{k!} z^k\right]$$

$$= \mathbb{E}_S\left[e^{-\lambda(1-z)S}\right]$$

$$= \mathcal{L}_S(\lambda(1-z)) \qquad \left(\mathcal{L}_S(s) = \mathbb{E}[e^{-sS}]\right)$$

Q: Transition probabilities (P_{ij}) for the embedded Markov chain?

A: $P_{ij} \equiv Pr[$ j jobs in system at next departure given i jobs after current departure $]$

Case: $i \geq 1$: $P_{ij} = Pr[(j-i+1)$ arrivals during one service time$]$

$= a_{j-i+1}$

Case: $i = 0$: $P_{0j} = Pr[$ j arrivals during the service time of the job (yet to arrive) $]$

$= a_j$

so $P = [P_{ij}] =$

$$\begin{bmatrix} a_0 & a_1 & a_2 & a_3 & \cdots \\ a_0 & a_1 & a_2 & a_3 & \cdots \\ & a_0 & a_1 & a_2 & \cdots \\ & & a_0 & a_1 & \end{bmatrix}$$

Balance equations

$\pi_j = Pr[N_d = j] = Pr[N = j]$

$$\begin{aligned} \pi_0 &= \pi_0 a_0 + \pi_1 a_0 \\ \pi_1 &= \pi_0 a_1 + \pi_1 a_1 + \pi_2 a_0 \\ \pi_2 &= \pi_0 a_2 + \pi_1 a_2 + \pi_2 a_1 + \pi_3 a_0 \\ &\vdots \\ \pi_j &= \pi_0 a_j + \pi_1 a_j + \pi_2 a_{j-1} + \cdots + \pi_{j+1} a_0 \end{aligned}$$

Now for the useful trick: Convert the system of equations into just one equation

- multiply b.e. for π_j by z^j
- add all the equations

$$\left[\quad \pi_0 = \pi_0 a_0 + \pi_1 a_0 \quad\right] \times z^0$$

$$\left[\quad \pi_1 = \pi_0 a_1 + \pi_1 a_1 + \pi_2 a_0 \quad\right] \times z^1$$

$$\left[\quad \pi_2 = \pi_0 a_2 + \pi_1 a_2 + \pi_2 a_1 + \pi_3 a_0 \quad\right] \times z^2$$

$$\vdots$$

$$\left[\quad \pi_j = \pi_0 a_j + \pi_1 a_j + \pi_2 a_{j-1} + \cdots + \pi_{j+1} a_0 \quad\right] \times z^j$$

$$\Rightarrow \sum_{j=0}^{\infty} \pi_j z^j = \pi_0 \sum_{j=0}^{\infty} a_j z^j + \pi_1 \sum_{j=0}^{\infty} a_j z^j + (\pi_2 z) \sum_{j=0}^{\infty} a_j z^j + \cdots + (\pi_j z^{j-1}) \sum_{j=0}^{\infty} a_j z^j + \cdots$$

$$\Rightarrow G_N(z) = G_{A_s}(z)\left(\pi_0 + \pi_1 + \pi_2 z + \pi_3 z^2 + \cdots + \pi_j z^{j-1} + \cdots\right)$$

$$= G_{A_s}(z)\left(\pi_0 - \frac{\pi_0}{z} + \frac{\pi_0}{z} + \pi_1 + \pi_2 z + \cdots + \pi_j z^{j-1} + \cdots\right)$$

$$= G_{A_s}(z)\left(\pi_0\left(1 - \frac{1}{z}\right) + \frac{1}{z} G_N(z)\right)$$

$$\Rightarrow z \cdot G_N(z) = G_{A_s}(z)\left(\pi_0(z-1) + G_N(z)\right)$$

$$= \mathcal{L}_s(\lambda(1-z))\left(\pi_0(z-1) + G_N(z)\right)$$

$$\Rightarrow G_N(z) = \frac{\pi_0 (1-z)\, \mathcal{L}_s(\lambda(1-z))}{\mathcal{L}_s(\lambda(1-z)) - z}$$

Q: $\pi_0 = ?$

A1: $(1-\rho)$ (duh!)

A2: $G_N(1) = 1 \Rightarrow \dfrac{\pi_0 \cdot (1-1)\, \mathcal{L}_s(\lambda(1-z))}{\mathcal{L}_s(\lambda(1-1)) - 1} = 0$ $\left(\frac{0}{0} \Rightarrow \text{use l'Hospital}\right)$

$$\Rightarrow \frac{\pi_0 (-1)}{(-\lambda)\, \mathcal{L}_s'(s)\big|_{s=0} - 1} = \frac{-\pi_0}{(-\lambda)(-E[S]) - 1} = 1$$

$$\Rightarrow \pi_0 = 1 - \rho$$

□

Finally

$$G_N(z) = \frac{(1-\rho)(1-z)\,\mathcal{L}_S(\lambda(1-z))}{\mathcal{L}_S(\lambda(1-z)) - z}$$

z-transform of $N^{M/G/1/FCFS}$

Sanity check: $M/M/1 \Rightarrow \mathcal{L}_S(\lambda(1-z)) = \frac{\mu}{\mu+\lambda-\lambda z}$

which gives $G_N(z) = \frac{(1-\rho)(1-z)\,\mu/(\mu+\lambda-\lambda z)}{\mu/(\mu+\lambda-\lambda z) - z}$

$$= \frac{(1-\rho)(1-z)\,\mu}{(\mu-\lambda z)(1-z)}$$

$$= \frac{1-\rho}{1-\rho z} = (1-\rho)\left(1+\rho z+\rho^2 z^2+\cdots\right)$$

(that is, $Pr[N=n] = (1-\rho)\rho^n$) □

From $G_N(z)$ to $\mathcal{L}_T(s)$

Distribution of N gives us some useful information about system but the most important performance metric is T_Q or T

customer's experience

Not easy to extend M/M/1 analysis

→ by PASTA, arrival sees N jobs in system but N & E are correlated.

↳ residual size of job at server

But N_d = # of arrivals during the system/sojourn time of the departing job

(because single server FCFS ⇒ jobs leave in the order of arrival)

Q: What is the number of Poisson Process (λ) arrivals during system time T?

A: $\mathcal{L}_T(\lambda(1-z))$

$$\Rightarrow \quad \mathcal{L}_T(\lambda(1-z)) = G_N(z) = \frac{(1-\rho)(1-z)\,\mathcal{L}_S(\lambda(1-z))}{\mathcal{L}_S(\lambda(1-z)) - z}$$

substitute $\lambda(1-z) = s \Rightarrow z = 1 - s/\lambda$

$$\Rightarrow \quad \boxed{\mathcal{L}_T(s) = \frac{(1-\rho)\,s\,\mathcal{L}_S(s)}{s - \lambda(1-\mathcal{L}_S(s))}}$$

Laplace transform of M/G/1/FCFS response time.

and since:

$$T = T_Q + S$$

(T_Q: time in queue; S: job size) ; T_Q and S are independent

$$\Rightarrow \quad \mathcal{L}_T(s) = \mathcal{L}_{T_Q}(s)\,\mathcal{L}_S(s)$$

$$\Rightarrow \quad \boxed{\mathcal{L}_{T_Q}(s) = \frac{(1-\rho)\,s}{s - \lambda(1-\mathcal{L}_S(s))}}$$

Laplace transform of M/G/1 FCFS queueing/wait time/delay.

We can actually say more about T_Q if we do some rearranging

$$\mathcal{L}_{T_Q}(s) = \frac{1-\rho}{1-\lambda\left(\frac{1-\mathcal{L}_S(s)}{s}\right)} = \frac{1-\rho}{1-\rho\left(\frac{1-\mathcal{L}_S(s)}{s\,\mathbb{E}[S]}\right)}$$

$$\boxed{\mathcal{L}_{T_Q}(s) = \frac{1-\rho}{1-\rho\cdot\mathcal{L}_{S_e}(s)} = (1-\rho)\sum_{n=0}^{\infty}\left(\rho\,\mathcal{L}_{S_e}(s)\right)^n}$$

($\mathcal{L}_{S_e}(s)$ = Laplace t. of stationary excess of S)

Imp: $\left(T_Q = \sum_{i=0}^{N} S_{e_i}\right)$ where $N \stackrel{d}{=} \text{Geom}(1-\rho) - 1$ & S_{e_i} are iid. samples of stationary excess of S. !

Remark:

* We have already seen this result when we showed that "M/GI/1 under symmetric scheduling policy has stationary # jobs distributed like an M/M/1 ... ... and conditioned on N=n, the residual sizes of the N jobs are iid samples from S_e"

* For FCFS: T_Q = unfinished workload ahead of an arrival
  = unfinished workload for any "work conserving" policy
  = unfinished workload for a work conserving ($\phi(n)=1$) symmetric policy.

□

Now that we have $\mathcal{L}_T(s)$ & $\mathcal{L}_{T_Q}(s)$ we can:

(1) invert it to find the distribution (or spectral representation) of T, T_Q
(2) calculate moments
(3) study it asymptotically (e.g. Heavy traffic: $\rho \uparrow 1$)

Example

(1) M/M/1

$$\mathcal{L}_T(s) = \frac{(1-\rho)\, s\, \frac{\mu}{\mu+s}}{s - \lambda\left(1 - \frac{\mu}{\mu+s}\right)} = \frac{(\mu-\lambda)}{(\mu-\lambda)+s}$$

$$\Rightarrow T \sim Exp(\mu-\lambda)$$

(2) Variance of $T_Q^{M/G/1/FCFS}$:

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

(3) As $\rho \to 1$, $(1-\rho)\, T_Q \sim$ Exp with mean $\mathbb{E}[S_e] = \frac{\mathbb{E}[S^2]}{2\mathbb{E}[S]}$

Alternate derivation of $\mathcal{L}_T(s)$

We will write 4 relationships between $\mathcal{L}_T(s)$, $\mathcal{L}_{T_\mathcal{B}}(s)$, $G_N(z)$, $G_{N_\mathcal{B}}(z)$ and eliminate

(1) $N \stackrel{d}{=} A_T \Rightarrow$

$$G_N(z) = \mathcal{L}_T(\lambda(1-z))$$

(2) $T = T_\mathcal{B} + S \; ; \; T_\mathcal{B} \perp S \Rightarrow$

$$\mathcal{L}_T(s) = \mathcal{L}_{T_\mathcal{B}}(s) \cdot \mathcal{L}_S(s)$$

(3) $N_\mathcal{B} \stackrel{d}{=} A_{T_\mathcal{B}} \Rightarrow$

$$G_{N_\mathcal{B}}(z) = \mathcal{L}_{T_\mathcal{B}}(\lambda(1-z))$$

(4) $N_\mathcal{B} = (N-1) + \mathbb{1}_{\{N=0\}} \Rightarrow$

$$\begin{aligned} G_{N_\mathcal{B}}(z) &= \pi_0 + \pi_1 + \pi_2 z + \pi_3 z^2 + \cdots \\ &= \pi_0 + \frac{1}{z}\left(\pi_1 z + \pi_2 z^2 + \cdots\right) \\ &= \pi_0\left(1 - \frac{1}{z}\right) + \frac{1}{z}\left(\pi_0 + \pi_1 z + \cdots\right) \\ &= \pi_0\left(1 - \frac{1}{z}\right) + \frac{1}{z} G_N(z) \end{aligned}$$

or,

$$G_N(z) = z\, G_{N_\mathcal{B}}(z) + \pi_0(1-z)$$

substitute (1), (3) into (4)

$$\mathcal{L}_T(\lambda(1-z)) = z\, \mathcal{L}_{T_\mathcal{B}}(\lambda(1-z)) + \pi_0(1-z)$$

eliminate $\mathcal{L}_T$ using (2)

$$\mathcal{L}_{T_\mathcal{B}}(\lambda(1-z)) \cdot \mathcal{L}_S(\lambda(1-z)) = z\, \mathcal{L}_{T_\mathcal{B}}(\lambda(1-z)) + \pi_0(1-z)$$

Substitute $s = \lambda(1-z)$; $\pi_0 = 1-\rho$

$$\boxed{\mathcal{L}_{T_\mathcal{B}}(s) = \frac{(1-\rho)s}{s - \lambda(1 - \mathcal{L}_S(s))}}$$

Approach #3 for $\mathcal{L}_T(s)$: Lindley's recursion

Recall: for GI/GI/1 FCFS

$$W_{n+1} = (W_n + S_n - A_{n+1})^+$$

- W_{n+1}: waiting time of job $(n+1)$
- W_n: waiting time of job n
- S_n: size of job n
- A_{n+1}: inter-arrival time of job $(n+1)$

W_n : discrete time, continuous state Markov process.

Therefore:

$$W \stackrel{d}{=} (W + S - A)^+$$

To complete:

(1) Write equation for $f_W(w)$

(2) take Laplace transform using same approach we used for $G_N(z)$

Beyond FCFS scheduling

M/GI/1 FCFS $\Rightarrow$ $\mathbb{E}[T]_{\$} = \rho \frac{\mathbb{E}[S]}{1-\rho} \cdot \left(\frac{C_S^2+1}{2}\right)$

so if $\rho = \lambda \mathbb{E}[S] < 1$, but $\mathbb{E}[S^2] = \infty$, then $\mathbb{E}[T]_{\$} = \infty$

e.g. $Pr[S > x] = \frac{1}{(1+x)^{3/2}}$

Designing queueing systems to mitigate the effect of job size variability is one of the core aims of queueing theory

Common performance tuning knobs

(1) choose # servers, speed (ie one fast vs. many slow)

(2) load balancing, routing (i.e., which server to send the job to)

(3) scheduling (ie. sequencing / prioritizing jobs at servers)

Scheduling policies

Characterized mainly along two axes

(1) Preemptive vs. Non-preemptive policies

(2) size-based vs. blind policies

Examples

	Non-preemptive	Preemptive
Blind	FCFS: First Come First Serve LCFS: Last Come First Serve ROS: Random Order of Service	PS: processor Sharing P-LCFS: Preemptive LCFS LAS: Least Attained Service
Size-based	SJF: Shortest Job First PRIO: class-based priority	SRPT: Shortest Remaining Processing Time PSJF: Preemptive Shortest Job First

M/GI/1 Scheduling : Non-preemptive, Blind policies

These are the simplest policies :

(1) FCFS : pick the earliest arrival to serve

(2) LCFS : pick the last arrival from queue to serve

(3) ROS : pick a random job to serve

Q: Which of the above has the smallest $\mathbb{E}[T]$?

A: Consider example: $S \sim Exp(\mu)$

Little's law $\Rightarrow \mathbb{E}[T] = \frac{1}{\lambda}\mathbb{E}[N]$

with $S \sim Exp(\mu)$, all blind policies (preemptive and non-prempt) have exactly the same distribution of N

(CTMC is identical to M/M/1/FCFS)

> **Theorem**: For GI/M/1
>
> $$\mathbb{E}[T^P] = \mathbb{E}[T^{FCFS}]$$
>
> for all blind scheduling policies P.

Now consider generally distributed S

$\Rightarrow$ N_d still forms an embedded DTMC with the same transition probabilities as FCFS

$P_{ij} = Pr[$ j jobs left behind by next departure given current departure leaves behind i jobs $]$

P_{ij} only depended on the fact that under FCFS job sizes were iid (independent of i) & scheduling non-preemptive

$\Rightarrow$ N_d has the same distribution for all blind non-preemptive policies

$\Rightarrow$ N has the same distribution for blind, non-preemptive

$\Rightarrow$ $\mathbb{E}[N]$ is the same for all blind, non-preemptive

Theorem: For M/GI/1

$$\mathbb{E}[T^P] = \mathbb{E}[T^{FCFS}]$$

for all blind, non-preemptive policies P.

So we don't gain anything in $\mathbb{E}[T]$ by messing with scheduling as long as it is blind and non-preemptive

Q How about $\text{var}(T)$?

Theorem: For any blind, non-preemptive policy P, under GI/GI/1

$$\text{var}(T^{FCFS}) \leq \text{var}(T^P) \leq \text{var}(T^{LCFS})$$

Therefore FCFS is the "optimal" blind, non-preemptive policy.

Proof (of $\text{var}(T^{FCFS}) \leq \text{var}(T^P)$):

We will couple two sample paths; one under P other under FCFS

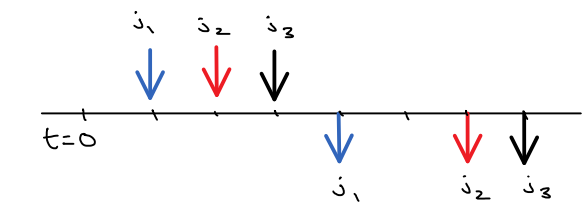
* couple arrival times $a_1 \leq a_2 \leq a_3 \leq \ldots$.

* couple service times

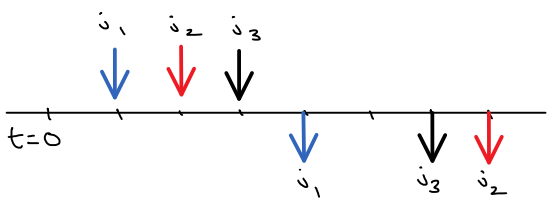
s_n = service time of n^{th} job to begin service

example: $P \equiv LCFS$, $a_1 = 1, a_2 = 2, a_3 = 3$; $s_1 = 3, s_2 = 2, s_3 = 1$

FCFS:



LCFS:



* from coupling of $\{a_i\}$ and $\{s_i\}$

the departure times are coupled under P & FCFS

$d_1 \le d_2 \le d_3 \le \dots$

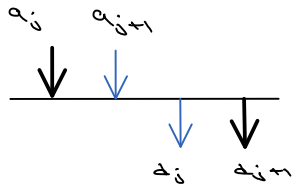
(Proof by induction)

T_j = system time of j^{th} arrival

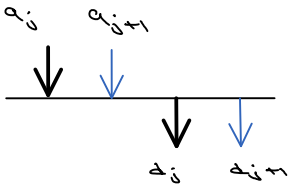
$T_j^{FCFS} = d_j - a_j$ (j^{th} departure is also the j^{th} arrival)

Assume: $\exists\ j$ s.t.

$T_j^P = d_{j+1} - a_j$ $\quad$ $T_{j+1}^P = d_j - a_{j+1}$ $\quad$ (P:)



$T_j^{FCFS} = d_j - a_j$ $\quad$ $T_{j+1}^{FCFS} = d_{j+1} - a_{j+1}$ $\quad$ (FCFS:)



$T_{j+1}^P \le T_j^{FCFS}$, $T_{j+1}^{FCFS} \le T_j^P$

and

$$T_j^P + T_{j+1}^P = T_j^{FCFS} + T_{j+1}^{FCFS}$$

$\Rightarrow$

$$\left(T_j^P\right)^2 + \left(T_{j+1}^P\right)^2 \ge \left(T_j^{FCFS}\right)^2 + \left(T_{j+1}^{FCFS}\right)^2$$

for arbitrary P, we can keep making pairwise exchanges to get from FCFS departure sequence to P departure sequence

each such swap increases $\sum_{j=1}^{n} \left(T_j^P\right)^2$

In summary: $\sum_{j=1}^{n} \left(T_j^P\right)^2 \ge \sum_{j=1}^{n} \left(T_j^{FCFS}\right)^2$

$$\Rightarrow \lim_{n\to\infty} \frac{1}{n} \sum_{j=1}^{n} \left(T_j^P\right)^2 \ge \lim_{n\to\infty} \frac{1}{n} \sum_{j=1}^{n} \left(T_j^{FCFS}\right)^2$$

$$\Rightarrow \quad \mathbb{E}\left[(T^P)^2\right] \geq \mathbb{E}\left[(T^{FCFS})^2\right]$$

together with $\mathbb{E}[T^P] = \mathbb{E}[T^{FCFS}]$

$$\Rightarrow \quad \text{var}(T^P) \geq \text{var}(T^{FCFS})$$

(a similar argument shows $\text{var}(T^P) \leq \text{var}(T^{LCFS})$

Note: We did not make any assumption about arrival process
But we coupled service times not at arrival but
deferring the decision until job reaches server
(since P was blind, this is kosher)

To argue a measure preserving bijection b/w sample path under P & FCFS, suffices if job sizes are i.i.d.

