Lecture 1: Course Intro and Hashing

Lecturer: *Christopher Musco*

1 Hashing: Preliminaries

Hashing can be thought of as a way to *rename* an address space. For instance, a router at the internet backbone may wish to have a searchable database of destination IP addresses of packets that are whizing by. An IP address is 128 bits, so the number of possible IP addresses is 2^{128}, which is too large to let us have a table indexed by IP addresses. Hashing allows us to rename each IP address by fewer bits. Furthermore, this renaming is done probabilistically, and the renaming scheme is decided in advance before we have seen the actual addresses. In other words, the scheme is *oblivious* to the actual addresses.

Formally, we want to store a subset S of a large universe U (where $|U| = 2^{128}$ in the above example). And $|S| = m$ is a relatively small subset. For each $x \in U$, we want to support 3 operations:

- *insert*(x). Insert x into S.
- *delete*(x). Delete x from S.
- *query*(x). Check whether $x \in S$.

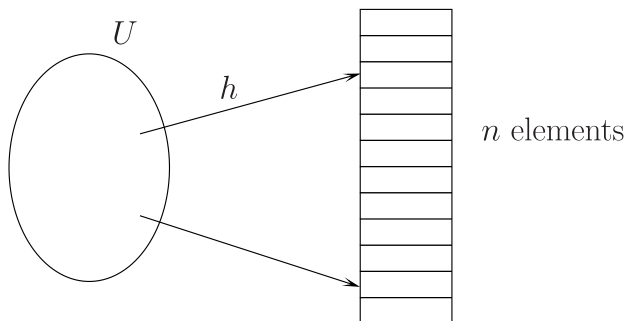

Figure 1: Hash table. x is placed in $T[h(x)]$.

A hash table can support all these 3 operations. We design a hash function

$$h : U \longrightarrow \{0, 1, \dots, n-1\} \tag{1}$$

such that $x \in U$ is placed in $T[h(x)]$, where T is a table of size n. Typically, we can assume that $m \leq n \ll |U|$, although we will talk about some applications where we hash to a set with size $n < m$.

Since $|U| \gg n$, multiple elements can be mapped into the same location in T, and we deal with these collisions by constructing a linked list at each location in the table.

One natural question to ask is: how long is the linked list at each location?

This can be analysed under two kinds of assumptions:

1. Assume the input is the random.
2. Assume the input is arbitrary, but the hash function is random.

Assumption 1 may not be valid for many applications.

Hashing is a concrete method towards Assumption 2. We designate a set of hash functions $\mathcal{H}$, and when it is time to hash S, we choose a random function $h \in \mathcal{H}$ and hope that on average we will achieve good performance for S. This is a frequent benefit of a randomized approach: no single hash function works well for every input, but the average hash function may be good enough.

2 Hash Functions

What do we want out of a random hash function? Ideally, we would hope that h "evenly" distributes the elements of S across the hash table. One option would be to map every element in U to a random value in $[n]$. However, constructing such a "fully random" hash function is very expensive: we would need to build a lookup table with $|U|$ rows, each storing $\log_2(n)$ bits to specify the value of $h(x) \in [n]$ for one $x \in U$. At this cost, we might as well have just stored our original data in a $|U|$ length array – it's often simply impossible.

The goal in hashing is to find a *cheaper* function (fast and space efficient) that's still *random enough* to evenly distribute elements of S into our table. For a family of hash functions $\mathcal{H}$, and for each $h \in \mathcal{H}$, $h : U \longrightarrow [n]$[1], what we mean by "random enough".

For any $x_1, x_2, \ldots, x_m \in S$ ($x_i \neq x_j$ when $i \neq j$), and any $a_1, a_2, \ldots, a_m \in [n]$, ideally a random $\mathcal{H}$ should satisfy:

- $\mathbf{Pr}_{h \in \mathcal{H}}[h(x_1) = a_1] = \frac{1}{n}$.
- $\mathbf{Pr}_{h \in \mathcal{H}}[h(x_1) = a_1 \wedge h(x_2) = a_2] = \frac{1}{n^2}$. Pairwise independence.
- $\mathbf{Pr}_{h \in \mathcal{H}}[h(x_1) = a_1 \wedge h(x_2) = a_2 \wedge \cdots \wedge h(x_k) = a_k] = \frac{1}{n^k}$. k-wise independence.
- $\mathbf{Pr}_{h \in \mathcal{H}}[h(x_1) = a_1 \wedge h(x_2) = a_2 \wedge \cdots \wedge h(x_m) = a_m] = \frac{1}{n^m}$. Full independence (note that $|U| = m$).

