Data Communication Networks

Lecture 5

Saad Mneimneh
Computer Science
Hunter College of CUNY
New York

Abstraction

- Customers arrive at random times to obtain service

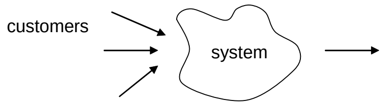

  - e.g. customers are packets assigned to a communication link
  - service time $=L/C$
- Questions of interest
  - what is the average number of customers in the system (typical number waiting in queue or undergoing service)?
  - what is the average delay per customer (typical time a customer waits in queue + service time)?
- These quantities are often obtained in terms of known information such as
  - customer arrival rate (typical number of customers entering the system per unit time)
  - customer service rate (typical number of customers the system serves per unit time when it's constantly busy)

Definitions

Let's work out what we mean by average or typical.

Define:

- $N(t) =$ number of customers in the system at time t
- $A(t) =$ number of customers who arrives in $[0, t]$
- $T_i =$ time spent in the system by customer i^{th} customer

Averages

- A notion of "typical" number of customers observed up to time t is the time average

$$N_t = \frac{1}{t} \int_0^t N(\tau) d\tau$$

  In many systems of interest N_t converges to a steady state

$$N = \lim_{t \to \infty} N_t$$

- Similarly, we define $\lambda_t = \frac{A(t)}{t}$ and and the time average arrival rate $\lambda = \lim_{t \to \infty} \lambda_t$ (assuming limit exists)
- We also define

$$T_t = \frac{\sum_{i=1}^{A(t)} T_i}{A(t)}$$

  and the time average customer delay (assuming limit exists)

$$T = \lim_{t \to \infty} T_t$$

Little's theorem

$$N = \lambda T$$

Little's theorem expresses the natural idea that crowded systems are associated with long delays

- Rainy day
  - traffic moves slower (large T)
  - streets are more crowded (large N)
- Fast food restaurant
  - fast service (small T)
  - requires small waiting area, e.g. drive through (small N)

Proof of Little's theorem

We will prove it under some simplifying assumptions:

- System is initially empty, i.e. $N(0) = 0$
- System is FIFO
- System becomes empty infinitely many times

Let $A(t)$ and $D(t)$ be the arrivals and departures respectively, then:

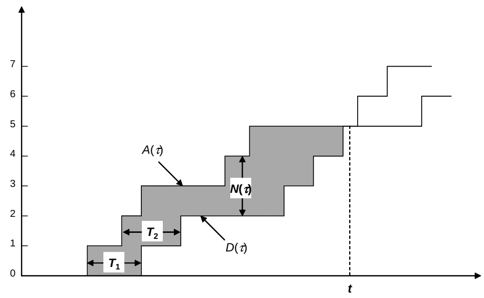

System empty at t: $\int_0^t N(\tau)d\tau = \sum_{i=1}^{A(t)} T_i = \frac{A(t) \sum_{i=1}^{A(t)} T_i}{A(t)}$. Dividing by t, we get $N_t = \lambda_t T_t$. Taking the $\lim_{t\to\infty}$ (assuming steady state), we get $N = \lambda T$.

Relaxing third assumption

Assuming that the system does not necessarily become empty infinitely many times, we can always write:

$$\sum_{i=1}^{D(t)} T_i \leq \int_0^t N(\tau)d\tau \leq \sum_{i=0}^{A(t)} T_i$$

Therefore,

$$\frac{D(t)}{t} \frac{\sum_{i=1}^{D(t)} T_i}{D(t)} \leq \frac{1}{t} \int_0^t N(\tau)d\tau \leq \frac{A(t)}{t} \frac{\sum_{i=1}^{A(t)} T_i}{A(t)}$$

If we only assume that $\lambda = \lim_{t\to\infty} \frac{A(t)}{t} = \lim_{t\to\infty} \frac{D(t)}{t}$ (arrival rate is equal to departure rate), and $\lim_{t\to} T_t = T$ then

$$\lambda T \leq \lim_{t\to\infty} \frac{1}{t} \int_0^t N(\tau)d\tau \leq \lambda T$$

$$N = \lambda T$$

We can also relax the initially empty and FIFO assumptions.