□

Corollary: $\text{var}(T_Q^{FCFS}) \leq \text{var}(T_Q^P) \leq \text{var}(T_Q^{LCFS})$

for all blind, non-preemptive policies P in G/GI/1.

□

If time permits, we will prove the following stronger optimality result for FCFS:

Thm: For G/GI/1, if the job size distribution has increasing failure rate (IFR), then

$$N^{FCFS} \leq_{st} N^P$$

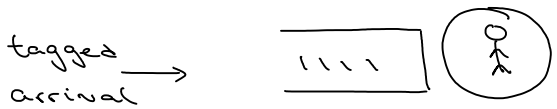
for any blind policy (preemptive or non-preemptive) P.

That is, once we start serving a job, it is optimal to finish serving it when don't know job sizes

M/G/1/LCFS analysis: T_Q^{LCFS}

We have argued that LCFS is the worst blind, non-preemptive policy, but it is still instructive to analyze T_Q^{LCFS}

"Tagged job analysis"



Case: arrival finds system idle (with prob. $(1-\rho)$)

$$T_Q^{idle} = 0 \Rightarrow \mathcal{L}_{T_Q^{idle}}(s) = 1$$

Case: arrival finds system busy (with prob. ρ)

↳ sees N_Q jobs in queue

LCFS ⇒ arrival jumps ahead of jobs in queue

↳ sees job in server with residual size $\sim S_e$

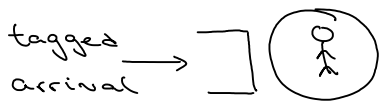
LCFS ⇒ must wait for this job to finish

+ (all jobs that arrive after me but before this job finishes)

+ (all jobs that arrive during ↑ jobs' service)

+ ...

In other words: it is as if



* arrival sees $N_Q = 0$, one job at server
* waits until the system idles

≡ busy period started by S_e !

M/GI/1 busy period analysis refresher:

define

B = busy p. started by S

B_x = busy p. started by deterministic size x

B_X = busy p. started by random size X.

recall:

$A_x \equiv$ # Poisson(λ) arrivals during random duration X

$$G_{A_X}(z) = \mathbb{E}[z^{A_x}] = \mathcal{L}_X(\lambda(1-z))$$

Therefore:

$$B_x = x + \sum_{i=1}^{A_x} B_i$$

where B_i = busy period started by i^{th} arrival during service of x

$$\Rightarrow \mathcal{L}_{B_x}(s) = \mathbb{E}[e^{-sB_x}]$$

$$= \mathbb{E}\left[e^{-s\left(x + \sum_{i=1}^{A_x} B_i\right)}\right]$$

$$= \mathbb{E}_{A_x}\left[e^{-sx} \cdot \left(\mathbb{E}[e^{-sB}]\right)^{A_x}\right]$$

$$= \mathbb{E}_{A_x}\left[e^{-sx}\left(\mathcal{L}_B(s)\right)^{A_x}\right]$$

$$= e^{-sx} G_{A_x}(\mathcal{L}_B(s))$$

$$= e^{-sx} e^{-(\lambda(1-\mathcal{L}_B(s)))x}$$

$$= e^{-(s+\lambda(1-\mathcal{L}_B(s)))x}$$

$$\Rightarrow \mathcal{L}_{B_x}(s) = \mathcal{L}_X(s+\lambda(1-\mathcal{L}_B(s)))$$

and

$$\mathcal{L}_B(s) = \mathcal{L}_{B_S}(s) = \mathcal{L}_S(s+\lambda(1-\mathcal{L}_B(s)))$$

since T_B^{busy} = b. period started by S_e

$$\mathcal{L}_{T_B^{busy}}(s) = \mathcal{L}_{B_{Se}}(s)$$

$$= \mathcal{L}_{S_e}\left(s + \lambda(1 - \mathcal{L}_B(s))\right)$$

$$= \frac{1 - \mathcal{L}_S(s+\lambda(1-\mathcal{L}_B(s)))}{(s+\lambda(1-\mathcal{L}_B(s)))\ \mathbb{E}[S]}$$

$$= \frac{1 - \mathcal{L}_B(s)}{(s+\lambda(1-\mathcal{L}_B(s)))\ \mathbb{E}[S]}$$

Finally,

$$\mathcal{L}_{T_B^{LCFS}}(s) = (1-\rho) + \rho \cdot \frac{1 - \mathcal{L}_B(s)}{\left(s+\lambda(1-\mathcal{L}_B(s))\right)\mathbb{E}[S]}$$

$\Rightarrow$

$$\boxed{\mathcal{L}_{T_B^{FCFS}}(s) = (1-\rho) + \frac{\lambda\left(1 - \mathcal{L}_B(s)\right)}{s + \lambda - \lambda\,\mathcal{L}_B(s)}}$$

Exercise:

Show (for M/GI/1):

$$\mathbb{E}\left[(T_B^{LCFS})^2\right] = \frac{\mathbb{E}\left[(T_B^{FCFS})^2\right]}{1-\rho}$$

The GI/M/1 queue

Service times $\sim Exp(\mu)$ i.i.d.

Interarrival times $A_i \sim$ Generally distributed with $\mathbb{E}[A] = \frac{1}{\lambda}$

$$\rho \doteq \frac{\lambda}{\mu} \qquad \rho = \text{load} = \text{fraction of time server is busy}$$

$N(t) = \#$ jobs at time t

Q: Is $N(t)$ a CTMC?

A: No. The transition rate $i \to i+1$ depends on time since last arrival, or time until next arrival.

So a complete Markovian description will need to keep track of some information about arrival process.

Idea: look at system only at instants of job arrivals

ie. a DTMC embedded at arrival instants.

Let $N_a(n) \doteq$ # jobs seen by n^{th} arrival

$Y(n) =$ # events of a Poisson process with rate μ during interarrival time of n^{th} arrival

then,

$$N_a(n) \stackrel{d}{=} \left(N_a(n-1) + 1 - Y(n) \right)^+$$

(Notice the similarity to Lindley recursion)

Let $Y \stackrel{d}{=} Y(1)$ denote the number of Poisson(μ) events during A.

$$d_i \doteq Pr[Y = i] \qquad i = 0, 1, 2, \ldots$$

Q: What are the transition probabilities for embedded DTMC?

A:

$P_{ij} = P_r[$ next arrival sees j jobs given current arrival sees i jobs$]$

$$P_{ij} = \begin{cases} 0 & j > i+1 \\ d_0 & j = i+1 \\ d_1 & j = i \\ \vdots & \\ d_i & j = 1 \\ \sum_{k=i+1}^{\infty} d_k & j = 0 \end{cases}$$

or,

$$P = \begin{bmatrix} d_{\geq 1} & d_0 & & \\ d_{\geq 2} & d_1 & d_0 & \\ d_{\geq 3} & d_2 & d_1 & d_0 \\ & & & \ddots \end{bmatrix} \qquad \left(d_{\geq j} = \sum_{k=j}^{\infty} d_k \right)$$

We have already encountered such transition matrices when we looked at Matrix Analytic Methods!

Here is the line of reasoning again:

(1) $\frac{\pi_{i+1}}{\pi_i} = \mathbb{E}[$ # visits to state $i+1$ b/w two consecutive visits to state i $]$

We assumed the above relation to be true. Here is the

Proof: recall renewal reward theorem

$$\text{avg. rate of reward} = \frac{\mathbb{E}[\text{reward per renewal cycle}]}{\mathbb{E}[\text{length of a renewal cycle}]}$$

By definition, DTMC regenerates on every visit to state i

$\Rightarrow$ renewal cycles are visits to state i

reward = 1 on every visit to state $i+1$

$$\Rightarrow \lim_{t\to\infty} \frac{R(t)}{t} = \pi_{i+1} = \frac{\mathbb{E}[\text{visits to } i+1 \text{ in a renewal cycle}]}{\mathbb{E}[\text{length b/w two visits to } i]}$$

$$= \frac{\mathbb{E}[\#\text{ visits to } i+1 \text{ during a renewal cycle}]}{1/\pi_i}$$

□

(2) claim : $\frac{\pi_{i+1}}{\pi_i}$ is indep. of i for $i \geq 0$

Proof : consider $i_1 < i_2$

(i) $P_{i_1 i_1+1} = P_{i_2 i_2+1} = d_0$

Therefore with probability $(1-d_0)$; the number of visits to $(i+1)$ are 0 , otherwise we transition to $(i+1)$.

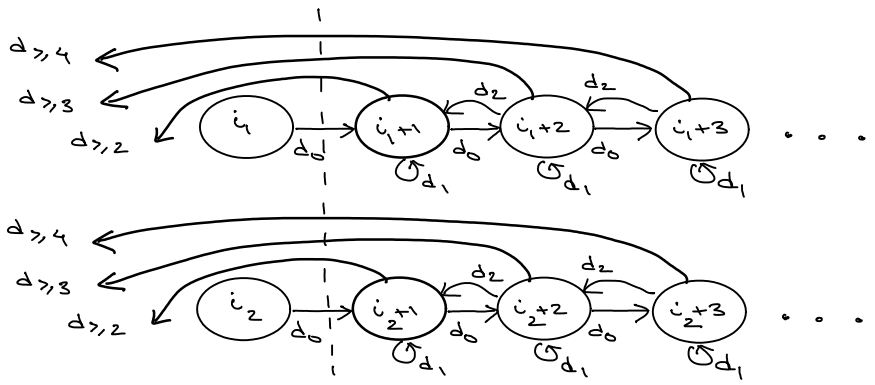
(ii) Since the DTMC is <u>skip-free to the right</u> , the excursion ends on a transition to state $j \leq i$.

But for both i_1, i_2 , the transient process on the aggregated state spaces

i_1 : $\{0,1,2,\ldots,i_1\}, i_1+1, i_1+2, i_1+3, \ldots$

i_2 : $\{0,1,2,\ldots,i_2\}, i_2+1, i_2+2, i_2+3, \ldots$

with initial state (i_1+1) , (i_2+1) are identical



□

(3) $\frac{\pi_{i+1}}{\pi_i} = \beta$ for $i = 0, 1, 2, \ldots$

$\Rightarrow \quad \pi_i = \pi_0 \beta^i = (1-\beta)\beta^i$

to solve for β, we can use any balance equation.

using B.E. for π_1:

$$\pi_1 = \pi_0 d_0 + \pi_1 d_1 + \pi_2 d_2 + \cdots$$

$$\Leftrightarrow \quad (1-\beta)\beta = (1-\beta)\sum_{i=0}^{\infty} d_i \beta^i$$

$$= (1-\beta)\,\mathbb{E}[\beta^Y]$$

$$= (1-\beta)\,\mathcal{L}_A(\mu(1-\beta))$$

$Y \triangleq$ # Poisson(μ) events in interarrival time A

$\Rightarrow G_Y(z) = \mathcal{L}_A(\mu(1-z))$

Therefore: β is the smallest non-negative solution of

$$\boxed{\beta = \mathcal{L}_A(\mu(1-\beta))}$$

Claim: $\beta = \mathcal{L}_A(\mu(1-\beta))$ has a unique soln in $(0,1)$ when $\rho < 1$.

Proof: $\mathcal{L}_A(\mu(1-z)) = \mathbb{E}[z^Y] = \sum_{i=0}^{\infty} d_i z^i$

(i) since $d_i \geq 0$ & z^i are convex, increasing

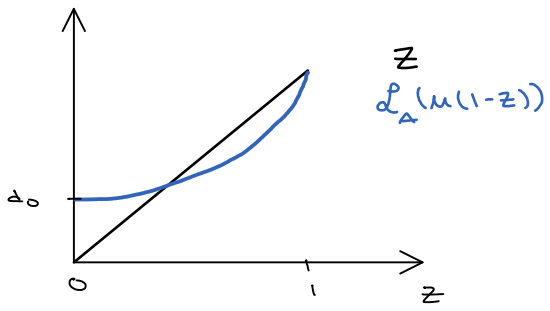
$\Rightarrow \mathcal{L}_A(\mu(1-z))$ is convex increasing

(ii) $\mathcal{L}_A(\mu(1-0)) = d_0 > 0$

(iii) $\mathcal{L}_A(\mu(1-1)) = \mathcal{L}_A(0) = 1$

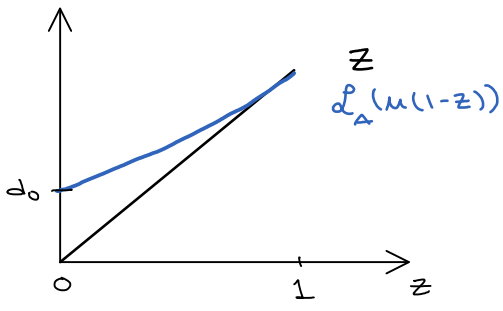
Therefore we have the following situations :

(i)



unique $\beta \in (0,1)$

(ii)



no solution for $\beta \in (0,1)$

To find which case we are in

$$\frac{\partial}{\partial z} \mathcal{L}_A(\mu(1-z)) \Big|_{z=1} = -\mu \frac{\partial}{\partial s} \mathcal{L}_A(s) \Big|_{s=0}$$

$$= (-\mu)(\mathcal{L}_A'(0))$$

$$= (-\mu)(-\mathbb{E}[A])$$

$$= \mu/\lambda$$

$$= 1/\rho$$

$\Rightarrow$ when $\rho < 1$, we are in the first case of unique $\beta \in (0,1)$.

□

To solve for β

(i) pick $\beta(0) \in (0,1)$

(ii) iterate $\beta(n+1) = \mathcal{L}_A(\mu(1-\beta(n)))$

Waiting time distribution

$$P_r[W=0] = P_r[N_a=0] = 1-\beta$$

But conditioned on non-zero wait time

$$[W \mid W>0] = \sum_{i=1}^{N_a} S_i$$

$$\stackrel{d}{=} \text{Exp}(\mu(1-\beta))$$

$[N_a \mid N_a > 0]$ has Geom$(1-\beta)$ distribution

Remarks :

* $E[N_q] = \frac{\beta}{1-\beta}$

Therefore GI/M/1 does not seem to exhibit $\left(\frac{1}{1-\rho}\right)$ behavior like the M/GI/1 queue

* EXAMPLE : D/M/1

$$\mathcal{L}_A(s) = e^{-s/\lambda} \Rightarrow \beta = \mathcal{L}_A(\mu(1-\beta)) = e^{-s(1-\beta)/\rho}$$

when $\rho = 0.5$, $\beta \approx 0.203$

$\Rightarrow$ Even though the server is busy 50% time, almost 80% jobs find the server idle on arrival.

This is one example of the broader principle: variability (in service or interarrival times) hurts performance

* If we fix A, but slow service so that

$$\mu \downarrow \lambda \Rightarrow \rho \uparrow 1$$

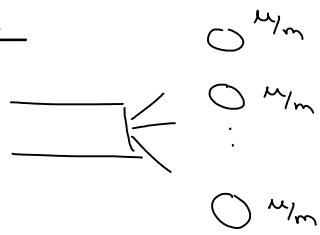
then,

(*HW) $$\lim_{\rho \uparrow 1} \frac{(1-\beta)}{(1-\rho)} = \frac{2}{C_A^2+1} \qquad \left(C_A^2 = \frac{\text{Var}(A)}{(E[A])^2}\right)$$

So, we do get back the $\left(\frac{1}{1-\rho}\right)$ asymptote in heavy traffic

---

GI/M/m



m servers of speed μ/m each.

Analysis almost identical

Let: $\pi_i \doteq P_r[N_a = i]$ (stat. dist of embedded chain)

then $\pi_{(m-1)+i} = \pi_m \beta^i$

where $\beta \in (0,1)$ solves $\beta = \mathcal{L}_A(\mu(1-\beta))$

$\Rightarrow \quad (N_a - (m-1) \mid N_a \geq m) \sim Geom(1-\beta)$

$\Rightarrow \quad [W \mid W > 0] \sim Exp(\mu(1-\beta))$

Theorem: Let $W^{G/M/m}$ denote the waiting time in the system with m servers of speed μ/m each.

Then

$$m_2 < m_1 \Rightarrow W^{G/M/m_2} \geq_{st} W^{G/M/m_1}$$

Proof (via coupling) is identical to M/M/1 vs. M/M/2 proof we covered earlier.

(Note that the proof did not require the Poisson assumption on the arrival process.)

□

GI/GI/1 queue - I : Lindley recursion, Kingman's bounds

With GI/GI/1 queue we enter the realm of queueing theory where exact analysis would be usually not possible and we will rely on bounds and approximations

This naturally entails a tradeoff b/w the accuracy of the bounds/ approxmations vs. the simplicity of their expressions.

GI/GI/1 : i.i.d. service times $S_i \sim F_S(\cdot)$; $\mathbb{E}[S] = 1/\mu$

i.i.d. interarrival times $A_i \sim F_A(\cdot)$; $\mathbb{E}[A] = 1/\lambda$

$$\rho \doteq \lambda/\mu \quad \Rightarrow \quad \lim_{t\to\infty} \frac{1}{t} \sum_{i=1}^{N(t)} S_i = \rho$$

$N(t)$ = # arrivals by t.

Defintions : W_0 = workload at time 0

u_n = arrival time of n^{th} customer (C_n)

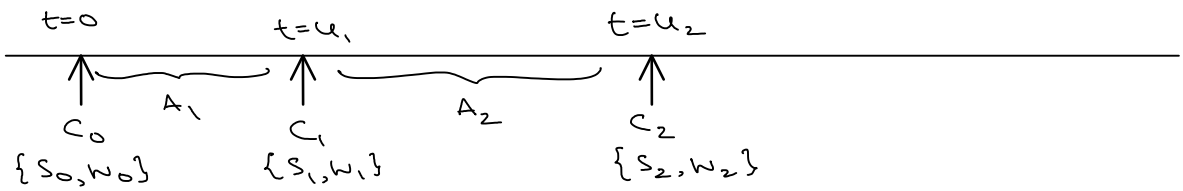
$u_0 = 0$

$A_n = u_n - u_{n-1}$ = interarrival time of C_n

S_n = service requirement of C_n $(n=0,1,2,\ldots)$

W_n = waiting time of C_n

= unfinished workload at $t = u_n$



Lindley recursion :

$$W_{n+1} = (W_n + S_n - A_{n+1})^+$$

I. Stability of GI|GI|1

Proving stability of GI|GI|1 queues (ie. $\lim_{n\to\infty} W_n < \infty$) is more challenging than M|G|1 or G|M|1 since we do not have the theory of Markov chains at our disposal.

