Appendix B

Leader Election: Rounds Needed with Uncorrelated Loss

For the Leader Election algorithm in Chapter 5, the difference between the delay metrics for the correlated versus uncorrelated case result from the difference in the expected number of rounds before everyone gets a message. In the correlated case, this is $1/(1-l)$. In this appendix, we calculate the uncorrelated case.

Since the losses are independent, we can think of each loss as a coin-toss experiment. Let $l = Pr[heads]$ and $(1-l) = Pr[tails]$. A process wants to toss the coin until it comes up tails, i.e., the process receives the message. Each process is performing an independent coin toss experiment and there are N processes in total.[1] The question is: what is the expected value of the maximum number of tosses before they all get tails?

Consider one process. It takes i tosses before a process gets a tails with probability $l^{(i-1)} * (1-l)$, just as in the correlated case (Section 5.3.1.2). So, the probability that it takes less than or equal to i tosses to get a tails,

$$\begin{aligned} Pr[\leq i \text{ tosses to get a tails}] &= (1-l) + l*(1-l) + l^2*(1-l) + ... + l^{(i-1)}*(1-l) \\ &= (1-l)(1 + l + l^2 + ... + l^{(i-1)} \\ &= 1 - l^i \end{aligned}$$

What is the probability that given N processes the **maximum** number of tosses is k? For the maximum to be equal to k, at least one must be k, and all the others must be less than or equal to k. However it is more complicated, because any process could be the maximum and we must be careful not to double count because others being less than or equal to k does not prevent them from being equal to k.

Because each process behaves independently, we know that the probability for all processes to

[1] Actually there is one less than the number of processes in LE, which we call N here for simplicity.

have received a message in less than or equal to k tosses is

$$\begin{aligned} Pr[all \leq k] &= Pr[\leq k \text{ tosses to get a tails}]^N \\ &= (1 - l^k)^N \end{aligned}$$

Therefore, the probability that the maximum number of tosses is exactly equal to k is

$$\begin{aligned} Pr[max = k] &= Pr[all \leq k] - Pr[all \leq (k-1)] \\ &= (1 - l^k)^N - (1 - l^{(k-1)})^N \end{aligned}$$

which represents the probability that all processes have received a message in less than or equal to k attempts minus the probability that all processes received the message in less than or equal to $k-1$ attempts. In essence, we subtract out the overlap between these two probabilities and are left with the probability that exactly k attempts are required. Thus, the expected value of the maximum number of tosses, given N processes, is

$$\begin{aligned} E[max] &= \sum_{k=1}^{\infty} k \times Pr[max = k] \\ &= \sum_{k=1}^{\infty} k \times ((1 - l^k)^N - (1 - l^{(k-1)})^N) \\ &= \sum_{k=1}^{\infty} k \times \sum_{i=0}^{N} (-l^k)^i \binom{N}{i} - (-l^{k-1})^i \binom{N}{i} \\ &= \sum_{i=0}^{N} \binom{N}{i} \sum_{k=1}^{\infty} k \times ((-l^k)^i - (-l^{k-1})^i) \\ &= \sum_{i=0}^{N} \binom{N}{i} \sum_{k=1}^{\infty} k \times (-1)^i l^{(k-1)^i} (l^i - 1) \\ &= \sum_{i=0}^{N} (-1)^i \binom{N}{i} (l^i - 1) \sum_{k=1}^{\infty} k \times l^{(k-1)^i} \\ &= \sum_{i=0}^{N} (-1)^i \binom{N}{i} (l^i - 1) \sum_{k=1}^{\infty} k \times (l^i)^{k-1} \\ &= \sum_{i=0}^{N} (-1)^i \binom{N}{i} (l^i - 1) \frac{1}{(1 - l^i)^2} \\ &= \sum_{i=0}^{N} (-1)^i \binom{N}{i} \frac{-1}{(1 - l^i)} \\ &= \sum_{i=0}^{N} (-1)^{i+1} \binom{N}{i} \frac{1}{(1 - l^i)} \end{aligned}$$

The result $E[max]$ can be recast as the expected number of announcements that must be made by the leader process before leadership is established.