Generally speaking, we encounter a tradeoff. The more random $\mathcal{H}$ is, the greater the number of random bits needed to generate a function h from this class, and the higher the cost of computing h. The challenge is to prove that, even when we use few random bits, the hash stable still performs well in terms of insert/delete/query time.

2.1 Goal One: Bound expected number of collisions

As a first step, we want to understand the expected length of a single linked list. Note that this is just the first step towards understanding the runtime of our desired operations. Assume that $\mathcal{H}$ is a pairwise-independent hash family.

[1] We use $[n]$ to denote the set $\{0, 1, \ldots, n-1\}$

Let L_x be the length of the linked list containing x; this is just the number of elements with the same hash value as x. Let random variable

$$I_y = \begin{cases} 1 & \text{if } h(y) = h(x), \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

So $L_x = 1 + \sum_{y \in S; y \neq x} I_y$, and

$$\mathbb{E}[L_x] = 1 + \sum_{y \in S; y \neq x} \mathbb{E}[I_y] = 1 + \frac{m-1}{n} \tag{3}$$

Usually we choose $n > m$, so this expected length is less than 2. Later we will analyse this in more detail, asking how likely is L_x to exceed say 100.

The expectation calculation above doesn't need full independence; pairwise independence would actually suffice. In fact, we don't even need pairwise independence, we just need the probability of a collision to be small. This motivates the next idea.

3 2-Universal Hash Families

Definition 1 (Carter Wegman 1979) *Family $\mathcal{H}$ of hash functions is 2-universal if for any $x \neq y \in U$,*

$$\mathop{\mathbf{Pr}}_{h \in \mathcal{H}}[h(x) = h(y)] \leq \frac{1}{n} \tag{4}$$

Exercise: Convince yourself that this property is weaker than pairwise independence – i.e. that every pairwise independent hash function also satisfies (4).

We can design 2-universal hash families in the following way. Choose a prime $p \in \{|U|, \ldots, 2|U|\}$,[2] and let

$$f_{a,b}(x) = ax + b \mod p \qquad (a, b \in [p], a \neq 0) \tag{5}$$

Then let

$$h_{a,b}(x) = f_{a,b}(x) \mod n \tag{6}$$

One this to note about $f_{a,b}$ is that $f_{a,b}(x_1) \neq f_{a,b}(x_2)$ if $x_1 \neq x_2$. Why? If $f_{a,b}(x_1) = f_{a,b}(x_2) = s$, then

$$a(x_1 - x_2) = 0 \mod p,$$

which can't be the case if $a \in 1, \ldots p-1$ and $(x_1 - x_2) \neq 0$. Can someone tell me why?

This isn't immediately useful, because we will still have a hash collision when $f_{a,b}(x_1) = f_{a,b}(x_2) \mod n$, but it's helpful to note.

[2] How do we know that such a prime exists? This is due to Bertrand's Postulate, which exactly states that such a prime exists. Second, how do we find such a prime? One option is to guess random numbers between $|U|$ and $2|U|$, check if they're prime, and continue until we find one. The Prime Number Theorem states that each guess is likely to be prime with probability roughly $1/\log(|U|)$. Also, the AKS primality test lets us test whether a number is in fact prime in time $\text{poly}(\log(|U|))$. Alternatively, one could imagine an online pre-computed database of primes that lie in the correct range.

Lemma 1
For any $x_1 \neq x_2$ and $s \neq t$, the following system

$$ax_1 + b = s \mod p \tag{7}$$
$$ax_2 + b = t \mod p \tag{8}$$

has exactly one solution (i.e. one set of possible values for a, b). In that solution, $a \neq 0$.