We will use Lindley recursion as our starting point, and in the process also gain some intuition about the sample paths of G|G|1.

> **Theorem:** If $\rho < 1$, then for the G|G|1 queue,
>
> (i) $W_n \xrightarrow{d} W$ as $n \to \infty$, where W satisfies
>
> (ii) $\Pr(W < \infty) = 1$

Before delving into the proof, we define some auxiliary stochastic processes.

Define:

$$X_n \doteq S_{n-1} - A_n \qquad n = 1, 2, 3, \ldots$$

$$Z_n = X_1 + X_2 + X_3 + \cdots + X_n$$
$$Z_0 \doteq 0$$

$$M_n = \max_{0 \le k \le n} Z_k$$
$$M = \sup_{0 \le k < \infty} Z_k$$

$\{X_n\}$ is often called the "netput process"
(Note: $\mathbb{E}[X_n] = \mathbb{E}[S] - \mathbb{E}[A] = \mathbb{E}[A](\rho - 1)$)

$\{Z_n\}$ is a simple unreflected random walk with i.i.d increments given by $\{X_n\}$, starting at 0.

$\{M_n\}$ is a non-decreasing; non-negative, running maximum process for the random walk $\{Z_n\}$

Since Lindley recursion involves "reflecting" the workload process whenever it hits 0, we will find it easier to study it using the unreflected $\{Z_n\}$ random walk.

Lemma: $W_n \stackrel{d}{=} \max\left(W_0 + Z_n,\ M_{n-1}\right)$

Proof: Unfolding Lindley recursion:

$$\begin{aligned}
W_n &= \max\left(0,\ W_{n-1} + X_n\right) \\
&= \max\left(0,\ \max\left(0, W_{n-2} + X_{n-1}\right) + X_n\right) \qquad \text{(L.R. to } W_{n-1}\text{)} \\
&= \max\left(0,\ W_{n-2} + X_{n-1} + X_n,\ X_n\right) \\
&= \max\left(0,\ W_{n-3} + X_{n-2} + X_{n-1} + X_n,\ X_{n-1} + X_n,\ X_n\right) \qquad \text{(rearranging)} \\
&\ \ \vdots \\
&= \max\left(0,\ W_0 + X_1 + X_2 + \cdots + X_n,\ X_2 + X_3 + \cdots + X_n,\ \ldots,\ X_{n-1} + X_n,\ X_n\right) \qquad (*)
\end{aligned}$$

Since $\{X_n\}$ are i.i.d., $X_n \stackrel{d}{=} X_1$, $X_{n-1} \stackrel{d}{=} X_2$, $\ldots$

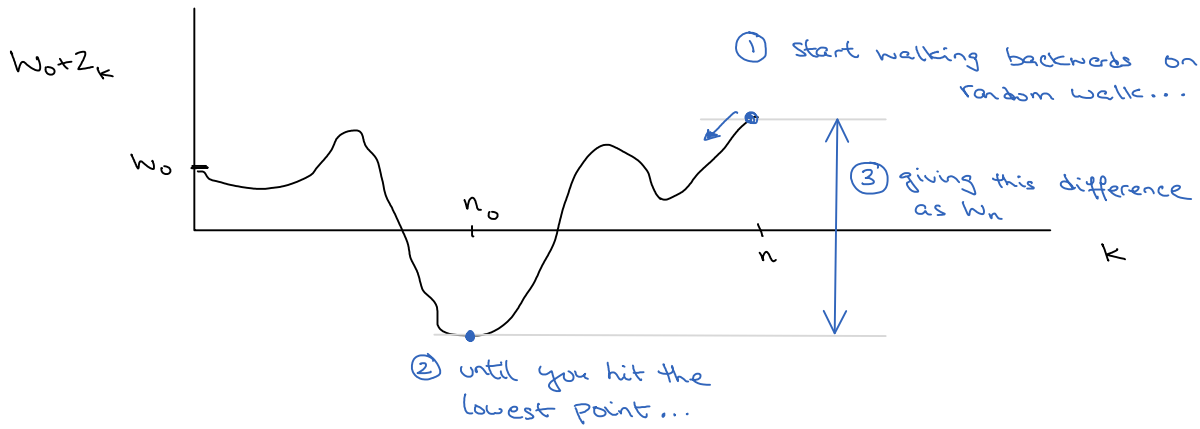
with this renaming:

$$\begin{aligned}
W_n &\stackrel{d}{=} \max\left(0,\ W_0 + X_n + X_{n-1} + \cdots + X_1,\ X_{n-1} + X_{n-2} + \cdots + X_1,\ \ldots,\ X_2 + X_1,\ X_1\right) \\
&= \max\left(0,\ W_0 + Z_n,\ Z_{n-1},\ Z_{n-2},\ \ldots,\ Z_1\right) \qquad \text{(note: the equality is only in distrib.)} \\
&= \max\left(W_0 + Z_n,\ M_{n-1}\right)
\end{aligned}$$

□

Let us understand (*) pictorially:

$$\begin{aligned}
W_n &= \max\left(X_n,\ X_n + X_{n-1},\ X_n + X_{n-1} + X_{n-2},\ \ldots,\ X_n + X_{n-1} + \cdots + W_0,\ 0\right) \\
&= W_0 + Z_n - \min\left(W_0 + X_1 + \cdots + X_{n-1},\ W_0 + X_1 + \cdots + X_{n-2},\ \ldots,\ 0,\ W_0 + X_1 + X_2 + \cdots + X_n\right) \\
&= W_0 + Z_n - \min\left(0,\ \min_{1 \le k \le n} W_0 + Z_k\right)
\end{aligned}$$



② n_0 represents the last customer to experience zero wait. $\Rightarrow \{W_k\}$ does not reflect after this point $\Rightarrow W_n = \sum_{n_0+1}^{n} X_i$

You can imagine this view would be a messy way to deal with W_n.
Which is why the statement of the Lemma is remarkable.

Suppose $W_0 = 0$. The Lemma then says:

$$W_n \stackrel{d}{=} \max(Z_n, M_{n-1}) = M_n$$

That is: W_n (an outcome of a reflected random walk) can simply be written as the maximum of an unreflected random walk!

Now we are ready to prove our stability theorem:

<u>Proof of theorem:</u> $\rho < 1 \Leftrightarrow \mathbb{E}[X] < 0$

Therefore,

$$\frac{Z_n}{n} \to \mathbb{E}[X] < 0 \quad \text{with prob. 1 by SLLN}$$

$$\Rightarrow \quad Z_n \to -\infty \quad \text{as } n \to \infty \text{ with prob 1}$$

$$\Rightarrow \quad \text{Prob}(M < \infty) = \text{Prob}\left(\sup_{0 \le k < \infty} Z_k < \infty\right) = 1$$

Now: $Z_n \to -\infty \Rightarrow W_0 + Z_n \to -\infty$ w.p. 1

$\Rightarrow$ for all n large enough

$$W_n \stackrel{d}{=} \max(W_0 + Z_n, M_{n-1}) = M_{n-1}$$

Since $\{M_n\}$ is non-decreasing, $M_n \uparrow M$

$$\Rightarrow \quad W_n \stackrel{d}{\to} M \quad \text{as } n \to \infty.$$

□

II Kingman's bounds for G/G/1 and G/G/m FCFS.

While exact analysis of G/G/1 is often impossible, Kingman obtained fairly sharp bounds on $\mathbb{E}[W]$ in a series of papers (1964-1970).

Theorem: For the GI/GI/1 queue

$$\frac{\lambda^2 \mathbb{E}[((S-A)^+)^2]}{2\lambda(1-\rho)} \leq \mathbb{E}[W] \leq \frac{\rho^2 c_S^2 + c_A^2}{2\lambda(1-\rho)}$$

Proof: denote $x^- = \max(0, -x)$ (the negative part of x)

so that $x = x^+ - x^-$

this gives: $\mathbb{E}[x] = \mathbb{E}[x^+] - \mathbb{E}[x^-]$ (*)

$\mathbb{E}[x^2] = \mathbb{E}[(x^+)^2] + \mathbb{E}[(x^-)^2]$ (**)

$\mathrm{var}(x^+) + \mathrm{var}(x^-) = \mathrm{var}(x) - 2\mathbb{E}[x^+]\mathbb{E}[x^-]$ (***)

From Lindley: $W \stackrel{d}{=} (W + S - A)^+$

for convenience: let $\hat{W} = (W+S-A)^+$ $\left(\hat{W} \stackrel{d}{=} W\right)$

$Y = (W+S-A)^+$

(*) gives

$$\mathbb{E}[W+S-A] = \mathbb{E}[\hat{W}] - \mathbb{E}[Y]$$

$$\Rightarrow \quad \mathbb{E}[W] + \mathbb{E}[S] - \mathbb{E}[A] = \mathbb{E}[\hat{W}] - \mathbb{E}[Y]$$

$$\Rightarrow \quad \boxed{\mathbb{E}[Y] = \mathbb{E}[A] - \mathbb{E}[S] = 1/\lambda - 1/\mu}$$

(***) gives:

$$\mathrm{var}(\hat{W}) + \mathrm{var}(Y) = \mathrm{var}(W+S-A) - 2\mathbb{E}[\hat{W}]\mathbb{E}[Y]$$

$$= \mathrm{var}(W) + \mathrm{var}(S) + \mathrm{var}(A) - 2\mathbb{E}[\hat{W}]\mathbb{E}[Y]$$

$\hat{W} \stackrel{d}{=} W \Rightarrow \text{var}(\hat{W}) = \text{var}(W)$

$$\Rightarrow \mathbb{E}[W] = \mathbb{E}[\hat{W}] = \frac{\text{var}(S) + \text{var}(A) - \text{var}(Y)}{2\mathbb{E}[Y]}$$

$$\Rightarrow \boxed{\mathbb{E}[W] = \frac{\text{var}(S) + \text{var}(A) - \text{var}(Y)}{2(\mathbb{E}[A] - \mathbb{E}[S])}}$$

Upper bound : since $\text{var}(Y) \geq 0$

$$\mathbb{E}[W] \leq \frac{\text{var}(S) + \text{var}(A)}{2(\mathbb{E}[A] - \mathbb{E}[S])}$$

$$= \frac{c_S^2/\mu^2 + c_A^2/\lambda^2}{2(1/\lambda - 1/\mu)}$$

$$= \frac{\rho^2 c_S^2 + c_A^2}{2\lambda(1-\rho)}$$

Lower bound : we want an upper bound on $\text{var}(Y)$

$$\text{var}(Y) = \mathbb{E}[Y^2] - (\mathbb{E}[Y])^2$$

$$= \mathbb{E}\left[\left((W+S-A)^-\right)^2\right] - (\mathbb{E}[A] - \mathbb{E}[S])^2$$

since $W \geq 0 \Rightarrow (W+S-A)^- \leq (S-A)^-$

$$\Rightarrow \text{var}(Y) \leq \mathbb{E}\left[\left((S-A)^-\right)^2\right] - (\mathbb{E}[A] - \mathbb{E}[S])^2$$

$$= \mathbb{E}[(S-A)^2] - \mathbb{E}[((S-A)^+)^2] - (\mathbb{E}[A] - \mathbb{E}[S])^2$$

$$= \text{var}(S-A) - \mathbb{E}[((S-A)^+)^2]$$

$$= \text{var}(S) + \text{var}(A) - \mathbb{E}\left[\left((S-A)^+\right)^2\right]$$

plugging into $\mathbb{E}[W]$:

$$\mathbb{E}[W] \geq \frac{\mathbb{E}\left[\left((S-A)^+\right)^2\right]}{2(\mathbb{E}[A] - \mathbb{E}[S])}$$

$$= \frac{\lambda^2 \mathbb{E}\left[\left((S-A)^+\right)^2\right]}{2\lambda(1-\rho)}$$

□

Remarks:

(1) For the upper bound we lower bounded

$$Var(Y) \geq 0$$

But $Y = (W+S-A)^-$ is only positive when a job has zero waiting time.

In heavy traffic ($\rho \nearrow 1$), we expect Y to be close to 0. and hence the effect of ignoring $Var(Y)$ to be small.

This is indeed true and Kingman's upper bound is asymptotically tight in heavy traffic.

(2) The lower bound is more delicate, involving $\left((S-A)^+\right)^2$ rather than only $\mathbb{E}[S]$, $\mathbb{E}[A]$, $Var(S)$, $Var(A)$.

It is easy this is necessary for a non-trivial lower bound:

If we pick distributions for S and A such that

$$\text{support}(S) < \text{support}(A)$$

then $W \equiv 0$ even though $Var(S), Var(A) > 0$.

Kingman's bounds for GI/GI/m/FCFS

For an m server GI/GI/m/FCFS system with arrival rate λ & $\rho = \left(\frac{\lambda}{m}\right)\mathbb{E}[S]$:

Theorem: $$\mathbb{E}\left[W^{GI/GI/m/FCFS}\right] \leq \frac{m\rho^2 c_S^2 + c_A^2}{2\lambda(1-\rho)}$$

Proof sketch: Consider the following alternate scheduling policy

CYCLE: each server has its own queue, and

$C_1, C_{m+1}, C_{2m+1}, \ldots C_{km+1} \ldots \longrightarrow$ server 1

$C_2, C_{m+2}, C_{2m+2}, \ldots, C_{km+2}, \ldots \longrightarrow$ server 2

$\vdots$

$C_j, C_{m+j}, C_{2m+j}, \ldots, C_{km+j}, \ldots \longrightarrow$ server j $\quad (j=1,2,3\ldots m)$

that is, under CYCLE customers are assigned to the m servers in a cyclic (or round-robin) fashion.

Lemma: $\mathbb{E}\left[W^{GI/GI/1/FCFS}\right] \leq \mathbb{E}\left[W^{GI/GI/1/CYCLE}\right]$

(We'll sketch the proof in a bit, you will prove this formally in HW)

To find $\mathbb{E}\left[W^{GI/GI/1/CYCLE}\right]$ we can focus on any single server, say server 1.

→ service times are still i.i.d. $\sim S$

→ interarrival times are m-fold convolution of A

$$\Rightarrow \quad A' \stackrel{d}{=} \sum_{i=1}^{m} A_i$$

$$\Rightarrow \quad \mathbb{E}[A'] = m/\lambda \quad , \quad C^2_{A'} = \frac{C^2_A}{m}$$

plugging into Kingman's upperbound for GI/GI/1/FCFS

$$\mathbb{E}\left[W^{GI/GI/m/FCFS}\right] \leq \mathbb{E}\left[W^{GI/GI/m/CYCLE}\right]$$

$$\leq \frac{\rho^2 C_S^2 + C^2_{A'}}{2\lambda'(1-\rho)}$$

$$= \frac{\rho^2 C_S^2 + C^2_A/m}{2\lambda/m\,(1-\rho)}$$

$$= \frac{m\rho^2 C_S^2 + C^2_A}{2\lambda(1-\rho)}$$

□

Proof sketch of Lemma:

We will couple sample paths under FCFS and CYCLE by "shuffling" the service requirements.

Let: Z_n = arrival time of n^{th} job ($n=1,2,\ldots$)

Y_n = time instant at which n^{th} service initiation takes place (ie. across all servers)

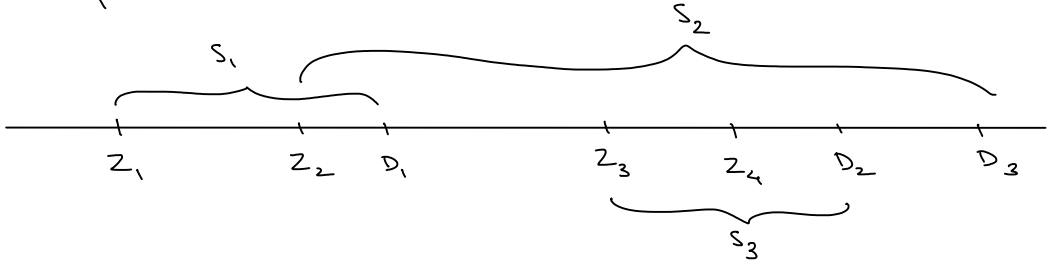
($Y_1 \le Y_2 \le Y_3 \le \ldots$)

S_n = service time associated with the n^{th} customer to enter service

(Note: for FCFS, S_n = service time of n^{th} arrival but not necessarily under CYCLE. Why?)

D_n = time of n^{th} departure from the system.

Here is a picture:



(Note: n^{th} departure need not be n^{th} job to enter service)

Claim 1: D_n is the n^{th} largest of (n^{th} order statistic)

$$\{ Y_i + S_i \}_{1 \le i \le n+m-1}$$

(ie., under both FCFS and CYCLE)

Claim 2: $$Y_n^{FCFS} = \max(Z_n, D_{n-m}^{FCFS})$$

$$Y_n^{CYCLE} \ge \max(Z_n, D_{n-m}^{CYCLE})$$

Claim 3: $\forall n \ge 1$: $Y_n^{CYCLE} \ge Y_n^{FCFS}$

(Hint: induction)

□

GI/GI/1 queue - II : Martingale bounds, Heavy traffic analysis

Last week : Kingman's bound for $\mathbb{E}[W^{GI/GI/1}]$

$$\mathbb{E}[W^{GI/GI/1}] \leq \frac{\rho^2 C_s^2 + C_A^2}{2\lambda(1-\rho)}$$

Usually we care about more fine grained metrics e.g. $\Pr[W > t]$

(1) can use Markov or Chebychev using $\mathbb{E}[W]$ or $\mathbb{E}[W^2]$ bounds

(2) can approximate G by Cox

(3) numerically solve $W \stackrel{d}{=} (W + S - A)^+$

(4) Asymptotics as approximations : TODAY

We will look at two asymptotics today

(i) $\Pr[W > t]$ for t large : Tail asymptotics

(ii) W distribution when $\rho \to 1$: Heavy traffic asymptotics

---

Recall our setup from last week :

(*) iid service times $\{S_i\}$ $i = 0, 1, 2 \ldots$

(*) iid. interarrival times $\{A_i\}$ $i = 1, 2, \ldots$ (C_0 arrives at $t=0$)

(*) Netput process : $X_i = S_{i-1} - A_i$

(*) W_n = waiting time of n^{th} customer

$W_0 = 0$ (for convenience)

Lindley Recursion: $W_n = (W_{n-1} + S_{n-1} - A_n)^+ = (W_{n-1} + X_n)^+$

(*) $Z_n \equiv$ unreflected random walk for $\{W_n\}$

$Z_0 = 0$; $Z_n = X_1 + X_2 + \cdots + X_n$

Unfolding Lindley recursion

$$W_n = \max\left(W_{n-1} + X_n, 0 \right)$$

$$= \max\left(\max(W_{n-2} + X_{n-1}, 0) + X_n, 0 \right)$$

$$= \max\left(W_{n-2} + X_{n-1} + X_n, X_n, 0 \right)$$

$$\vdots$$

$$= \max\left(X_1 + X_2 + \cdots + X_n, X_2 + \cdots + X_n, \cdots, X_{n-1} + X_n, X_n, 0 \right) \pm Z_n$$

$$= Z_n - \min\left(0, X_1, X_1 + X_2, X_1 + X_2 + X_3, \cdots, X_1 + X_2 + \cdots + X_n \right)$$

$$\boxed{W_n = Z_n - \min_{0 \le k \le n} Z_k}$$

Since, $X_1 \stackrel{d}{=} X_n$, $X_2 \stackrel{d}{=} X_{n-1}$, $\cdots$, we also get:

$$\boxed{W_n \stackrel{d}{=} \max_{0 \le k \le n} Z_k}$$

Kingman's martingale bounds

While a little weak, martingale bounds demonstrate the utility of this tool from probability theory to queueing systems

Let $\phi(\theta) = \mathbb{E}[e^{\theta X}]$ be the moment generating function of the iid $\{X_i\}$

If $\phi(\theta) < \infty$ for some $\theta > 0$; then define

$$\theta^* = \sup\{\theta : \phi(\theta) < 1\}$$

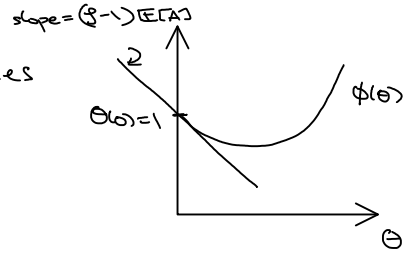
Q: When does θ^* exist?

