Contention Resolution

Instructor: Thomas Kesselheim

1 Waiting for One Success

Assume you are rolling a dice over and over again. How long to expect until it shows "6" for the first time? More generally, we observe a random process that is either "successful" or "not successful" in every round. All rounds are independent and in each we get the outcome "successful" with probability p. Let T denote the random variable indicating the number of rounds it takes until we see "successful" for the first time.

As it is a non-negative integer random variable, we can use the following formula to calculate the expectation.

$$\mathbf{E}\left[T\right] = \sum_{t=1}^{\infty} \mathbf{Pr}\left[T \geq t\right] \quad .$$

What is $\mathbf{Pr}\left[T \geq t\right]$? Note that we have $T \geq t$ if and only if in the first $t-1$ rounds we see "not successful". For an event $\mathcal{E}$, we let $\overline{\mathcal{E}}$ denote its complementary event. Regarding the probabilities, we have $\mathbf{Pr}\left[\overline{\mathcal{E}}\right] = 1 - \mathbf{Pr}\left[\mathcal{E}\right]$. Letting $\mathcal{S}_{t'}$ denote the event that we see "successful" in round t', this can be written as

$$T \geq t \text{ occurs} \qquad \text{if and only if} \qquad \bigcap_{t'=1}^{t-1} \overline{\mathcal{S}_{t'}} \text{ occurs} \quad .$$

And therefore

$$\mathbf{Pr}\left[T \geq t\right] = \mathbf{Pr}\left[\bigcap_{t'=1}^{t-1} \overline{\mathcal{S}_{t'}}\right] \quad .$$

As all rounds are independent, we have

$$\mathbf{Pr}\left[\bigcap_{t'=1}^{t-1} \overline{\mathcal{S}_{t'}}\right] = \prod_{t'=1}^{t-1} \mathbf{Pr}\left[\overline{\mathcal{S}_{t'}}\right] \quad .$$

Furthermore, $\mathbf{Pr}\left[\overline{\mathcal{S}_{t'}}\right] = 1 - p$. So overall

$$\mathbf{Pr}\left[T \geq t\right] = (1-p)^{t-1} \quad .$$

For the expectation, this gives us the following geometric sum

$$\mathbf{E}\left[T\right] = \sum_{t=1}^{\infty} \mathbf{Pr}\left[T \geq t\right] = \sum_{t=1}^{\infty} (1-p)^{t-1} = \frac{1}{p} \quad .$$

2 Collecting Coupons

In the previous example, we were rolling a dice until we saw one particular number for the first time. Next, we consider the number of rounds it takes until we have seen all of the numbers at least once. In the literature this question is known as the *coupon collector's problem*. There are n coupons one of which is drawn uniformly in every round with replacement. Let T be the number of rounds it takes until we have seen all coupons at least once. Clearly, T is never less than n but of course it can be much bigger. Figure 1 visualizes an experiment with $n = 100$.

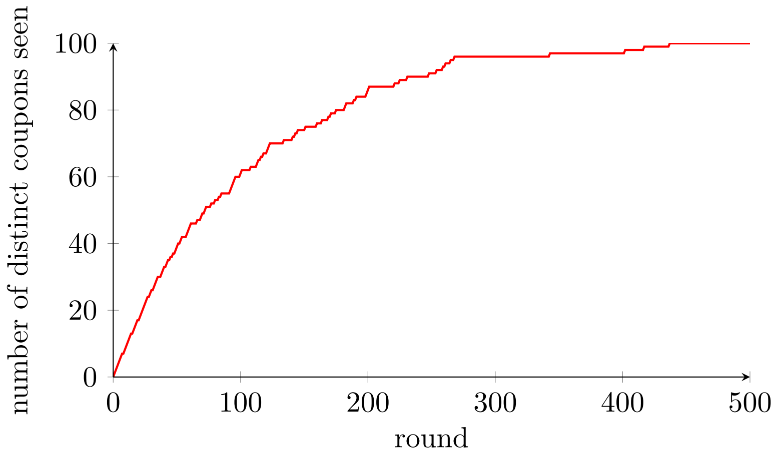

Figure 1: The progress made when collecting $n = 100$ coupons.