Proof: If you're familiar with modular arithmetic, this is clear. Since p is a prime, the integers $\mod p$ constitute a finite field. This implies that any element in $[p]$ has a multiplicative inverse $\mod p$, so we know that $a = (x_1 - x_2)^{-1}(s - t)$ and $b = s - ax_1$.

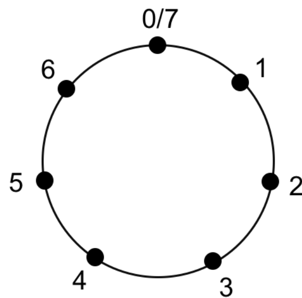

Figure 2: Modular arithmetic for prime $p = 7$.

It's not to hard to see this directly with a little thought. We want to claim that

$$a(x_1 - x_2) = (s - t) \mod p$$

has a unqiue solution a. Without loss of generality, assume that $x_1 > x_2$. When we multiply $(x_1 - x_2)$ by an integer, we're moving around the circle pictured in Figure 2 in increments of $(x_1 - x_2)$. Since p is prime, at each step before the p^{th} step, it better be that we hit a new element of $[p]$ on the circle. Otherwise, we would have found that $(x_1 - x_2)$ (which is $< p$) multiplies by some other number $< p$ to equal a multiple of p. This of course can't be true when p is prime.

So, as we multiply $(x_1 - x_2)$ by integers in $[p]$, we hit $(s - t) \mod p$ exactly once. □

By Lemma 1, since there are $p(p - 1)$ different possible choices of a, b:

$$\mathbf{Pr}_{a,b \leftarrow U(\{1,\ldots,p-1\} \times \{0,\ldots,p-1\})}[f_{ab}(x_1) = s \wedge f_{ab}(x_2) = t] = \frac{1}{p(p-1)} \tag{9}$$

Claim $\mathcal{H} = \{h_{a,b} : a, b \in [p] \wedge a \neq 0\}$ *is* 2*-universal.*

PROOF: For any $x_1 \neq x_2$,

$$\begin{aligned}
&\mathbf{Pr}[h_{a,b}(x_1) = h_{a,b}(x_2)] && (10)\\
=&\sum_{s,t\in[p],s\neq t} \mathbb{1}[s = t \mod n)] \cdot \mathbf{Pr}[f_{a,b}(x_1) = s \wedge f_{a,b}(x_2) = t] && (11)\\
=&\frac{1}{p(p-1)} \sum_{s,t\in[p],s\neq t} \mathbb{1}[s = t \mod n] && (12)\\
\leq&\frac{1}{p(p-1)} \frac{p(p-1)}{n} && (13)\\
=&\frac{1}{n} && (14)
\end{aligned}$$

where $\mathbb{1}$ is an indicator function (that is, $\mathbb{1}[x] = 1$ if statement x is true, and $\mathbb{1}[x] = 0$ otherwise). Equation (13) follows because for each $s \in [p]$, we have at most $(p-1)/n$ different t such that $s \neq t$ and $s = t \mod n$. $\square$

Can we design a collision free hash table then?

Solution 1: Collision-free hash table in $O(m^2)$ space.

Say we have m elements, and the hash table is of size n. Since for any $x_1 \neq x_2$, $\mathbf{Pr}_h[h(x_1) = h(x_2)] \leq \frac{1}{n}$, the expected number of total collisions is just

$$\mathbb{E}[\sum_{x_1\neq x_2} h(x_1) = h(x_2)] = \sum_{x_1\neq x_2} \mathbb{E}[h(x_1) = h(x_2)] \leq \binom{m}{2}\frac{1}{n} \tag{15}$$

Let's pick $n \geq m^2$, then

$$\mathbb{E}[\text{number of collisions}] \leq \frac{1}{2} \tag{16}$$

and so

$$\mathop{\mathbf{Pr}}_{h\in H}[\exists \text{ a collision}] \leq \frac{1}{2} \tag{17}$$

So if the size the hash table is large enough, we can easily find a collision free hash function. In particular, if we try a random hash function it will succeed with probability 1/2. If we see a collision when inserting elements of S into the table, we simply draw a new random hash function and try again. The expected function of this proceedure is:

$$\mathbb{E}[\text{time to insert } m \text{ items}] = m + \frac{1}{2}m + \frac{1}{4}m + \ldots = 2m.$$

Solution 2: Collision-free hash table in $O(m)$ space.