A: (1) If $\Pr[S > t] \le a e^{-\beta t}$ for some β then

$$\phi(\theta) < \infty \quad \forall \quad \theta < \beta.$$

(2) $$\left.\frac{d\phi(\theta)}{d\theta}\right|_{\theta=0} = \mathbb{E}[X] = \mathbb{E}[A](\beta - 1) < 0$$

which together with $\phi(0) = 1$ implies $\theta^* > 0$.



Thm: In a GI|GI|1|FCFS system, for the steady-state waiting time W,

$$\forall t: \quad P(W \geq t) \leq e^{-\theta^* t}$$

We begin with some basic definitions and inequality:

Defn: A (usually dependent) sequence M_n, $1 \leq n \leq N$ is called a supermartingale if $\mathbb{E}[|M_n|] < \infty$

$$\mathbb{E}[M_n | M_{n-1}, M_{n-2}, \ldots] \leq M_{n-1}$$

Example: $R_i \sim$ i.i.d. nonnegative with $\mathbb{E}[R_i] < 1$

Then: $M_n = \prod_{i=1}^{n} R_i$ is a supermartingale

Pf: $\mathbb{E}[M_n | M_{n-1}, M_{n-2}, \ldots] = \mathbb{E}[R_n M_{n-1} | M_{n-1}, \ldots, M_1]$

$$= \underbrace{\mathbb{E}[R_n]}_{\leq 1} \underbrace{M_{n-1}}_{> 0}$$

$$\leq M_n$$

□

(Sub / Super) Martingales are extremely powerful when combined with "maximal inequalities"

Theorem: Let X_n $(1 \leq n \leq N)$ be a submatingale. Then for $x > 0$

(1) $x \, P_r\left(\max_{1 \leq n \leq N} X_n \geq x\right) \leq \mathbb{E}[X_N^+]$

(2) $x \, P_r\left(\min_{1 \leq n \leq N} X_n \leq -x\right) \leq \mathbb{E}[X_N^+] - \mathbb{E}[X_1]$

Corollary: Let X_n $(1 \leq n \leq N)$ be a non-negative supermartingale. Then for $x > 0$

(3) $x \, P_r\left(\max_{1 \leq n \leq N} X_n \geq x\right) \leq \mathbb{E}[X_1]$

Proof of maximal inequalities (Source: "Prob. and its Applications" -Kallenberg)

(1) X_n, $1 \le n \le N$ is a submartingale

Define the following events

$$F = \{ \max_{1 \le n \le N} X_n \ge x \}$$

$$F_1 = \{ X_1 \ge x \}$$

$$F_k = \{ X_1 < x \} \cap \{ X_2 < x \} \cap \dots \cap \{ X_{k-1} < x \} \cap \{ X_k \ge x \}$$

Observe: F_n $n = 1, 2, \dots, N$ are disjoint events

and $F = F_1 \cup F_2 \cup F_3 \dots \cup F_N$

Since: $x \Pr(X_n \ge x) \le \mathbb{E}[X_n \mathbb{1}_{F_n}]$

we have $x \Pr(F_n) \le \mathbb{E}[X_n \mathbb{1}_{F_n}]$

$\le \mathbb{E}[X_N \mathbb{1}_{F_n}]$

Summing for $n = 1, \dots, N$:

$$x \Pr(F) = x \sum_{n=1}^{N} \Pr(F_n) \le \mathbb{E}[X_N \mathbb{1}_F] \le \mathbb{E}[X_N^+]$$

(2) Since $\{X_n\}$ is a submartingale; it can be decomposed as

$$X_n = M_n + A_n$$ "Doob decomposition"

where M_n is a martingale & A_n is a.s. non-decreasing, $A_1 = 0$

define $A_n = \sum_{k \le n} \mathbb{E}[X_k - X_{k-1} \mid \{X_i\}_{i \le k-1}]$

and $M_n = X_n - A_n$

Now: $x \Pr(\min_{1 \le n \le N} X_n \le -x) \le x \Pr(\min_{1 \le n \le N} M_n \le -x)$

$= x \Pr(\max_{1 \le n \le N} -M_n \ge x)$

$\le \mathbb{E}[M_N^-]$

$= \mathbb{E}[M_N^+] - \mathbb{E}[M_N]$

$\le \mathbb{E}[X_N^+] - \mathbb{E}[X_1]$

(applying part 1 to $-M_n$ which is a martingale $\Rightarrow$ also a submartingale

($\mathbb{E}[M_N] = \mathbb{E}[M_1]$; $M_1 = X_1$)

(3) X_n, $1 \le n \le N$ a non-negative supermartingale

Can perform Doob decomposition of X_n as

$$X_n = M_n - A_n$$

where M_n is a martingale and A_n is a.s. non-decreasing with $A_1 = 0$

Now:

$$\begin{aligned} x \Pr\left(\max_{1\le n\le N} X_n \ge x\right) &\le x \Pr\left(\max_{1\le n\le N} M_n \ge x\right) \\ &\le \mathbb{E}[M_N^+] \\ &= \mathbb{E}[M_N] + \mathbb{E}[M_N^-] \\ &= \underbrace{\mathbb{E}[M_1]}_{M_1 = X_1} + \mathbb{E}\Big[\underbrace{(X_N + A_N)^-}_{X_N \ge 0,\ A_N \ge 0}\Big] \\ &= \mathbb{E}[X_1] \end{aligned}$$

□

We are now ready to prove the tail bound on $W^{G/G/1}$

Proof of main theorem:

Since $W_n \stackrel{d}{=} \max_{0\le k\le n} Z_k$

$$\Pr(W_n \ge t) = \Pr\left(\max_{0\le k\le n} Z_k \ge t\right)$$

$$= \Pr\left(\max_{0\le k\le n} e^{\theta Z_k} \ge e^{\theta t}\right) \qquad \text{for } \theta > 0$$

for $0 < \theta < \theta^*$: $\mathbb{E}\left[e^{\theta X_i}\right] < 1 \Rightarrow \left\{e^{\theta Z_k}\right\}_{k=0}^{n}$ is a supermartingale

Applying Doob's maximal inequality

$$\begin{aligned} \Pr(W_n \ge t) = \Pr\left(\max_{0\le k\le n} Z_k \ge t\right) &\le \Pr\left(\max_{0\le k\le n} e^{\theta Z_k} \ge e^{\theta t}\right) \\ &\le e^{-\theta t}\, \mathbb{E}\left[e^{\theta Z_0}\right] \\ &= e^{-\theta t} \end{aligned}$$

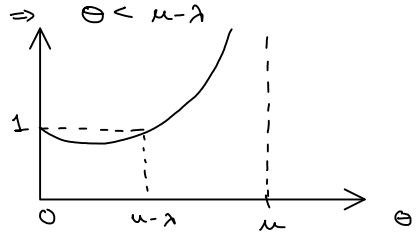
□

Example: consider M/M/1 : $\mathbb{E}[e^{\theta x}] = \mathbb{E}[e^{\theta s}]\,\mathbb{E}[e^{(-\theta)A}]$

$$= \left(\frac{\mu}{\mu-\theta}\right)\left(\frac{\lambda}{\lambda+\theta}\right)$$

$\frac{\mu}{\mu-\theta}\cdot\frac{\lambda}{\lambda+\theta} < 1 \Rightarrow \mu\lambda < \mu\lambda - \theta^2 + \theta(\mu-\lambda) \Rightarrow \theta < \mu-\lambda$

$(0 \le \theta < \mu)$



therefore $\theta^* = \mu-\lambda \Rightarrow$ Martingale bound $P_r(w_n > t) \le e^{-(\mu-\lambda)t}$

is asymptotically tight as $n, t \to \infty$

□

Heavy Traffic Theory

"Heavy traffic" is an overloaded term; means different things to different people. We are discussing today what I like to call "conventional" heavy traffic

This was the first heavy traffic regime investigated by researchers

Q: What does it mean?

A: We create a sequence of queuing system, indexed by n

n^{th} system: service distribution $S^{(n)}$; $\mathbb{E}[S^{(n)}] = 1/\mu^{(n)}$

interarrival distribution $A^{(n)}$; $\mathbb{E}[A^{(n)}] = 1/\lambda^{(n)}$

offered load : $\rho^{(n)} = \lambda^{(n)}\,\mathbb{E}[S^{(n)}]$

waiting time : $w^{(n)}$

Typically: $S^{(n)} = S$, but we scale the interarrival time distribution so that

(*) single server : $\rho^{(n)} \nearrow 1$

(*) m-server system: $\rho^{(n)} \nearrow m$

(*) ∞-server : $\rho^{(n)} \to \infty$

(*) queuing network : bottleneck station $\rho^{(n)} \to 1$

} as $n \to \infty$

Finally, we look at the asymptotic behavior of

single server, m-server } $W^{(n)}$, $W_q^{(n)}$

∞-server } $N^{(n)}$

queueing network } $\mathbf{X}^{(n)}$: the distrib. of # jobs at each station

For GI/GI/1, $W^{(n)} \to \infty$ as $\rho^{(n)} \to 1$ so this seems meaningless, but we expect

$$W^{(n)} = \Theta\left(\frac{1}{1-\rho^{(n)}}\right)$$

so hopefully

$$(1-\rho^{(n)})\, W^{(n)} \to \text{to some degenerate random variable. } W^*$$

This is exactly what will happen, and W^* gives us information about how the stochastic primitives (S, A) influence performance.

Q: So why look at heavy-traffic?

A: (1) we usually operate system under high-utilization where behavior is likely to resemble heavy traffic

(2) Exact analysis usually impossible; asymptotic limits usually "wash away" a lot of complicating details

(3) Much easier to go to a limiting regime, find the "asymptotically optimal" control policy, then argue it does well in the "real system".

↳ can be tricky; intuition goes wrong sometimes.

Examples:

We can already get non-trivial limits for two queueing systems:

(1) M/GI/1 with $S^{(n)} = S$; $\rho^{(n)} \nearrow 1/\mathbb{E}[S]$

$$(1-\rho^{(n)})\, W^{(n)} \to \text{Exp with mean } \mathbb{E}[S_e]$$

Intuition: we can write $W^{M|GI|1} = \sum_{i=1}^{N} Se_i$

where $N \sim Geom(1-\beta) - 1$

so $(1-\rho^{(n)}) N^{(n)} \to Exp(1)$

+ SLLN $\Rightarrow$ $(1-\rho^{(n)}) W^{(n)} \to Exp\left(1/\mathbb{E}[Se]\right) \stackrel{d}{=} W^*$

(2) $GI/M/1$ with $S^{(n)} = S \sim Exp(\mu)$, $A^{(n)} \stackrel{d}{=} \alpha^{(n)} \cdot A$ s.t. $\rho^{(n)} \uparrow 1$

we proved: $N_a \sim Geom(1-\beta) - 1$

$(\text{for } W^*)$ as $\rho^{(n)} \uparrow 1 \Rightarrow$ $(1-\beta^{(n)}) \sim (1-\rho^{(n)}) / \left((1+C_A^2)/2\right)$

$\Rightarrow$ $(1-\rho^{(n)}) N_a^{(n)} \to Exp\left(\frac{2}{1+C_A^2}\right)$

$\Rightarrow$ $(1-\rho^{(n)}) W^{(n)} \to Exp\left(\frac{2}{(1+C_A^2)\mu}\right) \stackrel{d}{=} W^*$

If we rewrite and stare at these:

$M|GI|1$: $(1-\rho^{(n)}) W^{(n)} \to W^* \sim$ Exp with mean $\left(\frac{C_S^2+1}{2}\right)\mathbb{E}[S]$

$GI/M/1$: $(1-\rho^{(n)}) W^{(n)} \to W^* \sim$ Exp with mean $\left(\frac{1+C_A^2}{2}\right)\mathbb{E}[S]$

we would be tempted to guess:

$GI|GI|1$: $(1-\rho^{(n)}) W^{(n)} \to W^* \sim$ Exp with mean $\left(\frac{C_S^2+C_A^2}{2}\right)\mathbb{E}[S]$

and we would be correct! (Note similarity with Kingman's bound)

Rest of the lecture we will look at the formal machinery for arriving at this result:

(1) Brownian motion / Wiener process

(2) Reflection map / Skorokhod map

(3) Continuous mapping theorem

Recall : $W_n = Z_n - \min_{0 \le k \le n} Z_k$

where $Z_n = X_1 + X_2 + \cdots + X_n$; $X_k = S_{k-1} - \Delta_k$

That is Z_n is a sum of iid random variables

⇒ this should remind you of Central limit theorem

ie. $\frac{Z_n - n\,\mathbb{E}[X]}{\sqrt{n}} \to N(0, \sigma^2)$

but this alone is not enough because we need information about the entire sequence $\{Z_k\}_{k=0}^{n}$ to talk about W_n.

Which leads us to the all important stochastic process : Brownian Motion.

Definition : (Standard) Brownian Motion

(Standard) Brownian Motion, $B(t)$, is a continuous time, continuous space stochastic process with

(1) $B(0) = 0$

(2) Stationary increments : $B(t+s) - B(t)$ depends on s but not t.

(3) Independent increments : for $t_1 < t_2 \le t_3 < t_4$

$B(t_2) - B(t_1) \perp B(t_4) - B(t_3)$

(4) Normally distributed increments

$B(t+s) - B(t) \sim N(0, s)$

↳ normal distribution with mean 0, variance s.

Note: property (4) implies that the sample paths of B.M. are continuous with probability 1.

In fact : (1)(2)(3)(4) ⇔ (1)(2)(3) ; continuous sample paths

That is : stationary independent increments + continuous sample paths implies that increments must have Normal distribution.

What are "non-standard" Brownian motions?

Definition: Brownian Motion with drift α and variance σ^2 are defined as

$$B_{(\alpha,\sigma^2)}(t) = \sigma B(t) + \alpha t$$

where $B(t)$ is a standard Brownian Motion.

That is: $B_{(\alpha,\sigma^2)}$ has stationary independent increments with

$$B_{(\alpha,\sigma^2)}(t+s) - B_{(\alpha,\sigma^2)}(t) \sim N(\alpha s, \sigma^2 s)$$

Properties of $B(t)$:

(1) Markov property: $B(t+s) = \big(B(t+s) - B(s)\big) + B(s)$

where $B(s)$ & $B(t+s) - B(s)$ are independent normal random variables

(2) Differential property: For any $s > 0$,

$$\{B_s(t) = B(t+s) - B(s) : t > 0\}$$

is also a Brownian motion, independent of $\{B(u) : u \le s\}$

(3) Scaling property: For any $c > 0$

$$\{\sqrt{c}\, B\left(\tfrac{t}{c}\right) : t \ge 0\} \overset{d}{=} \{B(t) : t \ge 0\}$$

(4) Symmetry: $-B(t) \overset{d}{=} B(t)$

(5) The Reflection Principle: for $a \in \mathbb{R}$ let

$$T_a = \inf\{t : B(t) = a\}$$

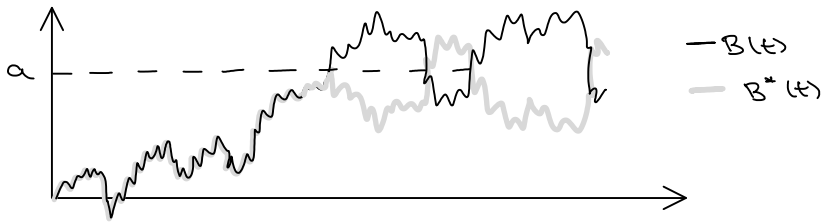
→ first hitting time of level a

and define

$$B^*(t) = \begin{cases} B(t) & t \le T_a \\ 2a - B(t) & t > T_a \end{cases}$$

then $B^*(t) \overset{d}{=} B(t)$

Pictorially:



(6) $B(t)$ is nowhere differentiable almost surely (!)

intuition: $|B(t+s) - B(t)| = |Y|$ where $Y \sim N(0, s)$

$= \sqrt{s}\left(\sqrt{\frac{2}{\pi}}\right)$

$\Rightarrow \lim_{s \to 0} \frac{|B(t+s) - B(t)|}{s} = \Theta\left(\frac{1}{\sqrt{s}}\right) \to \infty$

Also note that while $\mathbb{E}[B(t)] = 0$; $\mathbb{E}[|B(t)|] = \sqrt{\frac{2t}{\pi}}$

In fact:

Theorem (Law of the iterated Logarithm ; Khinchin)

For the standard Brownian motion, $B(t)$,

$$\limsup_{t \to \infty} \frac{B(t)}{\sqrt{2t \log\log t}} = 1 \quad \text{almost surely.}$$

(This result (appropriately scaled) also holds for sums of iid random vars.)

Donsker's Theorem (or) Functional Central Limit Theorem

CLT talks about centered sums of iid. random variables. for the "process level" analogue we start with a random walk:

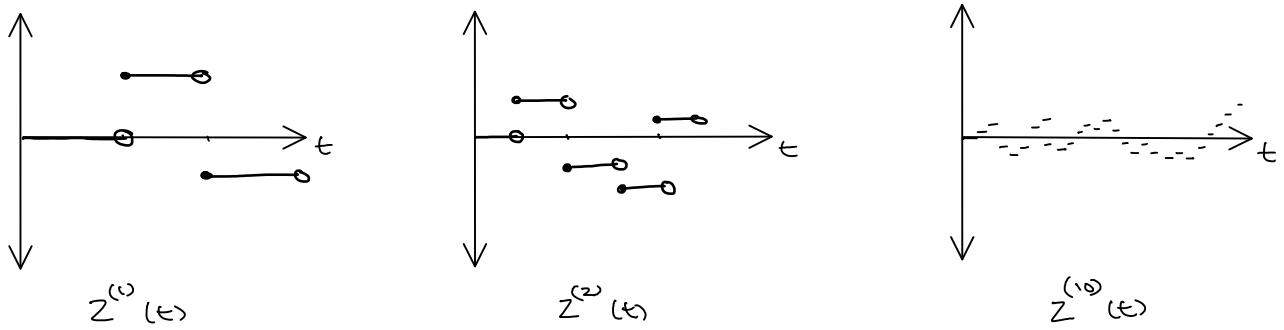
$$Z_0 = 0 \quad ; \quad Z_{n+1} = Z_n + X_{n+1}$$

where $X_i \sim$ iid with $\mathbb{E}[X] = \alpha$, $\text{var}(X) = \sigma^2$

and define the continuous time process:

$$\hat{Z}^{(n)}(t) = \frac{Z_{\lfloor nt \rfloor} - \lfloor nt \rfloor \alpha}{\sigma \sqrt{n}}$$

Q: What do sample paths of $\hat{Z}^{(n)}(t)$ look like?



