CS 547 Lecture 11: PASTA

Daniel Myers

Poisson Arrivals See Time Averages

The PASTA property refers to the expected state of a queueing system as seen by an arrival from a Poison process. An arrival from a Poisson process observes the system as if it was arriving at a random moment in time. Therefore, the expected value of any parameter of the queue at the instant of a Poisson arrival is simply the long-run average value of that parameter.

For example, at the instant of a Poisson arrival,

- the expected number of customers in the queue, including the one in service, is $\overline{Q}$
- the probability the server is busy is U
- the probability the server is idle is $1 - U$
- the expected number waiting and not being served is $\overline{Q} - U$

Proof

Here's a proof of a simplified version of PASTA. Consider a period of length t, divided into three non-overlapping regions with lengths a_1, b, and a_2. We'll show that probability that a Poisson process produces one arrival in the period of length b is the same as the probability of a randomly chosen point being in the interval b. This argument can be extended to a general case with any number of arrivals.

The probability that a single Poisson arrival on $[0, t]$ occurs in the interval of length b can be written as

$$\frac{P\{0 \text{ arrivals in } a_1 \textit{ and } 1 \text{ arrival in } b \textit{ and } 0 \text{ arrivals in } a_2\}}{P\{1 \text{ arrival in } t\ \}}$$

The denominator comes from conditioning on the probability of getting only one arrival during the entire period of length t. Split the numerator into its three independent components, then use the Poisson distribution to calculate the probability of getting 0 or 1 arrivals in each period.

$$\frac{e^{-\lambda a_1} \cdot \lambda b\, e^{-\lambda b} \cdot e^{-\lambda a_2}}{\lambda t\, e^{-\lambda t}}$$

Adding the exponents, we obtain

$$\frac{\lambda b\, e^{-\lambda(a_1+b+a_2)}}{\lambda t\, e^{-\lambda t}}$$

But $a_1 + b + a_2 = t$, so we can cancel the exponential terms and λ to give

$$P\{\text{one arrival in period of length } b\} = \frac{b}{t}$$

This is exactly the probability that a randomly chosen point in $[0, t]$ lands in the interval of length b. Thus, arriving from a Poisson process is statistically indistinguishable from arriving at a random moment in time.

Justification for the Poisson Process

At this point, we've shown that the Poisson process has several nice properties that aid in the analysis of queueing systems. Before actually analyzing any queues, we need to deal with one more fundamental question.

Is the Poisson process a good model for real systems?

There is, in fact, a theoretical justification for the existence of the Poisson process, which explains why it occurs in so many real-world systems. The justification rests on the fact that the behavior of large numbers of customers interacting independently can be closely approximated by the Poisson distribution.

Consider a system with huge universe of potential customers who *might* submit requests to the system. We can think of each customer as behaving like a Bernoulli trial: the customer flips an imaginary coin, and if the coin comes up heads with some small probability, the customer submits a request. In this model, the number of requests arriving to our system over a period of time depends on how many of these independent Bernoulli trials result in success.

$$P\{k \text{ successes}\} = P\{k \text{ successes out of } N \text{ Bernoulli trials}\}$$

The probability of getting k successes out of N independent Bernoulli trials with success parameter p is given by the *binomial distribution*.

$$P\{k \text{ successes out of } N \text{ Bernoulli trials}\} = \binom{N}{k}(1-p)^{N-k}p^k$$

If N is large and p is small – as they are in our model – the binomial distribution can be closely approximated by the Poisson distribution. Therefore, the probability of getting k arrivals in a period is approximately Poisson distributed, which implies that the arrival process is also nearly Poisson.

Therefore, in applications where subsequent arrivals tend to be mostly independent of each other, the Poisson process is often a realistic model for the arrival behavior.

Unfortunately, there are still many interesting problems that violate these assumptions. Particularly, many real-world applications – such as disk I/O – are capable of arrival sequences that are *burstier*[1] than the Poisson process.[2] As we move through our discussion of queueing systems, we'll discuss ways of adapating our models to deal with these cases.

[1] Is this a word? It is now.
[2] Though it should be noted that the Poisson process itself is fairly bursty