Probabilistic Little's

- In our analysis, we relied on a single sample function and computed averages over time (time averages)
- For almost every system of interest, we can replace time averages with ensemble averages, i.e.
  - N is replaced by $\bar{N} =$ expected number of customers in the system
  - T is replaced by $\bar{T} =$ expected delay per customer
  - λ is replaced by $\lim_{t\to\infty} \frac{\text{expected number of arrivals in } [0,T]}{t}$
- Usually λ is given as a property of arrivals, and $\bar{N}$ can be obtained by some simple analysis of p_n, the probability of having n customers in the system (later).

Example 1

Consider the following node where the arrival rate is λ packets per second and the link bandwidth is μ bps:

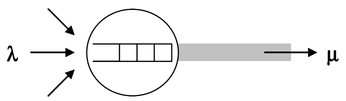

Little's theorem can be applied to any system or part of it.

- Looking at the node: $N = \lambda T$, where N is the average number of packets in the node and T is the average delay per packet
- Looking at the queue: $N_Q = \lambda W$, where N_Q is the number of packets in the queue and W is the average waiting time per packet
- Looking at the link: $\rho = \lambda L \frac{1}{\mu}$, where:
  - ρ is the number of bits currently being transmitted (served), also known as link utilization, efficiency, or throughput
  - L is the average packet length, and hence λL is the rate in bits (even this is an application of Little's theorem, multiply $N = \lambda T$ by L)
  - note that $N = N_Q + \rho$
- Looking at the link: $B = \lambda LD$, where B is the number of bits in transit and D is the propagation delay of the link

Example 2

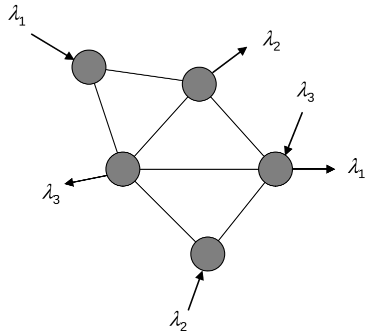

- For each subsystem, $N_i = \lambda_i T_i$
- For the whole system, $N = \lambda T$, where $N = \sum_i N_i$ and $\lambda = \sum_i \lambda_i$
- Therefore

$$T = \frac{\sum_i \lambda_i T_i}{\sum_i \lambda_i}$$

(weighted average by $\lambda_i's$)

Arrivals and Departures

- Packets arrive according to a random process typically modeled as Poisson
- A Poisson process is characterized by that interarrival times are independent and exponentially distributed

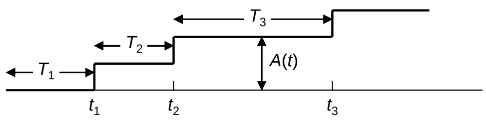

$$T_i = t_i - t_{i-1}$$

is exponentially distributed, i.e.

$$Pr(T \leq t) = 1 - e^{-\lambda t}$$

- Probability density function is $\lambda e^{-\lambda t}$ (derivative of $1 - e^{-\lambda t}$), i.e.

$$Pr(t_1 \leq T \leq t_2) = \int_{t_1}^{t_2} \lambda e^{-\lambda t}$$

$E[T]$ and $E[T^2]$

- Expected value of T

$$E[T] = \int_0^\infty t\lambda e^{-\lambda t} dt = \frac{1}{\lambda}$$

  Therefore, λ is interpreted as the rate of arrivals

- Expected value of T^2

$$E[T^2] = \int_0^\infty \lambda e^{-\lambda t} dt = \frac{2}{\lambda^2}$$

  Therefore, the variance is

$$\sigma^2(T) = E[(T - \frac{1}{\lambda})^2] = E[T^2] + E[(\frac{1}{\lambda})^2] - E[2T\frac{1}{\lambda}] = \frac{1}{\lambda^2}$$

  (linearity of expectation)

Poisson is memoryless

Given that I have waited for sometime t_0 and no arrival occurred, what is the probability that I have to wait for another t?