$Z^{(1)}(t)$ $Z^{(2)}(t)$ $Z^{(10)}(t)$

That is as n grows; we add more X_i's for each t; and to preserve the variance we divide by $\sqrt{n}$.

Q: What happens to $Z^{(n)}(1)$?

A: $Z^{(n)}(1) = \dfrac{Z_n - n\alpha}{\sigma\sqrt{n}} \to N(0,1)$ as $n \to \infty$

But now we are in a position to say more:

Theorem (Donsker's / Functional Central Limit Theorem):

Given an iid sequence $\{X_n : n \geq 1\}$ with mean α and variance σ^2, define $\{Z_n : n \geq 0\}$ where $Z_0 = 0$ and $Z_{n+1} = Z_n + X_{n+1}$

Let

$$\hat{Z}^{(n)}(t) = \frac{Z_{\lfloor nt \rfloor} - \lfloor nt \rfloor \alpha}{\sqrt{n}} \quad ; t \geq 0$$

then $\hat{Z}^{(n)} \xrightarrow{d} \sigma B$ as $n \to \infty$ (in space $\mathcal{D}$)

where $B(t)$ is the standard Brownian and $\xrightarrow{d}$ denotes convergence in distribution.

Remarks:

(1) if we plug in any fixed t, then FCLT gives us CLT. But FCLT says much more.

"The paths of scaled random walk (when viewed as functions) converge in distribution to the paths of the Brownian motion.

(2) What does convergence in function space mean?

It suffices for our purposes to say that for all finite collection of times $\{t_1, t_2, \ldots, t_k\}$

$$\{Z^{(n)}(t_1), Z^{(n)}(t_2), \ldots, Z^{(n)}(t_k)\} \xrightarrow{d} \{B(t_1), B(t_2), \ldots, B(t_k)\}$$

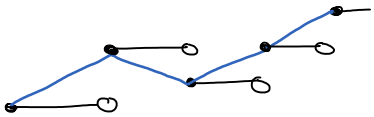
(3) What is the space $\mathcal{D}$?

$\mathcal{D}$ denotes the space of RCLL functions (also called "cadlag")

right continuous with left limits

note that $Z^{(n)}(t)$ does not have continuous sample paths.

This is a minor technical nuance which we will not bother with since we could have instead looked at the interpolated version:



with some cumbersome notation.

The Reflection Mapping (or) Skorokhod Map

We have one ingredient of our analysis already:

$$\hat{Z}^{(n)}(t) \doteq \frac{Z_{\lfloor nt \rfloor} - \lfloor nt \rfloor \mathbb{E}[X]}{\sqrt{n}} \to \sigma B(t)$$

where $\sigma^2 = \text{var}(X) = \text{var}(S-A) = \text{var}(S) + \text{var}(A)$

But we want to analyze W which we expressed as

$$W_n = Z_n - \min_{0 \le k \le n} Z_k$$

This map from $\{Z_n\}$ to $\{W_n\}$ is called the <u>reflection map.</u>

Lemma (Reflection mapping):

The reflection map $\phi : \mathbb{D} \to \mathbb{D}$ given by

$$\phi(x)(t) \doteq x(t) - \left(\inf_{0 \le s \le t} x(s) \right)^+$$

is continuous in the Skorohod J_1 topology

Remarks

(1) Recall $\mathbb{D}$ is the space of RCLL (right continuous w/ left limits) fns.

so given $x : [0,T] \to \mathbb{R}$, $x \in \mathbb{D}$

if $y = \phi(x)$;

then $y : [0,T] \to \mathbb{R}$, $y \in \mathbb{D}$

and $$y(t) \doteq (\phi(x))(t) = x(t) - \left(\inf_{0 \le s \le t} x(s) \right)^+$$

(2) Skorohod topology

This is a technical issue.

Suppose we were were working with $C[0,T]$ ← space of continuous fns. from $[0,T]$ to $\mathbb{R}$

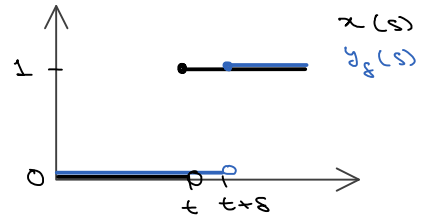
then we could use the uniform metric:

for $x, y \in C[0,T]$: $\|x - y\| = \sup_{0 \le t \le T} |x(t) - y(t)|$

But for processes with discontinuities this is too strong.

e.g.

$$x(s) = \begin{cases} 0 & s < t \\ 1 & s \ge t \end{cases} \qquad y_\delta(s) = \begin{cases} 0 & s < t+\delta \\ 1 & s \ge t+\delta \end{cases}$$



Then we want to say y_δ is close to x as $\delta \to 0$

but $\|x - y_\delta\| = 1 \quad \forall \delta > 0$

The Skorohod metric allows us to argue about such sample paths.

Definition (Skorokhod metric):

Let Λ be the space of strictly increasing continous functions $\lambda : [0,T] \to [0,T]$. The Skorokhod metric on $\mathcal{D}[0,T]$ is defined by:

$$\rho_s(x,y) = \inf_{\lambda \in \Lambda} \left(\|\lambda - I\| \vee \|x - y\lambda\| \right)$$

That is: $\rho_s(x,y) < \varepsilon$ if $\exists\ \lambda \in \Lambda$ s.t.

$$\sup_{0 \le t \le T} |\lambda(t) - t| < \varepsilon$$

and $$\sup_{0 \le t \le T} |x(t) - y(\lambda(t))| < \varepsilon.$$

So the reflection mapping lemma says

if $\{x_n\}$ is a sequence of functions in $\mathcal{D}$

and $x_n \to y \in \mathcal{D}$ in the Skorokhod metric

then $\phi(x_n) \to \phi(y)$ also in the Skorokhod metric.

Using the reflection map, we can define what is called the Reflected Brownian Motion:

Definition (Reflected Brownian Motion):

The reflected Brownian Motion with drift α and variance σ^2, denoted $R_{(\alpha,\sigma^2)}(t)$ is defined as

$$R_{(\alpha,\sigma^2)}(t) = \phi\left(B_{(\alpha,\sigma^2)}\right)(t)$$

where $B_{(\alpha,\sigma^2)}$ is a B.M. with drift α & variance σ^2

Q: Is $R_{(\alpha,\sigma^2)}$ a Markov process?

A: Yes

Q: Is $R_{(\alpha,\sigma^2)}$ a Brownian motion?

A: No. increments are not Normally distributed. ($R_{(\alpha,\sigma^2)}(t) \geq 0$).

Q: Are the increments of $R_{(\alpha,\sigma^2)}$ independent?

A: No. Increments in $[t_1, t_2]$ tells us something about $R_{(\alpha,\sigma^2)}(t_2)$ which influences future increments.

Remarks:

(1) For $\alpha = 0$, ie B.M. without drift

$$R_{(0,\sigma^2)} \overset{d}{=} |B_{(0,\sigma^2)}|$$

but NOT otherwise.

(2) Because of above, the name "reflected Brownian Motion" is a misnomer. A more apt name is "regulated Brownian Motion" and to play safe we will just use RBM.

(3) The map $\psi : \mathcal{D} \to \mathcal{D}$ defined as

$$\psi(x)(t) := -\left(\inf_{0 \leq s \leq t} x(s)\right)^+$$

is called the "regulator process" which prevents $\phi(x)$ from becoming negative.

For example: recall Lindley recursion

$$W_n = (W_{n-1} + X_n)^+ \quad \& \text{ denote } \quad Y_n = (W_{n-1} + X_n)^-$$

then
$$\begin{aligned} W_n &= W_{n-1} + X_n + Y_n \\ &= W_{n-2} + (X_{n-1} + X_n) + (Y_{n-1} + Y_n) \\ &\vdots \\ &= \sum_{i=1}^{n} X_i + \sum_{i=1}^{n} Y_i = Z_n + \sum_{i=1}^{n} Y_i \qquad \left(\Rightarrow \sum_{i=1}^{n} Y_i = -\min_{0 \leq k \leq n} Z_k \right) \\ &= \underbrace{\sum_{i=0}^{n-1} S_i}_{\text{total work brought in by arrivals}} - \Big(\underbrace{\sum_{i=1}^{n} A_i}_{\text{total elapsed time until arrival of } C_n} - \underbrace{\sum_{i=1}^{n} Y_i}_{\text{total idle time until arrival of } C_n} \Big) \end{aligned}$$

The Continuous Mapping Theorem

So far we have two ingredients:

(1) $\hat{Z}^{(n)}(t) \doteq \frac{Z_{\lfloor nt \rfloor} - \lfloor nt \rfloor \mathbb{E}[X]}{\sqrt{n}} \xrightarrow{d} \sigma B(t)$

(2) Using reflection map we can write

$$W_n = \phi(Z_n)$$

We now present the third ingredient that brings these two together:

Continuous Mapping Theorem:

Let (S_1, d_1) and (S_2, d_2) be two metric spaces

Let $\{X_n\}$, $\{X\}$ be S_1-valued random variables

Let $f : S_1 \to S_2$ be a function from S_1 to S_2

If: $X_n \xrightarrow{d} X$ in (S_1, d_1)

then: $f(X_n) \xrightarrow{d} f(X)$ in (S_2, d_2)

provided: f is X-almost everywhere continuous

[$Pr(X \in D_f) = 0$ where D_f is the set of discontinuity points of f]

(S_i, d_i) metric space $\Rightarrow$ S_i is the space

d_i is the metric

In words: Convergence in distribution of $X_n \xrightarrow{d} X$ is preserved by continuous functions.

Example: If $X_n \xrightarrow{d} X$ in $\mathbb{R}$

then $X_n^2 \xrightarrow{d} X^2$ in $\mathbb{R}$.

We are now ready to put our ingredients together

(0) $Z_n = X_1 + X_2 + \cdots + X_n$

$W_n = Z_n - \min_{0 \le k \le n} Z_k$

(1) $\hat{Z}^{(n)}(t) = \dfrac{Z_{\lfloor nt \rfloor} - \lfloor nt \rfloor \mathbb{E}[X]}{\sqrt{n}} \xrightarrow{d} \sigma B(t)$

(2) CMT $\Rightarrow \quad \phi(\hat{Z}^{(n)}(t)) \xrightarrow{d} \phi(\sigma B(t)) = R_{(0,\sigma^2)}(t)$

Now, we want to say

$$\frac{W_{\lfloor nt \rfloor}}{\sqrt{n}} \doteq \hat{W}^{(n)}(t) \stackrel{!}{=} \phi(\hat{Z}^{(n)})(t)$$

But the above $\stackrel{!}{=}$ is not correct

However if $\mathbb{E}[X] = 0$ (that is, $\rho = 1$) then

$$\hat{Z}^{(n)}(t) \doteq \frac{Z_{\lfloor nt \rfloor}}{\sqrt{n}} \xrightarrow{d} \sigma B(t)$$

and if we defined

$$\hat{W}^{(n)}(t) \doteq \frac{W_{\lfloor nt \rfloor}}{\sqrt{n}}$$

then we indeed have

$$\hat{W}^{(n)}(t) = \phi(\hat{Z}^{(n)})(t)$$

which when combined with the continuous mapping theorem gives

$$\hat{W}^{(n)}(t) \xrightarrow{d} \phi(B_{(0,\sigma^2)})(t) = R_{(0,\sigma^2)}(t)$$

Summarizing:

Theorem (FCLT for critically loaded GI/GI/1):

For the GI/GI/1 queue with $\rho = 1$ and $0 < \sigma^2 = \mathrm{Var}(S) + \mathrm{Var}(A) < \infty$, let

$$\hat{W}^{(n)}(t) \doteq W_{\lfloor nt \rfloor} / \sqrt{n}$$

Then,

$$\hat{W}^{(n)}(t) \xrightarrow{d} R_{(0,\sigma^2)} \stackrel{d}{=} \sigma |B|$$

RBM with drift = 0 and variance = σ^2.

So what does this say?

Recall $|B(t)| = \sqrt{\frac{2t}{\pi}} \Rightarrow |\sigma B(t)| = \sigma\sqrt{\frac{2t}{\pi}}$

So for the critically loaded GI|GI|1, waiting times grow linearly in $\sqrt{\text{var}(S)+\text{var}(A)}$ and as $\sqrt{t}$ in time.

That is neat but we would like to say something useful about "normal" queues

That is queues with $\rho < 1$.

Q: Why did we have to assume $\mathbb{E}[X] = 0$?

A: Recall $\hat{Z}^{(n)}(t) = \frac{Z_{\lfloor nt \rfloor} - \lfloor nt \rfloor \mathbb{E}[X]}{\sqrt{n}} \to \sigma B(t)$

$$\Rightarrow \frac{Z_{\lfloor nt \rfloor}}{\sqrt{n}} \to \sigma B(t) + \frac{\lfloor nt \rfloor}{\sqrt{n}} \mathbb{E}[X]$$

If $\mathbb{E}[X] = \alpha \neq 0$; then

$$\frac{Z_{\lfloor nt \rfloor}}{\sqrt{n}} \to \sigma B(t) + \frac{\lfloor nt \rfloor}{\sqrt{n}} \mathbb{E}[X]$$

$$\sim \sigma B(t) + t(\sqrt{n}\alpha)$$

So as n grows, $(\sqrt{n}\alpha)t$ "drowns" $\sigma B(t)$ and we get a <u>Functional Strong Law of Large Number (FSLLN)</u> result.

↳ also called <u>fluid scaling / fluid regime.</u>

To be able to say something interesting, we would like

$$\alpha\sqrt{n} \to \Theta \quad \Theta \text{ a constant}$$

i.e. $\alpha^{(n)} \sim \Theta/\sqrt{n}$

$$\Rightarrow \mathbb{E}[X^{(n)}] \sim \frac{\Theta}{\sqrt{n}} \to 0 \quad \text{as} \quad n \to \infty$$

We will do just this next!

Theorem (Donsker's Theorem / FCLT for doubly indexed sequences):

Suppose as $n \to \infty$

(1) $\{ X^{(n)} \} \xrightarrow{d} X$ where $\mathbb{E}[X^{(n)}] = \alpha^{(n)} \to \alpha = \mathbb{E}[X]$

(2) $\text{stddev}(X^{(n)}) = \sigma^{(n)} \to \sigma = \text{stddev}(X)$, $0 < \sigma < \infty$

As before define: $Z_0^{(n)} = 0$, $Z_{j+1}^{(n)} = Z_j^{(n)} + X_{j+1}^{(n)}$ (nth random walk)

Let $c_n > 0$ be a sequence : $c_n \to \infty$ as $n \to \infty$

For each n, define : $\hat{Z}^{(n)}(t) = \dfrac{Z^{(n)}_{\lfloor c_n t \rfloor} - \lfloor c_n t \rfloor \alpha^{(n)}}{\sigma^{(n)} \sqrt{c_n}}$: $t \geq 0$

Then as $n \to \infty$: $\hat{Z}^{(n)}(t) \xrightarrow{d} B(t)$

where $B(t)$ is standard Brownian Motion.

Remark: When $\alpha \neq 0$, the above theorem only adds little value over the vanilla FCLT.

The real power comes when $\alpha = 0$.

The way we will use it is as we motivated above:

$\sigma^{(n)} = \sigma$, $c_n = n$ $\Rightarrow$ $\dfrac{Z^{(n)}_{\lfloor nt \rfloor}}{\sqrt{n}} \to \sigma B(t) + \dfrac{\lfloor nt \rfloor}{\sqrt{n}} \alpha^{(n)}$

as n grows large $\dfrac{\lfloor nt \rfloor}{\sqrt{n}} \alpha^{(n)} \to \sqrt{n}\, t\, \alpha^{(n)}$: so we let $\alpha^{(n)} \sim \dfrac{\theta}{\sqrt{n}}$ as $n \to \infty$ $(\theta < 0)$

$\Rightarrow$ $\dfrac{Z^{(n)}_{\lfloor nt \rfloor}}{\sqrt{n}} \to \sigma B(t) + \theta t \overset{d}{=} B_{(\theta, \sigma^2)}$

Therefore: $\hat{W}^{(n)}(t) \doteq \dfrac{W^{(n)}_{\lfloor nt \rfloor}}{\sqrt{n}} = \phi\left(\dfrac{Z^{(n)}_{\lfloor nt \rfloor}}{\sqrt{n}} \right) \xrightarrow{d} \phi(B_{(\theta,\sigma^2)}) \overset{d}{=} R_{(\theta,\sigma^2)}$

(reflection mapping) (continuous mapping theorem)

Note: $\mathbb{E}[X^{(n)}] = \alpha^{(n)} \sim \frac{\theta}{\sqrt{n}}$ $\quad (\theta < 0)$

is equivalent to saying: $\mathbb{E}[S^{(n)}] - \mathbb{E}[A^{(n)}] \sim \frac{\theta}{\sqrt{n}}$

$$\Rightarrow \mathbb{E}[A^{(n)}]\left(\rho^{(n)} - 1\right) \sim \frac{\theta}{\sqrt{n}}$$

$$\Rightarrow \rho^{(n)} \approx 1 - \beta/\sqrt{n} \qquad \left(\beta = \frac{|\theta|}{\mathbb{E}[A]}\right)$$

Therefore to get the RBM limit

(1) $S^{(n)} \to S$; $A^{(n)} \to A$; $\rho^{(n)} \sim 1 - \beta/\sqrt{n}$

(2) we look at $\hat{W}^{(n)}(t) = \frac{W^{(n)}_{[nt]}}{\sqrt{n}}$

$\Rightarrow$ for the n^{th} $\hat{W}^{(n)}(t)$ = add (nt) many $X^{(n)}$

"scaling time by $1/n$"

"scaling time by $(1-\rho^{(n)})^2$"

$\equiv$ divide by $\sqrt{n}$

"scale space by $1/\sqrt{n}$"

"scale space by $(1-\rho^{(n)})$"

Summarizing:

Theorem (FCLT for stable GI/GI/1):

Suppose as $n \to \infty$

(1) $\{X^{(n)}\} \Rightarrow X$ where $\mathbb{E}[X^{(n)}] = \alpha^{(n)} \sim \theta/\sqrt{n}$

(2) $\text{stddev}(X^{(n)}) = \sigma^{(n)} \to \sigma = \text{stddev}(X)$ $\quad 0 \le \sigma < \infty$

then

$$\{\hat{W}^{(n)}(t)\} \equiv \left\{ \frac{W^{(n)}_{[nt]}}{\sqrt{n}} : t \ge 0 \right\} \Rightarrow R_{(\theta,\sigma^2)}(t)$$

in $\mathbb{D}$ as $n \to \infty$

More on RBM:

The reason for this rather tedious exercise of approximating waiting times by RBM was that

(i) we now only need to worry about $\mathbb{E}[S], \mathbb{E}[A], \text{var}(S), \text{var}(A)$

(ii) RBMs are tractable mathematical objects.