While we see many different coupons at the beginning, it takes more and more time until we see a new coupon.

To calculate the expectation of T, we use a trick. We divide the rounds into phases. Phase 1 is only the first round. Phase j starts after phase $j-1$ and ends in a round when we see a new coupon. Let X_j denote the length of phase j. By this definition, the time it takes until we see at least k different coupons is $\sum_{j=1}^{k} X_j$ and therefore $T = \sum_{j=1}^{n} X_j$.

When we are in phase j, we have already seen $j-1$ different coupons. So $n-(j-1) = n-j+1$ coupons are still new. The probability to see such a coupon in a particular round is $\frac{n-j+1}{n}$. By the result in Section 1, the expected length of phase j is therefore

$$\mathbf{E}\left[X_j\right] = \frac{n}{n-j+1} \quad .$$

To get the expectation of T, we use linearity of expectation.

Lemma 1.1 (Linearity of Expectation)**.** *Let $X_1, \ldots, X_n$ be (not necessarily independent) random variables. Then*

$$\mathbf{E}\left[\sum_{j=1}^{n} X_j\right] = \sum_{j=1}^{n} \mathbf{E}\left[X_j\right] \quad .$$

So, we get

$$\mathbf{E}\left[T\right] = \mathbf{E}\left[\sum_{j=1}^{n} X_j\right] = \sum_{j=1}^{n} \mathbf{E}\left[X_j\right] = \sum_{j=1}^{n} \frac{n}{n-j+1} = n \sum_{i=1}^{n} \frac{1}{i} = nH_n \quad .$$

H_n is the n-th harmonic number. It is easy to show that $\ln(n+1) < H_n < 1 + \ln n$ for all n and therefore $H_n = \Theta(\log n)$.

3 Contention Resolution

The coupon collector's problem may seem to be not much more than a mathematical puzzle. The following application shows that it actually many applications can be considered in a similar way.

We start with a typical example of a randomized algorithm. Assume there are n processes that each want to access a shared resource, which can only be accessed by one process at a time. For example, they each want to perform a radio transmission. If two processes attempt a transmission simultaneously, both fail. There is no way of coordinating processes, except for designing an algorithm.

In more detail, we have synchronized rounds. Our protocol starts in round 1. We want to minimize the time until all processes have successfully accessed the resource at least once. If the processes had IDs, we could have process i access the resource in round i and we would be done in n rounds. What can we do instead?

We let each process access the resource in each round with probability $p = \frac{1}{n}$ *independently*, even if it has already been successful once. Setting p this way seems reasonable for the time being; in the exercises, you will try out other values. Let T be the random variable representing the number of rounds until each process has been successful at least once. Whats T's distribution? Can we bound its expectation?

We define three kinds of events for every process i and every round t

- $\mathcal{A}_{i,t}$ is the event that process i *accesses* the resource in round t.
- $\mathcal{S}_{i,t}$ is the event that process i *successfully accesses* the resource in round t.
- $\mathcal{F}_{i,t}$ is the event that process i *does not successfully access* the resource in *any round* $1, \ldots, t$.

By these definitions, we have

$$T > t \text{ occurs} \quad \text{ if and only if } \quad \bigcup_i \mathcal{F}_{i,t} \text{ occurs} \enspace .$$

3.1 The Probability of a Single Success

We first start with the event $\mathcal{A}_{i,t}$. Naturally, we defined the probability of this event to occur to be p. So, formally, $\mathbf{Pr}\left[\mathcal{A}_{i,t}\right] = p$. The probability of the complementary event is therefore $\mathbf{Pr}\left[\overline{\mathcal{A}_{i,t}}\right] = 1 - p$.

