Randomized Online Algorithm
2017/03/29

Definitions in the randomized marking mouse (RMM) algorithm:

Locations where the mouse can hide $\equiv$ different pages 1, 2, 3, $\ldots$, m
Locations where mouse isn't are pages in cache
Cat probe sequence $\equiv$ page request sequence

Deterministic Mouse can achieve no better than (m − 1)-competitive

Random Marking Mouse (RMM) algorithm:
1. Start at a random location (cache entry)
2. When Cat probes a location, mark it
3. When Cat probes mouse's location, move to a random unmarked spot
4. If mouse is at last unmarked spot, clears marks (a new phase begins)

claim 1 $E[RMM_{cost}(p_1 p_2 \ldots p_n)] \leq O(\log m) OPT(p_1 p_2 \ldots p_n)$

Note   *"E" denotes "expected value", over possible values of mouse's choices. See note in next lecture for explanation*

Aside   *It is surprising that randomness gives us a probable gap in performance from deterministic algorithms. This seems to be a feature of online algorithms*

proof

Initially, the probability that mouse is at any spot is $\frac{1}{m}$

=> the first Cat probe finds mouse with probability $\frac{1}{m}$

Whether mouse is found or not in the first probe, the mouse is now at each of the m-1 unmarked spots with probability $\frac{1}{m-1}$, as the Cat now knows it is not in this initially probed spot.

Note: *that considering an intelligent cat that does not re-probe is equivalent to considering a cat that does re-probe, as any re-probes of a known empty spot will not contribute to the cost of RMM or OPT. Then we consider an intelligent cat here to simplify our argument.*

Then given that we have a smart cat that does not pick the same spot twice before the mouse moves, we get a similar increase in probability with each successive probe before the mouse is found:

probability for probe #2:
        Cat finds mouse with probability $\frac{1}{m-1}$
probability for probe #3:
        Cat finds mouse with probability $\frac{1}{m-2}$
probability for probe #4
        Cat finds mouse with probability $\frac{1}{m-3}$

etc.

To deal with these probabilities, we introduce indicator random variables, defined as follows:

Let $X_i \begin{cases} 0 \; if \; mouse \; is \; not \; found \; on \; probe \; i \\ 1 \quad if \; mouse \; is \; found \; on \; probe \; i \end{cases}$

Then we describe the number of times the mouse is found in a phase using an expectation:
note: the concept of expectation is described later

$$E[Number \; of \; times \; RMM \; is \; found \; during \; phase]$$
$$= E[\textstyle\sum_{i=1} X_i]$$
$$= \textstyle\sum_{i=1} E[X_i]$$
$$= \frac{1}{m} + \frac{1}{m-1} + \frac{1}{m-2} + \cdots + \frac{1}{1}$$
$$= H_m \approx \ln m$$

Where H_m is the sum of the first m terms of the harmonic series. Its approximate equality with ln(m) is a characteristic of the series.

Aside on Linearity of expectation:

Linearity of expectation means that for some random variables X and Y, and constant a, we have

1) $E[X + a] = E[X] + a$
2) $E[a*X] = a*E[X]$
3) $E[X+Y] = E[X] + E[Y]$

The first two are fairly intuitive. Let us consider the third.

For example, consider a game in which we flip a coin n times. Each time we flip heads, the payoff is $1 (and $0 for tails). Take X_i to be a random variable representing our ith flip, with a tails represented by $X_i = 0$ and a heads represented by $X_i = 1$

Given a fair coin:

$Prob[X_i = 1] = Prob[X_i = 0] = \frac{1}{2}$

Then it is simple to see that the expected payoff on a flip is given by

$E[X_i] = 1*Prob[X_i = 1] + 0*Prob[X_i = 0] = 1*1/2 + 0*1/2 = \frac{1}{2}$, or 50 cents

Further, if we consider the payoff for 2 rounds, call them X_i and X_{i+1} is given by

$E[X_i + X_{i+1}] = E[X_i] + E[X_{i+1}] = 1

It is fairly simple to see that linearity of expectation should apply here, as the 2 rounds are independent and do not affect each other, and so considering the addition of their expectations should not be different than considering the expectation of their additions. It is perhaps more surprising to note that this remains true in the case where the values are not independent (for addition)

Given that the cat probes all cache entries in each phase, we also have that the number of times OPT is found during a phase is at least one:

times OPT is found during a phase ≥ 1

Then we have $E[RMM_{cost}(p_1 p_2 \dots p_n)] \leq O(logm)OPT(p_1 p_2 \dots p_n)$ and claim 1 is true.

claim 2

For all mice A (deterministic or randomized)

$\exists p_1, p_2, \ldots, p_n \quad s.t. \quad E[A_{cost}(p_1 p_2 \ldots p_n)] \geq \log m \, OPT(p_1 p_2 \ldots p_n)$

proof

idea: show that a cat exists that will cause A to move $> \log m$ times more than OPT

If the Cat probes at random, then no matter what mouse A does, the Cat finds it with probability $\frac{1}{m}$.

Then the expected number of times A must move over a sequence of t probes is $\frac{t}{m}$

idea: optimal mouse will hide in the last spot that is probed

How many random cat probes does it take until the Cat examines all m spots? this is called the coupon collector problem, and an analysis of this problem gives expected number of tries as $m \ln m$

So opt moves once every $m \ln m$ probes while A moves $\frac{t}{m} = \frac{m \ln m}{m}$ times in this period

This gives: $\quad$ *Number of faults for A* $\geq \ln m * OPT$

Then we have that claim 2 holds.

Hash Functions

Universal Hash Functions

A set of hash functions $\mathcal{H}$ that map $\mathcal{U} \to \{0,1,\ldots,\text{m-1}\}$ is <u>universal</u> if, $\forall$ distinct keys k,l in $\mathcal{U}$ the number of hash functions $h \in \mathcal{H}$ s.t. $h(k) = h(l)$ is at most $\frac{|\mathcal{H}|}{m}$

Chaining using universal hash functions

Hash n keys into a table T of size m using hash function $h \in_R \mathcal{H}$
Where $\in_R$ signifies uniform random selection

Theorem

For key k
Let n_i = the number of items in bucket i
$\alpha = \frac{n}{m}$ = load factor

$$E\left[n_{h(k)}\right] \leq \begin{cases} \alpha + 1 & if\ key\ k\ \in T \\ \alpha & if\ key\ k\ \in T \end{cases}$$

Aside: $E[X] = \sum_y Prob(X = y) * y$

Where y is every possibly number in the universe
Most y's will have a probability of 0.
For example, consider rolling a 6-sided die. Then Prob(X = y) is 1/6 if y is one of 1,2,3,4,5,6 and 0 otherwise.

Let $X_{kl} \begin{cases} 1 & if\ h(k) = h(l) \\ 0 & othewise \end{cases}$

Continued in next lecture