Let us expand on the second point:

(1) We can show that

$$R_{(\alpha,\sigma^2)}(t) \stackrel{d}{=} M_t \doteq \sup_{0 \le s \le t} \{ \sigma B(s) + \alpha s \}$$

(using an argument we used to arrive at $W_n \stackrel{d}{=} \max_{0 \le k \le n} Z_n$).

which gives: $\Pr\left(R_{(\alpha,\sigma^2)}(t) \le x \right) = \Phi\left(\frac{x-\alpha t}{\sigma\sqrt{t}}\right) - e^{-2|\alpha|x/\sigma^2} \Phi\left(\frac{-x-\alpha t}{\sigma\sqrt{t}}\right), \quad x \ge 0$

where Φ is the standard normal cdf.

(2) When the drift term is negative, the stationary distribution of RBM is Exponentially distributed

$$R_{(\alpha,\sigma^2)}(\infty) \sim \text{Exp}\left(\frac{2|\alpha|}{\sigma^2}\right) \qquad (\alpha < 0).$$

Finally:

Theorem (Stationary Heavy Traffic Limit for GI/GI/1)

For a sequence of GI/GI/1 queues with $\rho^{(n)} \uparrow 1$ as $n \to \infty$, and:

(1) $\{ (S^{(n)}, T^{(n)}) \} \stackrel{d}{\to} (S, T)$ with

$$0 < \lim_{n\to\infty} \mathbb{E}[S^{(n)}] = \mathbb{E}[S] = 1/\mu = 1/\lambda = \mathbb{E}[A] = \lim_{n\to\infty} \mathbb{E}[A^{(n)}] < \infty$$

(2) $\text{Stddev}(S^{(n)}) = \sigma_S^{(n)} \to \sigma_S$, $\text{stddev}(A^{(n)}) = \sigma_A^{(n)} \to \sigma_A$, $0 < \sigma^2 = \sigma_S^2 + \sigma_A^2 < \infty$

then:

$$(1-\rho^{(n)}) W^{(n)} \stackrel{d}{\to} \text{Exp}\left(\frac{2}{\mu\sigma^2}\right)$$

Remark: In practice, given a system with $\rho < 1$ and some given $\mathbb{E}[S]$, $Var(S)$, $Var(A)$

you can approximate

$$W \approx Exp\left(\frac{2\mathbb{E}[S]}{Var(S) + Var(A)}\right)$$

Limit Interchange

FCLT only allowed us to argue $\hat{N}^{(n)}(t) \xrightarrow{d} B(t)$

for any interval $[0,T]$ as $n \to \infty$

Therefore we can only argue

$$\lim_{t\to\infty} \lim_{n\to\infty} \hat{W}^{(n)}(t) = \lim_{t\to\infty} R_{(\theta,\sigma^2)}(t) \sim Exp(2\theta/\sigma^2)$$

What we would like to say is

$$\lim_{n\to\infty} \lim_{t\to\infty} \hat{W}^{(n)}(t) = \lim_{n\to\infty} \hat{W}^{(n)}(\infty) = R_{(\theta,\sigma^2)}(\infty)$$

That is: the stationary distributions of the sequences of queues converges to the stationary distribution of the limiting RBM. (otherwise we cannot claim the last Theorem stated)

This is a delicate argument which uses tightness of $\hat{W}^{(n)}(\infty)$ and pre-compactness of the sequence $\{\hat{W}^{(n)}(\infty)\}$ and finally showing

$$\lim_{t\to\infty} \lim_{n\to\infty} Pr(\hat{W}^{(n)}(t) \geq x) = \lim_{n\to\infty} \lim_{t\to\infty} P(\hat{W}^{(n)}(t) \geq x)$$

which we omit.

GI|GI|1|PS queue

Since Processor Sharing is work conserving policy, for GI|GI|1|PS

as $\rho^{(n)} \uparrow 1$: $\quad (1-\rho^{(n)})\,\mathbb{E}[V^{(n)}] \to \left(\frac{c_s^2+c_A^2}{2}\right)\mathbb{E}[S]$

(Recall: V stands for unfinished workload)

If we now guess that (as in M|GI|1) the remaining size of jobs are iid samples of stationary excess

then combining with SLLN

$$(1-\rho^{(n)})\,\mathbb{E}[N^{(n)}] \to \frac{\left(\frac{c_s^2+c_A^2}{2}\right)\mathbb{E}[S]}{\left(\frac{c_s^2+1}{2}\right)\mathbb{E}[S]}$$

$$= \frac{c_s^2+c_A^2}{c_s^2+1}$$

and using Little's Law :

$$(1-\rho^{(n)})\,\mathbb{E}[T^{(n)}] \to \frac{1}{\lambda}\left(\frac{c_s^2+c_A^2}{c_s^2+1}\right) = \tau^*$$

Remarks :

(1) For $c_A^2 = 1$ (e.g. Poisson arrivals), τ^* is insensitive to S

(2) For any c_s^2, τ^* is increasing in c_A^2

$\Rightarrow$ variability in arrival process hurts

(3) Rewriting: $\tau^* = \frac{1}{\lambda}\left(1 + \frac{c_A^2-1}{c_s^2+1}\right)$

(a) if $c_A^2 > 1$, then τ^* is decreasing in c_s^2

(b) if $c_A^2 < 1$, then τ^* is increasing in c_s^2

$\Rightarrow$ <u>variable arrival process favors variable service distrib!</u>

(Our heuristic guess is in fact correct: See "GI|G|1 Processor Sharing Queue in Heavy Traffic" by S. Grishechkin. (1994) Adv. Appl. Probab.)

Multi-server heavy traffic regimes

Heavy-traffic for single server : $\rho^{(n)} = \lambda^{(n)} \mathbb{E}[S^{(n)}] \uparrow 1$

For multiserver systems, it can mean many things

(1) m fixed, $\rho^{(n)} = \lambda^{(n)} \mathbb{E}[S^{(n)}] \uparrow m$
↳ # servers

(2) $\rho^{(n)} = \lambda^{(n)} \mathbb{E}[S] \uparrow \infty$; $m^{(n)} \uparrow \infty$

i.e., the rate at which work comes in becomes large, and m grows with ρ

(so "load per server" $= \frac{\rho^{(n)}}{m^{(n)}}$ may not $\to 1$)

Note :

I will use: $\rho^{(n)} = \lambda^{(n)} \mathbb{E}[S^{(n)}]$

So $\rho^{(n)}$ = # servers (of speed 1) needed to serve the incoming demand

(Some people prefer to define $\rho^{(n)} = \frac{\lambda^{(n)} \mathbb{E}[S^{(n)}]}{m^{(n)}}$ = load per server

but this is less informative, as well as meaningless when trying to optimize $m^{(n)}$)

Conventional Heavy-traffic

Definition : Fix number of servers m

$S^{(n)} \to S$; $\mathbb{E}[S^{(n)}] \to \mathbb{E}[S] = 1/\mu$

$\lambda^{(n)} \uparrow m\mu \iff \rho^{(n)} = \lambda^{(n)} \mathbb{E}[S^{(n)}] \uparrow m$

- First multiserver heavy-traffic to be analyzed (Köllerström 1974)
- "Mathematician's point of view" = gives some insight into effect of S, A on response time;
but operationally not a desirable regime

In conventional heavy-traffic, GI/GI/m behaves like GI/GI/1

Theorem: If

(i) $\text{Var}(S) + \text{Var}(A) > 0$

(ii) $\mathbb{E}[S^3] < \infty$

(iii) $\mathbb{E}[A^{2+\delta}] < \infty$ for some $\delta > 0$

then

$$\left(1 - \frac{\rho^{(n)}}{m}\right) W^{(n)} \to \text{Exp}\left(\frac{2m\mu}{c_S^2 + c_A^2}\right) \quad \text{as: } \rho^{(n)} \uparrow m$$

The expectations of W, T, N, N_Q converge as long as $\mathbb{E}[S^2]$ & $\mathbb{E}[A^2]$ are finite:

Theorem: For GI/GI/m, as $\rho \uparrow m$:

$$\mathbb{E}[W] = \left(\frac{c_S^2 + c_A^2 + o(1)}{2\lambda\left(1 - \frac{\rho}{m}\right)}\right)$$

$$\mathbb{E}[N_Q] = \left(\frac{c_S^2 + c_A^2 + o(1)}{2\left(1 - \frac{\rho}{m}\right)}\right)$$

Many-server heavy traffic regimes

Modern multi-server heavy traffic theory takes a "system designer's" perspective:

Suppose you had an inexhaustible source of speed 1 servers

But it costs C_f to hire a server.

As $\rho^{(n)} \to \infty$, how many servers should you provision to balance cost-performance trade-off?

Q: Should the number of servers $(m^{(n)})$ grow as

(1) $m^{(n)} = (1+\delta)\rho^{(n)}$ for $\delta > 0$

(2) $m^{(n)} = \rho^{(n)} + g(\rho^{(n)})$ for $\frac{g(x)}{x} \to 0$ as $x \to \infty$

i.e. $g(x) = o(x)$

(3) $m^{(n)} = (1-\gamma)\rho^{(n)}$ for $\gamma > 0$

(4) ...

It turns out that depending on your "cost" and "performance" metric, any of these could be the correct choice.

Let us begin by examining $m^{(n)} = (1+\delta)\rho^{(n)}$

Also: for the rest of the lecture we will only look at $M/M/m^{(n)}$ systems.

Case 1: $m^{(n)} = (1+\delta)\rho^{(n)}$ for $\delta > 0$

Example: $\delta = 0.3$: $\rho = 100 \Rightarrow m = 130$ servers

$\rho = 1000 \Rightarrow m = 1300$ servers

Q: What cost-performance trade-off do you expect?

A: cost $= c \cdot m^{(n)} = c(1+\delta)\cdot\rho$

performance $= \mathbb{E}[N_Q]$? $(= \lambda \mathbb{E}[W])$

$P_Q = Pr[W>0]$?

For intuition: consider an $M/M/\infty$ system with load $\rho^{(n)} \to \infty$

$N_\infty^{(n)} \sim$ Poisson with mean $\rho^{(n)}$

$\sim$ sum of $\rho^{(n)}$ Poisson r.v.s. (each with mean 1)

by CLT: $\frac{N_\infty^{(n)} - \rho^{(n)}}{\sqrt{\rho^{(n)}}} \xrightarrow{\rho^{(n)}\to\infty} Z \sim N(0,1)$

So, $\quad Pr\left[N_\infty^{(n)} \geq (1+\delta)\, \rho^{(n)} \right] \sim Pr\left[\rho^{(n)} + \sqrt{\rho^{(n)}}\, Z \geq (1+\delta)\, \rho^{(n)} \right]$

$$= Pr\left[Z \geq \delta \sqrt{\rho^{(n)}} \right]$$

Useful fact: for $Z \sim N(0,1)$, $t \geq 0$

$$\frac{t}{t^2+1} \frac{1}{\sqrt{2\pi}} e^{-t^2/2} \leq Pr[Z \geq t] \leq \frac{1}{t} \cdot \frac{1}{\sqrt{2\pi}} e^{-t^2/2}$$

$$\sim \frac{1}{\sqrt{2\pi}} \quad \frac{1}{\delta \sqrt{\rho^{(n)}}} \quad e^{-\delta^2 \rho^{(n)}/2}$$

So, for $\delta > 0$; $Pr\left[N_\infty^{(n)} \geq (1+\delta)\, \rho^{(n)} \right] \to 0$ exponentially fast in $\rho^{(n)}$.

With overwhelming probability $1 - e^{-\theta(\delta)}$, an arrival finds an idle server

While the constants change, this intuition carries over to $M/M/m^{(n)}$ with $m^{(n)} = (1+\delta)\, \rho^{(n)}$

$$\Rightarrow \quad P_d = Pr[W^{(n)} > 0] \sim e^{-\theta(\rho^{(n)})}$$

and $\quad \mathbb{E}[W \mid W > 0] = \frac{1}{m^{(n)}\mu - \lambda^{(n)}} \sim \frac{1}{\mu \delta \rho^{(n)}}$

Excellent performance; poor cost (overstaffed system)

Hence also called "Quality Driven (QD) Regime"

Case 2: $m^{(n)} = \rho^{(n)} + o(\rho^{(n)})$

Q: What should $m^{(n)}$ be if we want

$$\Pr[W^{(n)} > 0] \to \alpha \quad \text{where } 0 < \alpha < 1 \text{ ?}$$

A: Recall: for $m^{(n)} = \rho^{(n)} + \delta\rho^{(n)}$

$$\Pr[N_\infty^{(n)} \geq m^{(n)}] \sim \frac{1}{\delta\sqrt{\rho^{(n)}}} e^{-\delta^2\rho^{(n)}/2}$$

$\Rightarrow$ if $\delta\sqrt{\rho^{(n)}} \to$ constant, then $\Pr[W^{(n)} > 0] \overset{?}{\to}$ constant

$$\Rightarrow \quad \delta \sim \frac{\text{constant}}{\sqrt{\rho^{(n)}}}$$

$$\Rightarrow \quad m^{(n)} = \rho^{(n)} + \beta\sqrt{\rho^{(n)}}$$

This is indeed true

- this heuristic was used by Erlang
- formalized by Halfin–Whitt

Theorem: (Halfin-Whitt 1981)

For $M/M/m^{(n)}$ with $\rho^{(n)} \to \infty$ as $n \to \infty$

$$\lim_{n\to\infty} P_d^{(n)} = \lim_{n\to\infty} \Pr[N^{(n)} \geq m^{(n)}] = \alpha \qquad (0 < \alpha < 1)$$

iff:

$$\lim_{n\to\infty} \frac{m^{(n)} - \rho^{(n)}}{\sqrt{\rho^{(n)}}} = \beta \qquad (\beta > 0)$$

in which case

$$\alpha = \left[1 + \beta\sqrt{2\pi}\,\Phi(\beta)\,e^{\beta^2/2}\right]^{-1}$$

($\Phi(\cdot)$ is the standard normal c.d.f.)

The regime $m = \rho + \beta\sqrt{\rho}$ is also called

- square-root staffing
- Halfin-Whitt regime
- Quality and Efficiency Driven (QED) regime.

Why QED?

Suppose we want to $\min_{m^{(n)}}$ (total cost) $= c(m^{(n)}) + \mathbb{E}[N_Q^{(n)}]$

$$= c(m^{(n)} - \rho^{(n)}) + P_d^{(n)} \frac{\rho^{(n)}/m^{(n)}}{1 - \rho^{(n)}/m^{(n)}} + c\rho^{(n)}$$

$$= c(m^{(n)} - \rho^{(n)}) + P_d^{(n)} \frac{\rho^{(n)}}{m^{(n)} - \rho^{(n)}} + c\rho^{(n)}$$

at optimality $\left(m^{(n)} - \rho^{(n)}\right) = \Theta\left(P_d^{(n)} \frac{\rho^{(n)}}{m^{(n)} - \rho^{(n)}}\right)$

$$\Rightarrow \quad m^{(n)} - \rho^{(n)} \sim \Theta\left(\sqrt{P_d^{(n)}\, \rho^{(n)}}\right)$$

$\Rightarrow$ square-root staffing

more precisely: $m = \rho + \beta\sqrt{\rho}$

where $\beta^* = \operatorname{argmin} c(m - \rho) + P_d \cdot \frac{\rho}{m - \rho}$

$$= \operatorname*{argmin}_{\beta} c\beta\sqrt{\rho} + \alpha(\beta) \frac{\rho}{\beta\sqrt{\rho}}$$

$$= \operatorname*{argmin}_{\beta} \left(c\beta + \frac{\alpha(\beta)}{\beta}\right)$$

□

Halfin-Whitt as an approximation for M/M/m

Sps: M/M/m with parameters λ, μ

Let $\beta = \dfrac{m - \lambda/\mu}{\sqrt{\lambda/\mu}}$

then approximate

$$P_d = P_r[N \geq m] \approx \alpha(\beta) \doteq \left[1 + \beta\sqrt{2\pi}\, \Phi(\beta)\, e^{\beta^2/2} \right]^{-1}$$

□

We can infact approximate the distribution of N.