$$Pr(T \leq t_0 + t | T > t_0) =?$$

$$Pr(T \leq t_0 + t | T > t_0) = \frac{Pr(T \leq t_0 + t, T > t_0)}{Pr(T > t_0)} = \frac{Pr(t_0 < T \leq t_0 + t)}{1 - Pr(T \leq t_0)} = \frac{\int_{t_0}^{t_0+t} \lambda e^{-\lambda\tau} d\tau}{1-(1-e^{-\lambda t_0})}$$

$$Pr(T \leq t_0 + t | T > t_0) = \frac{-e^{-\lambda\tau}|_{t_0}^{t_0+t}}{e^{-\lambda t_0}} = \frac{-e^{-\lambda(t_0+t)} + e^{-\lambda t_0}}{e^{-\lambda t_0}} = 1 - e^{-\lambda t} = Pr(T \leq t)$$

- Previous history does not help predicting the future
- Distribution of time until next arrival is independent of when last arrival occurred

Example

Suppose a bus arrives at a station according to a Poisson process with average interarrival time of 20 min (i.e. $1/\lambda = 20$)

- When a customer arrives at a station, what is the average amount of time until next bus?
  - 20 min, regardless of when previous bus arrived (memoryless)
- When a customer arrives at a station, what is the average amount of time since last bus departure?
  - 20 min, looking at the time in reverse, we will also see a Poisson process
- PARADOX: If an average of 20 min passed since last bus, and there is an average of 20 min until next bus, then we have an average of 40 min between buses! (how did that happen?)
  - there is conditioning on your arrival, you are likely to fall in a long interval

Intuition

Imagine throwing a ball at random

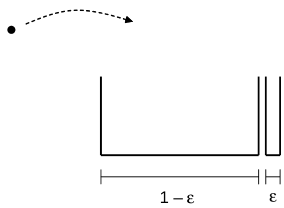

- The average bin size is
$$\frac{(1-\epsilon)+\epsilon}{2} = \frac{1}{2}$$
- The ball falling at random,
  - with probability $1-\epsilon$ will see a bin of size $1-\epsilon$
  - with probability ϵ will see a bin of size ϵ

  The expected bin size observed by the ball is
$$(1-\epsilon)(1-\epsilon) + \epsilon.\epsilon = (1-\epsilon)^2 + \epsilon^2 \approx 1 - 2\epsilon$$
- Large intervals have more weight

Mathematical interpretation

- Let $Y(t)$ be the time until next arrival

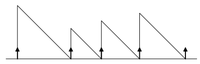

- Let $Z(t)$ be the time since last departure

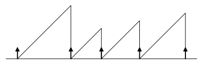

- Their sum is below

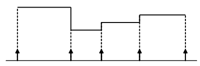

$$average = \frac{sum\ of\ areas}{length} = \frac{\sum_{i=1}^{n} T_i^2}{\sum_{i=1}^{n} T_i} = \frac{\sum_{i=1}^{n} T_i^2/n}{\sum_{i=1}^{n} T_i/n}$$

But $\sum_{i=1}^{n} T_i^2/n \to E[T^2]$ and $\sum_{i=1}^{n} T_i/n \to E[T]$

$$\frac{E[T^2]}{E[T]} = \frac{2/\lambda^2}{1/\lambda} = \frac{2}{\lambda}$$

Another characterization of Poisson

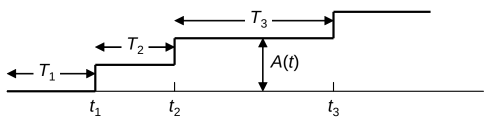

$$t_n = \sum_{i=1}^{n} T_i = T_1 + T_2 + \ldots + T_n$$

- It can be shown that t_n has the following probability density function

$$t_n \sim \frac{\lambda^n t^{n-1} e^{-\lambda t}}{(n-1)!}$$

- We are interested in obtaining the probability mass function for $A(t)$, i.e., what is $Pr(A(t) = n)$ for some integer n?

$$Pr(t_n \in (t, t+\delta]) = Pr(A(t) = n-1).Pr(T_n \leq \delta)$$

(we have used the memoryless property)

Another characterization of Poisson (cont.)

$$Pr(t_n \in (t, t+\delta]) = Pr(A(t) = n-1).Pr(T_n \leq \delta)$$

$$Pr(t_n \in (t, t+\delta]) = \int_t^{t+\delta} \frac{\lambda^n \tau^{n-1} e^{-\lambda\tau}}{(n-1)!} d\tau \approx \frac{\lambda^n t^{n-1} e^{-\lambda t}}{(n-1)!}\delta$$