What is the probability of the event $\mathcal{S}_{i,t}$? We can write it as the intersection of access events

$$\mathcal{S}_{i,t} = \mathcal{A}_{i,t} \cap \left(\bigcap_{j \neq i} \overline{\mathcal{A}_{j,t}} \right) \enspace .$$

This does not tell us anything about its probability yet. We need that the access attempts are independent between processes in a round. So therefore

$$\mathbf{Pr}\left[\mathcal{S}_{i,t}\right] = \mathbf{Pr}\left[\mathcal{A}_{i,t}\right] \cdot \prod_{j \neq i} \mathbf{Pr}\left[\overline{\mathcal{A}_{j,t}}\right] = p \cdot (1-p)^{n-1} \enspace .$$

The functions are visualized in Figure 2.

Fact 1.2.

(a) The function $x \mapsto \left(1 - \frac{1}{x}\right)^x$ *converges monotonically from* $\frac{1}{4}$ *up to* $\frac{1}{\mathrm{e}}$ *as* n *increases from* 2.

(b) The function $x \mapsto \left(1 - \frac{1}{x}\right)^{x-1}$ *converges monotonically from* $\frac{1}{2}$ *down to* $\frac{1}{\mathrm{e}}$ *as* n *increases from* 2.

So, using Fact 1.2, we get that with $p = \frac{1}{n}$ for all i and all t

$$\frac{1}{\mathrm{e}n} \leq \mathbf{Pr}\left[\mathcal{S}_{i,t}\right] \leq \frac{1}{2n} \enspace .$$

This bound is already enough to bound the expectation of T by $\Theta(n \log n)$ using the same calculations as for the coupon collector's problem. You will do this in the exercises.

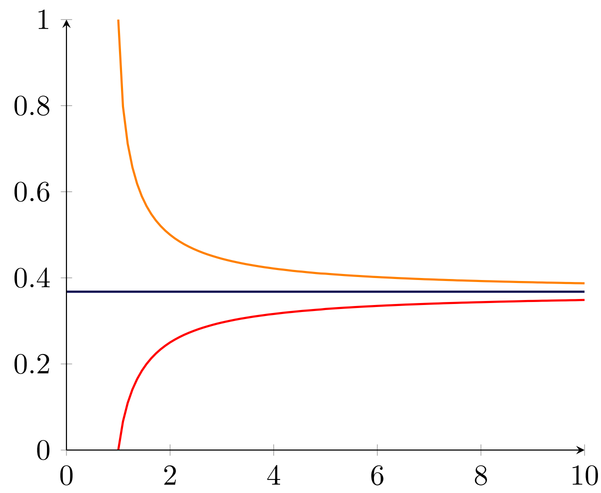

Figure 2: The functions bounded in Fact 1.2.

3.2 Bounding the Success Probability over Multiple Rounds and Multiple Processes

Bounding only the expectation does not tell us a lot about the actual distribution. Therefore, we have a closer look at the respective probabilities now. To this end, we will bound the probability of an event $\mathcal{F}_{i,t}$. It is defined as

$$\mathcal{F}_{i,t} = \bigcap_{r=1}^{t} \overline{\mathcal{S}_{i,r}} \ .$$

We can use independence between rounds to get

$$\mathbf{Pr}\left[\mathcal{F}_{i,t}\right] = \prod_{r=1}^{t} \mathbf{Pr}\left[\overline{\mathcal{S}_{i,r}}\right] \ .$$

We have already shown that $\mathbf{Pr}\left[\overline{\mathcal{S}_{i,r}}\right] \leq 1 - \frac{1}{\mathrm{e}n}$. So therefore

$$\mathbf{Pr}\left[\mathcal{F}_{i,t}\right] \leq \left(1 - \frac{1}{\mathrm{e}n}\right)^{t} \ .$$

Now let c be an arbitrary constant integer and set $t = c\lceil \mathrm{e}n \ln n \rceil$. We get

$$\mathbf{Pr}\left[\mathcal{F}_{i,t}\right] \leq \left(1 - \frac{1}{\mathrm{e}n}\right)^{t} \leq \mathrm{e}^{-c \ln n} = \frac{1}{n^c} \ .$$