Theorem: For $M/M/m^{(n)}$ with

(1) $\quad m^{(n)} \to \infty \quad$ as $\quad n \to \infty$

(2) $\quad \dfrac{m^{(n)} - m^{(n)}}{\sqrt{m^{(n)}}} \to \beta \quad$ as $\quad n \to \infty$

define $\hat{N}^{(n)} = \dfrac{N^{(n)} - m^{(n)}}{\sqrt{m^{(n)}}}$

then:

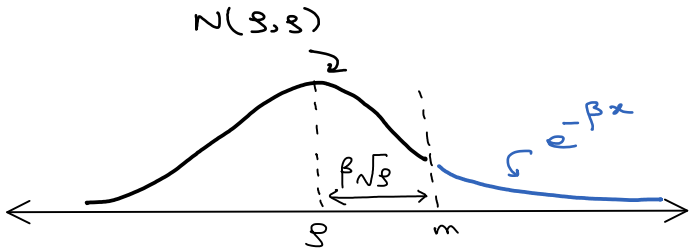
$$\lim_{n \to \infty} \hat{N}^{(n)} \xrightarrow{d} N^*$$

where

$$P_r[N^* \leq -x \mid N^* \leq 0] = \frac{\Phi(\beta - x)}{\Phi(\beta)} \qquad (x \geq 0)$$

$$P_r[N^* \geq x \mid N^* \geq 0] = e^{-\beta x} \qquad (x \geq 0)$$

Pictorially:



Remark: Halfin-Whitt also prove: for $\hat{Y}^{(n)}(t) = \dfrac{N^{(n)}(t) - m^{(n)}}{\sqrt{m^{(n)}}}$

$\hat{Y}^{(n)}(t) \xrightarrow{d} Y^*(t)$: a Brownian motion (diffusion) with variance 2μ

and state-dependant drift $\eta(x) = \begin{cases} -\mu\beta & x \geq 0 \\ -\mu(\beta + x) & x \leq 0 \end{cases}$ □

Case 3 : Efficiency Driven

Broadly anything with $\left(m^{(n)} - s^{(n)}\right)$ growing slower than $\sqrt{s^{(n)}}$ qualifies as Efficiency Driven

$$\Rightarrow \underbrace{\left(m^{(n)} - s^{(n)}\right)}_{\text{staffing cost}} << \underbrace{\mathbb{E}[N_Q]}_{\text{performance penalty}}$$

But a few are popular in literature

(1) $m^{(n)} = (1-\gamma)\, s^{(n)}$

Overloaded regime

- example: systems with abandonment
- example: systems with transient periods of overload

(2) $m^{(n)} = s^{(n)} + \underbrace{k}_{\text{a constant}}$

Non-Degenerate slowdown (NDS)

Recall: for a job of size x, slowdown $= \frac{\mathbb{E}[T(x)]}{x}$

Under QED : $\frac{\mathbb{E}[T(x)]}{x} = \frac{x + \mathbb{E}[W]}{x} = 1 + \frac{P_d \cdot \frac{1}{m\mu - s}}{x}$

$$= 1 + \frac{\alpha}{\mu \beta \sqrt{s}} \to 1$$

Under NDS : $\frac{\mathbb{E}[T(x)]}{x} = \frac{x + \mathbb{E}[W]}{x} = 1 + \frac{P_d \cdot \frac{1}{m\mu - s}}{x}$

$$= 1 + \frac{1}{k \mu x}$$

Example : $\min_{m^{(n)}} : c \cdot m^{(n)}$

subject to: $\mathbb{E}[W^{(n)}] \le 5 \cdot \mathbb{E}[S]$

References :

(1) Heavy-traffic limits for Queues with many Exponential Servers
S. Halfin, W. Whitt
Operations Research (1981)

(2) Dimensioning Large Call Centers
S. Borst, A. Mandelbaum, M. Reiman
Operations Research (2004)

Cμ rule; Conservation laws

"conservation law" : a characterization of system which is independent of the scheduling policy

(eg. conservation of energy in physical systems)

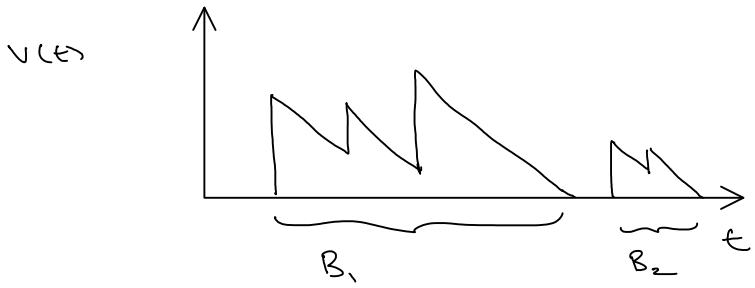
Examples :

(1) Little's law: $\mathbb{E}[N] = \lambda \mathbb{E}[T]$

(2) Utilization law : $\Pr(\text{server busy}) = \lambda \mathbb{E}[S]$

(3) Busy period duration (single server)

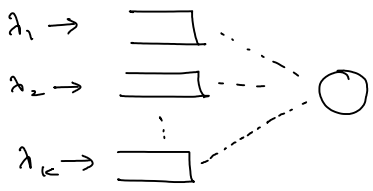
(4) Workload (V) (= virtual waiting time for FCFS)



In some cases, conservation laws allow deducing optimal scheduling policies , even though we can not solve for $\mathbb{E}[N]$ or $\mathbb{E}[T]$

Example : Multiclass G|M|1 priority system

* k - classes of jobs
* for class i : λ_i ; $S_i \sim Exp(\mu_i)$; cost = c_i



Q: What is the optimal pre-emptive (non-anticipatory) scheduling rule to

minimize $\mathbb{E}[C] = \sum_{i=1}^{K} c_i \mathbb{E}[N_i]$

Pre-emptive : can preempt a job currently in service and resume later

non-anticipatory : do not know future arrival stream, or the remaining sizes of jobs in system

Claim: Let the classes be numbered so that

$$c_1\mu_1 \geq c_2\mu_2 \geq \dots \geq c_K\mu_K$$

The optimal policy is to give preemptive priority to the class with largest $c_i\mu_i$ [$c\mu$-rule]

E.g. : $c_1 = c_2 = \dots = c_K$

$c\mu$ rule $\Rightarrow$ prioritize highest μ_i $\Rightarrow$ Shortest expected size

Proof: Denote V_i = total unfinished work due to class i jobs

$$V(t) = \sum_{i=1}^{K} V_i(t)$$

Note that $V(t)$ is conserved under any non-idling policy (also for G/G/1)

$\Rightarrow$ for all policies $\mathbb{E}[V] = \sum_{i=1}^{K} \mathbb{E}[V_i]$ is the same

Since our policies are non-anticipatory

$$\mathbb{E}[V_i(t)] = \frac{1}{\mu_i}\mathbb{E}[N_i(t)]$$

(each class-i job in system has expected residual size $\frac{1}{\mu_i}$)

$$\Rightarrow \mu_i \mathbb{E}[V_i] = \mathbb{E}[N_i]$$

$$\Rightarrow \mathbb{E}[C] = \sum_{i=1}^{K} c_i \mathbb{E}[N_i] = \sum_{i=1}^{K} c_i\mu_i \mathbb{E}[V_i]$$

Next we will employ an interchange argument:

Consider two preemptive priority polices A,B with priorities

A : $f_1 < f_2 < f_3 \dots < i < j < \dots < f_K$

B : $f_1 < f_2 < f_3 \dots < j < i < \dots < f_K$

(f_1: highest prio; f_K: lowest prio)

That is, A and B have the same relative priorities except $i<j$ is switched to $j<i$ for B.

Note: for $\ell \neq i,j$

$$\mathbb{E}[V_\ell^{(A)}] = \mathbb{E}[V_\ell^{(B)}]$$ (not true for non-preemptive)

Therefore

(1) $$\mathbb{E}[V_i^{(A)}] + \mathbb{E}[V_j^{(A)}] = \mathbb{E}[V_i^{(B)}] + \mathbb{E}[V_j^{(B)}]$$

(2) $$\mathbb{E}[C^{(A)}] - \mathbb{E}[C^{(B)}] = c_i\mu_i \left(\mathbb{E}[V_i^{(A)}] - \mathbb{E}[V_i^{(B)}]\right) + c_j\mu_j \left(\mathbb{E}[V_j^{(A)}] - \mathbb{E}[V_j^{(B)}]\right)$$

(1) $\Rightarrow$ $$\mathbb{E}[V_j^{(A)}] - \mathbb{E}[V_j^{(B)}] = -\left(\mathbb{E}[V_i^{(A)}] - \mathbb{E}[V_i^{(B)}]\right)$$

plugging into (2)

$$\mathbb{E}[C^{(A)}] - \mathbb{E}[C^{(B)}] = (c_i\mu_i - c_j\mu_j)\underbrace{\left(\mathbb{E}[V_i^{(A)}] - \mathbb{E}[V_i^{(B)}]\right)}$$

≤ 0 since A gives prio to i & B gives prio to j

$\Rightarrow$ $\mathbb{E}[C^{(A)}] \leq \mathbb{E}[C^{(B)}]$ iff $c_i\mu_i \geq c_j\mu_j$

For an arbitrary policy, we can keep making interchanges on sample path basis until we get to the $c\mu$-rule. □

References: For a general treatment of conservation laws:

(1) Conservation Laws, Extended polymatroids and Multiarmed Bandit Problems: a Polyhedral Approach
D. Bertsimas, J. Niño-Mora
Math of OR (1996)

(2) Fundamentals of Queueing Networks (Chap 11) (Springer)
H. Chen, D. Yao

Stochastic Stability and Theory of Lyapunov Functions

We were able to say a lot about systems we looked at so far

(*) M|M|1 type
(*) Jackson (product-form) Networks
} stationary distrib in closed form

(*) M|G|1, G|M|1 } stationary distrib in transform form (of Embedded DTMC)

(*) G|G|1 } $\rho < 1 \Rightarrow$ stat distrib exists; Kingman's bounds

(*) multiclass G|M|1 } - workload conservation $\Rightarrow$ ($\rho < 1 \Rightarrow$ stationary dist exists); - optimality of $c\mu$

But for general multiclass systems, things get much more complicated.

E.g.: under what conditions is a multiclass queueing network stable?

So far: $\underbrace{\rho < 1}$ $\Leftrightarrow$ "stability"
- for single-server
- for each server in Jackson networks

Defn: A queueing system is stable if $\lim_{t \to \infty} X(t)$ exists.

For general queing networks, $\rho < 1$ is not sufficient to guarantee stability. $\Rightarrow$ need tools to analyze this first order property

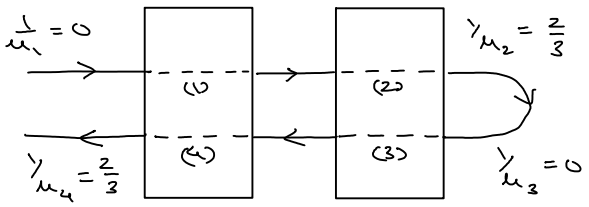
Popular tools: (1) Lyapunov functions
- allow designing stable policies
- allow cost optimization

(2) Fluid methods
(we will skip this topic in this course)

$\rho < 1$ not sufficient for stability

Consider the following multiclass 2-server system with static priority scheduling:

Lu-Kumar network



$$\rho_1 = \rho_2 = 2/3$$

Arrivals: deterministic at $t = 0, 1, 2, \ldots$

Service times: class dependent deterministic with mean $\frac{1}{\mu_i}$

Service: Server 1 prioritizes (4) over (1)
Server 2 prioritizes (2) over (3)

This kind of network is also called a reentrant line since jobs revisit some servers.

Initial state: M jobs of (1) at server 1 at $t = 0^-$

$t = 0^-$: M jobs of (2) at server 2
} (2) serve and become (3) but are stuck at S_2 due to priority

$t = 2M^-$: 3M jobs of (3) at server 2

$t = 2M^-$: 3M jobs of (4) at server 1
} (4) serve and leave system
(1) arrivals are stuck at server 1
no (2) or (3) jobs at server 2

$t = 4M^-$: 2M jobs of (1) at server 1

$\Rightarrow$ $t = 0^-$: M class (1) jobs

$t = 4M^-$: $2M$ class (1) jobs

$t = 12M^-$: $4M$ class (1) jobs

$t = 28M^-$: $8M$ class (1) jobs

$\vdots$

$\Rightarrow$ $N(t)$ diverges as $t \to \infty$ □

Remarks: (1) The deterministic assumption is for ease of exposition and not necessary

more generally: $\exists$ a finite time τ so that

$$N_2(t) = 0 \quad \text{or} \quad N_4(t) = 0 \qquad \forall\, t \geq \tau$$

$\Rightarrow$ server 1 and server 2 never process 'work' at the same time

$\Rightarrow$ $\frac{1}{\mu_2} + \frac{1}{\mu_4} > 1$ $\Rightarrow$ unstable Lu-Kumar network

(2) The reentrant property of routes is also not critical. Rybko-Stolyar provide an example of an unstable system without reentrant classes

(3) Bramson provides an example with a reentrant class which is unstable with FCFS scheduling
(See Section 3.2 of [3])

To summarize : we can not rely on $\rho_j < 1$ as a sufficient stability criterion

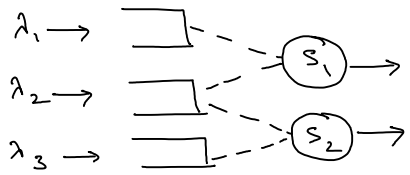
next we will see a sufficient stability condition called the Foster-Lyapunov criterion

---

Lyapunov functions: stochastic stability

We will see the use of Lyapunov fns. through a simple example:

Example: A 3-class 2-server flexible server system in discrete time



Slotted time system

arrivals: $A_i(t)$ arrivals of type i jobs during slot t

$A_i(t)$ are iid with mean λ_i

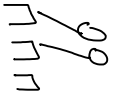
service: S_1 can serve either class 1 or 2 jobs (but not both) in a slot

S_2 can serve either class 2 or 3 jobs (but not both) in a slot

Only one server can serve class 2 jobs per slot

⇒ we can represent slot t scheduling decision as one of the following 3 vectors

$$\alpha(t) \in \left\{ \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \right\}$$

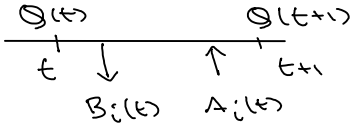
e.g. $\begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}$ ⇒ 

Class i work drains by $B_i(\alpha(t))$; $\mathbb{E}[B_i(\alpha(t))] = b_i(\alpha(t))$

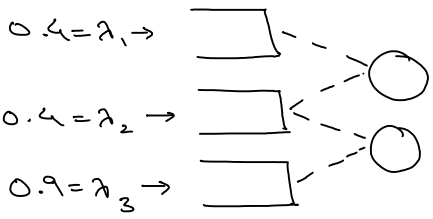
example: $b_i(\alpha(t)) = \alpha_i(t)$

⇒ on average drains by 1 unit if queue i is served

Queueing dynamics: $$Q_i(t+1) = \left(Q_i(t) - B_i(\alpha(t))\right)^+ + A_i(t)$$



Here is a concrete example with numbers:



$$\alpha(t) \in \left\{ \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \right\}$$

$$B(\alpha) = \begin{bmatrix} B_1(\alpha) \\ B_2(\alpha) \\ B_3(\alpha) \end{bmatrix} = \alpha$$

Q: What is a "good" scheduling rule?

↳ should at least stabilize queues

A possible schedule: $\begin{cases} \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} & \text{w. prob } 1/2 \\ \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} & \text{w. prob } 1/2 \end{cases}$

$\Rightarrow$ S_1 randomizes b/w Θ_1 & Θ_2

S_2 dedicated to Θ_3

Intuitively: in the long run $\begin{cases} \Theta_1 \text{ gets capacity } \bar{b}_1 = 0.5 > 0.4 = \lambda_1 \\ \Theta_2 \text{ gets capacity } \bar{b}_2 = 0.5 > 0.4 = \lambda_1 \\ \Theta_3 \text{ gets capacity } \bar{b}_3 = 1 > 0.9 = \lambda_3 \end{cases}$

So we expect queues to remain bounded.

To prove this we will use an appropriately defined Lyapunov function.

Definition: For a discrete-time Markov process $X(t)$ with state space S, a non-negative function $\Phi: S \to \mathbb{R}_+$ is a Lyapunov function if

(1) $\sup_{x \in S} \Phi(x) = \infty$

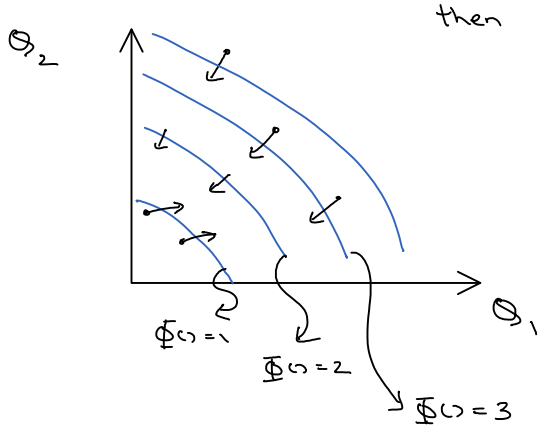
(2) Drift condition: $\exists \gamma > 0$, C_1, C_2

$$\mathbb{E}[\Phi(X(t+1)) \mid X(t) = x] - \Phi(x) \leq -\gamma + C_1 \mathbb{1}_{\{\Phi(x) \leq C_2\}}$$

Intuitively : Φ is a measure of how far you are from a "nice state" (e.g. empty system)

Therefore they are also called potential functions

Drift condition says : If far from nice state (those with $\Phi(x) \leq c_2$) then get closer by γ on average



"level sets" of $\Phi(\cdot)$

and average drift vectors

A suitably designed Lyapunov function helps in proving stability via the following theorem:

Theorem (Foster-Lyapunov criterion)

Let $\Phi : S \to \mathbb{R}_+$ satisfy the conditions mentioned in the defn.

Define hitting times of level sets as :

$$\tau_N(x) = \inf \{ t : \Phi(X(t)) \leq N \mid X(0) = x \}$$

Then, $\forall\ N > c_2$

$$\mathbb{E}[\tau_N(x)] \leq \frac{\Phi(x)}{\gamma} < \infty \quad \forall\ x \in S.$$

In words : for sets $B_N = \{ x : \Phi(x) \leq N \}$

the mean recurrence time is finite

$\Rightarrow$ B_N is a positive recurrent set of states

When S is countable, $|B_N|$ is finite $\Rightarrow$ $\exists$ some state x^* in B_N that is positive recurrent

$\Rightarrow$ all states x^* communicates with are positive recurrent

$\Rightarrow$ there exists a stationary distribution

(For general state spaces we have to worry about Harris recurrent sets, which we won't)

Proof (Foster-Lyapunov criterion):

Consider a fixed N and $x \in S$ with $\Phi(x) > c_2$; $N > c_2$

Define

$$\tau = \inf \{ t : \phi(x(t)) \le N \mid x(0) = x \}$$

and consider the process

$$Z(t) = \Phi(x(t)) + \gamma t - c_1 \sum_{u=0}^{t-1} \mathbb{1}_{\{\Phi(x(u)) \le c_2\}}$$

claim: $Z(t)$ is a supermartingale (i.e. decreasing in expectation)

Proof: (1) $\mathbb{E}[|Z(t)|]$ is finite

This follows if we show $\Phi(x(t))$ is finite

but $\mathbb{E}[\Phi(x(t+1))] \le \mathbb{E}[\Phi(x(t))] - \gamma + c_1 \Pr(\Phi(x(t)) \le c_2)$

$\le \mathbb{E}[\Phi(x(t))] - \gamma + c_1$

$\Rightarrow \mathbb{E}[\Phi(x(t))] \le \Phi(x) - \gamma t + c_1 t < \infty$

(2) $\mathbb{E}[Z(t+1) \mid \{x(0), \dots x(t)\}] \le Z(t)$

Pf:

$$\mathbb{E}[Z(t+1) \mid \{x(0) \dots x(t)\}] = \mathbb{E}[\phi(x(t+1)) \mid \{x(0), \dots x(t)\}] + \gamma + c_1 \mathbb{1}_{\{\phi(x(t)) \le c_2\}} + \gamma(t-1) + c_1 \sum_{u=0}^{t-1} \mathbb{1}_{\{\phi(x(u)) \le c_2\}}$$

$$\le \Phi(x(t)) + \gamma(t-1) + c_1 \sum_{u=0}^{t-1} \mathbb{1}_{\{\phi(x(u)) \le c_2\}}$$

$$= Z(t)$$

□

For a fixed integer M, this implies [via Optional Sampling Theorem

$$\mathbb{E}\left[Z(\tau \wedge M) \right] \leq Z(0) = \Phi(x)$$

but $\mathbb{E}\left[Z(\tau \wedge M) \right] = \Phi(x(\tau \wedge M)) + \gamma \mathbb{E}\left[\tau \wedge M\right]$

(since by assumption $N > c_2$
$\Rightarrow \phi(x(u)) > c_2$ for $u < \tau$
$\Rightarrow \mathbb{1}_{\{\phi(x(u)) \leq c_2\}} = 0$ for $u < \tau$

$\Rightarrow \gamma \mathbb{E}\left[\tau \wedge M\right] \leq \Phi(x)$ (since $\Phi(x(\tau \wedge M)) \geq 0$)

by monotone convergence theorem $\tau \wedge M \to \tau$ as $M \to \infty$

$$\mathbb{E}[\tau] \leq \frac{\Phi(x)}{\gamma}$$

□

Let's go back to our flexible queueing example

Step 1: Find a suitable Lyapunov function

This step is tricky, requiring an educated guess.