At this point, we have designed a hash table that has no collisions. The drawback is that it is that our table must be large: m^2 to store only m elements. But in reality, such a large table is often unrealistic. We may use a two-layer hash table to avoid this problem.

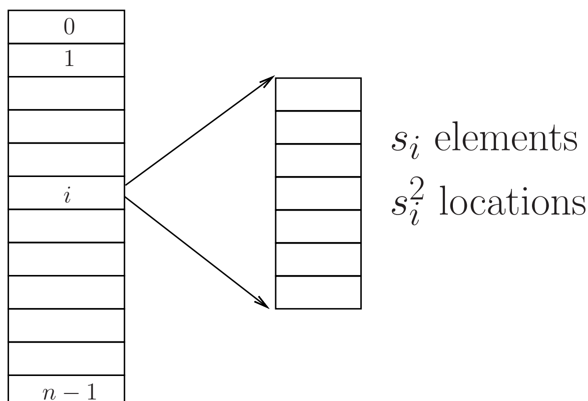

Figure 3: Two layer hash tables.

Specifically, let s_i denote the number of collisions at location i. If we can construct a second layer table of size s_i^2, we can easily find a collision-free hash table to store all the s_i elements. Thus the total size of the second-layer hash tables is $\sum_{i=0}^{m-1} s_i^2$.

To bound the expectation size of $\sum_{i=0}^{m-1} s_i^2$, we note that this sum is nearly equal to the total number of hash collisions, which we bound in Equation (15)! Specifically,

$$\mathbb{E}[\sum_i s_i^2] = \mathbb{E}[\sum_i s_i(s_i - 1)] + \mathbb{E}[\sum_i s_i] = \frac{m(m-1)}{n} + m \leq 2m \tag{18}$$

Including the first layer, we have now designed a hash table of expected size $3m$ to store m elements (so some overhead, but much less than before).

4 Preview to Lecture 3: Load Balancing

In our 2-level construction, we cared about limiting the *total size* of the hash table, and we were able to do so by bounding $\sum_{i=1}^{m} s_i^2$. However, we did not bound each s_i individually – it could be that some buckets of the first hash table are much larger than others. In some applications of hashing, this is something you want to avoid.

A simple example is when your hash table is distributed and each bucket (or a small set of buckets) is stored on a separate machine. The is a common architecture in large "no-SQL" databases like Amazon's DynamoDB or Apache Cassandra. In the distributed case, memory isn't a shareable resource across machines, so we care about showing that no s_i is too large (i.e no machine is overloaded).

Another example arises when hashing is used to distribute workload across multiple machines. As a toy example, suppose I look up directions from Princeton, NJ to Boston, MA on Google maps. Google has many different serves computing efficient driving routes and one potential strategy is to use a hash function to choose what server to send your request (i.e hash the start and end locations).

Question: Why is *hashing* a good strategy? Why not just send the request to an arbitrary or even randomly chosen server?

In Lecture 3 we will develop better tools for bounding the probability of random events and we will be able to establish effective bounds on $\max[s_i]$.

5 Other Considerations: Dealing with Adversaries

Other issues arise when hashing is used in a distributed way instead of simply to build data structures on a single machine. An interesting one is the issue of *resistance to adversarial attacks*. In particular, a user submitting requests to a centralized server may be able to *learn* the hash function being used by the server, even if that hash function is chosen randomly.

Question: How easy is it to learn our 2-universal hash function described in (6)?

This can open the door for denial of service attacks (DoS) attacks that intentionally issue "colliding" requests. Even if an adversary does not have the resources to take down a large web-service, they may have enough to take down one or several servers underlying that service. I'll post a blog post about these sorts of attacks and potential solutions on the course webpage.