This gives us a bound on the probability that *a fixed process* waits for more than t rounds for its first success. However, we are interested in the probability of the event that *there is a process* that waits for more than t rounds. Formally, we would like to bound the probability

$$\mathbf{Pr}\left[T > t\right] = \mathbf{Pr}\left[\bigcup_{i} \mathcal{F}_{i,t}\right] \ .$$

Unfortunately, the events $\mathcal{F}_{i,t}$ are not independent. Therefore, we cannot apply the techniques used so far. Instead, we rely on an apparently weak but very useful bound.

Lemma 1.3 (Union Bound)**.** *Given (not necessarily disjoint) events $\mathcal{E}_1, \mathcal{E}_2, \ldots, \mathcal{E}_n$, we have*

$$\mathbf{Pr}\left[\bigcup_{i=1}^{n} \mathcal{E}_i\right] \leq \sum_{i=1}^{n} \mathbf{Pr}\left[\mathcal{E}_i\right] \ .$$

With the union bound, we get that for $t = c\lceil en \ln n \rceil$

$$\mathbf{Pr}\left[T > t\right] = \mathbf{Pr}\left[\bigcup_i \mathcal{F}_{i,t}\right] \leq \sum_i \mathbf{Pr}\left[\mathcal{F}_{i,t}\right] \leq n \cdot \frac{1}{n^c} = \frac{1}{n^{c-1}} \ .$$

This bound on the probability gives us another way to bound the expectation of T as follows

$$\begin{aligned}
\mathbf{E}\left[T\right] &= \sum_{k=1}^{\infty} \mathbf{Pr}\left[T \geq k\right] \\
&\leq \sum_{c=0}^{\infty} \sum_{k=1}^{\lceil en \ln n \rceil} \mathbf{Pr}\left[T \geq c\lceil en \ln n \rceil + k\right] \\
&\leq \sum_{c=0}^{\infty} \mathbf{Pr}\left[T > c\lceil en \ln n \rceil\right] \lceil en \ln n \rceil \\
&= O(n \log n) \sum_{c=0}^{\infty} \mathbf{Pr}\left[T > c\lceil en \ln n \rceil\right] \\
&= O(n \log n) \left(1 + \sum_{c=0}^{\infty} \frac{1}{n^c}\right) \\
&= O(n \log n) \left(1 + \frac{1}{1 - \frac{1}{n}}\right) \\
&= O(n \log n) \ .
\end{aligned}$$

Indeed, we showed a much stronger guarantee, namely that T is bounded by $O(n \log n)$ with high probability.

Definition 1.4. *Let $f \colon \mathbb{N} \to \mathbb{R}_{>0}$ be a function. A sequence of random variables $(X_n)_{n \in \mathbb{N}}$ is bounded by $O(f(n))$* with high probability *if for every constant α there is a constant β such that*

$$\mathbf{Pr}\left[X_n \geq \beta f(n)\right] \leq \frac{1}{n^\alpha} \qquad \text{for all } n \in \mathbb{N} \ .$$

Definition 1.5. *Let $f \colon \mathbb{N} \to \mathbb{R}_{>0}$ be a function. A sequence of random variables $(X_n)_{n \in \mathbb{N}}$ is bounded by $\Omega(f(n))$* with high probability *if for every constant α there is a constant β such that*

$$\mathbf{Pr}\left[X_n \leq \beta f(n)\right] \leq \frac{1}{n^\alpha} \qquad \text{for all } n \in \mathbb{N} \ .$$

Definition 1.6. *Let $f \colon \mathbb{N} \to \mathbb{R}_{>0}$ be a function. A sequence of random variables $(X_n)_{n \in \mathbb{N}}$ is bounded by $\Theta(f(n))$* with high probability *if it is bounded by $O(f(n))$ with high probability and bounded by $\Omega(f(n))$ with high probability.*