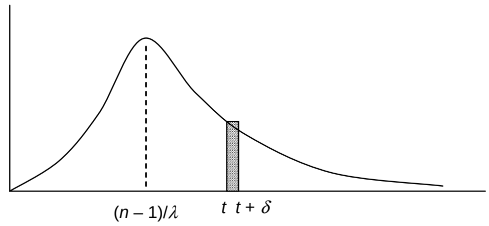

$$Pr(T \leq \delta) = 1 - e^{-\lambda\delta} = 1 - (1 - \lambda\delta + \frac{(\lambda\delta)^2}{2!} - \ldots) = \lambda\delta + o(\delta)$$

where $\lim_{\delta \to 0} \frac{o(\delta)}{\delta} = 0$

Another characterization of Poisson (cont.)

$$Pr(A(t) = n-1) = \frac{\lambda^n t^{n-1} e^{-\lambda t}\delta}{(n-1)!(\lambda\delta + o(\delta))} = \frac{\lambda^n t^{n-1} e^{-\lambda t}}{(n-1)!(\lambda + o(\delta)/\delta)}$$

Taking $\lim_{\delta \to 0}$, we get

$$Pr(A(t) = n-1) = \frac{\lambda^{n-1} t^{n-1} e^{-\lambda t}}{(n-1)!}$$

Poisson as a counting process $A(t)$:

$$Pr(A(t) = n) = \frac{(\lambda t)^n e^{-\lambda t}}{n!}$$

$$Pr(A(t') - A(t) = n) = \frac{[\lambda(t'-t)]^n e^{-\lambda(t'-t)}}{n!} \quad (stationary\ increment\ property)$$

$$E[A(t)] = \lambda t$$

$$E[A^2(t)] = \lambda t + (\lambda t)^2$$

$$\sigma^2(A(t)) = \lambda t$$

Yet another characterization...

Consider a small interval of time δ and define $\tilde{A}(\delta) = A(t+\delta) - A(t)$, then a Poisson process is defined as small increments:

- $Pr(\tilde{A}(\delta) = 0) = e^{-\lambda\delta} = 1 - \lambda\delta + o(\delta) \approx 1 - \lambda\delta$
- $Pr(\tilde{A}(\delta) = 1) = \lambda\delta e^{-\lambda\delta} = \lambda\delta + o(\delta) \approx \lambda\delta$
- $Pr(\tilde{A}(\delta) \geq 2) = o(\delta) \approx 0$

where $\lim_{\delta\to 0} \frac{o(\delta)}{\delta} = 0$

This characterization allows us to merge and split Poisson processes:

Merging

The sum of two Poisson processes with rates λ and μ is a Poisson process with rate $\lambda + \mu$

$$Pr(\tilde{A}(\delta) = 0) \approx (1 - \lambda\delta)(1 - \lambda\mu) \approx 1 - (\lambda + \mu)\delta$$

$$Pr(\tilde{A}(\delta) = 1) \approx \lambda\delta(1 - \lambda\mu) + \lambda\mu(1 - \lambda\delta) \approx (\lambda + \mu)\delta$$

$$Pr(\tilde{A}(\delta) \geq 2) \approx 0$$

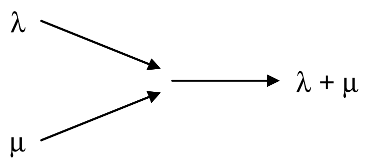

Splitting

A Poisson process with rate λ can be split into two independent Poisson as follows: Each arrival is independently sent to process 1 with probability p and to process 2 with probability $1-p$.

$$Pr(\tilde{A}_1(\delta) = 1) \approx p\lambda\delta = (p\lambda)\delta$$

$$Pr(\tilde{A}_1(\delta) \geq 2) \approx 0$$

$$Pr(\tilde{A}_1(\delta) = 0) =\approx 1 - (p\lambda)\delta$$

Similarly,

$$Pr(\tilde{A}_2(\delta) = 1) \approx (1-p)\lambda\delta = ((1-p)\lambda)\delta$$

$$Pr(\tilde{A}_2(\delta) \geq 2) \approx 0$$

$$Pr(\tilde{A}_2(\delta) = 0) =\approx 1 - ((1-p)\lambda)\delta$$