A good starting guess is the "Quadratic Lyapunov Function"

$$\Phi(Q(t)) = Q_1^2(t) + Q_2^2(t) + Q_3^2(t)$$

Step 2: Prove drift inequality.

Recall: $Q_i(t+1) = \left(Q_i(t) - B_i(t)\right)^+ + A_i(t)$

$\Rightarrow$

$$\Delta\Phi(t) \doteq \Phi(Q(t+1)) - \Phi(Q(t))$$

$$= \sum_{i=1}^{3} \left(\left(Q_i(t) - B_i(t)\right)^+ + A_i(t) \right)^2 - Q_i(t)^2$$

The following inequality will help us:

For $a, b > 0$: $((q-b)^+ + a)^2 \le q^2 + b^2 + a^2 + 2q(a-b)$

Proof:

$$\begin{aligned}((q-b)^+ + a)^2 &= ((q-b)^+)^2 + a^2 + 2a(q-b)^+ \\ &\le (q-b)^2 + a^2 + 2a(q-b)^+ \\ &\le (q-b)^2 + a^2 + 2aq \\ &= q^2 + b^2 + a^2 + 2q(a-b)\end{aligned}$$

□

Using the above inequality:

$$\Delta \Phi(t) \le \sum_{i=1}^{3} B_i^2(t) + A_i^2(t) + 2Q_i(t)\left(A_i(t) - B_i(t)\right)$$

$$\mathbb{E}\left[\Delta \Phi(t) \mid Q(t)\right] \le \sum_{i=1}^{3} \mathbb{E}\left[B_i^2(t) + A_i^2(t)\right] + 2\sum_{i=1}^{3} Q_i(t)\,\mathbb{E}\left[A_i(t) - B_i(t)\right]$$

since our policy is iid; $\mathbb{E}[A_i(t)] = \lambda_i$, $\mathbb{E}[B_i(t)] = \bar{b}_i$

and assume $\sum_{i=1}^{3} \mathbb{E}\left[A_i^2 + B_i^2\right] \le B$

$$\Rightarrow \mathbb{E}\left[\Delta \Phi(t) \mid Q(t)\right] \le B - 2\sum_{i=1}^{3} Q_i(\bar{b}_i - \lambda_i)$$

for the example: $\bar{b}_i - \lambda_i = 0.1$ for $i = 1, 2, 3$

$$\Rightarrow \mathbb{E}\left[\Delta \Phi(t) \mid Q(t)\right] \le B - 0.2\left(\sum_{i=1}^{3} Q_i(t)\right)$$

$$\Rightarrow \text{if } \sum_{i=1}^{3} Q_i(t) \ge \frac{B+\gamma}{0.2}$$

then $\mathbb{E}\left[\Delta \Phi(t) \mid Q(t)\right] \le -\gamma$

Or, $\mathbb{E}\left[\Delta \Phi(t) \mid Q(t)\right] \le -\gamma + B\,\mathbb{1}_{\left\{\Phi(Q(t)) \le \left(\frac{B+\gamma}{0.2}\right)^2\right\}}$

[Stability!!]

□

Quadratic Lyapunov is extremely useful because it also yields bounds on $\mathbb{E}[Q(t)]$.

Start with

$$\mathbb{E}[\Delta \Phi(t) \mid Q(t)] \leq B - 2\sum_{i=1}^{3} Q_i(t)\left(\bar{b}_i - \lambda_i\right)$$

$$\left(\text{let } \varepsilon = \min_i\left(\bar{b}_i - \lambda_i\right)\right)$$

$$\leq B - 2\varepsilon \sum_{i=1}^{3} Q_i(t)$$

Taking expectation wrt. $Q(t)$

$$\mathbb{E}[\Phi(Q(t+1))] - \mathbb{E}[\Phi(Q(t))] \leq B - 2\varepsilon \sum_{i=1}^{3} \mathbb{E}[Q_i(t)]$$

Summing for $t = 0, \ldots, T-1$ and rearranging

$$\frac{1}{T}\sum_{t=0}^{T-1}\sum_{i=1}^{3}\mathbb{E}[Q_i(t)] \leq \frac{B}{2\varepsilon} + \frac{\Phi(Q(0))}{T} - \frac{\mathbb{E}[\Phi(Q(T))]}{T}$$

$$\leq \frac{B}{2\varepsilon} + \frac{\Phi(Q(0))}{T}$$

letting $T \to \infty$

$$\sum_{i=1}^{3}\bar{Q}_i \doteq \limsup_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}\sum_{i=1}^{3}\mathbb{E}[Q_i(t)] \leq \frac{B}{2\varepsilon}$$

□

Remark: $\sum_i \bar{Q}_i \leq \frac{B}{2\varepsilon}$ ... $\sum A_i^2 + B_i^2$ ... $(\mu_i - \lambda_i)$

should remind you of Kingman's bound for G/G/m systems.

The min-drift / max-weight scheduling rule

Q: In the last example we happened to guess a stable policy. How can we do this for a general $\lambda = \{\lambda_1, \lambda_2, \lambda_3\}$?

A: We can write an LP and try to maximize $\varepsilon = \min_i (\bar{b}_i - \lambda_i)$

Denoting p_k = prob. of using k^{th} action $\alpha^{(k)}$

S-LP

$$\max_{p_k} \quad z$$

subject to:

$$z \le \bar{b}_i - \lambda_i \qquad \forall i$$

$$\bar{b}_i = \sum_k p_k b_i^{(k)} \qquad \forall i \qquad \left(\text{recall: } b_i^{(k)} = \mathbb{E}[B_i(\alpha^{(k)})]\right.$$

$$\sum_k p_k \le 1$$

$$p_k \ge 0 \qquad \forall k$$

Let $\varepsilon_{max}(\lambda)$ denote the optimal value of this LP.

$\Rightarrow$ if $\varepsilon_{max}(\lambda) > 0$ then the system can be stabilized by $\{p_k^*\}$

$$\sum_{i=1}^{3} \bar{Q}_i \le \frac{B}{2\varepsilon_{max}(\lambda)}$$

Q: What if we don't know λ or λ changes over time ?

Note that in our proof of stability, we only needed that Φ decreased

$\Rightarrow$ we can use Φ to design a scheduling policy

That is: choose $\alpha(t)$ as the one that minimizes $\mathbb{E}[\phi(Q(t+1))]$

Min-drift / Max-weight algorithm

At time t, choose $\alpha(t)$ as

$$\alpha(t) = \arg\max_{\alpha} \sum_i Q_i(t)\, b_i(\alpha)$$

Theorem: If $\forall t$, $\lambda(t)$ obeys

$$\varepsilon_{max}(\lambda(t)) \geq \varepsilon_0 > 0$$

then the min-drift/max-weight policy is stable and

$$\sum_i \overline{Q_i} \leq \frac{B}{2\varepsilon_0}$$

Proof:

Start with

$$\mathbb{E}[\Delta\Phi(t) \mid Q(t)] \leq B - 2\sum_{i=1}^{3} \mathbb{E}[Q_i(t)(B_i(t) - A_i(t)) \mid Q_i(t)]$$

$$= B + 2\sum_i Q_i(t)\mathbb{E}[A_i(t)|Q_i(t)] - 2\sum_{i=1}^{3} Q_i(t)\mathbb{E}[B_i(t) \mid Q_i(t)]$$

$$= B + 2\sum_i Q_i(t)\lambda_i(t) - 2\sum_{i=1}^{3} Q_i(t)b_i(t)$$

Since $\varepsilon_{max}(\lambda(t)) \geq \varepsilon_0$; there exists some randomized policy $\hat{P}_k(t)$ which yields $\hat{b}_i(t)$ such that

$$\hat{b}_i(t) - \lambda_i(t) \geq \varepsilon_0$$

but max-weight policy chooses the control which maximizes

$$\sum Q_i(t)b_i(t)$$

$$\Rightarrow \quad \sum_i Q_i(t)b_i(t) \geq \sum_i Q_i(t)\hat{b}_i(t)$$

$$\Rightarrow \mathbb{E}[\Delta\phi(t) \mid Q(t)] \leq B + 2\sum_i Q_i(t)\lambda_i(t) - 2\sum_i Q_i(t)\hat{b}_i(t)$$

$$\leq B - 2\varepsilon_0 \sum_i Q_i(t)$$

$$\varepsilon_0 > 0 \quad \Rightarrow \quad \mathbb{E}[\Delta\phi(t) \mid Q(t)] \leq -\gamma + B\,\mathbb{1}_{\{\phi(Q(t)) \leq \frac{B+\gamma}{2\varepsilon_0}\}}$$

and using a telescoping sum argument

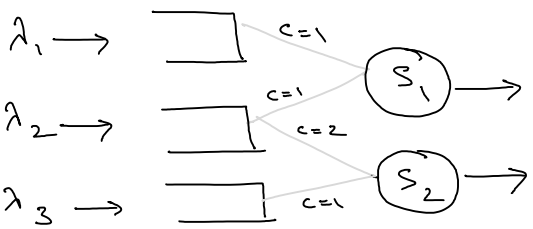
$$\limsup_{T\to\infty} \frac{1}{T}\sum_{\tau=1}^{T-1}\sum_i \mathbb{E}[Q_i(t)] \leq \frac{B}{2\varepsilon_0}$$

□

Cost optimization : drift + penalty algorithm

Lyapunov functions are even flexible, allowing optimizing some long run cost while maintaining stability.

Let's go back to our flexible server example:



Now we also have a cost associated with each arc

in particular: twice as expensive to serve class 2 using server 2

⇒ may want to idle server 2 even when work in Q_2

$$\alpha \in \begin{bmatrix}1\\1\\0\end{bmatrix}, \begin{bmatrix}1\\0\\1\end{bmatrix}, \begin{bmatrix}0\\1\\1\end{bmatrix}$$

$c(\alpha)$ 3 2 2

Goal: minimize $\bar{c} = \lim_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} c(\alpha(t))$

We can obtain a lower bound on $\bar{c}$ by solving an LP as before

C-LP

$$c^* = \min_{p_k} \sum p_k c(\alpha^k)$$

$$\text{s.t.} \quad \lambda_i \le \bar{b}_i = \sum_k p_k b_i^{(k)} \qquad \forall i$$

$$\sum_k p_k \le 1$$

$$p_k \ge 0$$

We can get quite close to c^* using a small twist on the min-drift algorithm.

Drift + Penalty algorithm

Let $V > 0$ be a weight parameter.

At time t choose the control $\alpha(t)$ as

$$\alpha(t) = \arg\min_{\alpha} \underbrace{V \cdot c(\alpha)}_{\text{penalty term}} - \underbrace{\sum_i Q_i(t)\, b_i(\alpha)}_{\text{drift term}}$$

Theorem: For the drift + penalty algorithm with weight parem V:

$$\bar{c} \le c^* + \Theta\left(\frac{1}{V}\right)$$

$$\sum \bar{Q}_i \le \Theta(V)$$

The above theorem highlights a cost-delay tradeoff.

→ as we increase $V \Rightarrow \bar{c}$ gets closer to c^*

but average queue length increase.

Proof: Let $\varepsilon_{max}(\lambda)$ denote the optimal value of S-LP corresp. to λ.

Consider an ε: $0 \le \varepsilon \le \varepsilon_{max}(\lambda)$

Let $\{\hat{p}_k\}$ be the optimal solution of C-LP for

$$\lambda' = (\lambda_1 + \varepsilon, \lambda_2 + \varepsilon, \lambda_3 + \varepsilon)$$

and let $c^*(\lambda')$ be the optimal value. ;

For the drift + penalty algorithm:

$$V \cdot \mathbb{E}[C(t) \mid Q(t)] + \mathbb{E}[\Delta \Phi(t) \mid Q(t)] = V \cdot C(\alpha(t) \mid Q(t) + \mathbb{E}[\Delta \Phi(t) \mid Q(t)]$$

$$\le B + 2\sum_i Q_i(t)\lambda_i + V \cdot C(\alpha(t) \mid Q(t)) - 2\sum_i Q_i(t)\, b_i(t)$$

Since drift + penalty minimizes this ;
it is smaller than the iid control $\{\hat{p}_k\}$

$$\leq \quad B + 2\sum_i Q_i(t)\lambda_i + V\cdot C^*(\lambda') - 2\sum_i Q_i(t)\underbrace{\hat{b}_i}_{=\sum_k \hat{P}_k b_i^{(k)}}$$

$$\geq (\lambda_i + \varepsilon)$$

$\Rightarrow$

$$V\cdot\mathbb{E}[C(t)|Q(t)] + \mathbb{E}[\Delta\phi(t)|Q(t)] \leq B + V\cdot C^*(\lambda') - 2\varepsilon\sum_i Q_i(t)$$

taking expectation and summing:

$$(*) \quad V\sum_{t=0}^{T-1}\mathbb{E}[C(t)] + \mathbb{E}[\phi(Q(T))] - \phi(Q(0)) \leq B\cdot T + VT\,C^*(\lambda') - 2\varepsilon\sum_{t=0}^{T-1}\sum_i \mathbb{E}[Q_i(t)]$$

Bounding cost

dividing $(*)$ by VT and rearranging

$$\frac{1}{T}\sum_{t=0}^{T-1}\mathbb{E}[C(t)] \leq C^*(\lambda') + \frac{B}{V} + \frac{\phi(Q(0))}{VT}$$

since this is true for all $0\leq\varepsilon\leq\varepsilon_{max}$; choosing $\varepsilon=0 \Rightarrow \lambda'=\lambda$

$$\Rightarrow \quad \frac{1}{T}\sum_{t=0}^{T-1}\mathbb{E}[C(t)] \leq C^*(\lambda) + \frac{B}{V} + \frac{\phi(Q(0))}{VT}$$

and $T\to\infty \Rightarrow$

$$\bar{C} \leq C^*(\lambda) + \frac{B}{V}$$

Bounding delay

dividing $(*)$ by $2\varepsilon T$ and rearranging:

$$\frac{1}{T}\sum_{t=0}^{T-1}\sum_i \mathbb{E}[Q_i(t)] \leq \frac{B}{2\varepsilon} + \frac{V}{2\varepsilon}\left(C^*(\lambda') - \frac{1}{T}\sum\mathbb{E}[C(t)]\right) + \frac{\phi(0)}{2\varepsilon T}$$

(true $\forall\ 0\leq\varepsilon\leq\varepsilon_{max}(\lambda)$)

Observe that $C^*(\lambda') - C^*(\lambda) \leq \eta\varepsilon$ for some constant η

since we only need to spend cost proportional to λ to serve it

(in our case $\eta\leq 3$)

$$\Rightarrow \quad C^*(\lambda') - \frac{1}{T}\sum_{t=0}^{T-1}\mathbb{E}[C(t)] \leq C^*(\lambda') - C^*(\lambda) \leq 3\varepsilon$$

$$\Rightarrow \frac{1}{T}\sum_{t=0}^{T-1}\sum_i \mathbb{E}[Q_i(t)] \le \frac{B}{2\varepsilon} + \frac{V\cdot 3\varepsilon}{2\varepsilon} + \frac{\phi(\Theta(0))}{2\varepsilon T}$$

Choose $\varepsilon = \varepsilon_{max}$ and let $T \to \infty$

$$\Rightarrow \sum_i \overline{Q_i} \le \frac{B}{2\varepsilon_{max}} + \frac{3V}{2}$$

□

There is an analogy between the drift + penalty algorithm and interior-point algorithm for C-LP.

Recall S-LP: $\min \sum_k p_k c(\alpha^k)$

s.t. $b_i = \sum_k p_k b_i^{(k)} \ge \lambda_i$

$\sum p_k \le 1$

$p_k \ge 0$

In an interior point solution, we simply move constraints to the objective function with a "penalty function"

* log-barrier for interior point

we will use a quadratic penalty

ie. $$\min \sum_k p_k c(\alpha^k) + \frac{1}{V}\sum_i \underbrace{\left(\left(\lambda_i - \sum_k p_k b_i(\alpha^k)\right)^+\right)^2}_{\text{penalty for } \lambda_i \ge b_i}$$

$\Rightarrow$ temporary backlog

$\Rightarrow Q_i$

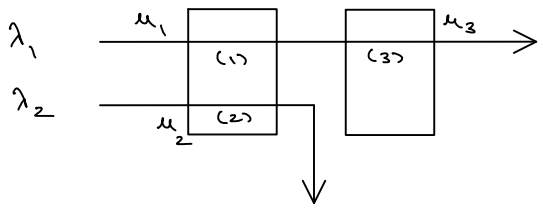
ie. $$\min \sum_k p_k c(\alpha^k) + \frac{1}{V}\sum_i (Q_i(t+1))^2$$

This is exactly what we minimize in drift+penalty.

Further, $\frac{2Q_i}{V}$ maps to the Lagrange multiplier for $(b_i \ge \lambda_i)$ constraint

Practice Exercise

Consider the following 2-server 3-class queueing network



The system operates in discrete slotted time.

Arrivals: $A_i(t)$ = # class i arrivals in slot t

$A_i(t) \in \mathbb{Z}_+$, iid ;

$\mathbb{E}[A_i(t)] = \lambda_i$, $\mathbb{E}[A_i^2(t)] < \infty$

Service process:

Server 1: at time t, chooses $\omega_1(t), \omega_2(t)$; $\omega_1(t) + \omega_2(t) \leq 1$

Server 2: at time t, chooses $\omega_3(t)$; $0 \leq \omega_3(t) \leq 1$

$B_i(t)$ = # potential class (i) departures in slot t

$\mathbb{E}[B_i(t)] = \omega_i(t)\mu_i$; $\mathbb{E}[B_i(t)^2] < \infty$

$B_i(t)$ $(i=1,2,3)$ are independent conditioned on $\omega_i(t)$ $(i=1,2,3)$

Work conserving policy: A policy $\{\omega_i(t)\}$ is work conserving if

(1) $\omega_1(t) + \omega_2(t) = 1$, $\forall t$ where $Q_1(t) + Q_2(t) > 0$

(2) $\omega_3(t) = 1$ $\forall t$

(3) $\omega_1(t) = 0$ if $Q_1(t) = 0$

(4) $\omega_2(t) = 0$ if $Q_2(t) = 0$

Prove that the system is stable under any work conserving policy if:

(1) $\lambda_1 + \lambda_2 < \min(\mu_1, \mu_2)$, and,

(2) $\mu_3 \geq \mu_1$.

(Hint: Use $\Phi(Q(t)) = (Q_1(t) + Q_2(t))^2 + Q_3^2(t)$)

References:

[1] ***Lyapunov Function Method (Lecture notes)***
Serguei Foss, Takis Konstantopoulos
http://www2.math.uu.se/~takis/L/StabLDC06/notes/SS_LYAPUNOV.pdf

[2] ***Stochastic Network Optimization with Application to Communication and Queueing Systems***
Michael J. Neely
http://www.morganclaypool.com/doi/abs/10.2200/S00271ED1V01Y201006CNT007

[3] ***Stability of Queueing Systems.***
Maury Bramson
http://www.math.umn.edu/~bramson/docs/Stability_of_Queueing_Networks.pdf

[4] ***Stability Properties of Constrained Queueing Systems and Scheduling Policies for Maximum Throughput in Multihop Radio Networks***
Leandros Tassiulas, Anthony Ephremides
IEEE Trans. Auto. Cont., Vol. 37, No. 12 (1992).