

Lecture 14

*Lecture date: February 28, 2024**Scribes: Ashwani Aggarwal, Gary Song*

1 Oblivious Ram Introduction

When Schwarzenegger was government of California he went skiing and broke his leg and was flown to St John. They kept it quiet that he was there. Imagine that they store the patient records in the cloud encrypted. About 300 people com to the ER every day. You can see all the records but don't know which one. If another day he breaks his leg again and goes to the ER his encrypted record will still be there. So you can tell his record is the overlap between the two. Another example, by seeing how many times a record is assessed reveals information. Encryption isn't enough.

You store records of the same size $x_1 \dots x_n$ and want to provide read and write. You want to provide that no matter what records are read and written the access pattern should be independent of your reads and writes. Even if the user reads and writes all files on the cloud it should be indistinguishable. This is called Oblivious Random Access Memory (ORAM). If it's a single user and single cloud we must encrypt.

As an abstraction, let's say that whatever value is stored is encrypted and we can only see what locations are read and written. Say that we have a sequence of the actual reads and write locations and we simulate another sequence that is independent, we would like these sequences of locations to be indistinguishable to any adversary. We can define overhead to be how many actual reads and writes we need to do with the cloud per virtual reads and writes that we want. It turns out that there is a provable $\log(n)$ lower bound for overhead assuming user has constant memory.

First Idea: We have n files and can safely store a permutation in the cloud. If we retrieve each file only once then it looks like we are picking random files. But if we retrieve a file again they know we are picking the same file and how many files. Thus, this would reveal information.

2 Sorting Networks

In a sorting network, the comparison gate takes in a, b and outputs the min of a, b and the max of a, b . If we have a bunch of comparison gates we can connect them so that if n numbers come in we can sort them. This is called a sorting network.

Two main sorting networks are as follows:

1. Batcher's sorting network with $O(n \log^2 n)$ gates and $O(\log^2 n)$ depth.
2. *AKS* network from Hungarian mafia. This is $O(n \log n)$ gates and $O(\log n)$ depth. However, the constant is in the billions so it is not usable.

The connection with ORAM is that this allows you to shuffle in the cloud according to random keys. Suppose you have a bunch of cards. You can take two cards and switch them and put them back. People see what two cards you touch, but do not know how they were switched.

From the perspective of a sorting network, we can do the following:

1. Download a, b from the cloud.
2. Compute the comparison gate according to some key the user chooses.
3. Encrypt and upload.
4. Repeat

This allows the user to permute the data in a way oblivious to the cloud.

3 Naive Oblivious Ram

Warm up solution: Why are sorting networks useful? Using the aforementioned procedure, I can permute all my data in the cloud. Say the user has local memory of size \sqrt{n} (which we call the buffer). We can store permuted data in the cloud with n real locations and \sqrt{n} dummy locations. If we want to read location i we check if it's in the buffer and if it's not we read it using its location computed with the permutation. If it's in the buffer we can read it from the buffer and instead read a random dummy location. This allows us to overcome the main problem of the first idea since files will be read in the same location at most once. The issue with this is that after \sqrt{n} we can't store things anymore. The solution is that after \sqrt{n} steps when the cache is full, we restart the process with a new permutation. This gives us a $O(\sqrt{n} \log^2 n)$ amortized overhead since sorting network needs to be run once every \sqrt{n} steps.

Next Solution: We put the buffer in the cloud. To make it secure we scan the entire buffer. There are two possibilities:

1. If the data is found in reading the buffer, we additionally read a random dummy location to obscure it.

2. If the data is not found in reading the buffer, we read the data from its permuted location and write it to the buffer.

The additional cost of this over the warm up solution is that for every read and write we need $O(\sqrt{n})$ steps. Similarly, when the buffer is full we need to start over with a new permutation requiring amortized $O(\sqrt{n} \log^2 n)$ time. Thus the overall overhead is $O(\sqrt{n} \log^2 n)$.

All of this assumes a honest, but curious cloud. For a malicious cloud, we have to take additional precautions. To start, pick a random s and you store $\alpha = (i, PRF_s(i) \oplus a)$. To prevent the cloud from changing it you can sign it or you put it into another PRF_s and store that value. Unless the cloud can also change this they are done. However, the cloud can still swap two values. You can fix this by also storing $PRF_s(\alpha, i)$ which gives context. This is still not good enough since you can take a value from the past and place it in the present. To address this requires more advanced time stamp methods that are out of scope.

4 Hash Tables

The first poly log solution was in 1990 and called the Hierarchical solution. We first talk about hash tables which are needed for it. Say that we have a hash function h . Suppose that we have n pairs of the form $(key, value)$. Our hash function can take in the key and output a number which taken modulo n gives us a position that we could store the pair in the table. This brings up the problem of collisions if two keys map to the same position. If this happens, we just store them into the same buffer at that position. If we make the buffer of size $O(\log^{1+\epsilon} n)$ for every position in the table then the probability of overflow is negligible.

There is a more recent data structure called a cuckoo hash table that can remove the need for these buffers. You pick two hash functions h_1, h_2 and a table from 1 to $O(n)$. For any key you compute $h_1(key)$ and try to put the pair in that bin. If we have a collision, then we kick out the pair currently in it, $(key_2, value_2)$, and put it in $h_2(key_2)$. This can create an avalanche where the item that is kicked out kicks out another item and so on. However, people showed that if you have a stash of size $\log(n)$ then it breaks with negligible probability.

5 Hierarchical Oblivious Ram

Back to ORAM. We create tables of sizes $2, 4, 8, \dots, 2^{\log n}$ as seen in Figure 1 in the cloud. The user chooses s_k randomly for every level from $1, \dots, \log n$. Let (i, a) be an arbitrary index value pair. When writing, the user would store a in the largest table at index $PRF_{s_{\log n}}(i) \bmod n$. When reading position i from the cloud, we have the following procedure:




Figure 1: Hierarchical Oblivious Ram Data Structure

1. For every level k , read from $PRF_{s_k}(i) \bmod 2^k$ until the desired item is found. Then read randomly from the rest of the levels.
2. When an item in the last level is read, move it into the first level, placing it at position $PRF_{s_1}(i) \bmod 2$.
3. When any level, k overflows, choose a new key s_{k+1} for the next level and transfer all of the contents of level k to the next level using s_{k+1} for computing the new positions.

The point of this is that every 2^k steps the shuffle on level k occurs. This amortized is polynomial in k and poly-log in n . In total, calculating gives us an overhead of $O(\log^{4+\epsilon}(n))$. Additionally, the user only requires constant memory since if done cleverly, everything can be accomplished with a single seed.

Recent Improvements: First off, instead of hash tables we can use cuckoo hash tables to improve the overhead. Looking in more detail at step 3 of the reading process when we have to shuffle also can improve the complexity. Take two levels of the hierarchical table. The key observation is that the items in each are already shuffled so all we need to do is shuffle between the two. Say that we have a vector of 0 and 1 and we need to sort it. This problem is called tight compaction. Running it backwards gives the desired shuffling. The question then becomes can we do tight compaction in linear time. It has been showed you can do it obliviously in linear time, but with a big constant.

In sum, these changes allow us to reduce to $O(\log(n))$, however even with some more improvements, the constant is 1400 which is impractical.