IEOR 151 – Lecture 22
Square Root Rule

1 M/M/∞ Queue

Recall that the M/M/∞ queue is a model with an infinite number of servers. In this model, the service rate is state-dependent and is given by $n\mu$ where n is the number of customers in line, and μ is the service rate for a single customer. Some calculations give that $L = \lambda/\mu$ and $W = 1/\mu$. Futhermore, the stationary distribution is given by a Poisson distribution with mean $\rho = \lambda/\mu$:

$$p_j = \frac{\rho^j \exp(-\rho)}{j!}, \qquad \forall j \geq 0.$$

2 Square Root Law

Suppose we are engineering the number of servers we should use in a queue. The equations for an M/M/s queue are difficult, but those of the M/M/∞ queue are much simpler. In the case where $\rho = \lambda/\mu$ is large, we can use approximate the equations for an M/M/∞ queue as a Normal distribution. In particular, we have that

$$\mathbb{P}(j \leq u) \approx \Phi((u - \rho)/\sqrt{\rho}),$$

where j is the number of people in the queue. If we would like to design the system so that the probability of waiting is α when we have hired s servers, then we have

$$\mathbb{P}(\text{waiting}) = \mathbb{P}(j \geq s) \Rightarrow \alpha \approx 1 - \Phi((s - \rho)/\sqrt{\rho}).$$

For example, suppose we would like to have a probability of waiting of 0.02, and $\rho = \lambda/\mu = 1000$. Then, we need

$$0.02 = 1 - \Phi((s - 1000)/\sqrt{1000}).$$

From a z-table, we get $\Phi(2.05) = 0.9798$. And so solving for s gives

$$(s - 1000)/\sqrt{1000} = 2.05 \Rightarrow s = 1064.8.$$

And so we should hire 1065 servers to ensure that the probability of waiting is 2%.

3 More Information and References

The material in these notes follows that of the course textbook "Service Systems" by Mark Daskin and of the Wikipedia article on "M/M/∞ queue" and "Poisson distribution".